



**Politecnico
di Torino**

ScuDo

Scuola di Dottorato ~ Doctoral School

WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Programme in Civil and Environmental Engineering (38th Cycle)

**From Space to Action:
Satellite Interferometric SAR for
Structural Health Monitoring**
Decision-ready diagnostics for existing structures

Raffaele TARANTINI

Supervisor:

Prof. Giuseppe Andrea Ferro

Doctoral Examination Committee:

Prof. Giuseppe Carlo Marano, Member, Politecnico di Torino

Prof. Filippo Ubertini, Referee, Università di Perugia

Prof. Giuseppina Uva, Referee, Politecnico di Bari

Politecnico di Torino
17 September 2026

This thesis is licensed under a Creative Commons License, Attribution - NonCommercial - NoDerivatives 4.0 International: see www.creativecommons.org. The text may be reproduced in unmodified form for non-commercial purposes only, provided that credit is given to the original author, and no derivative works are permitted.

I hereby declare that the contents and organisation of this dissertation constitute my own original work and do not compromise in any way the rights of third parties, including those relating to the security of personal data.

Turin, 30 April 2026

Raffaele Tarantini

A mia figlia, Zelda.

*Che tu possa avere
sempre il coraggio di
opporti ai prepotenti e
ai soprusi della vita,
custodendo intatta la
tua umanità.*

Papà

Acknowledgments

The author wishes to express his deepest gratitude to his supervisor, Professor Giuseppe Andrea Ferro, for the academic guidance, the methodological insight, and the intellectual mentorship that have shaped the present work. The trust placed in the research direction adopted in the final phase of the doctoral programme has been instrumental in the development of the methodological framework presented in this thesis.

The doctoral research has been supported by a fellowship jointly funded by the Department of Structural, Geotechnical and Building Engineering (DISEG) of Politecnico di Torino and by the Fondazione Cassa di Risparmio di Torino (CRT), to which the author expresses his sincere gratitude. The author thanks the DISEG community for providing the academic environment within which the present work has been conceived and conducted, and the Doctoral Programme in Civil and Environmental Engineering for the structured training opportunities offered throughout the three years of the programme.

Part of the research presented in this thesis was conducted during a visiting research period at the Jacobs School of Engineering of the University of California San Diego, beginning in August 2025, within the Cultural Heritage Engineering Initiative (CHEI) and ALERTCalifornia, both part of the Qualcomm Institute. The author is deeply indebted to Professor Falko Kuester for hosting the research stay and for the intellectual generosity with which the collaboration on the Palisades Fire application has been developed, and to Scott McAvoy for the ongoing methodological exchange on the airborne LiDAR processing component of the joint research. A preliminary phase of the doctoral research benefited from a visiting fellowship under the Horizon Europe Marie Skłodowska-Curie Staff Exchanges programme, project ReCharged (Grant ID 101086413), conducted at Factor Social in Lisbon during August and September 2023, under the guidance of Professor Gian Paolo Cimellaro, Professor José Manuel Palma-Oliveira, and Doctor Dalila Antunes. The author further acknowledges the technical training received through the Advanced Training Course on SAR and SAR Interferometry, organised by EO59 and held by Professor Daniele Perissin and Doctor Alfredo Rocca.

The author wishes to extend his sincere gratitude to all the professors and colleagues encountered during the doctoral programme, both within the Department of Structural, Geotechnical and Building Engineering of Politecnico di Torino and across the institutions visited during these years, whose presence has contributed in different ways to the academic and personal growth of the candidate. Within the same institution, the author warmly thanks Professor Giuseppe Carlo Marano for welcoming him into the Artificial Intelligence in Structural Engineering (ArtIStE) laboratory at Politecnico di Torino, and Doctor Raffaele Cucuzza, Doctor Jonathan Melchiorre, and Jana Candelaria Olivo Garcia for the advice and the interdisciplinary exchange of ideas. Sincere thanks are due to Professor Anna Reggio for the opportunity to contribute to her teaching in structural mechanics within the Architecture programme over successive academic years. The author also thanks Professor Giorgio Zavarise for entrusting him with the coordination of the international Spaghetti Bridge competition, an experience that brought together engineering, mentoring, and outreach. Warm thanks go as well to Francesco Pimpinella and Mattia Mairone, colleagues and friends with whom many conversations on engineering and on the road ahead have been shared.

A special thought goes to Doctor Stefania Coccimiglio, Doctor Gaetano Miraglia, and Professor Rosario Ceravolo, whose company and dialogue have made the doctoral experience richer in ways that go beyond any single contribution. A separate and heartfelt acknowledgement is owed to Engineer Maurizio Grassi, with whom, alongside Professor Giuseppe Andrea Ferro, the research on satellite-based structural monitoring has been developed and refined over time, and whose engagement and support have been essential to the consolidation of the line of investigation that culminates in the present thesis.

—

The following paragraphs switch to Italian, the author's mother tongue, in which the personal acknowledgements that close this section find their most sincere expression.

Le pagine che seguono cambiano lingua per accogliere un riconoscimento personale che trova nella lingua madre la sua espressione più sincera.

—

Un pensiero affettuoso va agli amici di Via Gambasca 11, la casa che ha fatto da sfondo a questi anni torinesi, per la quotidianità condivisa e per la leggerezza nei

giorni più intensi. E ai Giaguari Torino, la squadra di football americano di prima divisione in cui ho militato per tutti i miei anni a Torino: dentro e fuori dal campo, mi hanno ricordato che i traguardi più belli si raggiungono insieme.

Un ringraziamento profondo va alla mia famiglia di origine, alla quale devo tutto. A mia madre Dolores e a mio padre Lucio, per l'amore incondizionato, per la pazienza con cui hanno accompagnato ogni passo di questo percorso anche quando la distanza lo ha reso più difficile, e per i valori che mi hanno trasmesso e che continuano a guidarmi. A mio fratello Alessandro e a mia cognata Celeste, per la presenza costante e per l'esempio di vita che condividono con me. A mio zio Alessandro, punto di riferimento costante nei momenti che contano davvero.

Un ringraziamento sentito va alla famiglia di Dalila, che mi ha fin da principio accolto come uno dei suoi. A Katia, che ho avuto la fortuna di conoscere e che ci ha lasciato troppo presto: porto con me il ricordo della sua dolcezza e del suo sorriso, e sono profondamente grato per il tempo che abbiamo condiviso. A Francesco, per la generosità con cui ha aperto le porte della sua casa e del suo affetto. A Nico e Roberta, per la complicità e per l'affetto fraterno con cui mi hanno accolto.

A Dalila, mia compagna di vita, va il ringraziamento più importante. Tra miliardi di persone ho avuto la fortuna di incontrarla, e tutto ciò che è venuto dopo, compreso questo lavoro, porta il segno di quell'incontro. La sua presenza ha dato senso a questi anni, il suo sostegno ha reso possibili i traguardi raggiunti, e la sua forza nei momenti più difficili è stata la mia. In questi mesi, mentre la tesi prendeva forma, anche la nostra futura figlia Zelda cresceva con noi, e il pensiero di lei ha accompagnato ogni pagina di questo lavoro. A Zelda, che leggerà queste righe quando sarà grande, va l'augurio che la curiosità e la passione per la conoscenza l'accompagnino nella vita come hanno accompagnato me. A voi, Dalila e a Zelda, e alla famiglia che siamo diventati, va la promessa di una vita serena e piena: davanti a noi ci attendono grandi cose, e a ogni passo le costruiremo insieme.

Raffaele

Torino, 30 aprile 2026

Abstract

The translation of satellite multi-temporal Interferometric Synthetic Aperture Radar (MT-InSAR) measurements into actionable structural diagnostics is one of the principal open challenges in contemporary Structural Health Monitoring (SHM) of existing structures. Persistent Scatterer Interferometry (PSI) has matured rapidly as a source of millimetric displacement information. The gap between satellite displacement records and the diagnostic categories required by structural engineering practice has nevertheless remained largely unresolved, particularly at the building scale that underpins decision-making on existing assets.

The present thesis develops a coherence-aware methodological framework that articulates the diagnostic value of satellite InSAR observations across four progressively more demanding levels. The first level is screening. The second is displacement characterisation. The third is diagnostic interpretation. The fourth is continuous monitoring with automated anomaly detection. The framework rests on the principle that temporal coherence is a modelling variable rather than a quality filter, and that progression between levels depends on data conditions, complementary information sources, and explicit uncertainty quantification.

The framework is tested through three completed case studies and one ongoing research direction. The *Santissima Trinità* in Turin establishes that a null finding of structural stability is a meaningful monitoring outcome, and that the apparent point-density advantage of standardised European Ground Motion Service (EGMS) products on heritage masonry conceals a coherence cost that user-driven processing of raw Single-Look Complex (SLC) data systematically recovers. The retrospective forensic analysis of the *Himera Viaduct* and the *Albiano-Magra Bridge* formalises the Temporary Scatterer (TS) as a systematic diagnostic category. Lifecycle boundaries of TS encode the chronology of structural events at the resolution of the satellite revisit cycle. The *Sanctuary of Vicoforte* achieves the most complete realisation of the framework, identifying an application-specific coherence threshold of 0.56 through a Sparse Bayesian Learning (SBL) surrogate model. The same case operationalises the satellite displacement record as a stochastic oscillator monitoring statistic with continuous prediction capability. The ongoing application to the 2025 Palisades Fire in Los Angeles, conducted in collaboration with the Qualcomm Institute at the University of California San Diego, within the Cultural

Heritage Engineering Initiative (CHEI) and ALERTCalifornia, extends the framework into a post-event regime by fusing PS-InSAR scatterer lifecycle indicators with airborne Light Detection and Ranging (LiDAR) geometric change metrics at the per-building scale.

Four transferable methodological contributions emerge from the work. The first is the formalisation of the Temporary Scatterer as a systematic diagnostic category. The second is the coherence-aware SBL surrogate modelling framework. The third is the four-level diagnostic articulation that maps satellite observables onto decision-support categories. The fourth is the multi-modal fusion architecture for post-event building-scale damage assessment. The decision-ready diagnostic capability that gives the thesis its title is offered to the discipline as a contribution from which further developments in operational satellite-based SHM may proceed.

Keywords

Structural Health Monitoring; Satellite InSAR; Persistent Scatterer Interferometry; PS-InSAR; Multi-temporal Interferometry; EGMS; Sentinel-1; COSMO-SkyMed; Temporal Coherence; Sparse Bayesian Learning; Temporary Scatterer; Masonry Heritage; Bridge Monitoring; Infrastructure Collapse; LiDAR Fusion; Wildfire Damage Assessment; Decision-Ready Diagnostics

Table of contents

Acknowledgments	i
Abstract.....	iv
Table of contents.....	vi
List of figures.....	xi
List of tables.....	xxvii
Nomenclature and acronyms	xxx
Publications.....	xxxv
Research field and motivation	1
1.1. Introduction.....	1
1.2. Structural Health Monitoring of full-scale structures	3
1.2.1. Motivation for SHM	4
1.2.2. Challenges and constraints: scale, access, cost, and continuity...5	
1.2.3. Types of monitoring systems: contact, remote, and hybrid approaches.....	6
1.2.4. The role of satellite observations in SHM	8
1.3. Objectives and contributions of the thesis	10
1.4. Research questions and scope.....	12
1.5. Thesis outline and reading guide	19
Satellite data for SHM: state of the art	22
2.1. Satellite remote sensing basics	22
2.1.1. SAR sensors and acquisition modes	24
2.1.2. Geometry and resolution: wavelength, incidence angle, polarization, revisit time	29
2.1.3. Data types, DEMs, and auxiliary products	37
2.2. Satellite data in structural engineering.....	40

2.2.1.	Ground and structural displacements.....	41
2.2.2.	LOS measurements, decomposition issues, and thermal/seasonal effects	43
2.2.3.	Main limitations: decorrelation, geometric distortions, atmosphere, and referencing	47
2.3.	Interpretation frameworks and literature applications	53
2.3.1.	From geophysical origins to structural applications	54
2.3.2.	Cultural heritage and historical structures	54
2.3.3.	Bridges and transport infrastructure	56
2.3.4.	Urban buildings and underground construction effects	59
2.3.5.	Portfolio-level screening and prioritisation	60
2.3.6.	Integration with in-situ monitoring and numerical models	61
2.3.7.	Retrospective forensic analysis of structural failures	62
2.3.8.	The ReLUIIS guidelines, standardised products and the interpretation gap	66
2.3.9.	Remaining gaps and positioning of the thesis	67
	Interferometric SAR data and processing workflow	69
3.1.	Principles of InSAR and multi-temporal InSAR	69
3.1.1.	DInSAR	70
3.1.2.	PSI and PS-based approaches	80
3.1.3.	SBAS and related approaches	89
3.2.	Data sources and platforms.....	92
3.2.1.	Sentinel-1	93
3.2.2.	COSMO-SkyMed	94
3.2.3.	EGMS platform: products, access, and caveats	95
3.3.	User-driven processing workflow.....	100
3.3.1.	Dataset selection and acquisition strategy	101
3.3.2.	Co-registration and interferometric network construction.....	103
3.3.3.	DEM integration, geocoding, and export.....	104

3.3.4.	Small area processing	105
3.4.	Quality assurance and quality control.....	108
3.4.1.	Reflectivity, amplitude stability index, and sparse point selection 108	
3.4.2.	Reference point selection.....	116
3.4.3.	Temporal coherence and quality thresholds	118
3.5.	Validation issues and uncertainty sources	119
3.5.1.	Geocoding and localization	120
3.5.2.	Referencing strategies	121
3.5.3.	Precision, accuracy, and uncertainty propagation	122
	Methodology for satellite-based SHM of structures.....	124
4.1.	Point selection on structures	125
4.1.1.	Persistent scatterers on built assets	125
4.1.2.	Temporary scatterers and target lifecycle	129
4.1.3.	Geometric plausibility: roofs, facades, edges, and metallic elements 143	
4.2.	Time-series analytics for structural interpretation	146
4.2.1.	Linear trends	146
4.2.2.	Seasonal and thermal components.....	148
4.2.3.	Change-points and non-linear behaviour	155
4.3.	From LOS measurements to structural response metrics	159
4.3.1.	Relative motion and reference dependence	160
4.3.2.	Settlement patterns, gradients, and tilt-like behaviour	164
4.3.3.	Sanity checks for structural interpretation	165
4.3.4.	Retrospective event analysis and forensic constraints.....	167
4.4.	Handling heterogeneous datasets.....	169
4.4.1.	Sentinel-1 versus COSMO-SkyMed.....	169
4.4.2.	Harmonization and co-referencing	172
4.4.3.	Bias removal and interpretability issues	173

4.5.	Integration with complementary information	174
4.5.1.	Cross-validation with in-situ monitoring.....	174
4.5.2.	Integration with numerical models	175
4.5.3.	Fusion with non-SAR remote sensing data.....	176
4.6.	Toward decision-ready diagnostics: from displacement to engineering judgement	179
	Applications and case studies	184
5.1.	Screening and prioritization using EGMS products	184
5.2.	Case study I - Santissima Trinità in Turin: building-scale SHM of a masonry heritage structure	187
5.2.1.	Structural and historical context	188
5.2.2.	Available satellite datasets.....	196
5.2.3.	EGMS-based analysis	199
5.2.4.	User-driven Sentinel-1 analysis.....	204
5.2.5.	Comparison of point density, quality, and interpretability	208
5.2.6.	Implications for masonry heritage SHM	212
5.3.	Case study II - Retrospective forensic analysis of infrastructure collapses: the Himera Viaduct and the Albiano-Magra Bridge	215
5.3.1.	Himera Viaduct: EGMS and user-driven Sentinel-1 analysis	217
5.3.2.	Albiano-Magra Bridge: high-resolution COSMO-SkyMed interpretation	226
5.3.3.	Temporary scatterers, pre-/post-event analysis, and early-warning issues	237
5.4.	Case study III - The Sanctuary of Vicoforte: coherence-aware satellite monitoring and soil-structure interaction.....	240
5.4.1.	Structural and geotechnical context.....	243
5.4.2.	Available in-situ and satellite data.....	247
5.4.3.	PS-InSAR processing and temporal coherence assessment	253
5.4.4.	Integration with soil temperature and moisture variables.....	266
5.4.5.	Data-driven modelling of displacement behaviour.....	270

5.4.6.	Implications for integrated SHM of monumental buildings	280
5.5.	Synthesis for engineering practice	283
5.5.1.	A tiered framework for satellite-assisted SHM: from portfolio screening to structure-specific diagnostics	284
5.5.2.	Point density and spatial coverage on structures	289
5.5.3.	Coherence as a diagnostic variable	291
5.5.4.	Thermal and environmental effects on displacement interpretation 293	
5.5.5.	Temporary scatterers, retrospective inference, and the early-warning question	296
5.5.6.	Implications for decision-ready diagnostics	298
5.6.	Ongoing research: towards the fusion of PS-InSAR and airborne LiDAR for building-scale post-event assessment: the 2025 Palisades Fire (Los Angeles, California)	303
5.6.1.	Event context and motivation for multi-sensor assessment.....	305
5.6.2.	Study area, satellite datasets, and complementary LiDAR material 306	
5.6.3.	Preliminary PS-InSAR processing and scatterer lifecycle characterisation	310
5.6.4.	Conceptual framework for multi-modal fusion at the building-footprint scale.....	314
5.6.5.	Preliminary implications and directions for completion.....	315
	Discussion and conclusions	318
6.1.	The four methodological contributions.....	320
6.2.	Response to the research questions.....	323
6.3.	Limitations and threats to validity	325
6.4.	Recommendations for engineering practice	328
6.5.	Future research directions	331
6.6.	Closing of the thesis.....	332
	Bibliography	334

List of figures

Figure 2. 1. Passive remote sensing detection mechanisms. The target is illuminated by an external source (typically solar radiation) and the sensor records the reflected energy; measurement is constrained to daylight conditions and clear atmosphere (Coccimiglio, 2025).....	23
Figure 2. 2. Active remote sensing detection mechanisms. The platform emits a microwave pulse towards the target and records both amplitude and phase of the back-scattered signal; operation is independent of solar illumination and cloud cover, and phase coherence across repeated acquisitions enables interferometric displacement retrieval (Coccimiglio, 2025).....	23
Figure 2. 3. Timeline of the principal spaceborne SAR missions relevant to SHM, with indication of the frequency band, from 1992 to the present. Adapted from (Deslauriers et al., 2025).	24
Figure 2. 4. SAR side-looking geometry and synthetic aperture concept (Braun, 2019).	25
Figure 2. 5. Near polar and sun-synchronous low earth orbit configuration of spaceborne SAR missions. Adapted from (Types of Orbits, 2020).....	26
Figure 2. 6. Satellite acquisition geometry for ascending and descending orbits (Coccimiglio, 2025).....	27
Figure 2. 7. Comparison of the three principal SAR acquisition modes relevant to Structural Health Monitoring: Stripmap, ScanSAR, and TOPS. Adapted from (Boncori & Peter, 2019).....	28
Figure 2. 8. Main bands of electromagnetic spectrum.	29
Figure 2. 9. Differential penetration of microwave signals through vegetation as a function of the operating wavelength. Adapted from (Airbus Built Forest Monitoring Satellite Biomass Successfully Launched, 2025).	30
Figure 2. 10. Acquisition geometry of a spaceborne SAR system. Adapted from (Franceschetti & Lanari, 2018).	32
Figure 2. 11. Schematic representation of satellite radar acquisition in the case of Ascending orbit (left) and Descending orbit (right).	33

Figure 2. 12. Polarization: orientation of the transmitted and received electromagnetic field. Adapted from (Introduction to SAR—ArcGIS Pro Documentation, 2026).	34
Figure 2. 13. Range and azimuth resolution.....	36
Figure 2. 14. SLC vs GRD conceptual comparison. Adapted from (Colin, 2022).	38
Figure 2. 15. Definition of baseline (B), normal baseline ($B_{\perp}$) and parallel baseline ($B_{\parallel}$) in the plane orthogonal to the orbits.	39
Figure 2. 16. Decomposition of the displacement vector along the satellite LOS. Left panel: projection of the horizontal displacement components onto the satellite-looking direction, where α is the azimuth of the satellite heading vector (measured positive clockwise from north), SH is the satellite heading direction, SL is the satellite-looking direction, and d_E and d_N are the easting and northing components of the horizontal displacement, respectively. Right panel: projection of the vertical and horizontal displacement components onto the LOS direction, where d_U is the vertical displacement component and θ is the local incidence angle at the reflection point. Adapted from (Ng, 2010).	44
Figure 2. 17. Schematisation of InSAR satellite orbital tube.....	48
Figure 2. 18. Schematic of foreshortening, layover, and shadow distortions in side-looking SAR geometry. Ground points a, b, c are “seen” by radar as points a', b', c' in the slant range. (Alaska Satellite Facility, 2026a).....	49
Figure 2. 19. SAR image of the Eiffel Tower in Paris. Adapted from (Capella, 2024).	50
Figure 2. 20. Example of atmospheric artefacts on a differential interferogram of the Naples area, computed from a tandem pair with normal baseline of approximately -97 m. Left panel: differential tandem interferogram, in which the fringe pattern is dominated by atmospheric contributions. Right panel: Atmospheric Phase Screen estimated for the same image pair through the PS technique. As clearly visible, most of the interferometric fringes derive from atmospheric effects. Adapted from (Parizzi et al., 2006).....	51
Figure 2. 21. Distribution of COSMO-SkyMed PSs over the archaeological area of the Palatine Hill, Roman Forum, and Colosseum, colour-coded by mean LOS velocity over the observation period from March 2011 to June 2013. Panel (a) provides an overview of the archaeological site, while panels (b), (c), and (d) report detailed views of the arches of the Domitian aqueduct along Via di San Gregorio,	

of the Convento di San Bonaventura al Palatino, and of the Chiesa di Sant’Anastasia with the archaeological ruins in the south-western corner of the Palatine Hill, respectively. The figure illustrates the high PS density achievable on heritage masonry surfaces with X-band high-resolution sensors and the spatial selectivity required for element-level structural diagnosis.55

Figure 2. 22. PSs detected on the Consoli Palace in Gubbio from the ascending COSMO-SkyMed dataset through Persistent Scatterer Pair (PSP) analysis, colour-coded by mean LOS velocity. The left panel reports the displacement time series of two representative PSs located respectively on the roof and at the base of the palace, illustrating the differential motion between elevated and basement structural zones that the PSP analysis is designed to resolve. The figure exemplifies the building-scale interpretive use of high-resolution X-band PS data on a heritage masonry structure.56

Figure 2. 23. Surface deformation map (upper panel) and thermal dilation map (lower panel) of the Lupu Bridge, in Shanghai, China, derived from multi-temporal SAR interferometric processing.....58

Figure 2. 24. Distribution of the PSs identified in the case study area, colour-coded according to their temporal coherence in panel (a) and to their mean LOS displacement velocity over the full observation period in panel (b). The juxtaposition of the two panels illustrates how regions of low coherence may coincide with regions of apparent high velocity, highlighting the role of coherence as a quality filter that conditions the interpretive value of the velocity field.59

Figure 2. 25. Schematic representation of the MT-InSAR and relative stiffness method for the integrated assessment of tunnelling-induced damage on surface buildings, combining MT-InSAR displacement observations with the estimation of the relative stiffness between structure and ground.....60

Figure 2. 26. Highway network in the Rome metropolitan area analysed as case study, with interferometric displacement velocity maps of the PSs for two selected segments affected by subsidence. The two highlighted segments exhibit cumulative LOS displacements consistent with localised settlement, demonstrating the use of corridor-scale InSAR analysis for the prioritisation of infrastructure inspection along extended linear assets.....61

Figure 2. 27. Three-dimensional geological model of the Ponte della Musica subsoil, view from north-west (upper panel), and mean deformation velocity maps of the bridge for the ascending orbit (a) and descending orbit (b) acquisitions in false colours.62

Figure 2. 28. The Morandi Bridge before the 14 August 2018 collapse. Panel (A): two-dimensional structural scheme showing, from west to east, two single piers (1 and 2), six V-shaped piers (3 to 8), and three independent balanced systems (9 to 11), with dimensions in metres. Panel (B): three-dimensional view of the bridge.63

Figure 2. 29. Multi-temporal InSAR analysis of the Morandi Bridge combining COSMO-SkyMed and Sentinel-1A/B acquisitions. Panel (A): COSMO-SkyMed descending dataset, with PS height (top), cumulative LOS displacement at deck level over the period 2009-2018 (centre), and LOS displacement at ground-level points over the same period (bottom). Panel (B): Sentinel-1 descending dataset, with the same three layers over the period 2015-2018. Panel (C): COSMO-SkyMed ascending dataset, with the same three layers over the period 2011-2018.64

Figure 2. 30. Mt-InSAR analysis of the Morandi Bridge focused on the area of the collapsed pillar. Panels (a) and (b): SBAS and TomoSAR processing results from the ascending dataset. Panels (c) and (d): SBAS and TomoSAR processing results from the descending dataset.65

Figure 3. 1. Geometrical representation of the amplitude and phase of a SAR wave.70

Figure 3. 2. Panel (a): amplitude image of a representative scene, with brightness proportional to the backscattered intensity Panel (b): wrapped phase image of the same scene, with values in the interval from $-\pi$ to $+\pi$ and no immediate physical interpretability prior to interferometric processing. Adapted from (Balaji, 2024).71

Figure 3. 3. Geometry of the SAR interferometric baseline. Two satellite passes image the same target from orbital positions S1 and S2, separated by a baseline B decomposed into its perpendicular and parallel components with respect to the LOS. Adapted from (Alaska Satellite Facility, 2026b).72

Figure 3. 4. Descending (a) and ascending (b) interferograms of the 1999 Mw 7.1 Hector Mine earthquake in California. Each colour cycle represents ten centimetres of LOS displacement towards or away from the satellite. Purple plus and minus signs indicate areas of positive and negative LOS displacement; the mapped fault trace is shown as a thick line and the epicentre is marked with a star. Adapted from (Jónsson, 2002).73

Figure 3. 5. Effect of the flat-earth and topographic phase removal on a representative interferogram. Panel (a): Digital Elevation Model of the area of interest. Panel (b): raw interferogram before the removal of any systematic phase

contribution. Panel (c): interferogram after removal of the flat-earth phase, with residual fringes reflecting the topographic contribution. Panel (d): differential interferogram after removal of both the flat-earth and the topographic phase, retaining the displacement, atmospheric, and noise contributions.77

Figure 3. 6. Two interferograms of the Vesuvius area computed with different perpendicular baselines and presented without subtraction of the external DEM, in order to expose the dependence of the topographic fringe density on the baseline. Panel (a): interferogram with a 46 m perpendicular baseline, corresponding to a height of ambiguity of approximately 351.8 m. Panel (b): interferogram with a 178 m perpendicular baseline, corresponding to a height of ambiguity of approximately 93.4 m. The denser fringe pattern in panel (b), with approximately four times the fringe density of panel (a), reflects the linear dependence of the topographic phase sensitivity on the perpendicular baseline established by Equation 2. 4..... 79

Figure 3. 7. Phase unwrapping basic principle schematic. Adapted from (W. Yang et al., 2024).....80

Figure 3. 8. Effect of multi-temporal averaging on the radar reflectivity map of an urban area in Turin, illustrating the progressive reduction of speckle as the number of averaged acquisitions increases. Panel (a) reports the reflectivity map computed from five Sentinel-1 acquisitions, in which the salt-and-pepper texture characteristic of speckle dominates the visual appearance and obscures the structural detail of the urban fabric. Panel (b) reports the reflectivity map computed from 176 Sentinel-1 acquisitions over the same area, in which the temporal averaging suppresses the speckle and recovers the stable scattering signature of buildings, road surfaces, and built infrastructure.....83

Figure 3. 9. Geometric interpretation of the temporal coherence as a vector average in the complex plane. Panel (a): unit-length vectors aligned in a common direction, corresponding to small residual phases after the removal of the modelled topographic and deformation contributions, with the modulus of their average close to one. Panel (b): unit-length vectors distributed across the full interval from minus π to π , corresponding to large residual phases, with the modulus of their average close to zero.88

Figure 3. 10. Block diagram of the SBAS algorithm, showing the sequential processing steps from the generation of the differential interferograms through the SVD inversion to the retrieval of the deformation time series.91

Figure 3. 11. Revisit times of the satellites of the Sentinel-1 constellation. Adapted from (Prota et al., 2025)93

Figure 3. 12. Revisit times of the first generation (CSK) and second generation (CSG) COSMO-SkyMed constellation satellites. Adapted from (Prota et al., 2025)	95
Figure 3. 13. Calibrated product in the EGMS Data Explorer for a representative site, with the spatial distribution of the measurement points colour-coded by mean LOS velocity in the map view and the displacement time series of a selected point reported in the adjacent plot.	96
Figure 3. 14. Ortho product in the EGMS Data Explorer for a representative site, with the vertical and east-west displacement components resampled on a regular 100 m grid and the corresponding time series of a selected cell reported in the adjacent plot.	97
Figure 3. 15. EGMS product overview and the main product characteristics	97
Figure 3. 16. Summary of the principal applications of EGMS products in civil engineering and infrastructure monitoring, organised by category of engineering objective and by the level of EGMS product most appropriate to each category.	99
Figure 3. 17. Star-graph baseline plot of a representative dataset, with normal baseline reported on the vertical axis and acquisition date reported on the horizontal axis. The master image is in the centre, and each interferometric connection between master and slave is colour-coded according to the spatial coherence.	102
Figure 3. 18. Example of the meteorological record. From top to bottom: air temperature, atmospheric pressure, relative humidity, and daily precipitation.. The four variables are reported on a common time axis covering the observation period.	103
Figure 3. 19. Comparison between an optical image and a SAR amplitude image of the same scene, illustrating the distinct information content of the two acquisition modalities. Panel (a): SAR amplitude image of a representative area. Panel (b): optical image of the same area, acquired under standard daylight conditions, with partial cloud cover. The correspondence between visual prominence in the optical image and radar response is not direct. Adapted from (Gao et al. 2019).	110
Figure 3. 20. Comparison between an optical reference image and the SAR reflectivity map of the same urban coastal area. Panel (a): Google Earth optical view. Panel (b): mean amplitude image computed from the multi-temporal SAR stack, showing the characteristic backscatter signatures associated with water bodies, built-up areas, and vegetated surfaces.	111

Figure 3. 21. Radar signature of an isolated point-like reflector in an urban area. Panel (a): Google Earth optical view showing a light pole at the centre of the scene. Panel (b): SAR reflectivity map of the same area, in which the light pole appears as a bright pixel surrounded by attenuated side lobes in the range direction, the characteristic signature of a strong point-like scatterer. 112

Figure 3. 22. Cardinal sine function, which describes the impulse response of the SAR system in both the range and azimuth directions. The main lobe is centred at $x = 0$ with peak value of unity. The successive zeros occur at integer values of x , and the side lobes have progressively decreasing amplitude. The width of the main lobe between the first two zeros defines the nominal spatial resolution of the image, and the side lobe structure governs the multiplicative ambiguity of point scatterer responses in the range-azimuth plane. 113

Figure 4. 1. Example of an artificial corner reflector mounted on the *Sanctuary of Vicoforte* to enhance the radar visibility of selected structural points and to support the identification of stable measurement targets in the multi-temporal interferometric stack. 128

Figure 4. 2. Schematic representation of the four scatterer lifecycle configurations, expressed through the amplitude profile across the acquisition sequence. Panel (a): Temporary Scatterer of the first sub-type ($T_{on} > 1$, $T_{off} = N$), in which the scatterer appears at a datable epoch within the observation window and remains coherent until the end of the stack. Panel (b): Temporary Scatterer of the second sub-type ($T_{on} = 1$, $T_{off} < N$), in which the scatterer is present from the first acquisition and disappears at a datable epoch before the end of the stack. Panel (c): Temporary Scatterer of the third sub-type ($T_{on} > 1$, $T_{off} < N$), in which the entire coherence window of the scatterer falls within the observation period, bounded by datable T_{on} and T_{off} epochs on both sides. Panel (d): Persistent Scatterer ($T_{on} = 1$, $T_{off} = N$), with high and stable amplitude maintained throughout the entire stack. 132

Figure 4. 3. Multi-channel change detection through RGB amplitude composite over a representative urban port area. Panels (a), (b), and (c): mean amplitude maps computed from three SAR acquisitions at distinct epochs, with brightness proportional to the backscattered intensity of each scene. Panel (d): RGB composite obtained by assigning the amplitude maps of panels (a), (b), and (c) to the red, green, and blue channels respectively. Targets present in all three epochs appear as white or grey pixels, while targets present in only one or two epochs produce coloured features whose hue identifies the temporal interval of their appearance or disappearance. 134

Figure 4. 4. Coherence change detection composite on the same representative port area of Figure 4. 3. Panels (a), (b), and (c): coherence maps of three interferometric pairs computed on intervals of progressively later epochs within the observation window. Panel (d): RGB composite obtained by assigning panels (a), (b), and (c) to the red, green, and blue channels of a single image. Pixels with high coherence across all three pairs appear as white or light grey in the composite, while coloured pixels identify surfaces whose coherence is preserved in only a subset of the three intervals..... 136

Figure 4. 5. ASI of the candidate population over the same representative port area of Figure 4. 3. Panel (a): spatial map of the ASI, with the value at each candidate point reported on a colour scale ranging from zero for highly variable amplitude records to unity for stable persistent scatterers. Negative ASI values, which arise when the temporal standard deviation of the amplitude exceeds the temporal mean, are collapsed onto the zero level for the purpose of the colour mapping, and a distinct colour is assigned to identify their occurrence on the map. Panel (b): histogram of the ASI distribution across the same candidate population. 138

Figure 4. 6. Application of the TS workflow to a high-rise building under construction during the observation period. Panel (a): three-dimensional view on Google Earth of the PSs and TSs identified on the building, colour-coded by the estimated T_{on} , with PSs appearing at the first acquisition shown in dark blue and TSs appearing at successive epochs shown in progressively warmer colours. Panel (b): amplitude-temporal baseline plot of a representative TS selected from the upper portion of the new building, with the high-amplitude cluster identifying the live interval bounded by the estimated T_{on} and T_{off} . Panel (c): displacement time series of the same representative TS restricted to the live interval. In panel (b), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panel (c), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend..... 142

Figure 4. 7. Application of the parametric PSI estimation with thermoelastic coefficient on a representative high-rise building. Panel (a): reflectivity map of the area, computed as the temporal mean of the 72-image SAR amplitude stack. Panel (b): temporal coherence obtained with the two-parameter model that estimates residual height and linear velocity. Panel (c): temporal coherence obtained with the three-parameter model that adds the thermoelastic coefficient as a third unknown. Panel (d): spatial distribution of the estimated thermoelastic coefficient on the

building, expressed in radians per degree Celsius, exhibiting the characteristic vertical gradient that increases with the height of the PS above the base. 154

Figure 4. 8. Phase replica diagnostic and displacements reconstruction for a representative scatterer affected by non-linear deformation. Panel (a): displacement time series estimated by the periodogram inversion conditioned on the linear trend model, with the phase replicas. Panel (b): displacement time series of the same scatterer estimated without the imposition of the linear trend model, with the periodogram inversion restricted to the residual height parameter only. 159

Figure 4. 9. Residual height distribution on a high-rise building observed by X-band stack of 72 acquisitions, processed with three different reference point positions. Panel (a): reference point at the base of the building, on a stable ground surface adjacent to the building footprint. Panel (b): reference point at mid-height on the building. Panel (c): reference point at the top of the building. 163

Figure 4. 10. PSs identified over a motorway junction in Rome using medium-resolution C-band Envisat data acquired between 2002 and 2010 in panel (a) and high-resolution X-band COSMO-SkyMed data acquired between 2011 and 2014 in panel (b). Each PS is represented by a dot colour-coded by cumulative LOS displacement, with positive and negative values indicating motions towards and away from the satellite respectively. The contrast between the two panels illustrates the substantial gain in PS density achievable with X-band sensors on a built infrastructure asset, and the corresponding improvement in the spatial selectivity of the displacement characterisation at the element level.. 171

Figure 5. 1. Panel (a): Aerial view of the *Santissima Trinità* in Turin, embedded within the surrounding historic block. Panel (b): Plan of the urban block with the footprint of the church highlighted. 189

Figure 5. 2. *Santissima Trinità* in Turin: panel (a), plan; panel (b), longitudinal section of the dome and lantern. 190

Figure 5. 3. Close-range aerial view of the dome system of the *Santissima Trinità*. (a) “Cupolino” of the lantern, clad with metallic sheet. (b) Lantern shaft in exposed brick and tile cover of the dome supported by masonry false struts. (c) Metallic-clad drum at the base of the dome. 193

Figure 5. 4. Time series of selected points: (a) P₂₀, (b) P₂₄, (c) P₁₈₀₃, and (d) P₁₇₆₅. 202

Figure 5. 5. Quality attributes of the user-driven Sentinel-1 dataset: (a) star graph of the interferometric network centred on the master image with the

perpendicular baseline of each spoke and the colour of each spoke encoding the temporal coherence of the corresponding pair; (b) distribution of normal baselines across the 169 acquisitions; (c) distribution of Doppler centroid frequencies; (d) local temperature at the acquisition dates.205

Figure 5. 6. Small-area maps for the user-driven Sentinel-1 processing of the *Santissima Trinità*: mean amplitude across the 169-image stack.206

Figure 5. 7. Measurement points from the EGMS Calibrated product and from the user-driven Sentinel-1 processing over the *Santissima Trinità* in Turin, displayed on Google Earth in three-dimensional projection: panel (a) reports the area-wide distribution of all retained measurement points across the urban block centred on the church, while panel (b) reports the detail of the four selected points on the church itself, namely P20 and P24 from the EGMS Calibrated product and P1803 and P1765 from the user-driven processing. The four points are the basis for the comparative analysis of Section 5.2.4.....210

Figure 5. 8. The two infrastructure case studies of Section 5.3: (a) *Himera Viaduct*, collapsed deck (April 2015); (b) *Himera Viaduct*, rebuilt deck; (c) *Albiano-Magra Bridge*, historical view; (d) *Albiano-Magra Bridge*, collapse (April 2020).217

Figure 5. 9. User-driven Sentinel-1 processing of the *Himera Viaduct* over the period October 2014 to April 2016. Panels (a) and (b) report the star graphs of the interferometric networks in the ascending and descending orbits respectively, with each spoke representing one interferometric pair connecting a slave acquisition to the master, and with the colour of each spoke encoding the temporal coherence of the corresponding pair. Panels (c) and (d) report the reflectivity maps of the small-area processing window in the ascending and descending orbits respectively, with the position of the GCP selected to anchor the LOS solution highlighted in each map. The dual-orbit configuration provides the geometric diversity required to distinguish horizontal from vertical components of the displacement field.....218

Figure 5. 10. Slope-adjacent temporary scatterer identified on the *Himera Viaduct* in the ascending Sentinel-1 stack, interpreted as transient debris subsequently cleared from the site: (a) amplitude-temporal baseline plot; (b) displacement time series. In panel (a), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panel (b), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.220

Figure 5. 11. COSMO-SkyMed PSs over the *Himera Viaduct* and the adjacent Rocca di Sciara slope, processed in user-driven mode for the pre-collapse period from 2011 to 2014, colour-coded by mean LOS velocity. Panel (a) reports the ascending orbit, panel (b) the descending orbit. The dual-orbit comparison enables the qualitative attribution of velocity signatures to vertical or horizontal motion components on the structure and on the surrounding terrain.....222

Figure 5. 12. Relative displacement time series of two outliers on the *Himera Viaduct* identified by COSMO-SkyMed processing of the 2011-2014 archive: (a) ascending orbit, relative LOS velocity 6.11 mm/year, two-sigma exceedance from mid-2013; (b) descending orbit, relative LOS velocity 5.68 mm/year, two-sigma exceedance from early 2014. Adapted from (Macchiarulo, Milillo, et al., 2022).222

Figure 5. 13. EGMS Calibrated measurement points on the rebuilt deck of the *Himera Viaduct* (2018-2023): spatial distribution of the points on the reconstructed segments, with a panel in the lower left comparing the displacement time series of two representative measurement points.224

Figure 5. 14. User-driven COSMO-SkyMed processing of the *Albiano-Magra Bridge* (2011-2022): (a, b) coherence-coloured star graphs of the interferometric networks in ascending (a) and descending (b) orbits; (c, d) reflectivity maps of the Magra River crossing with the Ground Control Point indicated, ascending (c) and descending (d).....228

Figure 5. 15. Pre-collapse PSs over the *Albiano-Magra Bridge* and the surrounding terrain, processed from the COSMO-SkyMed Stripmap archive over the period 2011 to April 2020 in user-driven mode and colour-coded by mean LOS velocity. Panel (a) reports the ascending orbit, panel (b) the descending orbit. The Caprigliola hill on the eastern side of the bridge is included within the spatial extent of both panels, allowing the joint observation of bridge-scale displacement signatures and slope-scale ground motion. The pre-collapse Caprigliola cluster discussed in Section 5.3.2 is visible in both panels with opposite-sign LOS velocities consistent with a predominantly horizontal displacement field..230

Figure 5. 16. Amplitude-temporal baseline plots and displacement time series of a representative PS on a still-standing pier and a representative TS on a failed span of the *Albiano-Magra Bridge* in the ascending orbit (COSMO-SkyMed, 2011-2022): (a) PS amplitude-temporal baseline plot; (b) PS displacement time series; (c) TS amplitude-temporal baseline plot; (d) TS displacement time series. In panels (a) and (c), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panels (b) and (d), red triangles represent

individual LOS displacement values, and the solid red line interpolates the displacement trend.232

Figure 5. 17. Amplitude-temporal baseline plots and displacement time series of a representative PS on a still-standing pier and a representative TS on a failed span of the *Albiano-Magra Bridge* in the descending orbit (COSMO-SkyMed, 2011-2022): (a) PS amplitude-temporal baseline plot; (b) PS displacement time series; (c) TS amplitude-temporal baseline plot; (d) TS displacement time series. In panels (a) and (c), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panels (b) and (d), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.233

Figure 5. 18. Spatial distribution of persistent and temporary scatterers identified by user-driven COSMO-SkyMed processing on the *Albiano-Magra Bridge* and on the Caprigliola hill (2011-2022) on a Google Earth substrate, colour-coded by death date in both orbital geometries: (a) ascending orbit; (b) descending orbit.235

Figure 5. 19. Geometric and structural overview of the *Sanctuary of Vicoforte*: (a) aerial view of the Sanctuary and its immediate surroundings; (b) Internal view of the oval masonry dome.240

Figure 5. 20. Geometric and structural overview of the *Sanctuary of Vicoforte*: (a) plan of the structure, highlighting the oval dome footprint; (b) longitudinal section through the dome and the drum.241

Figure 5. 21. Schematic cross-section of the subsoil stratigraphy beneath the *Sanctuary of Vicoforte*, showing the variable thickness of the clayey-silt deposit (approximately 5 m on the south-western side, reducing to centimetres on the eastern side), the marlstone substratum, and the differential bearing conditions of the eight foundation buttresses (a). Structural plan of the *Sanctuary of Vicoforte* at foundation level, the eight foundation buttresses are labeled (b).244

Figure 5. 22. Survey of the cracks on the longitudinal section (a) and detail of the cracks on the openings (b).245

Figure 5. 23. Section of the dome: original iron circular ring and modern post-tensioning ring on the drum (left) and detail of the post-tensioning ring system (right).246

Figure 5. 24. Temporal correspondence between the water-table level measured by the piezometers and the surface soil moisture measured by the soil thermo-

hygrometer at the *Sanctuary of Vicoforte*. The close agreement between the two signals supports the use of surface moisture as the hydrological input to the data-driven model.248

Figure 5. 25. Layout of the soil monitoring sub-system at the *Sanctuary of Vicoforte*. The three piezometers PZ01, PZ02 and PZ03 monitor the groundwater level in the south-western sector of the Sanctuary, where the clayey-silt deposit reaches its maximum thickness. The soil thermo-hygrometer measures surface soil moisture and temperature at hourly sampling.....249

Figure 5. 26. Native time series of surface soil moisture and temperature recorded by the in-situ soil thermo-hygrometer at the *Sanctuary of Vicoforte* at one-hour sampling. The seasonal cycles and the inverse correlation between temperature peaks and moisture troughs are visible throughout the observation period.....250

Figure 5. 27. Spatial distribution of EGMS measurement points on and around the *Sanctuary of Vicoforte*: (a) Ortho dataset; (b) Calibrated dataset. The labels D1st and D2nd identify the descending acquisitions of the 2018-2022 and 2019-2023 windows respectively; A11st and A12nd identify the first ascending track for the same windows; A21st and A22nd identify the second ascending track.251

Figure 5. 28. Temporal coherence analysis of the EGMS Calibrated measurement points over the *Sanctuary of Vicoforte* in the descending orbit, with the 0.6 coherence threshold indicated by the red line: (a) histogram of the temporal coherence values; (b) spatial distribution of the points retained above and below threshold.....255

Figure 5. 29. Temporal coherence analysis of the EGMS Calibrated measurement points over the *Sanctuary of Vicoforte* in the two ascending orbital tracks across the 2018-2022 and 2019-2023 windows, with the 0.6 coherence threshold indicated by the red line: (a) and (b) ascending track 1, with histogram and spatial distribution; (c) and (d) ascending track 2, with histogram and spatial distribution.257

Figure 5. 30. Interferometric configuration of the user-driven COSMO-SkyMed Second Generation stack over the *Sanctuary of Vicoforte*: (a) histogram of the distribution of normal baselines across the 62 acquisitions; (b) star-shaped interferometric graph centred on the master, with edges coloured according to the average spatial coherence; (c) Copernicus DEM of the regional setting, with the area of interest indicated by the black crosses.258

Figure 5. 31. Mean reflectivity and spatial coherence of the user-driven COSMO-SkyMed Second Generation stack over the *Vicoforte* area: (a) mean

amplitude image with the Ground Control Point indicated in red; (b) spatial coherence map.259

Figure 5. 32. Finite element model of the *Sanctuary of Vicoforte*: (a) meshed geometry; (b) first mode of vibration; (c) second mode of vibration; (d) third mode of vibration. The model is implemented in Ansys®260

Figure 5. 33. User-driven small-area processing outputs for the *Sanctuary of Vicoforte*: (a) reflectivity map of the small area; (b) temporal coherence of the candidate PS, with empirical distribution; (c) mean LOS velocity, with empirical distribution; (d) cumulative LOS displacement over the observation period, with empirical distribution. The wide spread of the velocity and displacement distributions visible in panels (c) and (d), before coherence-based filtering, motivates the threshold selection procedure of Section 5.4.5.....262

Figure 5. 34. Estimated SBL predictive uncertainty as a function of the minimum temporal coherence threshold applied to the training subset for the *Vicoforte* basement PSs. The uncertainty is computed as the temporal mean of the predictive standard deviation σ normalised by its maximum value, and decreases monotonically with the threshold up to a breakpoint beyond which further exclusion of low-coherence points no longer reduces the predictive uncertainty. The vertical red line indicates the application-specific threshold of 0.56 derived from the breakpoint, which is adopted as the operational coherence threshold for the calibration of the SBL surrogate model in Section 5.4.6.....263

Figure 5. 35. Unfiltered LOS displacement time series for all candidate PS in the small area centred on the *Sanctuary of Vicoforte*. The wide spread of cumulative displacements, reaching approximately ± 400 mm by the end of the observation period, reflects the inclusion of low-coherence points affected by phase unwrapping errors, atmospheric residuals, and dielectric variability of the surrounding terrain, and motivates the coherence-based selection procedure developed in Section 5.4.5.264

Figure 5. 36. Spatial distribution and representative time series of the user-driven small-area PSs at the *Sanctuary of Vicoforte*: (a) two-dimensional Google Earth view of the spatial distribution; (b) three-dimensional Google Earth view; (c) example amplitude time series (upper panel) and LOS displacement time series (lower panel) for a representative basement PS, illustrating the high temporal coherence achieved by the user-driven processing and the seasonal thermoelastic component captured by the parametric inversion. In the upper panel of (c), black dots represent individual amplitude measurements, while red circles indicate the

mean value. In the lower panel of (c), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend. 266

Figure 5. 37. Surface soil moisture and surface temperature recorded by the in-situ thermo-hygrometer at the *Sanctuary of Vicoforte*, resampled to the temporal support of the satellite acquisitions (approximately 23 days). The seasonal cycles and the inverse correlation between temperature and moisture are preserved by the resampling.270

Figure 5. 38. LoS displacement time series predicted by the calibrated SBL model for the five retained basement PSs at the *Sanctuary of Vicoforte*: (a) point 1; (b) point 2; (c) point 3; (d) point 4; (e) point 5. The blue dots indicate the training subset, the magenta dots the validation subset, the cyan diamonds the testing subset; the solid red line is the SBL prediction; the dashed grey lines are the SBL uncertainty bands; the solid black line interpolates the satellite observations. ...275

Figure 5. 39. Average error metrics of the SBL displacement estimates at the *Sanctuary of Vicoforte*: Z-score (top), Normalised Root Mean Square Error (middle), and Normalised Mean Absolute Error (bottom), reported separately for the training, validation, testing, and total prediction subsets.276

Figure 5. 40. Stochastic oscillator time series for the five retained PSs at the *Sanctuary of Vicoforte*, computed from the Z-score of the SBL model residuals according to Equation 5. 11. Values close to unity correspond to a transient alert (Level 1, isolated exceedance), while persistent values close to unity correspond to a sustained alert (Level 2, prolonged exceedance), as defined in Section 5.4.6. The oscillator remains predominantly below 0.5 throughout the observation period, with occasional peaks during the warm summer months consistent with thermo-hydrological forcing not fully captured by the calibrated model.278

Figure 5. 41. Oversampled LoS displacement time series predicted by the calibrated SBL model for the five retained basement PSs at the *Sanctuary of Vicoforte*: (a) point 1; (b) point 2; (c) point 3; (d) point 4; (e) point 5. The prediction is computed at the one-hour sampling of the in-situ environmental inputs and extends beyond the last available satellite observation, with the SBL uncertainty bands in dashed grey indicating the time-varying confidence interval.279

Figure 5. 42. Integrated case-study workflow at the *Sanctuary of Vicoforte*: (a) PS-InSAR data processing chain, from data preparation to small-area processing and PS export; (b) SBL surrogate model and monitoring layer, from coherence threshold derivation to stochastic oscillator alert levels. Solid arrows indicate the

sequence of tasks within each macro-process; dashed arrows indicate transitions between the two macro-processes.....281

Figure 5. 43. Overview of the January 2025 Los Angeles fire complex, showing the perimeters of the Palisades Fire and of the concurrent Eaton, Hurst, Lidia, and other fires, with the regional context of the burn area in relation to the Greater Los Angeles area. Adapted from (NASA Scientific Visualization Studio, 2025).....304

Figure 5. 44. Pre-fire and post-fire airborne LiDAR point clouds the Pacific Palisades study area: (a) pre-fire point cloud, with the buildings geometry intact; (b) post-fire point cloud, with the substantial negative elevation change consistent with structural and vegetation removal. Processing and visualisation by the Cultural Heritage Engineering Initiative and ALERTCalifornia team at the Qualcomm Institute, University of California San Diego.309

Figure 5. 45. Per-parcel damage inventory for the Pacific Palisades community after the Palisades Fire of January 2025, showing the spatial distribution of buildings classified as destroyed and damaged by field inspection. Adapted from (Los Angeles County GIS, 2025).310

Figure 5. 46. Amplitude-temporal baseline plots and LOS displacement time series for representative scatterers in the Pacific Palisades study area: Fire-induced Temporary Scatterer on a destroyed building in ascending (panel a) and descending (panel b) geometry; Post-event Temporary Scatterer appearing in the debris field after the fire in ascending (panel c) and descending (panel d) geometry. In each amplitude-temporal baseline plot, black dots represent individual amplitude measurements, while red circles indicate the mean value. In each displacement time series, red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.....313

List of tables

Table 1. 1 Mapping of the six research questions to the specific objectives, to the chapters in which each question is developed methodologically, and to the case studies in which it is tested empirically.	17
Table 2. 1. Summary of the main limitations affecting satellite InSAR measurements, with their effect on the observation, the typical mitigation strategies, and their specific implications for building-scale SHM.	52
Table 4. 1. Synthesis of the four-level diagnostic framework introduced in Section 4.6, reporting for each level the information content delivered to the engineering user, the data and processing requirements, the conditions for reliability that must be satisfied, and an indicative application scenario. The table is the definitional reference for the framework and is recalled in the cross-case-study assessment of Section 5.5.	181
Table 5. 1. Summary of the two satellite datasets used in the case study of the <i>Santissima Trinità</i>	197
Table 5. 2. Quality attributes of the four selected MPs on the <i>Santissima Trinità</i> (P20 and P24 from EGMS Calibrated; P1803 and P1765 from user-driven Sentinel-1).	200
Table 5. 3. Comparative summary of the three InSAR datasets applied to the <i>Himera Viaduct</i>	224
Table 5. 4. Cross-case summary of the pre-collapse and peri-collapse InSAR evidence for the two infrastructure case studies of Section 5.3.	238
Table 5. 5. Numerical comparison of the EGMS Calibrated and Ortho measurement points on the <i>Sanctuary of Vicoforte</i>	251
Table 5. 6. Summary of the four satellite datasets used in the <i>Vicoforte</i> case study.	252
Table 5. 7. Quality characterisation of the EGMS Calibrated measurement points available on the <i>Sanctuary of Vicoforte</i> and its immediate surroundings, organised by orbital geometry and observation period. Three orbital tracks intersect the area, namely one descending and two ascending (Asc.1 and Asc.2), and two temporal releases of the product are considered, covering the periods 2018-2022	

and 2019-2023 respectively. For each combination, the table reports the total number of available MPs, the percentage of MPs whose temporal coherence exceeds 0.6, and the percentage classified as type 0 (high-quality permanent scatterers) according to the EGMS quality flag scheme. The full set of MPs satisfies the type 0 condition, and over 90 percent of the MPs satisfies the 0.6 coherence threshold across all six combinations, indicating that the EGMS Calibrated dataset on the Sanctuary is dominated by stable, persistent scatterers throughout the available observation window.254

Table 5. 8. Satellite soil moisture products relevant to the integration of environmental variables with InSAR displacement analysis for SHM. For each product, the table reports the source mission, the operating agency, the sensing principle (passive microwave, active C-band, radiometric), the native spatial resolution, the effective temporal resolution, the characterised soil depth (surface or root-zone), the access mode, and a concise indication of the typical SHM application regime. The table supports the data selection argument developed in Section 5.4.4, in which the Copernicus Global Land SWI is adopted as the environmental proxy for the *Vicoforte* case study.268

Table 5. 9. Satellite LST products relevant to the integration of thermal environmental variables with InSAR displacement analysis for SHM. For each product, the table reports the source mission, the operating agency, the thermal sensor and spectral domain, the native spatial resolution, the effective temporal resolution, the access mode, and a concise indication of the typical SHM application regime. The table complements the Soil Moisture products tabulation and supports the environmental variable selection argument developed in Section 5.4.4, in which the Sentinel-3 SLSTR LST is adopted as the thermal proxy for the *Vicoforte* case study.268

Table 5. 10. Estimated mean weight values of the basis functions of Equation 5. 10 for the five retained PSs at the *Sanctuary of Vicoforte*.276

Table 5. 11. Mapping between the operational stages of the satellite-assisted SHM framework (Section 5.5.1) and the diagnostic levels of the analytical framework (Section 4.6). For each stage, the table indicates the typical operational action, the data source, the diagnostic levels accessed, and the case studies that exemplify the stage.285

Table 5. 12. Comparative treatment of thermal and environmental effects across the three case studies of Chapter 5. The table maps the progression from passive recognition to surrogate modelling onto the diagnostic levels of Section 4.6.295

Table 5. 13. Synthesis of the diagnostic levels achieved by the three case studies of Chapter 5 against the four-level analytical framework of Section 4.6. For each level, the table indicates the case studies that achieved it, the enabling conditions, and the empirical limiting factors observed in the case study analysis.301

Table 5. 14. Summary of the satellite radar and airborne LiDAR datasets used in the Palisades Fire post-event assessment work of Section 5.6.307

Table 6. 1. Synthesis of the four methodological contributions developed in this thesis. For each contribution, the table reports the section in which it is introduced, the case study or case studies that test it empirically, the diagnostic level of the four-level framework to which it primarily relates, and the level of confidence with which the contribution is offered to the discipline.318

Nomenclature and acronyms

ALOS	Advanced Land Observing Satellite
ALS	Airborne Laser Scanning
AoI	Area of Interest
APS	Atmospheric Phase Screen
ASF	Alaska Satellite Facility
ASI	Italian Space Agency
ASI	Amplitude Stability Index
CHEI	Cultural Heritage Engineering Initiative
CLMS	Copernicus Land Monitoring Service
CNR	Consiglio Nazionale delle Ricerche
COSMO-SkyMed	Constellation of Small Satellites for Mediterranean basin Observation
CSG	COSMO-SkyMed Second Generation
CSK	COSMO-SkyMed First Generation
D_A	Amplitude Dispersion Index
DEM	Digital Elevation Model
DInSAR	Differential Interferometry Synthetic Aperture Radar

DISEG	Department of Structural, Geotechnical and Building Engineering
DPC	Dipartimento della Protezione Civile
DS	Distributed Scatterer
DSM	Digital Surface Model
DTM	Digital Terrain Model
EED Lab	Earthquake Engineering and Dynamics Laboratory
EGMS	European Ground Motion Service
EO	Earth Observation
ERS	European Remote Sensing satellite
ESA	European Space Agency
FE	Finite Element
FTS	Fire-induced Temporary Scatterer
GB-InSAR	Ground-Based Interferometric Radar
GCP	Ground Control Point
GNSS	Global Navigation Satellite System
GRD	Ground Range Detected
HH	Horizontal-Horizontal
HV	Horizontal-Vertical
ICP	Integrated Circuit Piezoelectric

InSAR	Interferometric Synthetic Aperture Radar
IREA	Istituto per il Rilevamento Elettromagnetico dell'Ambiente
IW	Interferometric Wide Swath
LARIAC	Los Angeles Regional Imagery Acquisition Consortium
LEO	Low Earth Orbit
LiDAR	Light Detection and Ranging
LOS	Line of Sight
LST	Land Surface Temperature
MP	Measurement Point
MST	Minimum Spanning Tree
MT-InSAR	Multi-Temporal InSAR
nDSM	normalised Digital Surface Model
NMAE	Normalised Mean Absolute Error
NRMSE	Normalised Root Mean Square Error
PolIMI	Politecnico di Milano
PolITO	Politecnico di Torino
PS	Persistent Scatterer
PSI	Persistent Scatterer Interferometry

PS-InSAR	Persistent Scatterer Interferometric Synthetic Aperture Radar
PSP	Persistent Scatterer Pair
PTS	post-event Temporary Scatterer
QPS	Quasi-PS
ReLUIs	Rete dei Laboratori Universitari di Ingegneria Sismica e Strutturale
RESORB	Restituted orbits
RMSE	Root Mean Square Error
RQ	Research Question
RVM	Relevance Vector Machine
SA	Small Area
SAR	Synthetic Aperture Radar
SBAS	Small BAseline Subset
SBL	Sparse Bayesian Learning
SHM	Structural Health Monitoring
SLC	Single Look Complex
SO	Stochastic Oscillator
SO	Specific Objectives
SRTM	Shuttle Radar Topography Mission

SSO	Sun-Synchronous Orbit
StaMPS	Stanford Method for PS
SVD	Singular Value Decomposition
SWI	Soil Water Index
TLS	Terrestrial Laser Scanner
TOPS	Terrain Observation by Progressive Scans
TS	Temporary Scatterer
UAV	Unmanned Aerial Vehicle
UNESCO	United Nations Educational, Scientific and Cultural Organization
VH	Vertical-Horizontal
VV	Vertical-Vertical
WUI	Wildland-Urban Interface

Publications

Tarantini, R., McAvoy, S., Ferro, G. A., & Kuester, F. (2026, manuscript in preparation). Building-scale wildfire damage assessment by fusing PS-InSAR time series and bi-temporal airborne LiDAR: The 2025 Palisades fire (Los Angeles, California).

Tarantini, R., Miraglia, G., Coccimiglio, S., Ceravolo, R., & Ferro, G. A. (2026). Soil displacement estimation from integrated sensing technologies in data-driven models biased by temporal coherence of PS-InSAR. *Land*, 15(2), Article 296. Special Issue “Ground Deformation Monitoring via Remote Sensing Time Series Data”. <https://doi.org/10.3390/land15020296>

Coccimiglio, S., Tarantini, R., Miraglia, G., Matteini, I., Scussolini, L., Ceravolo, R., & Ferro, G. A. (2026). Analysis of multi-source satellite-derived displacement data for structural monitoring of masonry heritage buildings. *Procedia Structural Integrity*, 78, 1032-1039. <https://doi.org/10.1016/j.prostr.2025.12.132>

Tarantini, R., Coccimiglio, S., Miraglia, G., Matteini, I., Grassi, M., Ceravolo, R., & Ferro, G. A. (2026). Pre- and post-collapse analysis of infrastructures using COSMO-SkyMed satellite data. In *Procedia Structural Integrity*, 84, 401–408. - III FABRE Conference: Existing bridges, viaducts and tunnels: research, innovation and applications. <https://doi.org/10.1016/j.prostr.2026.06.052>

Tarantini, R., Mairone, M., Ferro, G. A., Corrado, M., Givonetti, M., Gallina, G., & Masera, D. (2025). Seismic vulnerability assessment of RC bridge piers strengthened with GFRP rebars: A case study. In *Engineering materials, structures, systems and methods for a more sustainable future (pp. 1487-1492)*. CRC Press / Taylor & Francis. <https://doi.org/10.1201/9781003488644>

Tarantini, R., Coccimiglio, S., Miraglia, G., Grassi, M., Ceravolo, R., & Ferro, G. A. (2025, in press). EGMS and user-driven Sentinel-1 InSAR data for pre- and post-collapse analysis of the Himera viaduct. In *ARTISTE 2025 - First International Conference on Artificial Intelligence in Structural Engineering*.

Tarantini, R., Coccimiglio, S., Miraglia, G., Ceravolo, R., & Ferro, G. A. (2025). Leveraging satellite data for displacements optimization in structural health monitoring. *16th World Congress on Structural and Multidisciplinary Optimization (WCSMO 2025)*, 18-23 May 2025, Kobe Convention Center, Kobe, Japan. *International Society for Structural and Multidisciplinary Optimization (ISSMO)*.

Mairone, M., Tarantini, R., Ferro, G. A., Masera, D., Corrado, M., & Givonetti, M. (2025). Seismic behaviour at ultimate limit state of reinforced concrete structures retrofitted with glass fibre reinforced polymer rebars. *In Proceedings of the fib International Symposium on Concrete Structures: Extend lifetime, limit impacts (pp. 3697-3706)*. fib. ISSN 2617-4820.

Cardoni, A., Tarantini, R., Calidori, A., & Cimellaro, G. P. (2024). Seismic vulnerability and resilience of large-scale telecommunication networks. *In Proceedings of the 18th World Conference on Earthquake Engineering (18WCEE)*, Milan, Italy, 30 June - 5 July 2024. ISSN 3006-5933.

Cucuzza, R., Manuello Bertetto, A., Domaneschi, M., Tarantini, R., Cardoni, A., et al. (2024). Numerical and experimental dynamic analysis of 3D-printed pierced vault. *In Proceedings of the 2nd Italian Workshop on Shell and Spatial Structures (IWSS)*, Turin, Italy, 26-28 June 2023 (pp. 269-278). Springer. ISSN 2366-2557. ISBN 978-3-031-44327-5. https://doi.org/10.1007/978-3-031-44328-2_28

Tarantini, R., Cardoni, A., Cimellaro, G. P., Domaneschi, M., Pribaz, E., & Rupolo, G. (2023). Dynamic identification of a historical masonry building: The case study of Palazzo Rosciano. *In Proceedings of the 9th ECCOMAS Thematic Conference on Computational Methods in Structural Dynamics and Earthquake Engineering (COMPDYN 2023)*, Athens, Greece, 12-14 June 2023 (pp. 3700-3707). ISBN 978-618-5827-01-4. <https://doi.org/10.7712/120123.10673.22019>

Cardoni, A., Tarantini, R., & Cimellaro, G. P. (2023). Indoor tracking using unmanned aerial vehicles. *In Proceedings of COMPDYN 2023*, Athens, Greece, 12-14 June 2023 (pp. 4632-4640). ISBN 978-618-5827-01-4. <https://doi.org/10.7712/120123.10746.22021>

Tarantini, R., Cardoni, A., Marasco, S., Merlin, J., Pribaz, E., Rupolo, G., & Domaneschi, M. (2023). Dynamic characterisation and seismic vulnerability assessment of existing masonry port structures. *In Proceedings of the 17th World*

Conference on Seismic Isolation (WCSI 2022) (pp. 1032-1042). Springer. ISSN 2366-2557. https://doi.org/10.1007/978-3-031-21187-4_90

Cardoni, A., Cimellaro, G. P., Domaneschi, M., Merlin, J., Tarantini, R., Pribaz, E., & Rupolo, G. (2022). Seismic vulnerability assessment of existing port buildings: Parameters identification toward numerical simulations. *In Proceedings of the 11th International Conference on Structural Health Monitoring of Intelligent Infrastructure (SHMII 2022), Montreal, Canada, 8-12 August 2022 (pp. 470-473). ISSN 2564-3738.*

Chapter 1

Research field and motivation

1.1. Introduction

Structural Health Monitoring (SHM) has become a central component of infrastructure management and risk mitigation for civil structures whose safety and serviceability must be assessed over decades of operation (Balageas et al., 2010; Brownjohn, 2007; C. Farrar & Worden, 2013). In civil engineering, SHM refers to the process of observing structures through repeated or continuous measurements, with the aim of identifying changes in condition, tracking performance, and informing maintenance decisions. Within the SHM framework, the term damage denotes any departure from an established baseline state, whether through alteration of material properties, geometric configuration, member connectivity, or boundary conditions, that compromises current or future structural performance (C. Farrar & Worden, 2013).

Much of the existing built stock is now subjected to environmental loads, operational demands, and hazard scenarios that exceed what its designers originally envisaged. Its average age has become such that a substantial fraction of bridges and buildings has already reached or exceeded the service life assumed at the design stage (Biondini & Frangopol, 2016). Buildings, bridges, heritage monuments, and transport infrastructure routinely operate well beyond the service conditions assumed at the time of their construction. SHM contributes to this challenge not only through damage detection, but also by supporting the transition from time-based maintenance to condition-based strategies (Sohn, 2006).

The widespread implementation of SHM on full-scale civil structures, however, remains limited by practical constraints. Permanent monitoring systems require dedicated instrumentation, communication infrastructure, regular maintenance, and specialised personnel. Such instrumentation yields reliable data but entails considerable costs and logistical effort, qualities that make conventional sensor networks difficult to deploy at the scale of entire building portfolios (Brownjohn, 2007; C. Farrar & Worden, 2013). The difficulty is compounded in the case of historical and monumental structures, where physical access is often restricted, conservation protocols limit the placement of sensors, and the structural singularity of each edifice resists standardised monitoring approaches (Sohn, 2006). A second obstacle, epistemic in nature, lies in the separation of genuine structural change from the reversible fluctuations induced by temperature, humidity, soil conditions, and ordinary operational loading. Isolating damage-related variations from environmental and operational variability is one of the central problems of the discipline and remains an active area of methodological research (Sohn et al., 2002).

Within this context, satellite remote sensing has emerged as a complementary source of information for SHM. Among the available techniques, Synthetic Aperture Radar (SAR) interferometry is particularly relevant because it enables the retrieval of surface displacement from repeated radar acquisitions, independently of daylight and weather conditions (Bamler & Hartl, 1998; Massonnet & Feigl, 1998). Multi-temporal InSAR (MT-InSAR) techniques such as Persistent Scatterer Interferometry (PSI) reconstruct displacement time series at coherent radar targets with mean velocity precision of approximately 1 mm per year and per-epoch displacement standard deviations in the range of 2 to 4 mm, depending on the stack size, the wavelength, and the coherence of the target (Crosetto et al., 2016; Ferretti et al., 2001). These methods have progressively moved from geophysical applications toward engineering-oriented uses, including the monitoring of bridges, historical centres, and monumental structures (Calò et al., 2023; Farneti et al., 2023a; Milillo et al., 2018).

Extracting structural meaning from satellite-derived displacements remains, however, a non-trivial task. Interferometric Synthetic Aperture Radar (InSAR) measurements are inherently sparse in space and projected along the sensor Line of Sight (LOS). The measurement is referenced to a point assumed to be stable. The observable remains subject to decorrelation, atmospheric phase contamination, and the assumptions embedded in each processing chain (Crosetto et al., 2016; Prota et al., 2025). A recorded displacement, taken in isolation, does not constitute a structural diagnosis. Thermal cycles, ground-induced movement, transient

scatterers, and loss of coherence can each corrupt or mimic a genuine deformation signal. Translating satellite observations into reliable SHM information therefore demands a methodology that bridges measurement physics, processing strategy, and target identification. The methodology must further integrate uncertainty quantification with engineering judgement, within a single internally consistent framework.

The present thesis examines how interferometric satellite data can progress from raw LOS displacement records to decision-relevant input for the monitoring of existing buildings and infrastructure. Its guiding argument is that InSAR delivers its greatest diagnostic value when it is incorporated into a workflow that accounts explicitly for acquisition geometry, temporal coherence behaviour, measurement uncertainty, the structural plausibility of the observed patterns, and the complementary role of ground-based data. The argument rests on two premises that the remainder of the thesis develops systematically. The first is that the choice of processing strategy is not neutral with respect to the structural conclusions that can be drawn from the data: the granularity at which parameters are controlled, the stringency of the coherence threshold, and the spatial scale of the analysis all interact with the diagnostic question being asked, and the transition from standardised products to user-driven processing corresponds to a transition from population-level screening to structure-specific diagnosis. The second is that temporal coherence should be treated as a modelling variable in its own right, rather than as a binary filter applied before interpretation: the distribution of coherence values across measurement points carries information about the physical state of the target, and the systematic use of this information opens diagnostic possibilities that fixed-threshold workflows foreclose. Taken together, these premises delimit the methodological space within which the thesis operates and motivate the framework elaborated in the chapters that follow.

1.2. Structural Health Monitoring of full-scale structures

SHM of full-scale structures deals with systems exposed to actual operational loads, environmental variability, imperfect boundary conditions, and ageing processes. Unlike laboratory testing, full-scale monitoring must account for the fact that the measured response reflects not only the mechanical state of the structure but also temperature, humidity, soil conditions, use patterns, and other external actions. Changes in measured response do not automatically imply damage, and one of the persistent challenges of the discipline is the distinction between normal variability

and meaningful structural anomalies (Doebbling et al., 1996; C. R. Farrar & Worden, 2007; Sohn, 2006).

From an operational perspective, full-scale SHM may involve displacements, strains, rotations, accelerations, modal properties, crack openings, environmental parameters, and the long-term variation of all of these. No single sensing technology captures the full behaviour of a complex civil structure. Instrumented systems are often spatially sparse and intrusive, and remain expensive to maintain over long periods (Brownjohn, 2007; Lynch, 2006). For these reasons, the field has progressively evolved toward integrated frameworks that combine heterogeneous measurements, uncertainty-aware interpretation, and decision support for asset management (C. R. Farrar et al., 2025; Worden et al., 2015). The following sections discuss four aspects of this broader perspective: the motivations for SHM, the practical constraints that limit its adoption, the distinction among monitoring approaches, and the role of satellite observations.

1.2.1. Motivation for SHM

The primary motivation for SHM of full-scale civil structures lies in the need to preserve safety, serviceability, and durability throughout the life cycle of assets often expected to operate far beyond their original design horizon. Three converging factors define the challenges specific to civil infrastructure (Brownjohn, 2007). The first is the physical extent of the assets involved. The second is the heterogeneity of their structural systems. The third is the decades-long horizon over which they must remain serviceable. Much of the existing built stock now operates under environmental actions and hazard exposures more severe than those anticipated at the design stage. As these assets age, their response is collectively modified by progressive material deterioration, by shifts in boundary restraint, soil-structure coupling, altered loading regimes, and the cumulative interaction of all of these. The combination affects buildings, bridges, and monumental structures alike. For heritage assets, the imperative to monitor is strengthened by conservation constraints and by restricted physical access. The need to keep any intervention as unobtrusive as possible reinforces this imperative further. Monitoring, under these circumstances, functions both as a diagnostic instrument and as a means of accumulating durable knowledge about the structural system and its long-term trajectory (Brownjohn, 2007; Coccimiglio, 2025; C. R. Farrar & Worden, 2007).

Conventional maintenance philosophies that rely on periodic visual inspections and post-damage corrective action fall short when confronted with this degree of

complexity. The SHM community has repeatedly stressed that the principal benefit of continuous observation is a fundamental reorientation of engineering practice. The reorientation moves practice away from reactive repair cycles and toward decisions whose value can be quantified within a pre-posterior Bayesian framework (Thöns, 2018). Sustained or periodically repeated measurement enables three operational gains. The practitioner can follow the real evolution of structural behaviour. Departures from expected trends are detected at an early stage. Intervention choices are anchored to recorded evidence rather than predetermined schedules.

A third motivation, less frequently acknowledged yet no less significant, is epistemic in nature. Full-scale civil structures are characterised by uncertainty in geometry, material properties, boundary conditions, and actual loading history. Changes in measured response may derive from damage or from non-damage-related variability, and the central problem of SHM is to distinguish between the two (Doebeling et al., 1996; Sohn et al., 2002). Monitoring reduces this uncertainty progressively over time. At the same time, high-quality local monitoring systems cannot realistically be deployed on every existing structure (Di Mucci et al., 2024). This is one of the reasons why the field is moving toward multi-source frameworks that combine direct sensing, data analytics, and remote observations to achieve informative and scalable monitoring strategies.

1.2.2. Challenges and constraints: scale, access, cost, and continuity

The implementation of SHM on full-scale civil structures is constrained by four recurrent issues. The first is scale. Structures are spatially extended, mechanically heterogeneous, and embedded in complex environmental contexts. Measurements collected at a limited number of instrumented locations may not capture distributed phenomena such as differential settlements or spatially variable environmental effects (Di Mucci et al., 2024). The problem becomes even more severe when the goal is to screen entire stocks of buildings or infrastructures, where the same instrumental density cannot be replicated across all assets.

The second issue is access. Many structures are difficult to instrument because of height, geometry, heritage constraints, traffic exposure, or safety requirements during installation. Sensor placement is often dictated by practical compromise rather than by optimal measurement design, which means that some portions of the structure may remain only partially observable.

The third constraint is cost. A permanent monitoring system requires not only sensors but also acquisition hardware, power supply, communication infrastructure, data storage, calibration, and expert interpretation. Recent reviews confirm that deployment cost and long-term maintenance remain persistent obstacles to widespread adoption, particularly for small- and medium-scale assets (Elsisi et al., 2025).

The fourth constraint is continuity. Long-term systems are exposed to sensor drift, data gaps, communication failures, and evolving operational conditions. The longer the observation period, the more difficult it becomes to distinguish true structural changes from environmentally induced fluctuations. Continuous acquisition is not sufficient unless accompanied by rigorous procedures for contextualisation and validation (Thöns, 2018).

These four constraints interact. Restricted physical access drives up the cost of both installing and servicing sensor equipment. As costs rise, fewer measurement points can be justified, and spatial coverage diminishes accordingly. Where temporal continuity cannot be guaranteed, the data lose interpretive power and the long-term return on the monitoring investment erodes. The real difficulty therefore concerns the design of observation strategies that extract meaningful structural information within the practical constraints imposed by heritage and existing assets. The acquisition of additional data, in itself, has become a less binding limitation.

1.2.3. Types of monitoring systems: contact, remote, and hybrid approaches

Contact-based monitoring relies on transducers mounted directly on the structural member or in its immediate vicinity. Commonly deployed devices include accelerometers, strain gauges, displacement transducers, inclinometers, fibre-optic sensing cables, Global Navigation Satellite System (GNSS) receivers, crack meters, environmental probes, and a variety of specialised devices for site-specific applications (Brownjohn, 2007; C. Farrar & Worden, 2013). Among these, Distributed Fibre Optic Sensing (DFOS) has emerged as a particularly informative technology for civil structures, because a single cable can deliver strain or temperature measurements at metre or sub-metre spatial resolution along its entire length, and has been successfully deployed on bridges, tunnels, and heritage masonry (Bado & Casas, 2021; Barrias et al., 2016). The principal advantage of contact instruments, considered as a family, is their capacity to deliver spatially localised measurements of well-defined physical quantities, often sampled at high

temporal resolution. Contact sensing remains the backbone of most long-term monitoring systems for bridges and other major civil structures, including towers and monumental buildings. At the same time, these systems require direct access, installation works, maintenance, calibration, and data-management infrastructure. For existing and heritage structures, these requirements may prove prohibitive or incompatible with conservation constraints.

Remote monitoring systems acquire information without physical contact with the structure and encompass a heterogeneous set of technologies that share the property of decoupling observation from instrumentation. Terrestrial Laser Scanning (TLS) provides dense three-dimensional geometric reconstructions of built surfaces. The technique supports geometric survey, crack mapping, and bi-temporal change detection at the millimetre to centimetre scale (Alba et al., 2011; Riveiro et al., 2013). Airborne Laser Scanning (ALS), typically deployed from fixed-wing aircraft or helicopters, extends the same principle to the landscape scale. ALS captures building footprints, roof geometries, and terrain elevation over areas of tens of square kilometres. Point densities now routinely exceed tens of points per square metre. Vision-based monitoring relies on digital image correlation and photogrammetric reconstruction. The technique can resolve displacements, vibrations, and surface damage without any sensor attached to the structure. Its performance remains sensitive to lighting, image scale, and environmental variability (Dong & Catbas, 2021; Y. Yang et al., 2024). Unmanned Aerial Vehicles (UAVs) have progressively entered the SHM toolbox as flexible carriers for high-resolution optical imagery, photogrammetric block adjustment, thermography, and lightweight Light Detection and Ranging (LiDAR) payloads. The platform enables close-range inspection of façades, decks, and roofs that are difficult or hazardous to access by other means (Feroz & Abu Dabous, 2021). Ground-Based Interferometric Radar (GB-InSAR), typically operated in Ku or Ka band, provides dense displacement measurements of a radar-illuminated scene at sub-millimetre precision and sampling rates as high as hundreds of Hertz. Its spatial coverage is inherently bounded by the radar footprint. The technique offers the closest terrestrial analogue to satellite PS-InSAR and has been widely used for bridges, dams, landslides, and historical buildings (Monserrat et al., 2014; Pieraccini & Miccinesi, 2019). Satellite-based observations extend the scale of remote monitoring further, enabling the simultaneous acquisition of data over multiple assets and the surrounding territory, and constitute the specific focus of this thesis. Across the family of remote techniques, the principal advantages are reduced invasiveness and broader spatial coverage, together with the capacity to reach locations that are impractical to instrument directly. The spatial reach of these

techniques makes them particularly suitable for screening activities. They are also effective for observing phenomena distributed over areas larger than a single structure, such as ground settlements and slope movements at regional scale.

Hybrid monitoring architectures bring together contact and remote measurements within a unified SHM strategy. The underlying logic is one of complementarity: the weaknesses inherent in any single sensing family can be offset by the capabilities of another. Contact transducers supply temporally dense, physically grounded readings at discrete locations, whereas remote techniques extend the observable domain, eliminate the need for direct instrumentation, and furnish broader contextual information about the structural surroundings. Several recent surveys converge on the view that progress in the field depends less on advances within individual sensing technologies than on the capacity to merge heterogeneous data streams. This merging relies on principled fusion schemes, interpretation methods that propagate uncertainty explicitly, and decision-support workflows that translate combined observations into actionable engineering judgements (C. R. Farrar et al., 2025; Mardanshahi et al., 2025). For existing structures, hybrid approaches prove especially attractive because they allow the selective deployment of in-situ sensors where local accuracy is indispensable, while using remote observations to extend coverage, reconstruct historical trends, or screen multiple assets simultaneously.

Within this classification, satellite interferometric observations occupy a specific position. They do not replace direct structural monitoring, but they extend the spatial and temporal scope of observation in ways that are difficult to achieve with conventional instrumentation alone. The role of satellite data within SHM is examined in the following section.

1.2.4. The role of satellite observations in SHM

Satellite observations address several of the constraints discussed in the preceding sections. They provide spatial scalability, because a single dataset can cover multiple assets, urban sectors, and surrounding ground conditions simultaneously (Nettis et al., 2023; Talledo et al., 2022). They are non-invasive, requiring no sensor installation, site access, or interruption of service during acquisition. They offer historical depth, because archived radar acquisitions make it possible to reconstruct displacement histories long before the installation of any dedicated monitoring system (Crosetto et al., 2016). They also capture environmental context, because the observed scene includes not only the structure but also terrain, neighbouring

buildings, and ground deformation patterns that may influence structural behaviour. These characteristics position satellite observations as a viable tool both for screening large building portfolios and for interrogating individual structures in detail; the distinction between the two uses rests primarily on the processing strategy adopted and on the spatial scale of interest.

Satellite-derived measurements, however, do not yield a structural diagnosis on their own. What the sensor records is a single displacement component projected along the satellite LOS, rather than the complete three-dimensional movement of a point on the structure. Vertical and east-west components can be recovered only when ascending and descending acquisitions over the same target are jointly processed. The north-south component remains inaccessible because of the near-polar orbital geometry of current SAR missions. The measurement is defined with respect to a selected reference target. It remains susceptible to atmospheric phase contributions. It can be retrieved only at locations where the radar return preserves adequate coherence throughout the observation period. Temporal sampling is limited to the satellite revisit cycle, and spatial coverage on a single building depends on geometry, target visibility, and resolution. These constraints mean that the engineering value of satellite data depends entirely on how they are processed, filtered, and interpreted. User-driven processing remains essential when the objective is structure-specific analysis. Point quality, coherence thresholds, geocoding, and reference selection can each affect the engineering meaning of the results (Coccimiglio, Miraglia, et al., 2022; Coccimiglio, 2025; Prota et al., 2025).

Satellite interferometric observations can accordingly serve the SHM process along three distinct axes. The first axis is that of screening: satellite-derived velocity and displacement maps enable the identification of assets, façades, or zones whose kinematic behaviour departs from expectations and that may therefore warrant targeted field investigation. The second axis concerns interpretation: spatially distributed displacement records, accumulated over multi-year observation windows, provide evidence from which the analyst can reconstruct the progression of structural or geotechnical phenomena and distinguish persistent trends from transient fluctuations. The third axis involves integration: interferometric products, when combined with ground-based instrumentation, numerical modelling, and informed engineering judgement, strengthen the overall assessment by supplying independent, spatially extensive constraints that no single in-situ technique can replicate on its own. In this thesis, satellite InSAR is framed within this three-level logic and examined as a component of a broader monitoring and interpretation strategy for existing structures.

1.3. Objectives and contributions of the thesis

An engineering-oriented framework for satellite interferometric SHM of existing structures must bridge three distinct gaps. The first gap lies between the rapid expansion of freely available displacement datasets, delivered by Sentinel-1, COSMO-SkyMed, and the European Ground Motion Service (EGMS) (Ferretti et al., 2023; Kotzerke et al., 2022), and the interpretive requirements of structural engineering practice, for which displacement measurements become useful only when embedded in a context of structural knowledge, physical plausibility, and quantified uncertainty. The second gap concerns the level of analytical granularity: the regional and urban scales at which MT-InSAR has been most extensively validated do not transfer automatically to the scale of individual buildings or bridge decks. The third gap is operational and separates standardised data products from the diagnostic judgements that engineers and asset managers are asked to formulate. The overarching objective of this thesis is to develop a methodological framework that addresses these three gaps in an integrated manner, supported by four specific objectives (SO) that define the scope and direction of the research.

The first specific objective is to establish the conditions under which satellite InSAR can provide structurally meaningful information at the building scale. Existing applications have demonstrated the capacity of the technique on individual structures (Cavalagli et al., 2019; Milillo et al., 2018), yet the boundary between scenarios in which building-scale interpretation is feasible and those in which it is not has not been systematically mapped. This objective is addressed through a combination of methodological development, presented in Chapters 2, 3, and 4, and empirical testing across structures of different typology, geometry, and urban context, reported in Chapter 5.

The second specific objective is to formalise a tiered monitoring workflow that connects standardised EGMS products, user-driven InSAR processing, and integrated interpretation within a single decision framework. The workflow specifies what each analytical stage can and cannot resolve, what signals in the data justify escalation from one stage to the next, and how the satellite evidence relates to the diagnostic levels that a structural engineer requires. This framework is developed in Chapter 4 and in the synthesis presented in Section 5.5, where the tiered workflow is articulated in relation to the case studies.

The third specific objective is to investigate the role of temporal coherence, thermal and environmental effects, and Temporary Scatterer (TS) diagnostics as

sources of structural information, rather than as quality filters or nuisance parameters. The underlying proposition is threefold: coherence carries diagnostic content when treated as a continuous variable; satellite-derived environmental data can substitute for in-situ sensors in data-driven displacement models; and the lifecycle boundaries of TS encode the chronology of structural events. These propositions are developed methodologically in Chapter 4 and tested empirically through the case studies presented in Chapter 5.

The fourth specific objective is to extend the satellite monitoring framework into the domain of multi-modal sensor fusion for post-event structural assessment. The fusion of InSAR temporal information with airborne LiDAR geometric change data at the per-building scale has not been attempted in the published literature for wildland-urban interface fire events. The methodological framework for this fusion is developed in Section 4.5.3 and tested in Section 5.6 on the damage inventory produced by the 2025 Palisades Fire in Los Angeles.

On the basis of these four objectives, the main contributions of this thesis can be stated concisely. The work develops an engineering methodology for the structural reading of satellite InSAR measurements on buildings and infrastructure, taking as its foundation the diagnostic principles codified in the ReLUIIS guidelines (Prota et al., 2025) and extending them in four specific directions: quality control informed by coherence behaviour, harmonisation of displacement records derived from independent data sources, identification and treatment of TS, and fusion of interferometric products with complementary monitoring techniques. A tiered monitoring workflow is proposed and tested against real cases, progressing from EGMS-based portfolio screening through user-driven interferometric processing to integrated structural interpretation, with each tier corresponding to one of four progressive diagnostic levels.

All multi-temporal interferometric processing reported in the three completed case studies and in the ongoing Palisades application was conducted by the author through a fully user-driven workflow operating directly on SLC data, with explicit control over master image selection, interferometric network construction, reference point definition, amplitude and coherence thresholds, atmospheric phase screen estimation, and geocoding parameters. The processed datasets comprise the 169-scene Sentinel-1 descending stack over the *Santissima Trinità* in Turin, the dual-orbit Sentinel-1 peri-collapse archives over the *Himera Viaduct*, the unified 2011-2022 dual-orbit COSMO-SkyMed Stripmap archive over the *Albiano-Magra Bridge*, the 62-scene COSMO-SkyMed Second Generation small-area stack over

the *Sanctuary of Vicoforte*, and the dual-orbit Sentinel-1 archive comprising 70 ascending and 73 descending acquisitions over the Pacific Palisades study area. Standardised EGMS products are employed throughout the thesis as the primary instrument for Level 1 portfolio-scale screening and for cross-validation of the user-driven results, in accordance with the tiered framework developed in Chapter 4. The user-driven processing of raw SLC data is the primary instrument for Level 2 displacement characterisation. The same processing supports the higher diagnostic levels, where structure-specific control over coherence thresholds, reference frame, and processing parameters becomes a precondition for reliable interpretation. The two processing strategies are complementary within a single decision framework. Their joint use across the case studies allows the comparison of standardised and user-driven outputs on the same structures. The comparison provides one of the empirical contributions of the thesis. Four case studies are presented, spanning masonry heritage, infrastructure collapse, environmentally driven soil-structure interaction, and post-fire damage assessment. Together they delineate the conditions under which each diagnostic level is achievable. The cases also identify the factors that currently limit further progression.

The broader implications of these contributions for engineering practice are synthesised in Chapter 6, which closes the thesis with a discussion of limitations, transferability, and prospective developments. The research questions that operationalise the four objectives are formulated in the next section. The six research questions (RQ) that follow are not in one-to-one correspondence with the four objectives: RQ1 and RQ2 operationalise the first specific objective; RQ3 and RQ4 articulate the third; RQ5 addresses aspects of the second and third jointly; RQ6 operationalises the fourth. This non-bijective mapping reflects the interdependence of the analytical questions that arise at the interface of satellite remote sensing and structural engineering.

1.4. Research questions and scope

The objectives and contributions outlined in the preceding section respond to a set of open questions that arise at the intersection of satellite remote sensing and structural engineering. Although MT-InSAR has been applied to structures for more than two decades, the literature reveals persistent gaps between the maturity of the interferometric technique and the interpretive needs of the structural engineering community. The research questions formulated below are designed to target these gaps directly, and each is addressed through a combination of methodological development and empirical investigation within the subsequent chapters.

RQ1. Under what conditions can satellite InSAR observations provide structurally meaningful information, and what are the boundaries of their interpretive validity for building-scale SHM?

The capacity of Persistent Scatterer InSAR (PS-InSAR) to resolve millimetre-scale LOS displacements has been demonstrated extensively at the urban and regional scale (Crosetto et al., 2016; Ferretti et al., 2001). When the analysis moves to the scale of individual buildings, three measurement attributes govern the structural value of the analysis. The first is the number of usable measurement points. The second is their spatial arrangement across the structure. The third is the physical surfaces from which they originate. The regional-scale literature has seldom examined these attributes in a systematic manner. Whether a given surface produces a detectable scatterer depends on the relationship between the geometry of the building envelope and the incidence angle of the satellite. The three-dimensional configuration of the surrounding urban fabric further conditions the detectability. The scatterer population that results from this interplay is frequently sparse, irregularly clustered, and tied to architectural elements whose load-bearing function or structural relevance may not be immediately apparent. Several studies have identified these constraints in individual cases, including applications to infrastructure assets (Milillo et al., 2018), heritage masonry (Cavalagli et al., 2019), and tunnelling-induced building damage (Giardina et al., 2019), yet a structured assessment of the conditions under which building-scale interpretation becomes feasible, and of the conditions under which it does not, remains absent from the literature. The present thesis addresses RQ1 primarily through Chapters 2 and 4, where the interpretive framework is developed, and through the case studies of Chapter 5, which test that framework against structures of different typology, geometry, and urban context.

RQ2. How do standardised EGMS products and user-driven InSAR processing compare in terms of point density, measurement quality, and interpretability at the building scale, and under what circumstances is each approach sufficient or insufficient?

The EGMS has made satellite displacement data freely accessible to non-specialist users across European territory, and its potential for infrastructure screening is increasingly recognised in the structural monitoring literature (Ferretti et al., 2023; Kotzerke et al., 2022; Prota et al., 2025). The validation of EGMS products, however, has focused primarily on regional-scale consistency and on agreement with GNSS reference networks, rather than on the fitness of EGMS data

for structural interpretation at the scale of individual buildings or bridge decks. The temporal coherence thresholds applied in the EGMS processing chain differ from those that a structural engineer would impose when selecting points for a building-specific analysis. Point density on a given structure may appear adequate in the EGMS product yet prove insufficient for resolving displacement gradients. The comparison between standardised and user-driven processing has been discussed qualitatively in the literature (Talledo et al., 2022), but a quantitative evaluation across the three axes of density, quality, and structural interpretability, conducted on the same structures with both approaches, has not been published. RQ2 is addressed through Case studies I and II in Chapter 5, where EGMS and user-driven processing results are compared on a heritage building and on two collapsed infrastructure elements, and through the decision workflow formalised in Section 5.5.1.

RQ3. Can temporal coherence be treated not merely as a quality filter for discarding unreliable points, but as a diagnostic variable that carries information about the scattering environment and the structural condition?

Within the conventional PS-InSAR processing chain, temporal coherence acts as the sole gate through which candidate measurement points must pass. A point whose coherence falls below a predefined threshold is excluded from the solution; one that exceeds the threshold is accepted. The threshold value itself is, in most implementations, fixed by convention or selected through manual inspection of the coherence distribution. That the spatial pattern of coherence across a building, or its evolution over the observation period, might itself encode information of structural or diagnostic relevance is a possibility that the literature has so far explored only marginally. Coherence degradation on a structural element may indicate physical change in the scattering surface, such as material loss, restoration works, or collapse, rather than mere measurement noise. Conversely, the dependence of model uncertainty on the coherence of the input data has not been formalised in the context of data-driven displacement modelling. RQ3 is addressed in Chapter 4, where the interpretation of coherence patterns on built assets is developed, and in Case study III (Section 5.4), where a coherence-aware framework for surrogate modelling is proposed and empirically tested using the *Sanctuary of Vicoforte*.

RQ4. To what extent can satellite-derived environmental variables replace or supplement in-situ environmental sensors within data-driven models of thermally and hydrologically influenced structural displacement?

The influence of temperature and soil moisture on the displacement of structures founded on compressible soils has been extensively documented through in-situ monitoring (Ceravolo et al., 2021; Sohn, 2006). In-situ sensors, however, are expensive, spatially sparse, and often installed only after a problem has been identified. Satellite-derived Land Surface Temperature (LST) and Soil Water Index (SVI) are available as free, continental-scale products with temporal densities comparable to those of many in-situ networks. Whether these satellite products can serve as effective proxies for the environmental inputs required by data-driven displacement models, particularly for structures where in-situ environmental data are absent or limited, has not been investigated systematically. The *Vicoforte* case study (Section 5.4) provides an empirical test of this proposition by comparing surrogate models trained on satellite-derived environmental inputs against the in-situ temperature and piezometric records maintained by the permanent monitoring system.

RQ5. What is the diagnostic value of TS in the retrospective analysis of structural failures, and can their lifecycle be used to constrain the chronology and spatial extent of collapse events?

TS, defined as measurement points that meet the phase stability criterion within a limited time window yet lose coherence before the end of the full observation span, have been reported in several retrospective analyses of infrastructure failures (Macchiarulo, Giardina, et al., 2022; Milillo et al., 2019). Despite these documented occurrences, no systematic effort has been made to formalise the TS as a distinct diagnostic category within the PS-InSAR processing and interpretation chain. The sudden emergence of a TS on a structural element, or its disappearance from the scatterer population, may signal an alteration of the reflecting surface attributable to construction activity, partial demolition, a change in surface material, or the onset of structural distress. In the context of collapse analysis, the temporal boundaries of a TS lifecycle can, in principle, constrain the date of a physical event with a precision determined by the satellite revisit time. Whether this information can be systematically extracted, and whether it adds diagnostic value beyond what the displacement time series alone can provide, are open questions. RQ5 is addressed through Case study II (Section 5.3), where TSs are identified on the *Himera Viaduct* and the *Albiano-Magra Bridge* in both pre- and post-collapse configurations, and through the ongoing research presented in Section 5.6, where TS lifecycle analysis is combined with LiDAR-derived building change metrics for post-fire damage classification.

RQ6. Can satellite InSAR and airborne LiDAR be fused at the building scale to produce damage classifications that neither sensor can achieve independently in a post-event assessment scenario?

Multi-sensor fusion for post-event structural assessment has been explored primarily in the earthquake domain and predominantly at the urban block scale (Dong & Catbas, 2021). The fusion of InSAR-derived temporal information with LiDAR-derived geometric information at the scale of individual building footprints represents an approach that has not been tested in the wildland-urban interface fire context. The complementarity between the two sensors is straightforward at a conceptual level: LiDAR captures volumetric change with high spatial fidelity but offers no historical depth, whereas InSAR accumulates a displacement record over months or years but cannot resolve the full three-dimensional geometry of a damaged element. The question that RQ6 poses is twofold: whether a per-building feature vector that draws on both data sources yields better classification performance than either source alone, and under what acquisition and processing conditions the gain proves statistically meaningful. The ongoing research presented in Section 5.6 examines this proposition against the damage inventory produced by the 2025 Palisades Fire in Los Angeles, which serves as a preliminary empirical testbed.

Four deliberate delimitations bound the scope of the present work. On the processing side, all interferometric analyses are carried out within the Persistent Scatterer Interferometry (PSI) formulation of the Politecnico di Milano (POLIMI) school; Small Baseline Subset (SBAS)-based approaches are discussed where they provide useful context but are not applied directly, and alternative multi-temporal formulations such as the Quasi-Persistent Scatterer (QPS) approach are presented in Chapter 3 as methodological extensions of PSI that the present work does not exploit operationally. Regarding structural typology, the investigation concentrates on masonry heritage buildings and on reinforced concrete bridges and viaducts; steel frames, timber construction, and geotechnical assets such as embankments or retaining walls lie outside the perimeter of the study. With respect to numerical modelling, the thesis neither develops nor validates independent finite element representations of the monitored structures; finite element analysis enters the discussion only as a complementary interpretive resource in Case study III. Finally, the work does not address the operational deployment of real-time or near-real-time satellite monitoring architectures; every analysis presented is retrospective in character and relies on archived SAR acquisitions.

The six research questions formulated above, the four specific objectives defined in Section 1.3, the chapters in which each question is developed methodologically, and the case studies in which each question is tested empirically are mapped onto a single reference grid in Table 1. 1, which also indicates the primary literature gap that motivates each question. The grid serves as a persistent reference throughout the thesis and supports the alternative reading itineraries suggested in Section 1.5.

Table 1. 1 Mapping of the six research questions to the specific objectives, to the chapters in which each question is developed methodologically, and to the case studies in which it is tested empirically.

RQ	Short title	Primary literature gap	Specific objective addressed	Methodological development	Empirical testing
RQ1	Conditions and boundaries of structurally meaningful InSAR information at the building scale	Lack of systematic assessment of the scatterer-population conditions under which building-scale interpretation becomes feasible, as opposed to demonstrated individual-case feasibility at the urban and regional scale	SO1	Chapters 2 and 4	Case studies I, II and III (Chapter 5)
RQ2	Quantitative comparison of EGMS and user-driven processing at the building scale	Absence of a paired quantitative evaluation across density, quality and structural interpretability of standardised and user-driven products on the same structures	SO1	Chapter 3	Case studies I and II (Sections 5.2 and 5.3); synthesis in Section 5.5.1
RQ3	Temporal coherence as a	Limited exploration of	SO3	Chapter 4	Case study III

RQ	Short title	Primary literature gap	Specific objective addressed	Methodological development	Empirical testing
	diagnostic variable rather than a binary quality filter	the spatial and temporal patterns of coherence as carriers of structural and diagnostic information, beyond their role as threshold-based filters			(Section 5.4)
RQ4	Substitution potential of satellite-derived environmental variables for in-situ sensors in data-driven displacement models	Absence of systematic investigation of LST and SWI products as proxies for in-situ thermal and hydrological inputs in surrogate displacement models	SO3	Chapter 4	Case study III (Section 5.4)
RQ5	Diagnostic value of TS for the retrospective chronology and spatial extent of structural failures	Lack of formal consolidation of TS as a distinct diagnostic category within the PS-InSAR interpretation chain, despite documented occurrences in forensic analyses	SO2 and SO3	Chapter 4	Case study II (Section 5.3); Section 5.6
RQ6	Building-scale fusion of InSAR and airborne LiDAR for post-event damage classification	First systematic test of per-building InSAR and LiDAR fusion in a wildland-urban interface fire context, a	SO4	Section 4.5.3	Section 5.6 (Palisades Fire, ongoing research)

RQ	Short title	Primary literature gap	Specific objective addressed	Methodological development	Empirical testing
		regime not previously addressed in the published literature			

1.5. Thesis outline and reading guide

The thesis is organised in six chapters whose internal logic proceeds from general framing, through progressively more specialised methodological development, to empirical testing and synthetic discussion. The sequence reflects the ambition to construct a single coherent argument in which each chapter supplies a necessary component of the overall framework, and no case study can be correctly interpreted without the methodological apparatus developed beforehand.

Chapter 1, which the reader is now concluding, establishes the context of the research. It motivates the need for SHM of existing full-scale structures, reviews the spectrum of monitoring technologies currently deployed, locates satellite interferometric observations within that spectrum, and formulates the research questions that the remainder of the thesis addresses. The boundaries of the study and the interpretive standpoint from which satellite data are read throughout the work are set in this chapter.

Chapter 2 reviews the state of the art of satellite data for SHM. The chapter begins with the physical principles of SAR acquisition, the geometric parameters that control the spatial resolution and sensitivity of the observations, and the auxiliary products that support the interpretation of displacement records. It then surveys the progressive migration of SAR interferometry from its geophysical origins toward structural engineering applications, distinguishing the domains in which the technique has reached methodological maturity from those in which its use remains exploratory. The chapter closes with a discussion of the ReLUI guidelines as the first systematic codification of structural interpretation principles for InSAR data, and identifies the interpretive gap that the remainder of the thesis seeks to narrow.

Chapter 3 presents the InSAR data and the processing workflow adopted in the thesis. The chapter opens with the principles of conventional Differential

Interferometric SAR (DInSAR) and extends the treatment to multi-temporal approaches, with particular emphasis on the PSI formulation adopted in the present work; the QPS and SBAS approaches are discussed as complementary lines of the methodological literature. The three data sources used throughout the thesis, namely Sentinel-1, COSMO-SkyMed, and the EGMS, are described together with their respective acquisition characteristics and product levels. The chapter then details the user-driven processing workflow, from dataset selection and co-registration through Digital Elevation Model (DEM) integration, geocoding, and small-area processing. Quality assurance procedures, reference point selection criteria, and sources of uncertainty are discussed in the final sections of the chapter.

Chapter 4 develops the methodological core of the thesis. Point selection on structures is treated in relation to the geometric plausibility of scatterers on roofs, façades, and architectural details, and the notion of TS is introduced as a distinct diagnostic category. Time-series analytics are presented as the bridge between raw displacement records and structural interpretation, with dedicated treatment of linear trends, seasonal and thermal components, and change-point analysis. The translation of LOS measurements into structural response metrics is then addressed, together with the handling of heterogeneous datasets, co-referencing strategies, and the integration of satellite products with complementary sources of information, including in-situ monitoring, numerical models, and non-SAR remote sensing data. The chapter closes with a discussion of what a decision-ready diagnostic framework requires and how the satellite evidence can be organised to meet that requirement.

Chapter 5 reports the empirical work of the thesis. The empirical work opens with a section on EGMS-based screening that demonstrates the use of standardised products for portfolio-level prioritisation. Three case studies follow. Case study I (Section 5.2) addresses the *Santissima Trinità* church in Turin and compares EGMS and user-driven processing at the building scale on a masonry heritage structure. Case study II (Section 5.3) examines the retrospective forensic analysis of the *Himera Viaduct* and the *Albiano-Magra Bridge*, and introduces TS as diagnostic markers of structural events. Case study III (Section 5.4) investigates the *Sanctuary of Vicoforte* and develops a coherence-aware framework for soil-structure interaction monitoring. The framework integrates PS-InSAR observations with in-situ environmental variables through a Sparse Bayesian Learning (SBL) surrogate model. A synthesis section (Section 5.5) then proposes a tiered framework for satellite-assisted SHM. The synthesis discusses the implications of the case studies for engineering practice and articulates the conditions under which each diagnostic level becomes achievable. The chapter closes with a section on ongoing research

(Section 5.6). The section presents the preliminary results and methodological framework of an investigation on the 2025 Palisades Fire in Los Angeles. PS-InSAR time series are combined with airborne LiDAR geometric change metrics for post-event damage assessment at the building scale. The investigation is conducted in collaboration with the Qualcomm Institute at the University of California San Diego, within the Cultural Heritage Engineering Initiative (CHEI) and ALERTCalifornia.

Chapter 6 draws together the findings of the preceding chapters, discusses their limitations and threats to validity, formulates recommendations for engineering practice, and outlines prospective developments. Particular attention is devoted to the integration of multi-sensor fusion, the use of digital twins as predictive frameworks, and the future evolution of building-scale decision-ready diagnostics as an emerging research programme.

The chapters are designed to be read sequentially, but alternative itineraries may serve different readers. A reader whose primary interest lies in structural engineering applications may proceed from Chapter 1 directly to Chapter 5, consulting Chapters 3 and 4 as methodological support when specific analytical questions arise. A reader whose primary interest lies in remote sensing and interferometric processing may concentrate on Chapters 2 and 3 and consult Chapter 5 selectively, focusing on the case studies most aligned with their research field. A reader interested primarily in the synthesis of the research may proceed from the research questions of Chapter 1 to Section 5.5 and then to Chapter 6, returning to the methodological chapters for the details that support specific claims. Regardless of the reading order chosen, the six research questions formulated in Section 1.4 serve as a persistent reference throughout the thesis and are addressed in full by the end of Chapter 5, with their closure consolidated in Section 6.2. The three itineraries correspond, respectively, to the structural engineer and the asset manager concerned with the diagnostic use of satellite data, to the remote sensing specialist interested in the methodological adaptation of InSAR to the structural engineering domain, and to the reviewer or programme coordinator seeking a rapid orientation to the contributions of the thesis.

Chapter 2

Satellite data for SHM: state of the art

2.1. Satellite remote sensing basics

Satellite remote sensing refers to the acquisition of information about the Earth's surface without direct physical contact, through sensors that record electromagnetic radiation reflected, emitted, or backscattered by natural and artificial targets (Thenkabail et al., 2018). A fundamental distinction is that between passive and active sensors. Passive sensors (Figure 2. 1) measure naturally available radiation, typically reflected sunlight or thermal emission. Active sensors (Figure 2. 2) emit their own signal and record the returned response after interaction with the target (Rodriguez & Martin, 1992).

SAR belongs to the family of active microwave systems. Because microwave signals do not depend on solar illumination and can propagate through clouds, SAR acquisitions are possible day and night and under most atmospheric conditions. This operational independence is one of the main reasons why SAR has become a reference technology for deformation monitoring. The observational capability that supports satellite-based SHM today is the cumulative outcome of several decades of mission design and methodological development. The timeline of the principal remote sensing systems is shown in Figure 2. 3.




Figure 2. 1. Passive remote sensing detection mechanisms. The target is illuminated by an external source (typically solar radiation) and the sensor records the reflected energy; measurement is constrained to daylight conditions and clear atmosphere (Coccimiglio, 2025).




Figure 2. 2. Active remote sensing detection mechanisms. The platform emits a microwave pulse towards the target and records both amplitude and phase of the back-scattered signal; operation is independent of solar illumination and cloud cover, and phase coherence across repeated acquisitions enables interferometric displacement retrieval (Coccimiglio, 2025).

The usefulness of any satellite observation for engineering purposes depends on the combination of spatial, temporal, spectral, and radiometric resolution offered by the sensing system. Spatial resolution defines the ground area represented by a pixel. Temporal resolution determines how frequently the same area is revisited. These dimensions of resolution do not vary independently of one another; any real acquisition system reflects a set of engineering compromises among them rather than an unconstrained maximisation of each. From the standpoint of SHM, the decisive criterion is whether the spatial granularity, the revisit frequency, and the electromagnetic properties of a given dataset are adequate for the scale and the nature of the phenomenon under investigation. The sections that follow examine these factors with specific reference to SAR sensors, concentrating on acquisition modes, geometric parameters, and the data products of greatest relevance to the monitoring of structures.




Figure 2. 3. Timeline of the principal spaceborne SAR missions relevant to SHM, with indication of the frequency band, from 1992 to the present. Adapted from (Deslauriers et al., 2025).

2.1.1. SAR sensors and acquisition modes

SAR sensors are active microwave systems that illuminate the Earth's surface with a transmitted electromagnetic pulse and record the backscattered response returned to the antenna (Meyer, 2019; Rodriguez & Martin, 1992). The interaction between the radar signal and the target depends on the wavelength of the emitted pulse, the geometric and dielectric properties of the surface, and the local incidence angle.

Part of the incident energy is scattered away from the sensor, part is absorbed, and part is returned toward the antenna. The amplitude of the returned signal carries information about the reflectivity and roughness of the target, while the phase encodes the distance between the sensor and the scatterer. Both quantities are preserved in the complex SAR image and are essential for interferometric applications.

SAR systems operate in a side-looking geometry. The antenna is oriented at an angle with respect to the vertical, so that the radar beam illuminates a strip of terrain to one side of the satellite ground track. The term “synthetic aperture” (Figure 2. 4) refers to the fact that the motion of the platform along the flight direction is exploited to simulate an antenna much longer than the physical one, thereby improving azimuth resolution beyond what a real antenna of practical size could achieve (Meyer, 2019; Rodriguez & Martin, 1992). Image formation is based on two orthogonal directions. Range resolution is controlled by the bandwidth of the transmitted pulse and determines the ability to distinguish targets at different distances from the sensor. Azimuth resolution is controlled by the synthetic aperture length and determines the ability to distinguish targets along the flight direction.




Figure 2. 4. SAR side-looking geometry and synthetic aperture concept (Braun, 2019).

Spaceborne SAR missions for Earth Observation (EO) operate along near-polar, sun-synchronous orbits (SSO) at altitudes typically ranging from approximately 600 to 800 km, within the Low Earth Orbit (LEO) band. The near-polar inclination, shown in Figure 2. 5, provides global coverage including the polar regions, while the sun-synchronous configuration ensures that each location on the Earth surface is observed at nearly the same local solar time on every revisit. The operational altitude of each mission is a design compromise between radiometric budget, coverage, and orbital stability, and its specific value determines both the ground footprint and the revisit interval. The implications of the revisit interval on the temporal sampling of displacement time series are discussed in the paragraphs that follow.




Figure 2. 5. Near polar and sun-synchronous low earth orbit configuration of spaceborne SAR missions. Adapted from (Types of Orbits, 2020).

Each orbital cycle comprises an ascending arc, in which the platform travels broadly from south to north, and a descending arc that reverses this trajectory (Figure 2. 6); in both cases the antenna is oriented to illuminate the terrain to the right of the velocity vector, a convention adopted by nearly all current missions. The strip of ground that falls within the antenna beam during a single pass defines the swath, while the elapsed time between two consecutive acquisitions over the same target area constitutes the revisit time. These parameters are not merely orbital characteristics. They directly determine the temporal sampling of the displacement history, the availability of ascending and descending viewing geometries for a given site, and therefore the feasibility of LOS decomposition into vertical and horizontal components.




Figure 2. 6. Satellite acquisition geometry for ascending and descending orbits (Coccimiglio, 2025).

SAR data can be acquired in different modes, each designed as a trade-off between spatial resolution, swath width, and radiometric performance (Moreira et al., 2013). In Stripmap mode the antenna maintains a constant pointing angle relative to the ground while the platform advances along its orbit, sweeping out a continuous ground strip that offers a practical compromise between spatial resolution and areal coverage. Stripmap acquisition has long represented the standard operating mode for interferometric studies, and the bulk of the published PS-InSAR literature rests on data collected in this configuration. Spotlight mode adopts a different strategy: the beam is progressively steered so as to keep the same patch of terrain within the footprint for a longer fraction of the overpass, thereby sharpening the azimuth resolution at the cost of a narrower imaged corridor. This gain in spatial detail makes Spotlight geometry especially well-suited to analyses conducted at the scale of individual structures. In ScanSAR mode, the antenna scans multiple adjacent sub-swaths in elevation, achieving wider coverage at the expense of azimuth resolution. ScanSAR is therefore useful for regional-scale observation rather than local characterisation.

A notable evolution is the Terrain Observation by Progressive Scans (TOPS) concept, which can be regarded as a burst-mode evolution of ScanSAR (Grandin, 2015). In TOPS, the antenna beam is steered in azimuth during each burst, achieving wide swath coverage with improved radiometric uniformity compared to conventional ScanSAR. Sentinel-1 acquires its standard Interferometric Wide Swath (IW) products in TOPS mode, which produces three sub-swaths covering approximately 250 km on the ground. The TOPS acquisition strategy has direct

consequences for interferometric processing, because co-registration must account for the burst structure and the azimuth-varying Doppler centroid. Three of the main acquisition modes adopted by the SAR missions and the geometric principle of each mode are illustrated in Figure 2. 7.




Figure 2. 7. Comparison of the three principal SAR acquisition modes relevant to Structural Health Monitoring: Stripmap, ScanSAR, and TOPS. Adapted from (Boncori & Peter, 2019).

The choice of acquisition mode has practical consequences for SHM applications. Acquisition modes that afford finer spatial resolution, notably Spotlight and Stripmap, lend themselves to the detection of discrete scatterers on individual buildings and bridges, a context in which the number of coherent targets per unit area directly conditions the depth of any structural reading. Broad-coverage modes such as the IW configuration of Sentinel-1 serve a different purpose: they are better matched to territorial screening campaigns or to the surveillance of extended infrastructure corridors where completeness of coverage outweighs point-level detail. Of the satellite missions of greatest relevance to the present work, Sentinel-1 warrants particular mention for its systematic C-band acquisitions in TOPS mode, which deliver approximately 5 m resolution in the range direction and 20 m in azimuth. COSMO-SkyMed and TerraSAR-X operate in X-band and offer Stripmap and Spotlight configurations with spatial resolutions of approximately 3 m in Stripmap HIMAGE mode and sub-metre values in Spotlight configuration, which proves particularly advantageous for building-scale and bridge-scale analyses (Covello et al., 2010; Werninghaus & Buckreuss, 2010). The engineering use of SAR data must therefore always consider the acquisition mode together with the sensor band, because resolution, coverage, burst structure, and interferometric behaviour all depend on this choice.

2.1.2. Geometry and resolution: wavelength, incidence angle, polarization, revisit time

The physical meaning and engineering usefulness of SAR data are governed by a small set of acquisition parameters that control the interaction between the radar signal and the observed target. The radar wavelength governs the spatial scale at which the electromagnetic wave (Figure 2. 8) interacts with the illuminated surface. X-band systems, operating at a wavelength of approximately three centimetres, respond preferentially to fine geometric detail and tend to resolve individual architectural features on built targets with greater precision. C-band signals at approximately 5.55 centimetres (specifically 5.547 cm for Sentinel-1) and L-band signals at approximately 23.50 centimetres penetrate more deeply into vegetation canopies and loosely consolidated media, becoming more responsive to surface characteristics that vary over larger spatial extents.




Figure 2. 8. Main bands of electromagnetic spectrum.

Figure 2. 9 shows the differential penetration of microwave signals through vegetation as a function of the operating wavelength. In X-band the incident signal is scattered primarily by leaves and fine branches of the upper canopy, resulting in a dominant surface return. In C-band the signal penetrates through the canopy and interacts with mid-sized branches, producing a backscattering response that combines contributions from multiple vertical levels. In L-band the longer wavelength allows the signal to penetrate the full canopy and reach major branches, trunks, and in certain conditions the underlying surface. The differential penetration has direct implications for the identification of coherent scatterers on infrastructure surrounded by vegetation, as discussed in Sections 3.2 and 4.1.3.




Figure 2. 9. Differential penetration of microwave signals through vegetation as a function of the operating wavelength. Adapted from (Airbus Built Forest Monitoring Satellite Biomass Successfully Launched, 2025).

The practical consequence for engineering applications is that X-band missions such as COSMO-SkyMed and TerraSAR-X typically yield a higher density of identifiable scatterers on man-made structures, while C-band platforms such as Sentinel-1, with their wider swath and systematic revisit, are better adapted to sustained regional observation over extended time spans. Wavelength enters the measurement process in a second, more fundamental way. The interferometric phase, expressed as a function of the LOS displacement, obeys the relationship given in Equation 2. 1:

$$\varphi_{disp} = \frac{4\pi}{\lambda} \cdot d_{LOS} \quad 2.1$$

where φ_{disp} is the displacement-induced interferometric phase in radians, λ is the radar wavelength, and d_{LOS} is the LOS displacement. The factor of two arises from the two-way travel path of the radar signal between sensor and target. A complete phase cycle of 2π therefore corresponds to a LOS displacement equal to

half the wavelength, that is, approximately 2.78 centimetres for C-band and 1.56 centimetres for X-band. This quantity defines the unambiguous displacement range accessible within a single interferogram before phase wrapping becomes an issue, and ties the geometric sensitivity of the technique directly to the frequency band of the sensor.

The incidence angle is the angle between the incoming radar beam and the local vertical at the target location, and varies across the swath from near range to far range. For Sentinel-1 IW mode it spans from approximately 29° in the near range of sub-swath IW1 to more than 46° in the far range of sub-swath IW3. For COSMO-SkyMed, the incidence angle can be selected between approximately 20° and 60° depending on the acquisition configuration. Smaller incidence angles favour stronger returns from horizontal surfaces and increase sensitivity to vertical displacements. Larger incidence angles improve the ability to observe vertical features such as façades, but also increase the likelihood of foreshortening, layover, and shadow, phenomena discussed in Section 2.2.3, on steep terrain or tall structures (Hanssen, 2001). The incidence angle therefore controls not only backscatter intensity, but also which portions of a building are visible to the radar and how reliably displacement information can be attributed to a specific structural element. The geometric projection of a displacement vector onto the LOS direction depends explicitly on the incidence angle, and the full three-dimensional projection is developed in Section 2.2.2.

The definitions of slant range, ground range, and incidence angle used throughout this thesis refer to the side-looking acquisition geometry common to all spaceborne SAR systems, in which the antenna beam is directed laterally with respect to the direction of flight rather than towards the nadir point below the platform. The geometric configuration in the plane perpendicular to the flight direction is shown in Figure 2. 10. The orbital platform moves along the azimuth direction at altitude h above the reference surface, and the antenna illuminates a swath on the ground extending from the near range, where the incidence angle is smaller, to the far range, where it is larger. For a given target, the slant range r denotes the distance from the sensor measured along the LOS, the ground range is the corresponding distance measured on the reference surface, and the local incidence angle θ is the angle between the LOS and the vertical at the target location.




Figure 2. 10. Acquisition geometry of a spaceborne SAR system. Adapted from (Franceschetti & Lanari, 2018).

The side-looking geometry of Figure 2. 10 is realised operationally through two complementary orbital configurations, namely the ascending and descending passes, whose geometry is shown in Figure 2. 11. During the ascending pass the satellite moves from south-south-west to north-north-east and the antenna beam is directed laterally to the east, whereas during the descending pass the trajectory is reversed, and the antenna beam is directed westwards.




Figure 2. 11. Schematic representation of satellite radar acquisition in the case of Ascending orbit (left) and Descending orbit (right).

Polarisation describes the orientation of the transmitted and received electromagnetic field and is expressed through combinations such as Vertical-Vertical (VV), Horizontal-Horizontal (HH), Horizontal-Vertical (HV), and Vertical-Horizontal (VH), as shown in Figure 2. 12. Co-polarised channels (VV or HH) are the standard choice for deformation monitoring and interferometric applications because they maximise the signal returned from surface and dihedral scatterers typical of built environments. Cross-polarised channels (HV, VH) carry information on volume scattering and are more relevant to land cover characterisation than to structural monitoring. Sentinel-1 IW mode acquires in VV+VH as standard, whereas COSMO-SkyMed Stripmap commonly uses HH. The polarisation channel influences target stability and coherence over time, and therefore affects the suitability of a scatterer for multi-temporal interferometric analysis.




Figure 2. 12. Polarization: orientation of the transmitted and received electromagnetic field. Adapted from (Introduction to SAR—ArcGIS Pro | Documentation, 2026).

Revisit time is the interval between successive acquisitions over the same area. It determines the temporal sampling of the displacement history and affects the ability to reconstruct trends, seasonal cycles, and event-related changes. Sentinel-1 has a nominal revisit time of twelve days per satellite. The constellation originally operated with two satellites, Sentinel-1A launched in April 2014 and Sentinel-1B launched in April 2016, which provided a combined revisit of six days. Sentinel-1B ceased operations on 23 December 2021 following an on-board power anomaly, with the end-of-mission formally declared by European Space Agency (ESA) on 3 August 2022. The effective revisit interval over European territory was reduced to twelve days until Sentinel-1C was successfully launched in December 2024 and Sentinel-1D in November 2025, restoring the nominal six-day revisit with the two-satellite configuration.

COSMO-SkyMed can achieve revisit times as short as four days with its four-satellite constellation, although access to these data is not open. Shorter revisit times improve temporal coherence, reduce the risk of phase aliasing for fast deformation,

and allow finer resolution of seasonal behaviour. Longer revisit intervals limit the observation of rapidly evolving phenomena and increase temporal decorrelation, particularly over vegetated or changing surfaces (Zebker & Villasenor, 1992).

Spatial resolution in the range direction (Figure 2. 13), defined as the minimum separation at which two point targets along the LOS produce distinguishable echoes after pulse compression, is controlled by the bandwidth of the transmitted pulse according to Equation 2. 2:

$$\rho_r = \frac{c}{2B} \quad 2.2$$

where ρ_r is the slant-range resolution, c is the speed of light, and B is the pulse bandwidth. The factor of 2 accounts for the two-way path. Ground-range resolution is obtained by dividing ρ_r by the sine of the local incidence angle, and therefore becomes finer at larger incidence angles for the same pulse bandwidth.

Azimuth resolution (Figure 2. 13), defined as the ability to distinguish between two closely spaced targets located at the same range but at different angular positions along the flight path, is controlled by the length of the synthetic aperture and is given, for focused SAR processing, by Equation 2. 3:

$$\rho_a \approx \frac{L}{2} \quad 2.3$$

where ρ_a is the azimuth resolution and L is the physical antenna length along the azimuth direction.




Figure 2. 13. Range and azimuth resolution.

The remarkable feature of Equation 2. 3 is that ρ_a is independent of both the radar wavelength and the slant range to the target. This result follows from the cancellation between the linear increase of the synthetic aperture length with range, equal to $\lambda R/L$, and the inverse dependence of the focused resolution on the synthetic aperture dimension itself. The approximation symbol reflects the assumption of uniform antenna illumination and of a focused processing scheme without azimuth windowing; in operational systems, the weighting filters applied to suppress sidelobes of the impulse response broaden the main lobe by a factor typically comprised between 1.1 and 1.3, depending on the specific window adopted. The nominal figure $L/2$ therefore represents a theoretical lower bound rather than the operational resolution delivered by the processing chain. For Sentinel-1 IW mode, the nominal single-look resolution is approximately 5 m in range and 20 m in azimuth. For COSMO-SkyMed Stripmap HIMAGE, typical values are approximately 3 m in both range and azimuth. TerraSAR-X Spotlight achieves sub-metre resolution in both directions.

Equations 2. 1 and 2. 3 involve two distinct parameters of the radar system that must not be confused, although both relate to electromagnetic quantities. The bandwidth B is a property of the transmitted pulse and determines the spatial resolution of the system: a wider bandwidth produces a narrower main lobe of the impulse response in the range direction and therefore a finer ability to distinguish

adjacent targets. The wavelength λ is a property of the carrier frequency and determines the interferometric sensitivity of the system: a shorter wavelength produces a larger phase change per unit of physical displacement along the LOS, and therefore a finer ability to detect small movements. The two parameters are physically independent. A sensor may have high resolution (large bandwidth) but low displacement sensitivity (long wavelength), or the converse. For structural monitoring purposes, the practical implication is that the choice of sensor involves a trade-off. X-band systems such as COSMO-SkyMed offer both high resolution, with bandwidth up to approximately 300 MHz in Spotlight mode, and high displacement sensitivity, corresponding to 1.56 centimetres per phase cycle, but at the cost of reduced coverage and higher susceptibility to phase ambiguity on fast-moving targets. C-band systems such as Sentinel-1 offer wider coverage and a larger unambiguous displacement range of 2.78 centimetres per phase cycle, but at coarser resolution, with bandwidth of approximately 56 MHz in IW mode, and with fewer PS on compact structural targets. These differences have direct consequences for building-scale SHM, because finer resolution increases the probability of identifying individual coherent targets on structural elements, while coarser resolution may result in mixed scatterers whose physical attribution is ambiguous. The engineering relevance of resolution cannot be assessed in isolation, however. A dataset with fine spatial resolution but poor temporal coherence or unfavourable incidence angle may be less useful for structural interpretation than a coarser dataset with better temporal coverage and more stable phase behaviour.

2.1.3. Data types, DEMs, and auxiliary products

A SAR dataset is never defined by the image alone. Product type, elevation information, orbit files, and metadata all influence what can be extracted from the data and how reliably it can be used in interferometric workflows. The most fundamental distinction is between complex and detected products. For Sentinel-1, the two main Level-1 product families are Single Look Complex (SLC) and Ground Range Detected (GRD) (European Space Agency, 2016). SLC products are delivered in the native acquisition geometry and preserve both amplitude and phase for each pixel. They represent the standard input for interferometric and multi-temporal processing. GRD products are intensity images projected in ground range, multi-looked and detected, and therefore do not retain the phase information required for interferometry. A conceptual comparison between SLC and GRD is shown in Figure 2. 14.

A single Sentinel-1 IW SLC product covers approximately 250 km in range and 250 km in azimuth and contains six measurement files corresponding to the three sub-swaths (IW1, IW2, IW3) and two polarisation channels (VV and VH). Each file is typically several gigabytes. For interferometric SHM applications, only the co-polarised channel (VV for Sentinel-1, HH for COSMO-SkyMed) and the relevant sub-swath are normally used. An analogous distinction between complex and detected products exists for TerraSAR-X and COSMO-SkyMed, where interferometric workflows likewise require SLC-equivalent products that preserve phase.




Figure 2. 14. SLC vs GRD conceptual comparison. Adapted from (Colin, 2022).

Digital Elevation Models (DEMs) are required at multiple stages of SAR and InSAR processing. They support geocoding, which relates radar coordinates (range and azimuth) to geographic coordinates (latitude and longitude). They are used to compute and remove the topographic phase component from interferograms, which would otherwise mask the displacement signal. They also enable the simulation of geometric distortions such as foreshortening and layover, allowing the user to predict which parts of a structure are visible in a given acquisition geometry. The topographic phase contribution for a given interferometric pair is expressed by 2. 4, which establishes the proportionality between residual topography and measured phase in the differential interferogram. The form of this relationship, adopted here in the convention of Hanssen 2001, is given by:

$$\varphi_{topo} = \frac{4\pi}{\lambda} \cdot \left(\frac{B_{\perp}}{R \cdot \sin \theta} \right) \Delta h \quad 2. 4$$

where φ_{topo} is the topographic phase contribution in radians, λ is the radar wavelength, $B_{\perp}$ is the normal baseline between the two acquisitions (Figure 2. 15),

R is the slant range from sensor to target, θ is the local incidence angle, and Δh is the elevation of the target above the reference surface used during interferogram formation.




Figure 2. 15. Definition of baseline (B), normal baseline ($B_{\perp}$) and parallel baseline ($B_{\parallel}$) in the plane orthogonal to the orbits.

The quantity that governs the sensitivity of the interferometric phase to residual topography is the ratio $B_{\perp}/R \sin \theta$, often expressed as the height of ambiguity (Equation 2. 5):

$$h_a = \frac{\lambda \cdot R \cdot \sin \theta}{2B_{\perp}} \quad 2. 5$$

which represents the elevation change corresponding to one complete phase cycle. Short baselines produce high heights of ambiguity and reduce the contamination of the displacement signal by residual topography, which is one of the reasons why modern interferometric missions such as Sentinel-1 maintain a tightly controlled orbital envelope during repeat passes. The full decomposition of the interferometric phase into its topographic, displacement, atmospheric, orbital, and noise components is developed in Section 3.1.1.

The quality of the DEM has a direct impact on the separation between topographic and displacement-related phase components. Residual elevation errors

propagate into the estimated displacement and may complicate the association between measurement points and structural elements, particularly at building scale. The Copernicus GLO-30 DEM at 30 m resolution has progressively become the default reference for large-area interferometric processing, owing to its near-global coverage, open access, and vertical accuracy within a few metres over most terrestrial surfaces (Airbus Defence and Space, 2022). The older Shuttle Radar Topography Mission (SRTM) DEM at 90 m resolution remains available as a default option in many interferometric processing chains and is adequate for regional-scale analyses. When higher accuracy is required, for instance in urban areas with tall buildings or in mountainous terrain, user-supplied DEMs at finer resolution, for example from airborne LiDAR, can be imported into the processing workflow.

Auxiliary products complete the information required for accurate pre-processing. Precise orbit files define the satellite position and attitude during each acquisition and contribute directly to co-registration quality, baseline computation, and flat-earth phase removal, that is, the suppression of the phase trend due to Earth curvature that is addressed in detail in Section 3.1.1. For Sentinel-1, orbit products are distributed through the Copernicus Precise Orbit Determination service and are available at different timeliness levels. Restituted (or operational) orbits (RESORB) are available within hours of the acquisition and are sufficient for most applications; precise orbits, available after approximately twenty days, provide the highest accuracy. The automated retrieval of orbit files is a standard feature of modern interferometric processing chains, as described in detail in Section 3.3. Acquisition metadata, including timing, burst boundaries, calibration parameters, and instrument configuration, are embedded in the product annotation files and are read automatically by the processing software. Engineering users should therefore be aware that SAR scenes are not self-contained images, but structured datasets embedded in a broader framework of geometric and auxiliary information, and that the quality of the final interferometric result depends on the completeness and correctness of all these inputs.

2.2. Satellite data in structural engineering

InSAR has progressively entered the practice of structural engineering as a source of spatially distributed displacement information that complements conventional sensor-based monitoring systems. The fundamental appeal of the technique rests on three properties that no ground-based instrument can fully replicate. The first is the capacity to cover large areas without any physical access to the structure. The

second is the availability of satellite archives extending over one or more decades, which allows retroactive application to structures for which no monitoring programme was active at the time. The third is a sensitivity to range change that, under favourable conditions, reaches the sub-millimetre level. For Sentinel-1, the freely accessible observation record covers most of Europe from 2014 onwards, providing a temporal baseline of more than a decade at the time of writing. Such a record cannot be matched by any instrument deployed after the onset of a damage process. However, translating satellite displacement data into structural engineering information requires a careful understanding of three fundamental aspects. The measurement captures motion only along the satellite LOS, not in any structural reference direction. Displacement signals at the millimetre level may be masked or distorted by atmospheric artefacts and thermal effects. Not all structural surfaces generate usable coherent scatterers, and the spatial density of detectable points varies widely with building geometry, material, and orientation relative to the satellite. These constraints define the conditions under which satellite data can provide reliable diagnostic input for structural assessment. The following sections address the physical content of InSAR displacement measurements on built assets, the decomposition problems arising from the single-geometry observation, and the systematic environmental effects that must be separated from any structural signal before interpretation can proceed.

2.2.1. Ground and structural displacements

Satellite interferometric observations can provide displacement information at both ground level and the level of the structure itself. The distinction between ground displacement and structural displacement is essential for the engineering interpretation of satellite-derived data. The motion of a structural target may derive from several distinct physical causes. These causes include structural deformation, differential settlement of the supporting ground, terrain instability, and environmental loading. In many practical cases, the engineering problem requires distinguishing between these contributions rather than simply recording the total observable displacement. The conflation of ground and structural signals is a recurrent source of interpretative ambiguity in SHM applications (Burland et al., 1977; Prota et al., 2025). Accordingly, the following paragraphs examine each class of displacement in turn, with attention to their physical meaning and their role in satellite-based structural interpretation.

Ground displacements represent the first and most extensively studied class of motions detectable through satellite interferometry. They may be associated with

subsidence, landslide activity, consolidation of compressible soils, groundwater drawdown, or long-term settlement processes. In structural engineering, such movements are relevant because they frequently act as external kinematic drivers of structural response. A bridge or masonry building may deform not from internal deterioration alone, but from foundation movement induced by terrain instability. Several studies have demonstrated that satellite interferometry can capture coupled ground-structure mechanisms in infrastructure monitoring (Macchiarulo, Giardina, et al., 2022; Talledo et al., 2022). The retrospective identification of pre-failure displacement signatures in satellite time series has been demonstrated for several bridge collapse events and is examined in detail in Case study II (Section 5.3). Structural displacements, while related to ground movements in several important cases, are associated more directly with the response of built assets.

Structural displacements are motions that can be attributed more directly to the behaviour of the structure or of specific structural elements. Satellite observations may reveal long-term subsidence trends, relative movements between structural zones, or thermally induced cyclic behaviour. In bridges, interferometry has been applied to reconstruct displacement trends for decks, piers, and abutments (Farneti et al., 2023a; Tonelli, Valentini, et al., 2023). For heritage buildings, satellite data have been used to observe façade displacements and thermal-cycle-related behaviour (Coccimiglio et al., 2026). However, the resulting point density depends on coherence, target visibility, and resolution, and may be insufficient for localised damage detection. In practice, the distinction between ground and structural displacements is often blurred, owing to the complexity of soil-structure interaction.

The boundary between ground and structural displacement is particularly difficult to define in cases of soil-structure interaction. A radar target on a structure may reflect both structural response and foundation movement simultaneously. Conversely, a terrain point near the asset may carry information directly relevant to structural interpretation. Monumental buildings on deformable ground and bridges affected by slope instability illustrate this coupling well. A further complication arises because satellite radar systems measure only the displacement component projected along the sensor's LOS and not the full displacement vector. As a result, the same physical motion may produce markedly different LOS signatures depending on incidence angle and orbit direction (Lorenz et al., 2024). This geometric dependency is examined in the following section, which addresses LOS measurements, decomposition approaches, and their implications for structural interpretation.

2.2.2. LOS measurements, decomposition issues, and thermal/seasonal effects

The displacement measured by satellite interferometry is not the full three-dimensional vector of a target's motion. Rather, it represents the component projected along the sensor-to-target line of sight. The LOS displacement d_{LOS} is related to the target's east, north, and vertical motion components through a projection defined by the satellite geometry, as expressed in Equation 2. 6:

$$d_{LOS} = d_N \cdot \sin \theta \cdot \sin \alpha - d_E \cdot \sin \theta \cdot \cos \alpha + d_U \cdot \cos \theta \quad 2. 6$$

where d_E , d_N , and d_U are the east, north, and vertical (up) components of the displacement vector in the local ENU reference frame, θ is the local incidence angle, and α is the satellite heading angle measured clockwise from north, for a right-looking acquisition geometry. The geometric projection that links the three-dimensional displacement vector to the scalar LOS measurement is illustrated in Figure 2. 16.

The sign convention in Equation 2. 6 corresponds to a LOS measured positive in the direction from target to sensor, which is the convention adopted throughout this thesis. As a result, a motion away from the sensor produces a negative LOS value. A purely vertical subsidence of one centimetre observed at an incidence angle of 40° produces a LOS displacement of approximately -7.66 millimetres, computed as ten millimetres of downward motion multiplied by the cosine of 40° . An eastward displacement of one centimetre observed in ascending geometry, with heading angle approximately equal to 350° and incidence angle equal to 40° , produces a LOS displacement of approximately -6.33 millimetres when the full expression of Equation 2. 6 is evaluated, confirming that the east sensitivity of near-polar missions is comparatively high in ascending geometry at moderate incidence angles. For descending geometry with heading angle approximately 190° , the same eastward displacement produces a LOS value of approximately +6.33 millimetres, with opposite sign. The north component of motion maps onto LOS through a coefficient $\sin(\theta) \cdot \sin(\alpha)$, which for near-polar missions with heading close to 0° or 180° remains far smaller than $\sin(\theta)$, and which vanishes in the limit of exactly polar orbits. For operational near-polar missions such as Sentinel-1, with heading angles close to 350° (ascending) or 190° (descending), the north sensitivity is small but not zero, whereas the east and vertical sensitivities dominate the observed LOS signal.




Figure 2. 16. Decomposition of the displacement vector along the satellite LOS. Left panel: projection of the horizontal displacement components onto the satellite-looking direction, where α is the azimuth of the satellite heading vector (measured positive clockwise from north), SH is the satellite heading direction, SL is the satellite-looking direction, and d_E and d_N are the easting and northing components of the horizontal displacement, respectively. Right panel: projection of the vertical and horizontal displacement components onto the LOS direction, where d_U is the vertical displacement component and θ is the local incidence angle at the reflection point. Adapted from (Ng, 2010).

The physical meaning of a LOS measurement varies considerably with sensor geometry and structural orientation. A vertical ground settlement and an out-of-plane façade movement produce very different LOS signatures for the same incidence angle. In practice, the same physical phenomenon may be clearly observable in one geometry and virtually invisible in another (Cernuto et al., 2026). The selectivity of the LOS projection is especially consequential for structural assets whose principal deformation direction is weakly aligned with the radar sensitivity axis. Building-scale geometry often prevents a direct correspondence between the LOS signal and the engineering variable of interest (Prota et al., 2025). Decomposing LOS measurements into physically interpretable displacement components is therefore a desirable but challenging step in structural interpretation.

Decomposition into vertical and east displacement components requires at least two independent viewing geometries, typically an ascending orbit and a descending orbit acquisition. Under the assumption of negligible north motion, the LOS measurements from ascending and descending tracks can be combined as expressed in Equation 2. 7:

$$\begin{pmatrix} d_{LoS,asc} \\ d_{LoS,desc} \end{pmatrix} = A \cdot \begin{pmatrix} d_E \\ d_U \end{pmatrix} \quad 2.7$$

where the two-by-two design matrix A contains the projection coefficients for the two orbital geometries (Equation 2. 8):

$$A = \begin{pmatrix} -\sin \theta_{asc} \cdot \cos \alpha_{asc} & \cos \theta_{asc} \\ -\sin \theta_{desc} \cdot \cos \alpha_{desc} & \cos \theta_{desc} \end{pmatrix} \quad 2.8$$

The elements of A are the projection coefficients of the easting and vertical motion components onto the LOS direction for each acquisition geometry, where θ_{asc} and θ_{desc} are the local incidence angles and α_{asc} and α_{desc} are the heading angles of the ascending and descending orbits respectively. The condition number of A increases as the two LOS directions become more parallel, with consequences for the stability of the inversion. Inversion of A yields an estimate of the east and vertical displacement components. The low north sensitivity of near-polar missions remains the principal source of residual uncertainty in this decomposition, and the result should be interpreted as a two-component approximation of a three-component motion rather than as a complete reconstruction of the displacement vector (Lorenz et al., 2024). Further practical limitations emerge at building scale: ascending and descending datasets rarely observe the identical point on a structure, since the scatterers that remain coherent under one viewing geometry are generally not the same as those coherent under the other (Di Carlo et al., 2021). In bridges, the measured LOS signal may reflect vertical settlement, thermal deformation, or combined support and deck motion simultaneously (Farneti et al., 2023b). Even when points are co-located, the retrieved components remain partial because the east and vertical sensitivities in Equation 2. 6 are not fully independent and the condition number of A depends on the incidence angles involved (Lorenz et al., 2024). Decomposition should therefore be regarded as an interpretative aid rather than as a reliable reconstruction of the structural displacement vector.

Beyond geometric limitations, displacement time series on structures are systematically affected by environmental and thermal variability. Thermal dilation causes structures to expand and contract cyclically in response to temperature changes, generating periodic LOS displacement oscillations whose amplitude depends on the thermal expansion coefficient of the structural material, the characteristic dimension of the dilating element, and the projection of the dilation

vector onto the LOS direction (Monserrat et al., 2011). A linear model relating LOS displacement to temperature change is expressed as Equation 2. 9:

$$d_{LOS,thermal}(t) = K \cdot \Delta T(t) \quad 2. 9$$

where K expressed in millimetres per degree Celsius is the effective LOS thermal sensitivity of the observed element, and $\Delta T(t)$ expressed in degrees Celsius is the temperature anomaly with respect to a reference epoch. For a predominantly vertical dilation, the parameter K takes the form $\alpha_T \cdot L \cdot \cos \theta$, where α_T is the coefficient of linear thermal expansion of the structural material and L is the characteristic vertical dimension. For a predominantly horizontal dilation along a direction whose azimuth β is measured clockwise from north, the parameter K takes the form $\alpha_T \cdot L \cdot \sin \theta \cdot \sin(\alpha_{sat} - \beta)$, where α_{sat} denotes the satellite heading angle and the term $\sin(\alpha_{sat} - \beta)$ expresses the alignment between the dilation direction and the horizontal LOS projection. The value of K therefore depends not only on the material and geometry of the structure but also on the acquisition geometry of the specific dataset (Tonelli, Caspani, et al., 2023). In addition to annual thermal cycles, seasonal effects from moisture variation and groundwater fluctuations may be superimposed on structural time series. These contributions may obscure genuine structural trends, particularly where the structural response is slow relative to the seasonal signal. Analytical approaches for separating these contributions have been proposed in the literature and are developed methodologically in Section 4.2.2.

From an engineering perspective, LOS measurements require a double interpretative effort. The analyst must first understand what component of motion is observable under the given geometric conditions, then separate structurally informative displacement contributions from those driven by thermal or seasonal variability. LOS observations should not be treated as direct structural response metrics. They are physically filtered measurements that require contextualisation through geometry, mechanical understanding, and external validation. In this thesis, LOS measurements are treated as the primary observational level from which structural interpretation begins, and their decomposition, contextualisation, and separation from environmental effects are examined as necessary steps for building a reliable SHM-oriented interpretation framework.

2.2.3. Main limitations: decorrelation, geometric distortions, atmosphere, and referencing

Beyond the interpretative challenges associated with LOS geometry and environmental variability, satellite InSAR observations are affected by a set of physical and methodological limitations that constrain the quality, spatial coverage, and reliability of the retrieved displacement information. These limitations originate from different stages of the measurement chain and must be understood before any structural interpretation is attempted. The following discussion addresses the four factors most relevant to engineering applications: decorrelation, geometric distortions, atmospheric delay, and reference dependence.

Decorrelation refers to the loss of statistical similarity between two SAR acquisitions and represents one of the most pervasive limitations of interferometric analysis. The interferometric coherence, defined as the normalised cross-correlation of the complex signals from two co-registered images, quantifies the degree to which the phase difference between the two acquisitions carries meaningful information (Hanssen, 2001). When coherence is high, the interferometric phase can be interpreted in terms of geometric and displacement contributions. When coherence drops, the phase becomes increasingly dominated by noise and loses its diagnostic value. Decorrelation has two main components.

Temporal decorrelation occurs when the scattering properties of the observed surface change between acquisitions. Vegetated areas, agricultural fields, and water bodies are typical examples of surfaces where the radar return varies rapidly over time, producing low coherence even at short temporal baselines (Zebker & Villasenor, 1992). Temporal decorrelation is wavelength-dependent: shorter wavelengths such as X-band are more sensitive to small-scale surface changes than longer wavelengths such as L-band.

Geometric decorrelation, on the other hand, arises from differences in the observation geometry between two acquisitions and is controlled by the spatial (normal) baseline. When the baseline exceeds a critical value that depends on wavelength, incidence angle, and terrain slope, the spectral overlap between the two observations decreases and coherence is lost (Rosen et al., 2000). Modern satellite missions such as Sentinel-1 operate within a tightly controlled orbital tube, a virtual, narrow 3D corridor in space that a SAR satellite must strictly follow to ensure repeat-pass interferometry compatibility (Figure 2. 17). It keeps normal baselines small and reduces geometric decorrelation. Temporal decorrelation, however,

remains a significant constraint, especially in non-urban environments and at high spatial resolution.




Figure 2. 17. Schematisation of InSAR satellite orbital tube.

The acquisition geometry of side-looking SAR systems introduces three well-known distortion effects in the radar image: foreshortening, layover, and shadow. All three arise from the fact that the radar measures distances in slant range rather than on the ground, and that surfaces with different orientations relative to the LOS are projected differently into the image. Foreshortening (Figure 2. 18) occurs when a slope faces the sensor, so that the slant-range distance between the base and the top of the slope is compressed and the feature appears shorter in the image than it is on the ground. The radar return from the entire slope is concentrated into fewer resolution cells, resulting in anomalously bright pixels.

When the slope exceeds the local incidence angle, the compression becomes so severe that the top of the feature is imaged closer to the sensor than its base. This extreme case is known as layover (Figure 2. 18), and its visual consequences in the SAR image can be counterintuitive for a reader accustomed to optical imagery. The high-resolution X-band SAR amplitude image of the Eiffel Tower in Paris shown in Figure 2. 19 illustrates the phenomenon. In an orthorectified optical image viewed from above, the Tower appears as a compact footprint at ground level, with the upper structure projecting vertically and therefore invisible. In the SAR amplitude image, the Tower appears as an elongated bright feature extending from its true ground position toward the near range, because the radar signal reaches the top of the structure before it reaches the base. The tip of the Tower, being closer to

the satellite in slant range than the surrounding ground, is mapped to a position in the image that lies on the near-range side of the base, producing a geometric inversion in which the structure appears to have fallen toward the satellite. The phenomenon is not an artefact of the processing but a direct consequence of the side-looking acquisition geometry. For structural monitoring applications, the layover effect means that PS identified on tall buildings or towers may be geocoded at planimetric positions that differ substantially from their true ground-plan location, with the offset increasing with the height of the reflecting element above the reference DEM. The residual height estimated during the PS-InSAR inversion (Section 3.1.2) partially corrects this effect, but the correction is only as accurate as the height estimate itself, and a systematic positional bias may remain for elements whose elevation departs significantly from the DEM reference surface.




Figure 2. 18. Schematic of foreshortening, layover, and shadow distortions in side-looking SAR geometry. Ground points a, b, c are “seen” by radar as points a’, b’, c’ in the slant range. (Alaska Satellite Facility, 2026a)

Shadow (Figure 2. 18) is the complementary effect: when a slope faces away from the sensor more steeply than the incidence angle, the radar beam cannot illuminate the area behind the obstacle. No signal is returned from shadowed zones, and the corresponding pixels contain only noise. Shadow may render entire façades invisible to the satellite. The severity of all three effects depends on the incidence angle, the orbit direction, and the local geometry of the building and its surroundings. In practice, the combined use of ascending and descending acquisitions partially mitigates these limitations, because areas affected by layover in one geometry may be correctly imaged in the other (Crosetto et al., 2016).




Figure 2. 19. SAR image of the Eiffel Tower in Paris. Adapted from (Capella, 2024).

Atmospheric delay constitutes one of the most insidious sources of error in differential interferometry. A common misconception among users approaching InSAR from a general remote sensing background is that the ability of radar signals to penetrate clouds implies immunity of the interferometric measurement from atmospheric effects. The distinction that must be drawn is between the amplitude and the phase of the radar signal. The amplitude, which records the backscattered intensity from ground targets, is largely unaffected by the troposphere under most weather conditions: even in the presence of heavy cloud cover, the radar pulse reaches the ground and returns a usable amplitude measurement. The phase, however, records the two-way travel path with sub-wavelength sensitivity, and the refractive index of the troposphere introduces a path delay that varies spatially and temporally with temperature, pressure, and water vapour content (Hanssen, 2001; Massonnet & Feigl, 1998). A cloud layer that is invisible in the amplitude image can introduce a phase delay of several radians, equivalent to centimetres of apparent LOS displacement, into the interferogram. A typical tropospheric artefact may reach variances of 30 to 50 millimetres at length scales of about 100 kilometres, while at scales of about 1 kilometre the relative atmospheric error is typically below 8 millimetres, with the spatial structure function increasing approximately as the square root of the distance between observation points (Sandwell et al., 2025). In a single interferogram, the atmospheric contribution cannot be separated from the displacement signal without external information, and this is one of the principal reasons why single-pair DInSAR has limited applicability for slow deformation monitoring. MT-InSAR techniques address this problem by exploiting redundancy in large stacks of interferograms. Under the assumption that atmospheric delay is spatially correlated but temporally uncorrelated from one acquisition to the next,

the atmospheric phase screen can be estimated and partially removed during the inversion process (Ferretti et al., 2001). The effectiveness of this separation depends on the number of available images, the temporal sampling, and the spatial characteristics of the atmospheric signal relative to the deformation pattern of interest. At building scale, where the spatial extent of the observed structure is small compared to the correlation length of typical atmospheric artefacts, the atmospheric contribution tends to be spatially uniform across the structure. Relative displacement patterns within a single building are therefore less affected by atmospheric noise than absolute displacement values referred to a distant reference point.

The magnitude of the atmospheric contribution and its capacity to dominate the interferometric signal are illustrated in Figure 2. 20, which shows a differential interferogram acquired over the Naples area together with the Atmospheric Phase Screen (APS) estimated from the same image pair through the PSI technique. The fringe pattern visible in the left panel, despite the very short temporal baseline of the tandem acquisition and the negligible expected ground displacement between the two dates, is almost entirely accounted for by the atmospheric component shown in the right panel. The example establishes the order of magnitude of atmospheric residuals in single-pair interferometry and motivates the multi-temporal approaches discussed in Chapter 3, in which atmospheric screens are estimated and removed through stack redundancy rather than through single-pair processing.




Figure 2. 20. Example of atmospheric artefacts on a differential interferogram of the Naples area, computed from a tandem pair with normal baseline of approximately -97 m. Left panel: differential tandem interferogram, in which the fringe pattern is dominated by atmospheric contributions. Right panel: Atmospheric Phase Screen estimated for the same image pair through the PS technique. As clearly visible, most of the interferometric fringes derive from atmospheric effects. Adapted from (Parizzi et al., 2006).

The last limitation discussed here concerns the inherently relative nature of interferometric measurements. InSAR-derived displacements are differential in both time and space: the measurement represents the difference in path length between two acquisition dates, referenced to a spatial point assumed to be stable. The choice of the reference point directly affects the displacement values attributed to all other targets in the scene, and a moving reference introduces a systematic bias that cannot be corrected retrospectively. The formalisation of this limitation and the strategies available to mitigate it, including GNSS anchoring and the use of EGMS Calibrated products as external constraints, are developed in Sections 3.5.2 and 4.3.1. The ReLUIIS guidelines for the use of interferometric satellite data in structural applications explicitly recommend that users pay attention to the reference point strategy and assess its potential impact on the interpretation of displacement patterns (Prota et al., 2025).

The five limitations described above and summarised in Table 2. 1 are not independent of each other. Decorrelation reduces the number of measurable points and may leave entire portions of a structure unobserved. Geometric distortions determine which parts of a building are visible to the satellite and whether the returned signal can be unambiguously attributed to a specific structural element. Atmospheric delay affects the accuracy of the retrieved displacement, especially when the reference point is far from the target. The reference strategy, in turn, determines the frame within which all other quantities are expressed. For building-scale SHM, these limitations interact with the small spatial extent and the geometric complexity of structures to define a practical observability envelope: the set of conditions under which satellite InSAR can provide information that is both reliable and structurally meaningful. Understanding this envelope is a prerequisite for any responsible use of interferometric data in engineering diagnostics.

Table 2. 1. Summary of the main limitations affecting satellite InSAR measurements, with their effect on the observation, the typical mitigation strategies, and their specific implications for building-scale SHM.

Limitation	Effect on measurement	Typical mitigation	Implication for building-scale SHM
Temporal decorrelation	Loss of phase information over vegetated or changing surfaces.	PS and DS target selection; adoption of short temporal baselines.	Fewer measurable points on non-rigid or temporary elements.

Limitation	Effect on measurement	Typical mitigation	Implication for building-scale SHM
Geometric decorrelation	Loss of spectral overlap at large normal baselines.	Controlled orbital tubes (e.g., Sentinel-1).	Generally minor for modern missions.
Foreshortening and layover	Geometric compression or superposition of scatterers.	Joint use of ascending and descending acquisitions.	Tall buildings affected; ambiguity between facade and roof attribution.
Shadow	No radar returns from areas hidden behind steep obstacles.	Use of complementary orbit geometry.	Entire facades may remain invisible to the satellite.
Atmospheric delay	Spatially smooth phase artefacts that may mimic deformation.	MT-InSAR temporal filtering; Atmospheric Phase Screen estimation.	Relative patterns within a building are less affected than absolute values.
Reference dependence	All displacements are expressed relative to an assumed-stable point.	Careful reference selection; anchoring to GNSS or external networks.	A distant reference amplifies residual errors on the structure.

2.3. Interpretation frameworks and literature applications

The use of satellite InSAR for structural engineering purposes has developed along two converging trajectories. On one side, interferometric techniques that were originally designed for geophysical applications have been progressively applied to individual buildings, bridges, and heritage assets. On the other, a parallel effort has begun to codify the resulting experience into interpretation frameworks that a structural engineer can apply without specialist remote sensing training. Neither trajectory has reached a stable endpoint. The application literature continues to expand, while the codification effort remains largely confined to the Italian context. The present section traces both lines of development, identifies the points at which

they intersect, and locates the gaps that the subsequent chapters of this thesis are designed to address.

2.3.1. From geophysical origins to structural applications

The early technical references for InSAR interpretation were produced within the remote sensing community and addressed a geoscience audience. The ESA training manual on InSAR principles (Ferretti et al., 2007) remains one of the most widely referenced introductions to the subject. It covers the physical basis of SAR, interferogram formation, error sources, and processing strategies, but it does not discuss how interferometric products should be translated into structural engineering quantities. The methodological literature on MT-InSAR processing and validation (Crosetto et al., 2016; Hanssen, 2001) similarly provides the algorithmic and statistical foundations of the technique without extending to the engineering interpretation of displacement time series on individual structures.

The transition toward structure-specific investigations became feasible with the increasing availability of high-resolution X-band data from COSMO-SkyMed and TerraSAR-X, and with the accumulation of free C-band archives from Sentinel-1 starting in 2014. The combination of improved spatial resolution, shorter revisit times, and longer temporal baselines made it realistic to identify individual coherent scatterers on built structures and to reconstruct their displacement histories over periods of several years (Bürgmann et al., 2000; Massonnet & Feigl, 1998). The resulting body of application studies can be organised around four thematic domains, each of which has contributed specific insights to the emerging interpretation framework.

2.3.2. Cultural heritage and historical structures

Monuments, churches, towers, and archaeological sites have constituted one of the earliest and most productive domains of application. These structures are particularly suitable targets for InSAR because they often contain rigid masonry elements with stable radar reflectivity, they are located in urban environments with high scatterer density, and their long-term preservation generates sustained institutional interest. Cigna et al. (2014) demonstrated the use of PS-InSAR to map subsidence-related deformation patterns across the UNESCO World Heritage Site of the Historic Centre of Rome, linking satellite-derived velocity fields to geological and geotechnical information (Figure 2. 21).




Figure 2. 21. Distribution of COSMO-SkyMed PSs over the archaeological area of the Palatine Hill, Roman Forum, and Colosseum, colour-coded by mean LOS velocity over the observation period from March 2011 to June 2013. Panel (a) provides an overview of the archaeological site, while panels (b), (c), and (d) report detailed views of the arches of the Domitian aqueduct along Via di San Gregorio, of the Convento di San Bonaventura al Palatino, and of the Chiesa di Sant'Anastasia with the archaeological ruins in the south-western corner of the Palatine Hill, respectively. The figure illustrates the high PS density achievable on heritage masonry surfaces with X-band high-resolution sensors and the spatial selectivity required for element-level structural diagnosis.

Cavalagli et al. (2019) combined satellite radar interferometry with in-situ tiltmeter and extensometer measurements for static monitoring of historical monuments in Gubbio, Italy, using both Sentinel-1 and COSMO-SkyMed data (Figure 2. 22). The study established an influential precedent for the joint use of remote and contact sensing on heritage structures, showing that the two data sources capture complementary aspects of structural behaviour. The group has subsequently extended this approach to multi-span bridges, with explicit treatment of measurement uncertainty, in the pre-collapse analysis of the *Albiano-Magra Bridge* (Farneti et al., 2023a). Tang et al. (2016) applied advanced MT-InSAR techniques to the Summer Palace in Beijing, detecting millimetric subsidence patterns that affected the site over a multi-year observation window.




Figure 2. 22. PSPs detected on the Consoli Palace in Gubbio from the ascending COSMO-SkyMed dataset through Persistent Scatterer Pair (PSP) analysis, colour-coded by mean LOS velocity. The left panel reports the displacement time series of two representative PSPs located respectively on the roof and at the base of the palace, illustrating the differential motion between elevated and basement structural zones that the PSP analysis is designed to resolve. The figure exemplifies the building-scale interpretive use of high-resolution X-band PS data on a heritage masonry structure.

These studies confirmed that satellite data can provide useful displacement information on heritage structures, but they also highlighted recurring interpretive challenges. Target identification is complicated by the geometric complexity of historic buildings: domes, vaults, and irregular rooflines produce layover and shadow zones that reduce the number of usable scatterers. Coherence filtering must be more conservative than in regional-scale analyses, because the consequences of retaining a noisy point on a building are more severe than in a statistical average over thousands of ground targets. Cross-validation against independent measurements remains essential but is rarely available for structures that have not been instrumented prior to the satellite analysis.

2.3.3. Bridges and transport infrastructure

Interest in bridge monitoring through InSAR has grown substantially in the wake of several high-profile structural failures. The collapse of the Morandi Bridge in Genoa in 2018 and the failure of the *Albiano-Magra Bridge* in 2020 prompted retrospective analyses of pre-existing satellite data aimed at identifying potential precursory displacement signals (Macchiarulo, Giardina, et al., 2022; Milillo et al.,

2019). These studies generated considerable attention but also provoked an epistemological debate, summarised below, about the limits of retrospective inference.

Independent of the collapse-driven interest, a body of methodological work on bridge-scale InSAR has accumulated over the past decade. The potential of PSI for monitoring bridges in the Czech Republic was investigated by Lazecky et al. (2015), who demonstrated the ability to detect seasonal thermal deformation cycles on concrete viaducts. A comprehensive time-series analysis framework for bridge deformation using Sentinel-1 PSI was subsequently proposed by Schlögl et al. (2021), incorporating seasonal trend decomposition and spatial clustering of scatterers along the bridge alignment; the approach was validated against in-situ measurements on an integral concrete bridge in Vienna. Working on a major highway bridge in Ontario, Canada, Cusson et al. (2018) showed that bridge railings act as effective natural reflectors for X-band InSAR, and that the thermal sensitivity derived from satellite observations compared well with numerical modelling results. A rare direct comparison between InSAR, total station measurements, and finite element simulation for a curved highway bridge has more recently been provided by Lorenz et al. (2024), confirming the technique's potential for blind monitoring of bridge clusters while emphasising that reliable interpretation requires structural pre-knowledge. The explicit use of structural geometry to constrain the interferometric analysis itself was introduced by Qin et al. (2018), who developed a structure-driven multi-temporal DInSAR approach to map surface deformation and thermal dilation of arch bridges (Figure 2. 23).




Figure 2. 23. Surface deformation map (upper panel) and thermal dilation map (lower panel) of the Lupu Bridge, in Shanghai, China, derived from multi-temporal SAR interferometric processing.

For the Italian context, the sufficiency of X-band missions such as COSMO-SkyMed to support bridge-scale analysis has been confirmed by Bonaldo et al. (2023) and Biondi et al. (2020), who established that the scatterer density on man-made structural elements is adequate for the retrieval of deformation signals at the scale of individual spans. An automated clustering algorithm for PS time series on bridges has been developed by Caspani et al. (2026) and tested on the Po River Bridge and the Colle Isarco Viaduct using user-driven COSMO-SkyMed processing, demonstrating that structurally relevant deformation patterns can be identified with minimal user input (Figure 2. 24). Metrological validation has been provided by Giordano et al. (2025) through the comparison of InSAR-derived thermal displacements with an on-site static monitoring system on a multi-span prestressed concrete bridge, with the reconstructed structural displacements closely matching the in-situ record.




Figure 2. 24. Distribution of the PSs identified in the case study area, colour-coded according to their temporal coherence in panel (a) and to their mean LOS displacement velocity over the full observation period in panel (b). The juxtaposition of the two panels illustrates how regions of low coherence may coincide with regions of apparent high velocity, highlighting the role of coherence as a quality filter that conditions the interpretive value of the velocity field.

2.3.4. Urban buildings and underground construction effects

Milillo et al. (2018) presented one of the most frequently cited building-scale InSAR studies: the monitoring of structural damage caused by the Crossrail tunnelling project in London. Using high-resolution COSMO-SkyMed data, the study demonstrated that MT-InSAR could detect differential settlements at the scale of individual buildings and identify structures affected by the excavation before visible damage appeared (Figure 2. 25). The methodology bridged MT-InSAR with the relative stiffness method for tunnel-induced subsidence damage assessment, showing that remote sensing and structural engineering techniques can be combined to quantify building-specific vulnerability. Delo et al. (2022) applied a similar approach to monitor the effects of metro line construction in Rome, confirming that InSAR can capture ground and building deformation induced by tunnelling. Weissgerber et al. (2017) investigated the three-dimensional monitoring capabilities of high-resolution TerraSAR-X data for individual buildings in Paris. The study combined InSAR, DInSAR, and polarimetric SAR to extract height, displacement, and scattering information at the building scale.




Figure 2. 25. Schematic representation of the MT-InSAR and relative stiffness method for the integrated assessment of tunnelling-induced damage on surface buildings, combining MT-InSAR displacement observations with the estimation of the relative stiffness between structure and ground.

2.3.5. Portfolio-level screening and prioritisation

At a larger spatial scale, several studies have explored the use of InSAR for screening and prioritisation of infrastructure assets, with parallel methodological developments at Politecnico di Torino addressing the use of interferometric satellite data for the structural health monitoring of bridges and infrastructures (Coccimiglio et al., 2024). The approach is attractive because it exploits the wide-area coverage of satellite data to identify structures that may require closer investigation, without dedicating instrumentation to each individual asset. Calò et al. (2023) and Nettis et al. (2023) developed methodologies for combining InSAR-derived velocity maps with structural inventory data and fragility-based indicators. The methodologies were applied to screen large bridge portfolios in southern Italy (Figure 2. 26). Calò et al. (2024) further extended this approach through the BMS-InSAR platform, integrating satellite data with automated classifiers for portfolio-level displacement analysis. Guzmán-Acevedo et al. (2024) proposed a risk-based management framework incorporating satellite products into decision-support pipelines for bridge portfolios. Talledo et al. (2022) contributed a systematic analysis of the

potential and limitations of satellite radar interferometry for structural assessment and monitoring. The study used COSMO-SkyMed SBAS products over the city of Rome and discussed practical approaches for both preliminary and detailed evaluation of settlement-related structural damage.




Figure 2. 26. Highway network in the Rome metropolitan area analysed as case study, with interferometric displacement velocity maps of the PSs for two selected segments affected by subsidence. The two highlighted segments exhibit cumulative LOS displacements consistent with localised settlement, demonstrating the use of corridor-scale InSAR analysis for the prioritisation of infrastructure inspection along extended linear assets.

2.3.6. Integration with in-situ monitoring and numerical models

A growing number of studies have investigated the combination of InSAR data with contact-based monitoring systems and numerical structural models, moving beyond the use of satellite data as a standalone diagnostic tool. Satellite-derived displacement data have been combined with on-site vibrational measurements for the monitoring of the Ponte della Musica in Rome by Ponzo et al. (2021), showing that the two approaches capture complementary aspects of structural response (Figure 2. 27). A methodology for monitoring building performance through finite element analysis and machine learning, which uses InSAR-acquired displacement data as input for model calibration on the Vittoriano monument in Rome, has been proposed by Dabiri et al. (2025). The metrological performance of InSAR for road bridges was quantified by Tonelli, Caspani, et al. (2023), who also proposed SHM frameworks that explicitly incorporate satellite products into decision-support pipelines. Within the broader question of reliable engineering use, a systematic discussion of the problems and perspectives of satellite interferometric data for

SHM in seismic areas has been provided by Lenticchia et al. (2021), identifying the principal open challenges. The integration of satellite geophysical data with continuous on-site dynamic measurements for the monitoring of dynamic parameters of civil structures has been demonstrated by Coccimiglio, Miraglia, et al. (2022) on a monitored monumental building, providing a methodological reference for the joint use of remote and in-situ observables in heritage SHM that the present thesis builds upon in the *Vicoforte* case study of Section 5.4.




Figure 2. 27. Three-dimensional geological model of the Ponte della Musica subsoil, view from north-west (upper panel), and mean deformation velocity maps of the bridge for the ascending orbit (a) and descending orbit (b) acquisitions in false colours.

2.3.7. Retrospective forensic analysis of structural failures

A distinct and rapidly growing domain of application concerns the retrospective forensic analysis of structural failures using archival SAR data. The proposition underlying this use is that the continuously accumulating satellite radar archive

constitutes, in effect, an unintentional displacement record of every structure within the sensor's coverage area. When a structure fails, this archive can be retrieved and processed interferometrically to reconstruct the displacement history of the failed element over the years or decades preceding the event, without the structure having been instrumented or monitored at the time. Milillo et al. (2019) applied this approach to the *Polcevera Viaduct* (*Morandi Bridge*, Figure 2. 28) in Genoa after its collapse in August 2018. The study reported increasing relative displacements near the collapsed pier from COSMO-SkyMed data spanning 2003-2018 (Figure 2. 29). Lanari et al. (2020) reprocessed the same dataset with different techniques and found no evidence of the reported pre-collapse anomalies. The resulting scientific exchange remains unresolved and exposes the methodological sensitivity of forensic InSAR to processing choices (Figure 2. 30).





Figure 2. 28. The Morandi Bridge before the 14 August 2018 collapse. Panel (A): two-dimensional structural scheme showing, from west to east, two single piers (1 and 2), six V-shaped piers (3 to 8), and three independent balanced systems (9 to 11), with dimensions in metres. Panel (B): three-dimensional view of the bridge.






Figure 2. 30. Mt-InSAR analysis of the Morandi Bridge focused on the area of the collapsed pillar. Panels (a) and (b): SBAS and TomoSAR processing results from the ascending dataset. Panels (c) and (d): SBAS and TomoSAR processing results from the descending dataset.

Beyond bridges, forensic retrospective analysis has been applied to dam and tailings facility failures with stronger results. Grebby et al. (2021) demonstrated that the failure date of the Brumadinho tailings dam in Brazil could have been predicted approximately forty days in advance from Sentinel-1 data alone. The prediction used the inverse velocity method on accelerating deformation trends detected in the months preceding the January 2019 collapse. Tavakkoliestahbanati et al. (2024) detected pre-collapse deformation at the Kakhovka Dam in Ukraine beginning two years before its June 2023 destruction. For buildings, the collapse of Champlain Towers South in Surfside, Florida, in June 2021 became a prominent test case: earlier work by Fiaschi & Wdowinski (2020) had detected localised subsidence at the building site from 1993-1999 ERS data, but post-collapse reanalysis of Sentinel-1 data from 2016-2021 found no detectable displacement signal, illustrating the geometric and resolution constraints that limit building-scale forensic analysis with C-band sensors. The NIST investigation of the Surfside collapse formally incorporated InSAR processing but concluded that no sinking or settling was detectable in the five years before the event. Despite this growing body of work,

the forensic use of InSAR on civil structures lacks a systematic methodological framework, has not been tested in adversarial legal proceedings, and remains constrained by the epistemological gap between measured displacement and structural causation. The present thesis contributes to this emerging domain through the retrospective collapse analyses of Chapter 5, which develop the TS concept as a forensic dating tool and address the limits of pre-collapse attribution explicitly.

2.3.8. The ReLUIIS guidelines, standardised products and the interpretation gap

The accumulated experience described above has remained dispersed across the remote sensing, geotechnical, and structural engineering literatures. The most comprehensive attempt to date to consolidate this experience into an operational framework for structural engineers is the set of guidelines produced by the Italian ReLUIIS network in collaboration with the Italian Department of Civil Protection and the CNR-IREA (Prota et al., 2025). The document, titled “*Linee guida per l'utilizzo dei dati interferometrici satellitari ai fini dell'interpretazione del comportamento strutturale delle costruzioni*”, represents the first systematic guideline addressed specifically to the structural engineering community.

The guidelines are organised around four main topics. The first part presents the fundamentals of SAR interferometry that a structural engineer needs to understand before handling the data: image formation, available sensors and constellations, together with the differential and multi-temporal interferometric techniques. The second part addresses data processing aspects relevant to the end user, including the use of single and dual-geometry datasets, temporal and spatial resampling, spatial interpolation techniques, frequency-domain analysis, alongside the integration of InSAR data with GNSS observations. The third part, which constitutes the most original contribution of the document, deals directly with the structural interpretation of satellite-derived displacements. It proposes a classification of the actions that can produce measurable displacements on buildings: soil-structure interaction effects such as foundation settlements and consolidation, including bearing capacity problems where applicable; ground movements independent of the structure, including subsidence, landslides, seismic secondary effects, and mining-related deformations; water-related actions such as infiltration, scouring, and groundwater fluctuations; thermal and environmental effects; and ageing and degradation processes. The classification is significant because it provides the structural engineer with an explicit framework for formulating hypotheses about the origin of observed displacements, rather than

treating the InSAR output as a generic deformation signal. The fourth part of the guidelines illustrates these concepts through six case studies covering reinforced concrete buildings, masonry structures, bridges, and heritage assets, with contributions from several Italian universities including the research group at Politecnico di Torino.

A distinctive feature of the ReLUIIS guidelines is the emphasis on the limitations of the interpretation process itself. The document explicitly warns that the analysis of satellite data, if conducted incorrectly or over an inadequate time horizon, can lead to erroneous assessments of structural behaviour. It recommends that the interpreter always consider the geological and environmental context in which the structure is embedded, and that displacements observed in time series be attributed to specific physical causes only after excluding alternative explanations.

Beyond the Italian context, a comparable level of codification has not yet been achieved at the European or international scale. The EGMS, operated within the Copernicus Land Monitoring Service, provides standardised interferometric products over the entire European territory (Capes et al., 2023; Costantini et al., 2021). The EGMS product user manual describes data formats, reference frames, and quality indicators, but it does not include guidance on how the displacement products should be interpreted for structural engineering purposes. The responsibility for interpretation remains entirely with the user, and the manual explicitly states that the products are not intended for structure-specific assessment without additional processing and validation. The gap between data availability and interpretation capability that the ReLUIIS guidelines begin to address at the Italian national level therefore persists at the continental scale and constitutes one of the motivations for the present thesis.

2.3.9. Remaining gaps and positioning of the thesis

Despite the extensive and growing body of work reviewed in this section, several gaps remain in the literature. Most published studies focus on either regional-scale screening or single-structure case studies, and few propose a transferable methodology that links data acquisition, processing, quality control, and structural interpretation within a unified engineering framework. The separation between thermal, seasonal, and structural displacement components is often addressed through generic trend fitting, without explicit consideration of structural plausibility or uncertainty propagation. The role of temporal coherence as a modelling variable, rather than as a simple binary quality filter, has received limited attention. The

comparison between standardised products and user-driven processing for the same structure has rarely been performed quantitatively across the three axes of density, quality, and interpretability. The integration of satellite-derived environmental variables into data-driven displacement models, which could extend the applicability of satellite monitoring to structures lacking in-situ environmental sensors, has been explored in only a handful of studies. The diagnostic potential of TS for constraining the chronology and spatial extent of structural failures has been observed in individual case studies but has not been formalised as a systematic analytical category. The fusion of InSAR time-series information with three-dimensional geometric change data from LiDAR for post-event structural assessment at the building scale has not been attempted in the published literature.

The present thesis builds on this evolving landscape. The methodological framework developed in Chapters 4 and 5 draws on the principles established by the ReLUIIS guidelines and on the experience of the research groups cited in this section, while extending the approach toward a more explicit integration of coherence-aware quality control, multi-source data harmonisation, TS diagnostics, multi-sensor fusion, and decision-oriented interpretation at the building and infrastructure scale.

All multi-temporal interferometric processing reported in the three completed case studies and in the ongoing Palisades application was conducted by the author through a fully user-driven workflow operating directly on SLC data, with explicit control over master image selection, interferometric network construction, reference point definition, amplitude and coherence thresholds, atmospheric phase screen estimation, and geocoding parameters. The processed datasets comprise the 169-scene Sentinel-1 descending stack over the *Santissima Trinità* in Turin, the dual-orbit Sentinel-1 peri-collapse archives over the *Himera Viaduct*, the unified 2011-2022 dual-orbit COSMO-SkyMed Stripmap archive over the *Albiano-Magra Bridge*, the 62-scene COSMO-SkyMed Second Generation small-area stack over the *Sanctuary of Vicoforte*, and the dual-orbit Sentinel-1 archive comprising 70 ascending and 73 descending acquisitions over the Pacific Palisades study area. Standardised EGMS products are employed throughout the thesis exclusively as a comparative baseline and as a portfolio-level screening layer, never as a primary diagnostic source for structure-specific interpretation.

Chapter 3

Interferometric SAR data and processing workflow

3.1. Principles of InSAR and multi-temporal InSAR

Interferometric measurement of surface displacement from satellite SAR data rests on two complementary principles. The first is the extraction of range-change information from the phase difference between two co-registered SAR images acquired at different times, which constitutes the basic operation of differential SAR interferometry addressed in Section 3.1.1. The second is the use of multi-temporal image stacks to separate the displacement signal from atmospheric and geometric artefacts that contaminate single-pair interferograms, which motivates the persistent-scatterer and small-baseline formulations developed in Sections 3.1.2 and 3.1.3. This single-pair approach is subject to two fundamental limitations: temporal and geometric decorrelation, which degrades the phase quality over surfaces that change between acquisitions or are observed from significantly different orbital positions, and the atmospheric delay difference between the two acquisition epochs, which can mimic or obscure ground displacement signals. Both limitations motivated the development of MT-InSAR methods, which exploit stacks of tens to hundreds of images to separate the displacement signal from competing phase contributions and to produce temporally dense deformation time series.

Two principal families of methods have emerged from this effort. The PSI family, described in Section 3.1.2, targets individual stable coherent pixels and applies parametric estimation over the full temporal stack to recover displacement, residual height, and atmospheric phase simultaneously. The SBAS approach, described in Section 3.1.3, applies spatial multi-looking to interferograms formed from image pairs with limited orbital separation, recovering displacement time series on Distributed Scatterers (DS) by network inversion. Both families are relevant to structural monitoring applications, and their relative merits depend on the resolution, coherence, and geometry of the target structure.

3.1.1. DInSAR

The SLC product delivered by modern SAR missions represents the observed scene through two coupled fields, the amplitude and the phase, whose distinct informational content is illustrated in Figure 3. 1 and Figure 3. 2. The amplitude image reproduces the backscattered intensity of the scene in a form that is readily interpretable by visual inspection, as shown in the bottom-left panel. The phase image, by contrast, appears as a quasi-random pattern of values wrapped within the interval from $-\pi$ to $+\pi$, with no directly interpretable spatial structure, as displayed in the bottom-right panel. The top panel provides the geometrical interpretation of the SAR signal as an oscillating wave characterised by amplitude and phase, which underlies the interferometric processing developed in the following sections.




Figure 3. 1. Geometrical representation of the amplitude and phase of a SAR wave.




Figure 3. 2. Panel (a): amplitude image of a representative scene, with brightness proportional to the backscattered intensity Panel (b): wrapped phase image of the same scene, with values in the interval from $-\pi$ to $+\pi$ and no immediate physical interpretability prior to interferometric processing. Adapted from (Balaji, 2024).

Repeat-pass SAR interferometry, commonly referred to as InSAR, estimates surface deformation by comparing the phase of two complex SAR acquisitions acquired over the same area from similar orbital positions at different times. The formation of a SAR interferogram requires two acquisitions of the same target from orbital positions that are close but not coincident. The spatial separation between the two positions is known as the interferometric baseline, and its geometric decomposition is illustrated in Figure 3. 3. The baseline vector B connects the two orbital positions $S1$ and $S2$, and can be decomposed into a component perpendicular to the LOS, denoted $B_{\perp}$, and a component parallel to the LOS, denoted $B_{//}$. The distinction between the two components is not merely geometric but carries direct physical significance for the interferometric measurement. The perpendicular component $B_{\perp}$ controls the sensitivity of the interferometric phase to topographic relief, and its value affects the precision with which residual height and displacement can be separated in the subsequent inversion. The parallel component $B_{//}$ contributes instead to a phase term that is constant over the scene and is

typically absorbed in the reference phase. The along-track direction Y and the cross-track direction X shown in Figure 3. 3 correspond respectively to the azimuth direction, along which the synthetic aperture is formed, and to the range direction, along which the side-looking geometry develops the ground swath.




Figure 3. 3. Geometry of the SAR interferometric baseline. Two satellite passes image the same target from orbital positions S1 and S2, separated by a baseline B decomposed into its perpendicular and parallel components with respect to the LOS. Adapted from (Alaska Satellite Facility, 2026b).

The complex interferogram is formed by multiplying the SLC image of the reference acquisition by the complex conjugate of the second, as expressed in Equation 3. 1.

$$I(x, y) = C_1(x, y) \cdot C_2^*(x, y) \quad 3. 1$$

where C_1 and C_2 denote the co-registered SLC images of the reference (master) and repeat (slave) acquisitions, the asterisk indicates complex conjugation, and $I(x,y)$ is the resulting interferogram at pixel location (x,y) . The phase of the interferogram encodes the difference in sensor-to-target range between the two acquisition dates, combined with several additional contributions whose separation

is the central problem of interferometric processing. A representative example of an interferogram is shown in Figure 3. 4, which reports the ascending and descending interferograms of the 1999 Hector Mine earthquake in California. The concentric coloured bands visible around the epicentre are known as fringes, and correspond to successive cycles of the wrapped interferometric phase between $-\pi$, $+\pi$. Each fringe represents a LOS displacement equal to half the wavelength of the SAR signal, and the local density of fringes therefore quantifies the spatial gradient of the ground deformation between the two acquisition dates.




Figure 3. 4. Descending (a) and ascending (b) interferograms of the 1999 Mw 7.1 Hector Mine earthquake in California. Each colour cycle represents ten centimetres of LOS displacement towards or away from the satellite. Purple plus and minus signs indicate areas of positive and negative LOS displacement; the mapped fault trace is shown as a thick line and the epicentre is marked with a star. Adapted from (Jónsson, 2002).

Before examining the structure of the interferometric phase in detail, it is useful to clarify the physical meaning of the phase measurement itself and the reason why interferometric products are intrinsically differential. The phase of a radar echo encodes the fractional part of the two-way path length between the antenna and the target, expressed in units of the radar wavelength. For a target located at range R_T from the sensor, the one-way phase accumulation along the path is shown in Equation 3. 2:

$$\varphi_{ow} = \frac{2\pi}{\lambda} \cdot R_T \quad 3.2$$

and the two-way phase accumulation for the round trip of the radar signal is (Equation 3. 3):

$$\varphi_{tw} = \frac{4\pi}{\lambda} \cdot R_T \quad 3.3$$

The phase actually measured by the receiver, however, is confined to the interval $[0, 2\pi]$ or equivalently $[-\pi, +\pi]$, because only the fractional part of the complete phase cycle can be recovered from the sampled complex signal. The integer number of complete cycles accumulated along the path is lost. For a Sentinel-1 observation at C-band with $\lambda = 5.55$ cm and a slant range of approximately 800 km, the two-way path contains on the order of 2.9×10^{10} complete phase cycles; the measured phase retains no information about which of these billions of cycles is being observed. The physical situation is analogous to reading the angular position of a single blade on a rapidly spinning fan: when the fan stops, the blade position within one revolution is unambiguous, but the total number of revolutions completed during the observation is unknown. Two direct consequences follow from this limitation. First, interferometric displacement measurements are inherently differential, because only the change in fractional phase between two acquisitions carries useful information and the absolute phase is not recoverable. Second, the maximum unambiguous displacement that can be measured between two consecutive acquisitions is $\lambda/2$, which corresponds to a full 2π phase cycle along the two-way path; displacements exceeding this bound produce phase wrapping, as discussed in relation to Equation 3. 5 below.

The total interferometric phase at a generic pixel, in the wrapped form actually observed from $I(x,y)$, can be decomposed into five physically distinct components, each governed by a different set of acquisition parameters. The decomposition adopted throughout this thesis follows the standard formulation of Hanssen (2001) and is expressed by Equation 3. 4:

$$\varphi = \varphi_{flat} + \varphi_{topo} + \varphi_{def} + \varphi_{atm} + \varphi_{noise} \quad 3.4$$

where φ_{flat} is the flat-earth phase, φ_{topo} is the topographic phase, φ_{def} is the deformation phase, φ_{atm} is the atmospheric phase, and φ_{noise} is the residual noise phase including orbital residuals and decorrelation effects. Each of these five components has an analytical form that depends on specific physical quantities, as detailed below.

The flat-earth phase φ_{flat} arises from the curvature of the Earth and from the geometric difference between the two orbital positions at acquisition. For two targets at the same elevation but at different positions along the swath, the path length difference between master and slave varies systematically with the across-track coordinate because the two antennas observe each target from slightly different angles. This systematic variation generates a regular fringe pattern that would dominate the interferogram if not removed. The flat-earth phase is entirely determined by the orbital state vectors and the reference ellipsoid, and is removed during the interferogram formation using the precise orbital information described in Section 3.3.1.

The topographic phase φ_{topo} is proportional to the perpendicular baseline between the two acquisitions and to the target elevation above the reference surface, as established in Equation 2. 4 of Chapter 2. After flat-earth removal, the topographic contribution remains visible as a fringe pattern correlated with the local topography. The component is removed by simulating the expected phase contribution from an external DEM and subtracting the simulated field from the interferogram, as described in Section 3.3.3. The residual topographic phase that remains after DEM-based correction is proportional to the DEM error and represents one of the contributions to be estimated in the multi-temporal inversion of Section 3.1.2.

The deformation phase φ_{def} is proportional to the LOS displacement of the target between the two acquisition dates and constitutes the signal of interest for structural monitoring. The relationship between deformation phase and LOS displacement is given by Equation 3. 5, consistent with the sign convention introduced in Equation 2. 1 of Chapter 2:

$$\varphi_{def} = \frac{4\pi}{\lambda} \cdot d_{LOS} \quad 3.5$$

where φ_{def} is the deformation-induced phase in radians, λ is the radar wavelength, and d_{LOS} is the LOS displacement measured positive in the direction from target to sensor. The factor $4\pi/\lambda$ reflects the two-way travel path of the radar signal. One complete phase cycle of 2π in the wrapped interferogram therefore corresponds to a LOS displacement equal to $\lambda/2$, which defines the ambiguity interval of the measurement. For Sentinel-1 operating at C-band with $\lambda \approx 5.55$ cm, one interferometric fringe corresponds to approximately 2.78 cm of LOS motion;

for COSMO-SkyMed at X-band with $\lambda \approx 3.1$ cm, one fringe represents approximately 1.56 cm.

The atmospheric phase φ_{atm} arises from the spatially and temporally variable refractive index of the troposphere along the signal propagation path, as discussed in Section 2.2.3. Differential tropospheric delay between the two acquisition epochs introduces a phase pattern that is spatially correlated but cannot be distinguished from genuine displacement in a single interferogram without external information. In MT-InSAR approaches, the atmospheric contribution is estimated and removed by exploiting the redundancy of multiple interferograms, as developed in Sections 3.1.2 and 3.1.3.

The noise phase φ_{noise} collects all residual contributions that do not fit the four previous categories, including residual orbital errors not absorbed by the flat-earth correction, thermal noise of the radar receiver, and temporal and geometric decorrelation effects. Residual orbital errors produce long-wavelength trends across the scene and are typically corrected in post-processing by referencing the displacement field to ground truth, as described in Section 3.5.2. Decorrelation noise is intrinsically random and is suppressed statistically by the ensemble averaging that underlies multi-temporal approaches.

Once the flat-earth and topographic contributions have been removed through orbital information and the external DEM, the residual phase contains the deformation signal, the atmospheric screen, residual orbital and decorrelation noise, and any topographic error induced by DEM inaccuracies. In the single-pair DInSAR configuration, these residual contributions cannot be separated from one another without additional observations or external constraints. A fringe pattern in the residual phase is equally consistent with a surface displacement of $\lambda/2$ per fringe, with a DEM error equal to the height of ambiguity, or with a localised atmospheric anomaly of comparable phase magnitude. The analyst cannot determine which interpretation is correct without independent evidence.

The sequential effect of the flat-earth and topographic phase removal on a representative interferogram is illustrated in Figure 3. 5. Panel (a) shows the DEM of the area, which provides the topographic reference against which the corresponding phase contribution is simulated and removed. Panel (b) shows the raw interferogram, in which the fringe pattern is dominated by the flat-earth component and no structure of the underlying deformation can be recognised. Panel (c) shows the same interferogram after removal of the flat-earth phase, where the

residual fringes are correlated with the topographic relief. Panel (d) shows the differential interferogram obtained after removal of both systematic contributions, which isolates the displacement signal together with the atmospheric and noise components discussed in the following sections. The progression from panel (b) to panel (d) demonstrates that the interpretability of the interferometric signal in terms of ground deformation depends critically on the quality of the ancillary information, namely the orbital state vectors that define the flat-earth phase and the DEM that defines the topographic phase.




Figure 3. 5. Effect of the flat-earth and topographic phase removal on a representative interferogram. Panel (a): Digital Elevation Model of the area of interest. Panel (b): raw interferogram before the removal of any systematic phase contribution. Panel (c): interferogram after removal of the flat-earth phase, with residual fringes reflecting the topographic contribution. Panel (d): differential interferogram after removal of both the

flat-earth and the topographic phase, retaining the displacement, atmospheric, and noise contributions.

The fundamental property that enables the multi-temporal separation of these contributions is that each component depends on the interferometric geometry in a different way. The residual topographic phase is proportional to the perpendicular baseline: the same physical relief produces more fringes when observed with a larger baseline and fewer fringes when observed with a smaller one. The deformation phase depends only on the wavelength and on the LOS displacement magnitude and is independent of the perpendicular baseline. The atmospheric contribution varies randomly from one interferometric pair to another but remains spatially correlated within each pair. These three distinct dependencies, respectively on baseline geometry, on time, and on spatial position, constitute the physical basis for the multi-temporal approaches developed in Sections 3.1.2 and 3.1.3.

The dependence of the topographic phase sensitivity on the perpendicular baseline is illustrated in Figure 3. 6, which reports two interferograms of the Mount Vesuvius area obtained with substantially different perpendicular baselines. We deliberately retained the topographic phase contribution in both interferograms, without subtracting the external DEM, in order to make the dependence of the fringe density on the baseline directly visible. Had the topographic phase been removed, the fringes associated with the relief of the volcano would have been suppressed, leaving only the residual signal associated with surface deformation and atmospheric contributions. The interferogram with a 46 m perpendicular baseline produces a sparse fringe pattern, because a single cycle of the wrapped phase corresponds to a height difference of approximately 351.8 m. The interferogram with a 178 m perpendicular baseline produces a substantially denser fringe pattern, because the same cycle corresponds to a height difference of approximately 93.4 m. The near-fourfold increase in fringe density between the two panels is the direct visual counterpart of the linear scaling of the topographic phase with perpendicular baseline established by Equation 2. 4, and underlies the multi-temporal separation strategy developed in the following sections.




Figure 3. 6. Two interferograms of the Vesuvius area computed with different perpendicular baselines and presented without subtraction of the external DEM, in order to expose the dependence of the topographic fringe density on the baseline. Panel (a): interferogram with a 46 m perpendicular baseline, corresponding to a height of ambiguity of approximately 351.8 m. Panel (b): interferogram with a 178 m perpendicular baseline, corresponding to a height of ambiguity of approximately 93.4 m. The denser fringe pattern in panel (b), with approximately four times the fringe density of panel (a), reflects the linear dependence of the topographic phase sensitivity on the perpendicular baseline established by Equation 2. 4.

Phase measurements in a differential interferogram are inherently wrapped in the interval $[-\pi, +\pi]$, because the interferometric phase is the argument of a complex number and is physically indistinguishable from the same value plus any integer multiple of 2π . For spatial patterns whose true phase gradient between adjacent pixels exceeds $\pm\pi$, the observed interferogram shows abrupt jumps of 2π that do not correspond to real discontinuities of the underlying displacement field but to the wrapping operation. The recovery of the continuous unwrapped phase from the wrapped observation is known as phase unwrapping and is a non-trivial operation whose reliability depends on the spatial sampling of the phase gradient and on the quality of the interferogram. Several algorithms have been developed to address the two-dimensional unwrapping problem, including the branch-cut method of Goldstein et al. (1988) and the network-flow formulation of Chen & Zebker (2000). In DInSAR applications to slow deformation processes, phase unwrapping is often unnecessary because the total phase between acquisitions does not exceed the $\pm\pi$ bound, but the limitation becomes critical for rapid events such as earthquake surface ruptures or for structures with strong displacement gradients over short

distances. A schematic illustration of the phase unwrapping problem in interferometric SAR processing is shown in Figure 3. 7.




Figure 3. 7. Phase unwrapping basic principle schematic. Adapted from (W. Yang et al., 2024).

Single-pair DInSAR is most reliable for large, rapid deformation events such as co-seismic ruptures or volcanic unrest, where the displacement signal amplitude greatly exceeds the atmospheric and noise contributions. Atmospheric delays alone can produce apparent displacements of several millimetres to more than a centimetre, as discussed in Section 2.2.3. For the slow, millimetre-scale processes typical of structural and urban monitoring addressed in this thesis, single-pair DInSAR is rarely sufficient. Multi-temporal approaches address this limitation by exploiting the statistical properties of large image stacks to separate the deformation signal from the competing phase contributions. The two principal families of multi-temporal methods are presented in the following sections: PSI in Section 3.1.2 and SBAS in Section 3.1.3.

3.1.2. PSI and PS-based approaches

The differential interferometric measurement described in Section 3.1.1 assumes that the phase at each pixel can be meaningfully compared between two acquisitions. In practice this assumption holds only under restrictive conditions on the scattering environment within the resolution cell, and understanding those conditions requires a closer examination of how multiple scattering contributions combine within a single pixel.

Each resolution cell in a SAR image covers a ground area determined by the range resolution (Equation 2. 2) and the azimuth resolution (Equation 2. 3). For Sentinel-1 IW mode, the single-look resolution cell measures approximately 5 m in

range by 20 m in azimuth; for COSMO-SkyMed Stripmap HIMAGE, approximately 3 m by 3 m. Within each cell, the recorded signal is the coherent sum of the complex contributions from all scattering elements present. For a cell containing two scatterers of amplitude A_1 and A_2 located at ranges R_{T1} and R_{T2} along the LOS, the total signal is expressed by Equation 3. 6:

$$S = A_1 e^{j\varphi_1} + A_2 e^{j\varphi_2} \quad 3.6$$

Where (Equation 3. 7):

$$\varphi_k = \frac{4\pi}{\lambda} R_{Tk} \quad for \quad k = 1,2 \quad 3.7$$


The magnitude and phase of this sum depend on the phase difference, as shown in Equation 3. 8):

$$\Delta\varphi = \varphi_1 - \varphi_2 = \frac{4\pi}{\lambda} \cdot (R_{T1} - R_{T2}) \quad 3.8$$

which is governed by the difference in LOS distance between the two scatterers. When $\Delta\varphi$ equals an even multiple of π , the two vectors align constructively, and the resultant amplitude reaches its maximum $|S| = A_1 + A_2$. When $\Delta\varphi$ equals an odd multiple of π , the vectors oppose, and the amplitude reaches its minimum $|S| = |A_1 - A_2|$. For two scatterers of equal amplitude A , the two extrema correspond respectively to $2A$ and zero. For C-band at $\lambda \approx 5.55$ cm, the transition from constructive to destructive interference occurs over an LOS distance difference of just 2.78 cm. Within a resolution cell that extends over several metres of ground, the LOS distance difference between two scattering elements is determined by their three-dimensional positions relative to the satellite viewing geometry and cannot be known a priori.

The resulting pixel amplitude and phase are therefore deterministic but unpredictable, a condition that produces the characteristic salt-and-pepper texture of SAR imagery known as speckle. Speckle is not noise in the conventional sense. It is the deterministic consequence of coherent summation within the resolution cell, and it varies between acquisitions because the satellite flies along slightly different positions within its nominal orbital envelope, altering the LOS projection of each scatterer and therefore modifying the vector sum of Equation 3. 6. The

temporal variability of speckle is the physical origin of the decorrelation discussed in Section 2.2.3: when the vector sum within a cell changes unpredictably between acquisitions, the phase difference between the two images carries no stable information about surface displacement. The visual impact of speckle, and the corresponding gain that multi-temporal averaging affords on the interpretability of the radar reflectivity, is illustrated by comparing two reflectivity maps of the same urban area in Turin computed from a small subset and from the full Sentinel-1 stack. The salt-and-pepper texture that dominates the single-pair reflectivity map progressively dissolves as additional acquisitions are averaged into the temporal mean, revealing the stable backscattering structure of the urban fabric beneath, as shown in Figure 3. 8.



(a)



(b)

Figure 3. 8. Effect of multi-temporal averaging on the radar reflectivity map of an urban area in Turin, illustrating the progressive reduction of speckle as the number of averaged acquisitions increases. Panel (a) reports the reflectivity map computed from five Sentinel-1 acquisitions, in which the salt-and-pepper texture characteristic of speckle dominates the visual appearance and obscures the structural detail of the urban fabric. Panel (b) reports the reflectivity map computed from 176 Sentinel-1 acquisitions over the same area, in which the temporal averaging suppresses the speckle and recovers the stable scattering signature of buildings, road surfaces, and built infrastructure.

PSI was originally developed at Politecnico di Milano and has since become the standard technique for millimetre-scale deformation monitoring in urban and infrastructural environments (Crosetto et al., 2016; Ferretti et al., 2001). PSI targets resolution cells that maintain phase coherence across the entire image stack, irrespective of perpendicular baseline. These cells, termed Persistent Scatterers (PS), are dominated by a single stable electromagnetic reflector, such as a corner reflector, a metallic structure, or a rigid building element, whose backscatter properties do not change substantially over time. Man-made structures are particularly rich in PS candidates, which is one of the principal reasons for the suitability of PSI to the structural monitoring context addressed in this thesis.

For each PS candidate p and each differential interferogram i , the wrapped differential phase $\Delta\varphi_i(p)$ relative to a reference point p_0 follows the decomposition

of Equation 3. 4 introduced in Section 3.1.1. After the removal of the flat-earth contribution through orbital information and of the bulk topographic contribution through the external DEM, the double-differential phase, that is the phase difference between p and p_0 for interferogram i , retains four contributions expressed by Equation 3. 9:

$$\Delta\varphi_i(p) = \varphi_i^H(p) + \varphi_i^D(p) + \varphi_i^{APS}(p) + \varphi_i^{noise}(p) \quad 3. 9$$

where φ_i^H is the residual topographic phase due to residual height error with respect to the DEM, φ_i^D is the deformation phase accumulated between master and slave acquisitions, φ_i^{APS} is the differential atmospheric phase screen between the two epochs, and φ_i^{noise} is the residual noise phase that collects orbital residuals, decorrelation effects, and thermal noise of the receiver. Equation 3. 9 is the specialisation of Equation 3. 4 to the double-differential formulation adopted in the POLIMI PS approach (Perissin et al., 2011; Perissin & Rocca, 2006).

An operational distinction between the differential interferometry described in Section 3.1.1 and the PS-based estimation described here must be stated explicitly to avoid a common conceptual confusion. In conventional DInSAR, an interferogram is a two-dimensional spatial map, in which the phase difference is computed for every pixel in the co-registered image pair, producing a raster product that can be filtered, unwrapped, and interpreted as a continuous deformation or topographic field. In the PS-based approach, by contrast, the phase difference is computed only at the locations of the selected candidate points. For each PS candidate p and each slave image i , the PS-based estimation proceeds in three steps. The algorithm reads the complex value of the pixel at the range-azimuth position of p in the slave image. It then computes the phase difference with respect to the corresponding pixel at the same position in the master image. The phase of the reference point p_0 , computed in the same way for the same interferogram i , is finally subtracted. The operation is performed pixel by pixel in the temporal dimension defined by the interferometric network, which connects each slave to the master through the star graph introduced in Section 3.3.1, and produces a set of M scalar double-differential phases $\Delta\varphi_i(p)$ for each candidate point. No spatial interferogram is generated as a two-dimensional raster, because the PS-based approach operates on the time series of phase values at each candidate point rather than on the continuous phase field across the scene. No spatial interferogram is generated: the operation produces a set of M scalar phase values for each candidate point, not a two-dimensional raster. The absence of spatial filtering in this process

is deliberate, because spatial filtering would destroy the point-like information that distinguishes a PS from the surrounding DS. The PS estimation preserves the full spatial resolution of the original SAR data at the cost of restricting the measurement to discrete points rather than continuous fields.

Under the standard small-deformation, long-wavelength approximation, the residual topographic and linear deformation contributions to $\Delta\phi_i(p)$ can be modelled parametrically as functions of the interferometric baselines. The residual topographic phase scales linearly with the perpendicular baseline according to Equation 3. 10:

$$\phi_i^H(p) = \frac{4\pi}{\lambda} \cdot \frac{\Delta h_p}{R_0 \sin \theta} \cdot B_{ni} \quad 3. 10$$

where Δh_p is the residual height of the scatterer at point p with respect to the reference DEM, B_{ni} is the perpendicular baseline of interferogram i , R_0 is the reference slant range, θ is the local incidence angle, and λ is the radar wavelength. The linear deformation phase scales linearly with the temporal baseline according to Equation 3. 11:

$$\phi_i^D(p) = \frac{4\pi}{\lambda} \cdot \Delta v_p \cdot B_{ti} \quad 3. 11$$

where Δv_p is the mean LOS velocity at point p relative to the reference, B_{ti} is the temporal baseline of interferogram i . The sign of Δv_p follows the convention of Equation 3. 5, with positive values corresponding to a range-shortening displacement (motion toward the sensor) and negative values corresponding to a range-lengthening displacement (motion away from the sensor, typically associated with subsidence). Equations 3. 10 and 3. 11 together constitute the two-parameter linear model inverted in PSI processing to jointly estimate Δh_p and Δv_p for each candidate scatterer.

The structural difference between Equations 3. 10 and 3. 11 is the basis of the multi-temporal estimation. The residual height phase scales linearly with the perpendicular baseline B_{ni} : a PS with a given height error Δh_p produces a larger phase contribution in interferograms formed with distant slave images, and a smaller contribution in pairs with nearby slaves. The linear deformation phase, conversely, scales linearly with the temporal baseline B_{ti} : a PS moving at constant

velocity Δv_p accumulates more phase in interferograms spanning a longer time interval, regardless of the orbital geometry. When the differential phases from the full set of M interferograms are sorted by perpendicular baseline, the slope of the resulting trend is proportional to Δh_p . When the same phases are sorted by temporal baseline, the slope is proportional to Δv_p . Because the two independent variables B_{ni} and B_{ti} are generally uncorrelated across a well-designed interferometric network, the two parameters can be estimated jointly from the overdetermined system of observations.

The joint estimation of Δh_p and Δv_p from the M observed double-differential phases $\Delta \varphi_i(p)$ is performed through the maximisation of a complex-valued functional known as the periodogram. The periodogram is defined by Equation 3. 12:

$$\gamma_p(\Delta h, \Delta v) = \frac{1}{M} \cdot \sum_{i=1}^M e^{j[\Delta \varphi_i(p) - \varphi_i^H(\Delta h) - \varphi_i^D(\Delta v)]} \quad 3. 12$$

where the sum extends over the M interferograms, the residual phases in the exponent are computed by subtracting the modelled topographic and deformation contributions at trial values Δh and Δv , and the modulus of γ_p is the objective function to be maximised. The estimation algorithm performs a search over the two-dimensional parameter space $(\Delta h, \Delta v)$, evaluating the modulus of γ_p at each candidate pair, and selects the pair $(\Delta h_p, \Delta v_p)$ that maximises the modulus. The modulus of γ_p at the optimal pair is called the temporal coherence of the point and is denoted γ_p , as expressed by Equation 3. 13:

$$\gamma_p = |\gamma_p(\Delta h, \Delta v)| \quad 3. 13$$

The temporal coherence γ_p takes values between zero and one. Values close to unity indicate a phase-stable, reliable measurement point for which the two-parameter model explains the observed phase variability well. Values close to zero indicate that the model fails to account for the observed phase variability, either because the point is not a genuine PS, or because the physical deformation deviates substantially from the linear velocity assumption, or because the scattering properties of the target have changed during the observation period.

The formulation of the temporal coherence as the modulus of a complex average, rather than as a conventional goodness-of-fit statistic in the real domain, reflects a constraint of the interferometric measurement itself. The wrapped phase $\Delta\varphi_i$ is confined to the interval $[-\pi, \pi]$ regardless of the true accumulated phase, which may span many complete cycles over the observation period. A conventional least-squares regression in the real domain would attempt to fit a linear model to these wrapped phase values, producing erroneous results whenever the true phase trajectory crosses the $\pm\pi$ boundary. The crossing manifests as an apparent discontinuity in the observed phase, which a real-valued estimator would interpret as a residual, degrading the fit. The complex-domain formulation avoids this problem entirely, because the estimation operates on the unit-magnitude complex exponentials $e^{j\Delta\varphi_i}$ rather than on the phase values directly. The residual vectors in the exponent of Equation 3.12 rotate freely on the unit circle without encountering any boundary, and no unwrapping of the input phases is required.

The temporal coherence γ_p has a direct geometric interpretation that clarifies its role as an objective function. Each term in the sum of Equation 3.12 is a unit-length complex vector whose angle equals the phase residual after removal of the modelled topographic and deformation contributions. When the trial parameters $(\Delta h, \Delta v)$ are close to the true values, the residual phases are small and all M vectors point approximately in the same direction; their vector average has a modulus close to one. When the trial parameters are far from the true values, the residual phases are large and distributed across the full $[-\pi, \pi]$ interval; the corresponding vectors point in many different directions and their average has a modulus that tends toward zero (Figure 3.9). The periodogram therefore measures the degree of alignment of the residual phase vectors, and the search for the parameter pair that maximises γ_p is equivalent to finding the pair that best aligns all residuals in the complex plane.

Candidates with temporal coherence below a user-defined threshold are discarded from the final PS dataset. The PS selection in the POLIMI formulation additionally employs the Amplitude Dispersion D_A index of Ferretti et al. (2001), defined as the ratio of the temporal standard deviation to the temporal mean of the pixel amplitude, as a computationally inexpensive proxy for phase stability that does not require phase unwrapping. Points satisfying both the D_A threshold and the temporal coherence criterion constitute the final PS dataset used for structural interpretation. The Amplitude Dispersion and its relation to the ASI are presented formally in Section 3.4.1. The calibration of the temporal coherence threshold against the size of the interferometric stack is addressed in Section 3.4.3, where the finite-sample bias of the periodogram modulus is discussed explicitly.




Figure 3. 9. Geometric interpretation of the temporal coherence as a vector average in the complex plane. Panel (a): unit-length vectors aligned in a common direction, corresponding to small residual phases after the removal of the modelled topographic and deformation contributions, with the modulus of their average close to one. Panel (b): unit-length vectors distributed across the full interval from minus pi to pi, corresponding to large residual phases, with the modulus of their average close to zero.

The parametric model of Equations 3. 10 and 3. 11 assumes that the residual height Δh_p remains constant throughout the observation period and that the deformation can be represented by a single mean LOS velocity Δv_p . Both assumptions may be violated in practice. A scatterer whose physical height changes at some epoch within the stack, for example because of construction, demolition, or structural event, does not satisfy the height stationarity assumption and produces anomalous parameter estimates; the diagnostic implications of this condition are developed in the TS framework of Section 4.1.2. A deformation history that departs from the linear velocity assumption, such as a thermally driven periodic component or a non-linear creep, is addressed by enriching the phase model of Equation 3. 9 with additional deterministic basis functions, a strategy exploited in the time-series analytics developed in Section 4.2.

Several extensions of the original PSI algorithm have been proposed to address scenarios in which dominant point scatterers are absent or sparse. Hooper et al. (2004) developed the Stanford Method for PS, known as StaMPS, which selects pixels based on spatial phase continuity rather than on amplitude statistics alone, broadening applicability to non-urban and semi-natural environments. The QPS approach introduced by Perissin & Wang (2012), relaxes the single-master constraint of the original PSI formulation by allowing each interferogram to have its own master, effectively combining the PS and Small Baseline philosophies within a unified inversion framework. These extensions are methodologically

relevant to the broader context of multi-temporal InSAR, but the present thesis adopts the standard PSI formulation throughout the case studies, because the structural targets of interest provide abundant stable scatterers that satisfy the strict PS criterion and because the single-master star network preserves the full interpretability of the parametric inversion in terms of residual height and linear velocity. Non-linear deformation models, thermal expansion terms, and seasonal components can be incorporated by enriching the phase model of Equation 3.9 with deterministic basis functions, as developed methodologically in Section 4.2 and exploited in the case studies of Chapter 5.

The following section discusses the Small Baseline Subset approach, which addresses the multi-temporal inversion problem from a complementary perspective based on distributed scatterers rather than on dominant point reflectors.

3.1.3. SBAS and related approaches

The Small Baseline Subset approach, abbreviated as SBAS and introduced by Berardino et al. (2002), addresses the multi-temporal inversion problem from a perspective complementary to the PS paradigm. While PSI targets individual dominant scatterers in high-resolution cells, SBAS is designed to increase the density of coherent measurement points by working with distributed scatterers rather than with point reflectors. The fundamental design choice in SBAS is to restrict interferogram formation to pairs with small perpendicular and temporal baselines, thereby limiting both geometric and temporal decorrelation. Interferograms satisfying user-defined baseline thresholds are selected and organised into connected subsets. The resulting network is then inverted to recover the displacement time series at each coherent pixel, as illustrated in Figure 3.4.

The SBAS inversion is formulated as a linear system that relates the observed interferometric phase values to the unknown incremental LOS displacements between consecutive acquisitions. The formulation is expressed in matrix form by Equation 3.14:

$$d = G \cdot m \quad 3.14$$

where $d \in \mathbb{R}^N$ is the vector of N observed interferometric phase values, $m \in \mathbb{R}^S$ is the vector of S unknown incremental LOS displacements between consecutive acquisitions, and $G \in \mathbb{R}^{N \times S}$ is the design matrix that encodes the interferometric connections. Each row of G corresponds to one interferogram, with

entries equal to one for the slave acquisition and minus one for the master acquisition of that pair, and zero elsewhere; each column corresponds to one acquisition epoch.

The system is generally overdetermined, with N greater than S , and is solved using the Moore-Penrose pseudoinverse of the design matrix $G^\dagger$. The solution is expressed by Equation 3. 15:

$$\hat{m} = G^\dagger \cdot d \quad 3. 15$$

In the general SBAS configuration, the interferometric network may decompose into isolated subsets of acquisitions that are not mutually connected through short-baseline pairs, producing a rank-deficient design matrix for which the closed-form left inverse (Equation 3. 16) is not defined.

$$(G^T \cdot G)^{-1} \cdot G^T \quad 3. 16$$

The pseudoinverse is therefore computed through Singular Value Decomposition (SVD), as shown in Equation 3. 17:

$$G^\dagger = V \cdot \Sigma^\dagger \cdot U^T \quad 3. 17$$

with V , Σ , U the matrices of the decomposition (Equation 3. 18):

$$G = U \cdot \Sigma \cdot V^T \quad 3. 18$$

and $\Sigma^\dagger$ obtained by inverting the non-zero singular values of Σ . The SVD formulation accommodates both the full-rank and the rank-deficient case, the latter being typical of multi-subset SBAS networks.

A temporal smoothness regularisation term is typically added to suppress residual atmospheric noise and to stabilise the inversion when isolated interferogram clusters exist (Berardino et al., 2002). The output is a time series of cumulative LOS displacement at each coherent pixel, together with a mean annual velocity map.

The operational structure of the SBAS algorithm, as originally proposed by Berardino et al. (2002), is summarised in Figure 3. 10. The processing chain

comprises four sequential steps. The first is the generation of a set of differential interferograms selected according to a dual constraint on spatial and temporal baselines. The second is the coherence-based selection of the measurement points. The third is the inversion of the interferometric phase through Singular Value Decomposition of the design matrix that relates the observed phases to the unknown deformation increments between successive acquisitions. The fourth is the final retrieval of the displacement time series for each coherent pixel. The small-baseline constraint on the interferometric network is the distinctive feature of the approach and is adopted to minimise the geometric decorrelation that would affect pairs with large perpendicular baselines.




Figure 3. 10. Block diagram of the SBAS algorithm, showing the sequential processing steps from the generation of the differential interferograms through the SVD inversion to the retrieval of the deformation time series.

SBAS generally achieves higher measurement point density than PSI in areas of moderate coherence, such as agricultural land, forested hillslopes, and heterogeneous peri-urban environments. The spatial multi-looking inherent in the computation of coherence for distributed scatterers reduces the effective spatial resolution but improves the phase signal-to-noise ratio, making SBAS particularly effective for large-scale analyses of slow geohazards and urban subsidence over extended territories. PSI, by contrast, preserves full spatial resolution and achieves higher measurement precision at individual stable targets, at the cost of lower point density in low-coherence areas. For the building-scale and infrastructure-scale applications addressed in this thesis, the PSI formulation is preferred over SBAS, because the structural targets of interest, including facades, parapets, metallic

elements, and bridge decks, provide abundant stable corner reflectors with high temporal coherence, and the spatial resolution of the point-wise approach is essential for attributing measurements to specific structural elements.

Hybrid techniques such as SqueeSAR (Ferretti et al., 2011) jointly process PS and DS candidates within a unified covariance-based framework, partially reconciling the strengths of both paradigms and representing an active area of methodological development. In the user-driven implementation adopted throughout this thesis, the PSI formulation of Ferretti et al. (2001) is applied directly, without the intermediate relaxations admitted by the QPS extension. The choice is motivated by the nature of the structural targets considered, which exhibit the strongly dominant scatterer behaviour for which the original PSI formulation was designed, and by the priority given to the interpretability of the parametric inversion in the structural engineering context of the thesis.

3.2. Data sources and platforms

The quality and interpretability of a MT-InSAR analysis depend directly on the characteristics of the underlying SAR dataset. Four parameters are of particular relevance for structural monitoring applications. The wavelength determines the phase sensitivity to LOS displacement and the ambiguity interval. The spatial resolution governs the density of detectable coherent scatterers on built assets and the capacity to localise measurement points within the structural geometry. The revisit interval sets the minimum temporal sampling of the deformation signal and limits the detection of rapidly evolving processes. The length and continuity of the observation record control the ability to identify slow trends and to separate seasonal cycles from progressive movement with statistical confidence. No single sensor simultaneously optimises all four parameters. The choice of data source must therefore be adapted to the spatial scale and temporal characteristics of the monitoring problem. This section describes the three data sources employed in the work presented in this thesis. Section 3.2.1 covers the Sentinel-1 constellation operated by the ESA. Section 3.2.2 describes the COSMO-SkyMed constellation operated by the Italian Space Agency (ASI) and the Italian Ministry of Defence. Section 3.2.3 introduces the EGMS, a pre-processed InSAR product derived from Sentinel-1 data and covering the European continent. The complementary characteristics of these sources, and the challenges arising from their joint use in a single monitoring application, are discussed further in Section 4.4.

3.2.1. Sentinel-1

The Sentinel-1 constellation comprises two identical C-band SAR satellites, Sentinel-1A launched in 2014 and Sentinel-1B launched in 2016, developed and operated by the ESA within the Copernicus Earth Observation Programme. Sentinel-1 operates at a centre frequency of 5.405 GHz, corresponding to a wavelength of approximately 5.55 cm. The standard land acquisition mode is the IW mode, which employs the TOPS strategy to achieve a 250 km ground swath with a nominal spatial resolution of 5 m in range and approximately 20 m in azimuth. The swath is divided into three contiguous sub-swaths, each acquired in burst-by-burst fashion, and the SLC product delivers a Single Look Complex image in which all bursts within a sub-swath maintain phase coherence. The constellation was initially composed of Sentinel-1A, launched in April 2014, and Sentinel-1B, launched in April 2016. Sentinel-1B ceased operations on 23 December 2021 following an on-board power anomaly, with the end-of-mission formally declared by ESA on 3 August 2022, reducing the effective revisit interval over European territory from six to twelve days. The launch of Sentinel-1C in December 2024 restored the two-satellite configuration and re-established the nominal six-day revisit in most European regions, while Sentinel-1D, launched in November 2025, is currently entering operational transition and is intended to take over Sentinel-1A duties in 2026 (Figure 3. 11).




Figure 3. 11. Revisit times of the satellites of the Sentinel-1 constellation. Adapted from (Prota et al., 2025)

All Sentinel-1 data are freely accessible without charge through the Copernicus Data Space Ecosystem and the Alaska Satellite Facility (ASF), with a global archive extending back to 2014 for Sentinel-1A. This combination of free availability, dense global coverage, and a growing multi-year archive has established Sentinel-1 as the primary data source for large-scale InSAR applications worldwide. For

structural monitoring, the strengths of Sentinel-1 lie in its long time series, which enables the identification of slow trends and annual cycles at the urban and district scale (Crosetto et al., 2016). Its moderate spatial resolution, however, limits PS density on small or geometrically complex structures, and the co-registration between the TOPS sub-swaths requires specific processing procedures to avoid azimuth phase ramps at burst boundaries. When the area of interest spans the boundary between two frames in azimuth, the processing workflow described in Section 3.3 performs an automatic frame stitching operation that merges the adjacent frames into a single continuous image before interferometric processing begins.

3.2.2. COSMO-SkyMed

COSMO-SkyMed, whose full name is Constellation of Small Satellites for Mediterranean Basin Observation, is the primary source of X-band SAR data used in this thesis and the constellation from which the SLC stacks analysed in Chapters 4 and 5 are principally derived. The system was developed by the Italian Space Agency in cooperation with the Italian Ministry of Defence and constitutes the first Earth Observation mission conceived for dual civil and military purposes (Covello et al., 2010). The first generation, designated COSMO-SkyMed First Generation (CSK), consisted of four identical satellites equipped with SAR sensors operating at X-band ($\lambda \approx 3.1$ cm), launched progressively between June 2007 and November 2010; the third satellite, CSK3, was decommissioned in October 2022 following a deorbiting procedure. A second generation, designated COSMO-SkyMed Second Generation (CSG), was initiated with two satellites launched in December 2019 and January 2022 respectively, ensuring operational continuity and improving image quality relative to the first generation. The revisit interval for each individual satellite is 16 days, but the phased orbital configuration of the full constellation allows observation of the same scene approximately every four days under optimal conditions (Figure 3. 12).

The two acquisition modes most relevant to interferometric applications are the Stripmap HIMAGE mode, which provides a nominal spatial resolution of approximately $3 \text{ m} \times 3 \text{ m}$ and a swath of about 40 km, and the Spotlight mode, which achieves sub-metre resolution over a scene of approximately $10 \text{ km} \times 10 \text{ km}$. For the structural and geotechnical applications described in this thesis, Stripmap SLC data were used, as this mode offers the best combination of spatial coverage, phase quality, and PS density on built assets. Unlike Sentinel-1, COSMO-SkyMed data are not freely distributed; access for academic research is granted through ASI

or authorised channels, with the data request and licensing process adding a constraint to operational deployment that is absent for Sentinel-1. The following section discusses how Sentinel-1-derived displacement products are made available through the EGMS platform, which constitutes a third and qualitatively distinct data access pathway.




Figure 3. 12. Revisit times of the first generation (CSK) and second generation (CSG) COSMO-SkyMed constellation satellites. Adapted from (Prota et al., 2025)

3.2.3. EGMS platform: products, access, and caveats

The EGMS is a component of the Copernicus Land Monitoring Service (CLMS) that provides continental-scale, standardised, and freely accessible ground displacement products derived from MT-InSAR processing of Sentinel-1 data (Capes et al., 2023; Ferretti et al., 2023). The service was made publicly available in 2022 and currently covers the full territory of the European Economic Area, offering displacement time series for the period 2018-2022 in the first release and 2019-2023 in the updated release of October 2024. Unlike the raw SLC data described in Sections 3.2.1 and 3.2.2, EGMS products are pre-processed displacement outputs distributed at three successive levels of elaboration, designated Basic, Calibrated (Figure 3. 13), and Ortho (Figure 3. 14), each with a distinct referencing strategy and intended use (Kotzerke et al., 2022).

The three EGMS product levels address different user needs and carry different interpretative constraints, summarised in Figure 3. 15. EGMS product overview and the main product characteristics. The Basic product (Level 2A) provides LOS displacement time series and mean velocity estimates referenced to a virtual

reference point internal to each processed tile. The local referencing of this product means that measurements cannot be compared directly across tile boundaries, which limits its use for wide-area velocity mapping but does not impair intra-tile pattern recognition. The Calibrated product (Level 2B) introduces absolute referencing through GNSS-based harmonisation. The calibration enables cross-tile comparison of velocity fields and makes the product suitable for linear infrastructure assessment, where the coherent behaviour of multiple spans across a corridor must be evaluated jointly. The Ortho product (Level 3), decomposes the calibrated LOS measurements into vertical and east-west components on a regular 100 m grid, offering the most immediately interpretable displacement representation, but at the cost of spatial resolution and point density. The resampling to a regular grid eliminates many measurement points that fall on structural elements, which significantly reduces the usefulness of the Ortho product for structure-specific assessment. For bridge decks in particular, the shift to a 100 m grid can result in the loss of the majority of points located on the deck, leaving only points attributable to surrounding terrain. The Calibrated product therefore represents the most appropriate choice for infrastructure applications, while the Ortho product may suffice for coarse territorial screening of urban areas and building stocks where spatial aggregation is acceptable.




Figure 3. 13. Calibrated product in the EGMS Data Explorer for a representative site, with the spatial distribution of the measurement points colour-coded by mean LOS velocity in the map view and the displacement time series of a selected point reported in the adjacent plot.




Figure 3. 14. Ortho product in the EGMS Data Explorer for a representative site, with the vertical and east-west displacement components resampled on a regular 100 m grid and the corresponding time series of a selected cell reported in the adjacent plot.




Figure 3. 15. EGMS product overview and the main product characteristics

This qualitative difference from raw satellite data has direct consequences for engineering applications: EGMS products require no processing infrastructure or expertise in InSAR, making them immediately accessible to practitioners, but they impose fixed processing choices, reference point definitions, and temporal sampling that cannot be modified by the end user. The non-modifiability of the EGMS products has a direct consequence for their diagnostic use. The EGMS pipeline applies a single set of processing choices uniformly across the European territory, and the end user inherits these choices without the possibility of inspecting how they propagate through the reconstruction of displacement at any specific site. The distinction between standardised products and user-driven processing is

epistemic in nature, beyond its operational dimension. The user-driven workflow allows the analyst to hold the processing chain open to inspection, to test the sensitivity of the result to each parametric choice, and to adapt the workflow to the scale and nature of the monitoring problem. These properties are essential when the observable is to be used as an input to structural assessment, where the analyst must be able to account for every step between the raw SAR acquisitions and the reported displacement. The complementary role that the EGMS products play in this thesis follows from this distinction. Standardised products are employed where the diagnostic question is compatible with fixed processing choices, that is, in portfolio-level screening, in baseline confirmation of structural stability, and in the identification of structures that merit closer examination. User-driven processing is employed where the diagnostic question requires control over the processing chain, that is, in the structure-specific analyses of Chapter 5. The two strategies are not in competition. Each is the appropriate instrument for the diagnostic level it serves, and their joint use across the case studies enables the empirical comparison of standardised and user-driven outputs on the same structures, which constitutes one of the methodological contributions of the thesis. The following paragraphs describe the three product levels, the access pathway, and the caveats that must be considered when using EGMS data for structural interpretation.

A technical characteristic of the EGMS processing chain that has direct consequences for building-scale structural interpretation concerns the treatment of the Sentinel-1 sub-swath structure. As described in Section 2.1.1, the IW mode acquires data across three adjacent sub-swaths (IW1, IW2, IW3), each covering a strip of approximately 80 km in ground range with a slightly different incidence angle. The EGMS processing chain analyses each sub-swath independently, applying the PS candidate selection, the multi-temporal inversion, and the coherence filtering separately within each sub-swath before merging the results into the final product. A measurement point that satisfies the coherence threshold in one sub-swath but falls within the overlap zone between two sub-swaths may be retained even if it does not satisfy the threshold in the adjacent sub-swath. The consequence is that the EGMS product may contain measurement points in sub-swath overlap zones whose quality is evaluated against different interferometric networks and different atmospheric phase screen estimates, potentially introducing subtle inconsistencies in the velocity and coherence values of adjacent points.

User-driven processing, by contrast, can perform a combined analysis across all three sub-swaths, constructing a single interferometric network and a single atmospheric model over the merged scene. Points that appear in one sub-swath but

not in another are excluded from the combined solution, because their absence in the second sub-swath indicates insufficient phase stability under the combined assessment criterion. The result is typically a smaller number of retained points but with more uniform quality, because every point in the solution has been evaluated against the full interferometric network rather than against a sub-swath-specific subset.

For building-scale structural monitoring, this difference has a practical impact on the apparent point density and on the coherence statistics of the measurement points attributable to a given structure. A compact building located near a sub-swath boundary in the EGMS product may exhibit a higher apparent point count than the same building processed in a combined user-driven analysis, but the additional points may include lower-quality measurements that were retained because they satisfied the threshold in one sub-swath only. The comparative analysis presented in Case study I (Section 5.2.5) provides empirical evidence of this effect for the *Santissima Trinità* in Turin, where the EGMS product yields a higher apparent density than the user-driven solution but at lower average coherence. The principal applications of EGMS products in civil engineering and infrastructure monitoring are summarised in Figure 3. 16.




Figure 3. 16. Summary of the principal applications of EGMS products in civil engineering and infrastructure monitoring, organised by category of engineering objective and by the level of EGMS product most appropriate to each category.

3.3. User-driven processing workflow

The MT-InSAR analyses presented in this thesis are performed within a fully user-driven processing framework. This choice is methodological, not merely operational. A user-driven workflow exposes every stage of the reconstruction of displacement from SAR data, from the definition of the interferometric network to the selection of the reference point, the estimation of the topographic contribution, and the filtering of the retained measurement points. Each of these stages is governed by parametric decisions whose consequences propagate through the remainder of the chain. The control that the analyst retains over these decisions constitutes the substantive difference between a user-driven processing workflow and a pre-processed standardised product such as those delivered through the EGMS platform. Standardised products apply a uniform set of processing choices across wide geographic areas, and the end user inherits these choices as fixed inputs to the interpretation stage, without the possibility of inspecting how any single decision has shaped the reported displacement at a specific site. The user-driven workflow, by contrast, holds the entire processing chain open to inspection, allowing the analyst to test the sensitivity of the result to each parametric choice and to adapt the workflow to the scale and nature of the monitoring problem. These properties are preconditions for the use of InSAR displacement as an input to structural assessment, where each step between the raw SAR acquisitions and the reported displacement must be accountable to the diagnostic question being asked. The user-driven workflow described in this thesis is implemented within SARPROZ (EO59 Ltd., 2026), a MATLAB-based platform developed by Perissin and collaborators that realises the POLIMI PSI formulations through a modular, sequential processing architecture (Perissin et al., 2011, 2011; Perissin & Wang, 2012). The implementation deliberately exposes each parameter of the processing chain through an interactive interface, enabling the analyst to inspect intermediate results and to revise earlier choices in the light of the outputs of subsequent stages. The granularity at which the workflow is described below reflects the granularity at which decisions must be taken in practice; documenting these decisions explicitly is a precondition for reproducibility and for the critical discussion of the results. The remainder of this section proceeds in three phases. Section 3.3.1 addresses the selection and ingestion of the image dataset, including the definition of the master image, the area of interest, and the acquisition configuration. Section 3.3.2 describes the co-registration of slave images to the master reference and the construction of the interferometric network. Section 3.3.3 covers the integration of an external DEM, the generation of geocoded outputs, and the preparation of the

project for multi-temporal analysis. These three phases constitute the mandatory preparatory sequence that must be completed before the quality assurance operations described in Section 3.4 and the time-series estimation can proceed.

3.3.1. Dataset selection and acquisition strategy

The selection of an appropriate SAR dataset for structural monitoring requires decisions on four interdependent parameters: the sensor and acquisition mode, the temporal extent of the observation window, the orbit geometry, and the spatial extent of the area of interest. The sensor choice determines the spatial resolution and wavelength, as discussed in Section 3.2. For Sentinel-1, only the SLC product in IW mode is suitable for MT-InSAR processing. The SLC format preserves the full amplitude and phase information required for interferometric analysis, in contrast to the multi-looked GRD product, which is adequate for change detection but not for interferometry.

When the area of interest falls entirely within a single Sentinel-1 sub-swath, only that sub-swath needs to be imported. When the area spans the boundary between adjacent frames in azimuth, the processing environment performs an automatic frame stitching operation that merges the two images into a single consistent dataset before extraction. The temporal window should maximise the stack length while maintaining a consistent acquisition configuration, because changes in the satellite operational mode, in the polarisation channel, or in the orbital track invalidate co-registration between images acquired under different conditions and require the stack to be split accordingly. Both ascending and descending geometries should be acquired when the analysis requires LOS decomposition, as discussed in Section 2.2.2. The minimum stack size recommended in the literature for reliable PS detection and atmospheric phase estimation lies in the range of 15 to 25 images (Crosetto et al., 2016; Ferretti et al., 2001), although stacks of 50 or more images substantially improve velocity precision and atmospheric correction robustness.

For the applications described in this thesis, we downloaded SLC data from two sources. We retrieved Sentinel-1 IW scenes in VV polarisation through the Copernicus Data Space Ecosystem API and through the ASF Vertex interface, both of which support automated batch download using Python scripts that allow unattended retrieval of complete stacks. COSMO-SkyMed HIMAGE scenes in HH polarisation were accessed through the ASI Research and University Programme,

which provides academic institutions with a quota of scenes over selected Italian test sites within the MAPITALY framework after approval of a project proposal.

The master image is selected through an optimisation algorithm that jointly minimises the sum of perpendicular baselines, temporal baselines, and Doppler centroid frequency differences across all acquisitions in the stack. Selecting a master near the temporal and geometric centre of the dataset reduces the average decorrelation across all interferometric pairs and improves the conditioning of the inversion system, as visualised in the star-graph baseline plot produced at this stage, an example of which is shown in Figure 3. 17. Poor master selection, by contrast, can introduce asymmetric decorrelation that degrades coherence for a systematic subset of acquisitions and biases the resulting time series. The co-registration of all secondary images to the selected master is addressed in Section 3.3.2.




Figure 3. 17. Star-graph baseline plot of a representative dataset, with normal baseline reported on the vertical axis and acquisition date reported on the horizontal axis. The master image is in the centre, and each interferometric connection between master and slave is colour-coded according to the spatial coherence.

Ancillary data complement the SAR stack and support two distinct quality-control functions at the dataset preparation stage. Meteorological records (Figure 3. 18) for the area of interest, obtained from local weather stations or from global reanalysis services such as Visual Crossing, enable the identification of acquisitions

that may be unusable or degraded. Images acquired during heavy precipitation, with snow cover on the ground, or during exceptional meteorological events may exhibit anomalous atmospheric phase contributions or surface decorrelation that reduce the quality of the interferometric solution, and such acquisitions can be excluded from the stack or flagged for cautious interpretation. Meteorological records also provide the air temperature time series aligned with the SAR acquisition dates, which serves as external input for the estimation of the thermoelastic expansion coefficient in the parametric displacement model developed in Section 4.2.2. For structures located on compressible soils or subject to hydrological forcing, piezometric records and satellite-derived soil moisture products offer a complementary external input that feeds into the integrated interpretation frameworks of Section 4.5 and of the Case study III of Section 5.4.




Figure 3. 18. Example of the meteorological record. From top to bottom: air temperature, atmospheric pressure, relative humidity, and daily precipitation.. The four variables are reported on a common time axis covering the observation period.

3.3.2. Co-registration and interferometric network construction

Co-registration aligns all images in the stack to a common reference, the master image, so that each pixel in every slave image corresponds to the same ground resolution cell as the corresponding master pixel. Precise geometric alignment is a prerequisite for all subsequent processing operations, including interferogram formation and time-series estimation, as sub-pixel misregistration introduces phase

errors that propagate into the final displacement products. The master image is typically selected as the acquisition closest to the temporal centroid of the stack and at a moderate perpendicular baseline with respect to the majority of the other images, in order to minimise the average geometric decorrelation across the network. In the adopted workflow, the slave images are initially extracted with a spatial buffer relative to the master footprint, providing additional margin to accommodate orbital offsets before resampling. For Stripmap data and COSMO-SkyMed, co-registration relies on amplitude cross-correlation between patches extracted from master and slave images, using a reference window of 64×64 pixels and a search radius that can be increased for datasets with large orbital offsets or complex terrain. For Sentinel-1, precise orbital state vectors are used directly when available, typically 20 days after acquisition, with the amplitude cross-correlation strategy applied only to images for which only restituted orbits are accessible. The appropriate strategy is selected automatically for each image based on the availability of precise orbits. After co-registration, all slave images are resampled to the master grid at exactly the same pixel dimensions, and the buffer region is removed. The quality of co-registration should be assessed by visual comparison of master and co-registered slave images, verifying that stable features, such as building outlines and road intersections, are aligned consistently. The interferometric network adopted in this thesis is the single-master star graph centred on the master image. The connectivity of the resulting graph is verified before proceeding to interferogram generation by inspecting the star-graph plot and confirming that no isolated clusters of acquisitions remain.

3.3.3. DEM integration, geocoding, and export

The integration of an external DEM into the processing chain serves two distinct and equally necessary purposes: the removal of the topographic phase contribution from each interferogram, as formalised by Equation 2. 4 and by the decomposition of Equation 3. 4, and the geometric transformation that allows all results to be expressed in geographic coordinates for cross-referencing with ancillary data. The DEM is ingested and reprojected from its native geographic coordinate system into the SAR image coordinate system, producing a height layer with exactly the same pixel dimensions as the master image. The reprojection uses the orbital state vectors, the sensor geometry, and the incidence angle model to map each geographic DEM cell to its corresponding range-azimuth position in the SAR grid.

The processing environment retains the SRTM DEM at 90 m resolution as a historical default option, which remains adequate for regional-scale analyses. For

the applications presented in this thesis, we adopted the Copernicus GLO-30 DEM at 30 m resolution, because it offers near-global coverage, open access, and vertical accuracy within a few metres over most terrestrial surfaces (Airbus Defence and Space, 2022), providing substantially improved height resolution for building-scale interpretation than the SRTM alternative. When higher accuracy is required, for instance in urban areas with tall buildings or in mountainous terrain, user-supplied DEMs at finer resolution, for example from airborne LiDAR, can be imported into the processing workflow.

Once projected into SAR coordinates, the DEM is used to simulate a synthetic amplitude image by modelling the geometric effect of terrain slope and aspect on radar backscatter intensity. The synthetic amplitude serves as a comparison layer against the observed reflectivity map, whose definition and computation are presented in Section 3.4.1, providing a visual quality check of the geometric alignment: close agreement between the two images confirms that the DEM has been correctly registered to the SAR data. Where the alignment shows residual offsets, the geocoding can be refined using a small number of ground control points (GCP) identified by cross-referencing the reflectivity map with an optical reference in Google Earth or a comparable platform. The analyst selects each GCP by matching a recognisable stable feature in the SAR image, such as a corner reflector, a road intersection, or the base of an isolated structure, to its known geographic position.

After geocoding refinement, all outputs can be exported in geographic coordinates in KMZ, GeoTIFF, and ASCII formats for integration with GIS software. The geocoded reflectivity map and the ASI map, both introduced formally in Section 3.4.1, are the primary quality-assessment products reviewed at this stage, because they provide the first evidence of the density and spatial distribution of PS candidates on the target structures before interferometric processing begins.

3.3.4. Small area processing

The processing workflow described in Sections 3.3.1-3.3.3 operates at the scale of the full project site, which may extend over tens of kilometres in both range and azimuth. The reflectivity map, the ASI, the co-registered stack, and the external DEM are computed for the entire scene. The multi-temporal parameter estimation, however, can be performed at two distinct spatial scales: the full site (large area processing) and a localised crop (small area processing).. The two modes differ in the treatment of the atmospheric phase screen and in the computational strategy for

the parameter inversion. Large area processing, described in the published literature as the standard approach for regional-scale ground motion analysis, estimates and removes the atmospheric contribution explicitly across the full scene. Small area processing, which is the mode most directly relevant to the building-scale structural monitoring applications of this thesis, adopts a different physical assumption that simplifies the estimation at the cost of restricting the spatial extent of the analysis.

The physical basis of the small area assumption rests on the spatial correlation structure of the atmospheric phase screen. The tropospheric phase delay varies across the scene as a function of the spatial distribution of temperature, pressure, and water vapour at the time of each acquisition. The characteristic correlation length of this variation is typically on the order of several kilometres in temperate climates: two points separated by a distance much smaller than this length experience nearly identical atmospheric conditions, and the differential atmospheric phase between them is negligible relative to the structural phase contributions. When the analysis is restricted to a small area centred on a single structure or a compact group of structures, the differential atmospheric contribution between any measurement point and the reference point is assumed to be sufficiently small to be absorbed into the phase noise term without introducing a systematic bias in the estimated residual height and velocity.

The maximum admissible spatial extent of the small area depends on the radar wavelength. For a given physical atmospheric path delay ΔL_{atm} , the corresponding differential phase is shown in Equation 3. 19:

$$\Delta\phi_{atm} = \frac{4\pi\Delta L_{atm}}{\lambda} \quad 3. 19$$

A shorter wavelength produces a larger phase contribution for the same atmospheric condition, and therefore requires a smaller analysis area to keep the atmospheric residual below the estimation noise floor. As a practical guideline, the small area radius is limited to approximately 500 m for X-band sensors ($\lambda \approx 3.1$ cm), 800 m to 1 km for C-band sensors ($\lambda \approx 5.55$ cm), and larger distances for L-band sensors ($\lambda \approx 24$ cm). These values are indicative and depend on the local atmospheric conditions, the topographic relief, and the temporal separation between acquisitions. In flat, climatically stable areas, larger small areas may be acceptable; in mountainous or tropical environments where atmospheric variability is higher, more conservative limits apply, consistent with the spatial correlation analyses of tropospheric delay reported by Hanssen (2001) and Fattahi & Amelung (2014).

Operationally, the analyst creates the small-area window as a spatial crop of the co-registered stack, defining the crop boundaries in range and azimuth coordinates centred on the structure of interest and extending to include a surrounding buffer that provides contextual measurement points and candidate reference locations. The software extracts the corresponding sub-volume from the full stack, together with the associated DEM, reflectivity map, and ASI, and creates a self-contained sub-project within the main project directory. The sparse point selection (Section 3.4.1), reference point selection (Section 3.4.2), and multi-temporal parameter estimation (Section 3.1.2) are then performed independently within this sub-project, operating only on the points contained in the crop. The results are stored in the sub-project memory and can be geocoded and exported without affecting the main site-level project.

The small area mode offers three operational advantages for building-scale structural monitoring. The first is computational efficiency: by restricting the candidate population to the immediate vicinity of the structure, the periodogram search (Section 3.1.2) is applied to hundreds or thousands of points rather than to hundreds of thousands, reducing processing time from hours to seconds and allowing rapid iterative exploration of parameter settings. The second is local reference selection: the reference point can be chosen on or immediately adjacent to the structure, minimising the differential atmospheric residual and ensuring that the displacement estimates are relative to a physically meaningful location. The third is parameter adaptability: the amplitude threshold, coherence criterion, and estimation model can be tailored to the specific deformation regime expected on the structure, rather than being constrained by the settings appropriate for a heterogeneous regional scene. These advantages are exploited systematically in the case studies of Chapter 5, where small area processing is applied to the *Santissima Trinità* (Section 5.2.4), to the *Sanctuary of Vicoforte* (Section 5.4.3), and to the Palisades Fire study area (Section 5.6.3).

The transition from small area processing to large area processing becomes necessary when the monitoring objective requires coverage of multiple structures distributed across a wider territory, or when the atmospheric phase screen must be estimated and removed to achieve the required accuracy over distances exceeding the small area limit. Large area processing, which is not the primary focus of this thesis, achieves atmospheric estimation through spatial filtering of the phase residuals across a dense network of PS, exploiting the assumption that the atmospheric contribution is spatially smooth while the structural displacement varies at shorter spatial scales. The interested reader is referred to Ferretti et al.

(2001) and Perissin & Wang (2012) for a detailed treatment of the large area atmospheric estimation strategy.

3.4. Quality assurance and quality control

Quality assurance in MT-InSAR processing is not a single check applied after results have been produced. It is a sequence of decisions integrated into the processing chain at three stages, each of which constrains the reliability of all operations that follow. The first stage, addressed in Section 3.4.1, concerns the generation of the reflectivity map and the ASI, and the use of these parameters to define a sparse candidate point set from which PS will subsequently be drawn. The second stage, described in Section 3.4.2, concerns the selection of the spatial reference point, which anchors the entire displacement field and whose stability cannot be assumed without explicit verification. The third stage, presented in Section 3.4.3, applies a temporal coherence threshold to the estimated PS set, retaining only those points whose phase time series is sufficiently consistent with the adopted displacement model to support reliable interpretation. Errors introduced at the first stage inflate computational load and degrade the atmospheric phase correction by including unstable candidates. Errors at the second stage introduce a systematic offset that contaminates every displacement estimate across the project area and cannot be corrected retrospectively. Errors at the third stage either compromise the spatial density of the final point set, if the threshold is too strict, or admit unreliable measurements into the diagnostic output, if it is too permissive. A coherent understanding of all three stages is therefore necessary before any structural interpretation of the multi-temporal results can proceed.

3.4.1. Reflectivity, amplitude stability index, and sparse point selection

The reflectivity map is computed as the temporal mean of the pixel amplitude across all images in the stack, assigning to each resolution cell the average backscattering intensity recorded over the observation period rather than the instantaneous response at any single acquisition date. The distinction between a single SAR image and the reflectivity map is fundamental to the correct interpretation of the preliminary processing output, yet it is frequently overlooked by users who approach radar data with expectations formed by optical imagery. A single SAR image records the backscattering conditions of the scene at one specific acquisition epoch, and the amplitude of each pixel reflects the instantaneous scattering configuration, including the positions of all mobile objects (vehicles, vessels,

construction equipment), the moisture state of surfaces, and the specific orbital position of the satellite at that moment. The reflectivity map, by contrast, collapses the entire temporal dimension of the stack into a single statistical quantity. An object present in only one acquisition out of a hundred contributes one hundredth of its backscatter intensity to the corresponding pixel of the reflectivity map. Conversely, an object present in all acquisitions with consistent amplitude dominates the map value at its location.

The reflectivity map is therefore an indicator of temporal persistence of backscatter, not a representation of the physical state of the site at any given moment. A bright pixel in the reflectivity map indicates a location where strong backscatter was consistently recorded across many acquisitions; it does not indicate that a specific object was present at that location during the most recent acquisition. For structural monitoring, the practical consequence is that transient elements such as construction scaffolding, temporary shoring, parked vehicles, or debris may appear as faint contributions in the reflectivity map despite being prominently visible in individual images. The analyst who uses the reflectivity map to identify structural elements must therefore verify, through inspection of individual acquisition amplitudes or through the ASI introduced below, that the features observed in the map correspond to physically permanent elements of the structure.

The reflectivity map generated by the SAR amplitude stack differs fundamentally from an optical image of the same scene, and a direct comparison between the two modalities provides the most immediate entry point to the physics of SAR target response. Figure 3. 19 presents such a comparison over a representative area. Panel (a) shows the SAR amplitude image. Panel (b) shows the optical image of the same area, acquired under standard daylight conditions, with partial cloud cover, in which brightness is dominated by reflected solar radiation in the visible spectrum.




Figure 3. 19. Comparison between an optical image and a SAR amplitude image of the same scene, illustrating the distinct information content of the two acquisition modalities. Panel (a): SAR amplitude image of a representative area. Panel (b): optical image of the same area, acquired under standard daylight conditions, with partial cloud cover. The correspondence between visual prominence in the optical image and radar response is not direct. Adapted from (Gao et al. 2019).

The relationship between the physical scattering environment and the reflectivity map can be illustrated through a representative example (Figure 3. 20). A mean amplitude image computed from a multi-temporal stack over an urban coastal area reveals a characteristic pattern of backscatter intensities that maps directly onto the physical properties of the observed targets. Water bodies such as rivers, harbours, and lakes register as uniformly dark patches in the reflectivity map, because at the incidence angles typical of spaceborne SAR a calm water surface behaves as a specular mirror and nearly all incident energy is deflected away from the receiving antenna.




Figure 3. 20. Comparison between an optical reference image and the SAR reflectivity map of the same urban coastal area. Panel (a): Google Earth optical view. Panel (b): mean amplitude image computed from the multi-temporal SAR stack, showing the characteristic backscatter signatures associated with water bodies, built-up areas, and vegetated surfaces.

Urban areas present the opposite behaviour, appearing as dense aggregations of bright pixels, because building façades, metallic roof elements, and infrastructure components generate dihedral and trihedral reflection geometries that redirect a substantial fraction of the transmitted energy back toward the satellite. Within these bright aggregations, a single object becomes individually recognisable when its backscattered power is high enough to produce an impulse response that stands clear of the surrounding clutter floor. A light pole, to take a concrete example, appears in the reflectivity map as an isolated bright point flanked in the range direction by a sequence of progressively attenuated side lobes, the characteristic signature of a strong point-like reflector. The correspondence between the visual appearance of such a target and its radar response is shown explicitly in Figure 3. 21, where an optical view of a representative urban area is compared with the reflectivity map computed from the SAR amplitude stack. The origin of this response is electromagnetic rather than visual: the metallic composition of the pole and its vertical cylindrical profile form, together with the adjacent ground plane, a corner-reflector geometry that concentrates the return signal efficiently. This last observation carries a direct implication for any structural engineer approaching SAR imagery for the first time. The visual prominence of an object in the radar scene bears no correspondence to its optical appearance or its physical size. A large building covered by a flat polymer membrane may return less energy to the sensor than a compact metallic ventilation duct on a neighbouring rooftop, because the

amplitude of the radar response is dictated by three properties of the target: its electrical conductivity, the orientation of its surfaces relative to the LOS, and its roughness at the scale of the operating wavelength.




Figure 3. 21. Radar signature of an isolated point-like reflector in an urban area. Panel (a): Google Earth optical view showing a light pole at the centre of the scene. Panel (b): SAR reflectivity map of the same area, in which the light pole appears as a bright pixel surrounded by attenuated side lobes in the range direction, the characteristic signature of a strong point-like scatterer.

The physical shape of the response generated by a single point-like scatterer after SAR image formation deserves a brief clarification, because it underlies the local-maxima selection strategy described below. Each scattering target within the observed scene generates a response whose spatial form, after range compression and azimuth focusing, approximates a two-dimensional cardinal sine function (Figure 3. 22), commonly abbreviated as sinc, centred at the target's position in the range-azimuth plane. The peak amplitude of this function corresponds to the backscattered intensity of the target, while the width of the main lobe is determined by the system bandwidth in range, as expressed by Equation 2. 2 in Chapter 2, and by the synthetic aperture length in azimuth, as expressed by Equation 2. 3. Side lobes extend symmetrically on both sides of the peak, with progressively decreasing amplitude. The sinc response is a direct consequence of the finite bandwidth of the radar system, which produces a band-limited impulse response in the range direction, and of the finite synthetic aperture, which produces an analogous band-limited response in the azimuth direction.

The side-lobe structure of the cardinal sine response has direct implications for the identification of isolated coherent scatterers on built structures. A single physical scatterer generates a response that extends across several adjacent pixels in both the range and the azimuth direction, so that the unambiguous identification of a dominant target requires the isolation of the peak of the main lobe from the surrounding side-lobe structure and from the impulse responses of neighbouring scatterers. This geometric property motivates the local-maxima selection strategy adopted in the processing chain, which retains only the peak of each sinc response to avoid the spatial redundancy that would otherwise result from accepting all pixels within the main lobe of a single target as independent candidates.




Figure 3. 22. Cardinal sine function, which describes the impulse response of the SAR system in both the range and azimuth directions. The main lobe is centred at $x = 0$ with peak value of unity. The successive zeros occur at integer values of x , and the side lobes have progressively decreasing amplitude. The width of the main lobe between the first two zeros defines the nominal spatial resolution of the image, and the side lobe structure governs the multiplicative ambiguity of point scatterer responses in the range-azimuth plane.

The averaging operation that generates the reflectivity map substantially reduces the speckle noise inherent in individual SAR images. For a stack of N images with statistically independent speckle realisations, the effective number of looks increases proportionally, and the spatial texture of the map progressively approaches the true variability of the surface scattering properties. The product is therefore a map in the strict sense, not an image, because it does not correspond to

any particular physical state of the scene. On built assets, permanent structures with consistent corner-reflection geometry, including building façade edges, metallic roofing elements, and exposed infrastructure, produce high mean amplitude values that appear consistently bright and are readily identifiable as candidate PS locations. However, high mean reflectivity is a necessary but not sufficient condition for phase stability, because a target may reflect strongly yet exhibit temporally variable phase if multiple scatterers within the same resolution cell interfere differently from one acquisition to the next. A complementary index capturing the temporal stability of the amplitude, rather than its magnitude, is therefore required before a reliable candidate set can be defined.

The ASI is defined through Equation 3. 20:

$$ASI(p) = 1 - \frac{\sigma_A(p)}{\mu_A(p)} \quad 3. 20$$

where $\sigma_A(p)$ is the temporal standard deviation of the pixel amplitude at point p computed across the N images in the stack, expressed by Equation 3. 21:

$$\sigma_A = \sqrt{\frac{1}{N} \sum_{i=1}^N (A_i(p) - \mu_A(p))^2} \quad 3. 21$$

and $\mu_A(p)$ the temporal mean of the pixel amplitude, expressed by Equation 3. 22:

$$\mu_A = \frac{1}{N} \sum_{i=1}^N A_i(p) \quad 3. 22$$

with $A_i(p)$ denoting the amplitude of pixel p in image i . Values of ASI close to one indicate high temporal stability, and values approaching zero indicate strong variability. The ASI is algebraically related to the amplitude dispersion D_A introduced by Ferretti et al. (2001), defined by Equation 3. 23:

$$D_A(p) = \frac{\sigma_A(p)}{\mu_A(p)} \quad 3. 23$$

The relation between the two indices is therefore expressed by Equation 3. 24:

$$ASI(p) = 1 - D_A(p) \quad 3. 24$$

The two indices carry equivalent information, but the ASI convention makes the thresholding direction intuitive, because candidate pixels are retained above a minimum stability value. Typical thresholds lie between 0.7 and 0.85 depending on the coherence of the study area and the desired candidate density (Crosetto et al., 2016). Applying this threshold reduces the candidate set to pixels with demonstrably stable amplitude behaviour, limiting the computational load of the subsequent multi-temporal phase analysis and improving the robustness of the atmospheric phase correction by excluding unstable scatterers from the estimation network.

The ASI as defined by Equation 3. 20 is bounded between zero and one for pixels whose amplitude is positive in all acquisitions and whose temporal variability does not exceed the mean. In practice, however, the index can assume negative values when the temporal standard deviation of the amplitude exceeds the temporal mean. A negative ASI indicates that the amplitude at that pixel location varies so dramatically across the stack that the concept of a mean reflectivity ceases to be physically meaningful. The most common physical cause is the presence of a strong scatterer in a small number of acquisitions and its absence in the remainder. A vessel that occupies a harbour berth on two dates out of a hundred, a construction crane present for three months within a five-year stack, or a structural element destroyed during the observation period all produce this pattern: the few high-amplitude acquisitions inflate the standard deviation relative to the low mean computed over the full stack. Pixels with negative ASI are therefore candidate locations for TS whose lifecycle is short relative to the observation window, and a systematic survey of negative-ASI pixels on or near the structure can provide an early indication of physical changes in the scattering environment before the formal TS identification through the multi-temporal phase analysis of Section 4.1.2 is performed. The spatial distribution of negative ASI pixels also serves as a quality check: if negative values cluster in areas known to be physically stable, the cause is more likely to be processing artefacts or isolated atmospheric anomalies than genuine scattering changes.

Alongside the amplitude-based threshold, a sparse point selection mask is applied that retains only the pixels corresponding to the peak of the sinc response along the range direction, an operation known as local-maxima selection. As

established above, a single physical scatterer generates a sinc response extending across several adjacent pixels. Retaining only the peak of each response avoids spatial redundancy and ensures that the selected set contains one representative measurement per physical object rather than multiple correlated realisations of the same target. In physical terms, the selection of local maxima in the reflectivity map corresponds to the identification of the peak of each individual sinc response, that is, the position of the dominant scatterer within each resolution cell neighbourhood. Targets that produce strong, isolated sinc peaks, such as metallic fixtures, corner geometries, and bare rock outcrops, are retained as candidate PS. Targets whose sinc responses overlap with those of adjacent scatterers, or whose amplitude fluctuates substantially across the stack, produce broader and less distinct maxima that may fall below the selection threshold.

The density and spatial distribution of the resulting candidate set, visually inspected in conjunction with the reflectivity map, constitute the primary quality indicators at this stage. Where candidate density on the structure of interest appears insufficient, the analyst may relax the ASI threshold or revisit the master image selection before proceeding to the reference point definition addressed in the following section.

3.4.2. Reference point selection

All MT-InSAR displacement estimates are differential quantities: each LOS value represents the accumulated range change at a given PS relative to a single spatial reference, not an absolute geodetic measurement. Any undetected motion at the reference location is added as a systematic offset to every PS in the project, mimicking or partially cancelling genuine ground or structural movement across the entire study area. The selection of a stable reference point is accordingly one of the most consequential decisions in the processing workflow, and its effects cannot be corrected by any downstream quality filter.

A reliable reference point must satisfy three conditions simultaneously. The target must be located on a surface that is physically stable over the full observation period, free from long-term settlement, seasonal oscillation, or any deformation attributable to structural or geotechnical causes. The pixel must exhibit high amplitude stability, confirming that the PS response originates from a persistent dominant scatterer rather than a fluctuating distributed surface. In addition, the reference must be positioned at a sufficient distance from any known source of localised deformation, such as an active construction site, a zone of differential

compaction, or a fault-related subsidence field, to avoid contamination from spatially correlated signals that could be misinterpreted as reference stability. The reference point is operationally specified by declaring the range-azimuth coordinates of the selected pixel to the time-series estimation module. Once specified, the estimated displacement time series is expressed relative to the reference by construction, with the reference time series identically zero. All displacement time series subsequently expressed relative to that pixel, and the time series at the reference itself is identically zero by construction.

A fourth verification step complements the three statistical conditions listed above: visual inspection of the amplitude time series of the candidate reference point across all acquisitions in the stack. The reflectivity map and the ASI are summary statistics that compress the full temporal amplitude record into a single scalar. A point may satisfy all three statistical conditions, exhibiting high mean amplitude, low amplitude dispersion, and high spatial coherence, while nonetheless containing isolated acquisition epochs at which the amplitude departs substantially from the temporal mean. Such departures may result from transient interference by a nearby scatterer (a vehicle parked adjacent to the reference target on a single acquisition date), from temporary occlusion by construction equipment, or from anomalous atmospheric conditions that affect the backscatter intensity at the reference location. If an amplitude anomaly coincides with an epoch at which the interferometric phase is also disturbed, the differential phase computed for all other points with respect to the reference inherits the disturbance as a common-mode error that propagates through the entire displacement field at that epoch. Inspecting the amplitude time series before finalising the reference point selection allows the analyst to identify such anomalies and, where necessary, to exclude the affected acquisitions from the stack or to select an alternative reference point whose amplitude record is free of outliers. This verification is performed by plotting the amplitude series of the candidate point and comparing its temporal variability against the expected seasonal and orbital modulation pattern for a stable urban scatterer at the relevant wavelength.

Where no locally stable PS can be identified within the project area, two alternative strategies are available. The first consists of selecting a reference on bedrock outcrops or on deeply founded structures known from geotechnical records to be free of long-term settlement, even if such targets lie at the periphery of the study area. The second consists of importing an external velocity constraint from a GNSS permanent station or from a levelling benchmark and using it to anchor the InSAR reference frame to an absolute geodetic datum, a procedure discussed

further in Section 3.5.2. In both cases, the sensitivity of all displacement estimates to the reference choice should be assessed by repeating the analysis with at least one alternative reference location and comparing the resulting velocity maps. The degree of agreement between the two solutions provides an empirical bound on the referencing uncertainty, which feeds directly into the uncertainty propagation framework described in Section 3.5.3. With the reference point established, the temporal coherence of the estimated PS set can be evaluated, as addressed in the following section.

3.4.3. Temporal coherence and quality thresholds

The temporal coherence γ_p defined by Equation 3. 13 serves as the final quality filter applied to the PS dataset. Candidates with coherence below a user-defined threshold are discarded, and only points whose phase time series is sufficiently consistent with the two-parameter linear model of Equations 3. 10 and 3. 11 are retained for interpretation. The selection of the threshold value is therefore a decision that directly shapes the balance between spatial density and measurement reliability in the final output.

An important statistical property of the temporal coherence that affects the choice of quality thresholds is the positive bias introduced by the finite number of observations. The temporal coherence is computed as the modulus of the normalised vector sum of M complex residuals, as established by Equations 3. 12 and 3. 13. For a perfectly random set of M unit vectors uniformly distributed on the unit circle, the expected modulus of their average is not zero but approximately $\sqrt{\pi/4M}$, a result that follows from the two-dimensional random walk on the complex plane and is well documented in the interferometric phase statistics literature (Just & Bamler, 1994; Touzi et al., 1999).

The implication of this property becomes apparent at the lower end of the stack size range. With M interferograms equal to the number of parameters to be estimated, the residual phase is exactly zero at the solution of the normal equations and the temporal coherence is identically equal to one, regardless of whether the estimated parameters correspond to physically meaningful values. For the three-parameter model adopted in this thesis, which jointly estimates residual height, linear velocity, and a thermal or seasonal component, this degenerate condition occurs at M equal to three. The coherence estimator acquires discriminative power only when the number of interferograms exceeds the number of parameters by a

substantial margin, so that the residual phase carries information on the quality of the fit rather than reflecting the algebraic under-determination of the system.

For larger stack sizes the expected coherence under the null hypothesis of random residuals follows the previous approximation. With M equal to 25, the expected coherence for random data is approximately 0.18; with M equal to 71, approximately 0.11; with M equal to 200, approximately 0.06. The practical consequence is that the coherence threshold used to distinguish reliable estimates from noise must be calibrated against the number of interferograms in the stack rather than set as a universal constant. A threshold of 0.7 is conservative for a stack of 30 images, where the random-data expectation is approximately 0.16, but may be excessively stringent for a stack of 200 images, where the random-data expectation is approximately 0.06 and a threshold of 0.5 already provides strong discrimination. Conversely, a threshold of 0.5 applied to a stack of only 15 images, for which the random-data expectation is approximately 0.23, provides limited discriminative power because the threshold lies only about twice as far from the null expectation as from the unit maximum.

The coherence thresholds adopted in the case studies of Chapter 5 are interpreted in this context. The 0.6 threshold applied to the COSMO-SkyMed stack at *Vicoforte*, which contains 62 images and for which the random-data expectation is approximately 0.11, provides substantially stronger quality assurance than the same numerical threshold applied to the Sentinel-1 stack at the *Santissima Trinità*, which contains 169 images and for which the random-data expectation is approximately 0.07. The margin between the adopted threshold and the null expectation is what ultimately determines the discriminative power of the filter, and a fixed numerical threshold applied to stacks of different sizes therefore corresponds to different effective quality levels.

3.5. Validation issues and uncertainty sources

The quality assurance operations described in Section 3.4 ensure that the MT-InSAR processing chain produces internally consistent results. They do not, however, establish the relationship between those results and the physical quantities of interest to structural engineering practice, nor do they quantify the uncertainty budget that must be reported alongside any diagnostic conclusion. The reliability of MT-InSAR results for structural assessment therefore depends not only on the quality of the processing chain described in Sections 3.3 and 3.4, but also on an explicit characterisation of the uncertainty associated with each measurement.

Three distinct sources of uncertainty affect the final displacement products. The first is the geometric localisation of PS in geographic coordinates, which determines whether a measurement can be confidently attributed to a specific structural element. The second is the referencing of the displacement field to a stable datum, which controls the accuracy of absolute and relative velocity estimates. The third is the propagation of phase noise, atmospheric residuals, and model approximations into the precision and accuracy of the time series. These three sources are addressed in sequence in Sections 3.5.1 through 3.5.3. Their joint consideration is necessary before the structural interpretation framework of Chapter 4 can be applied, because a displacement gradient between two PSs on the same structure carries diagnostic meaning only if both the localisation error and the measurement precision are small relative to the gradient itself.

3.5.1. Geocoding and localization

The geographic position of each PS in the final output is determined by the geocoding transformation described in Section 3.3.3, which maps the range-azimuth pixel coordinates of the PS to a geographic coordinate pair using the orbital state vectors, the DEM, and any GCP correction applied during preprocessing. The localisation accuracy of this transformation depends on three factors: the precision of the orbital ephemeris, the vertical accuracy of the DEM used to account for topographic relief, and the quality of the GCP refinement. For Sentinel-1 with precise orbital state vectors and the Copernicus GLO-30 DEM, the expected horizontal localisation error is of the order of a few metres in flat terrain, increasing in areas of steep topography where DEM errors project more strongly onto the horizontal coordinate (Even & Schulz, 2020). For COSMO-SkyMed in Spotlight mode, the finer pixel spacing reduces the geometric uncertainty of the transformation itself, but the attribution of a PS to a specific structural position remains subject to the layover ambiguity discussed in Section 2.2.3. A PS detected at a given range-azimuth coordinate may correspond to any point along the layover column, from ground level to the top of the illuminated façade, and cannot be assigned to a specific floor or structural element without corroborating information from high-resolution optical imagery or field survey.

The practical consequence is that any attribution of a PS measurement to a structural element should be treated as a geometric hypothesis subject to verification, and localisation uncertainty should be reported explicitly alongside displacement estimates whenever PS measurements are used to infer displacement gradients over the height of a structure. External topographic knowledge can be

used to anchor the vertical position of individual PS where ground truth is available. In the *Sanctuary of Vicoforte* (Case study III of Section 5.4), we identified the apex cross of the dome as the highest scatterer of the structure through the residual height estimates. The known elevation of the apex cross above the base of the dome allowed the vertical rescaling of all other PSs associated with the monument, improving the attribution of measurement points to specific structural elements. This procedure illustrates how residual height information, when combined with architectural knowledge of the monitored structure, can mitigate the geocoding ambiguity inherent in the satellite geometry.

3.5.2. Referencing strategies

Whereas Section 3.4.2 addressed the operational selection of the reference point within the user-driven processing workflow, the present section focuses on the distinct question of how the resulting relative displacement field can be anchored to an external geodetic datum and of which strategies are available to quantify and reduce the residual referencing uncertainty. The InSAR measurement remains intrinsically differential, and the absolute accuracy of the results depends on the stability of the chosen reference. Three referencing strategies of increasing rigour are applicable depending on the available ground truth and the precision requirements of the application. The simplest strategy selects a reference PS on a surface assumed to be stable on the basis of geological or structural knowledge, without external validation. This approach is adequate for detecting relative displacement gradients within a compact study area, where the absolute velocity of the reference is irrelevant provided it remains constant over the observation period. A more rigorous strategy anchors the InSAR reference frame to one or more GNSS permanent stations by computing the difference between the InSAR LOS velocity at each station location and the GNSS-derived LOS projection of the three-dimensional velocity vector, then subtracting this offset from all PSs velocities across the project (Bähr & Hanssen, 2012). This correction removes any long-wavelength bias introduced by an imperfect reference choice and ties the displacement field to a geodetic datum, enabling comparisons with independent monitoring data and with results from other InSAR datasets processed with different references. The most complete strategy applies a spatially variable correction derived from a network of GNSS stations, accounting for large-scale orbital errors and broad atmospheric trends that a single-station correction cannot remove. In the absence of local GNSS stations, the EGMS Calibrated product described in Section 3.2.3 provides a pre-referenced velocity field that can serve as an external constraint for anchoring a user-processed dataset. The referencing

strategy adopted in each case study of Chapter 5 is documented explicitly, together with an estimate of the residual referencing uncertainty propagated into the final displacement products.

3.5.3. Precision, accuracy, and uncertainty propagation

The precision of a PS displacement time series, understood as the repeatability of the measurement under nominally identical conditions, is governed primarily by the phase noise at the PS location and by the quality of the atmospheric phase screen correction. For a PS with amplitude dispersion D_A , the standard deviation of the single-acquisition phase noise is approximately equal to D_A in radians under the high signal-to-noise ratio assumption adopted in (Ferretti et al., 2001). The corresponding LOS displacement noise is obtained through the phase-to-displacement relationship of Equation 3. 5, yielding a standard deviation equal to D_A multiplied by $\lambda/(4\pi)$ per epoch. For a well-selected PS with D_A of 0.1 radians and a Sentinel-1 wavelength of $\lambda \approx 5.55$ cm, this corresponds to approximately 0.44 mm of LOS displacement noise per acquisition. The noise accumulates in quadrature across the stack and yields a velocity precision of the order of 0.1 to 0.5 mm per year for a multi-year Sentinel-1 dataset, consistent with empirically reported values in the PSI literature (Crosetto et al., 2016). The atmospheric phase screen correction reduces the correlated component of this noise, but residual atmospheric artefacts remain the dominant source of uncertainty for individual epochs, particularly in areas with strong tropospheric gradients or near-field ionospheric disturbances (Fattahi & Amelung, 2014).

Accuracy, understood as the agreement between the InSAR estimate and the true displacement, is additionally affected by the referencing bias discussed in Section 3.5.2, by DEM errors projecting into residual topographic phase, and by any mismatch between the adopted displacement model and the actual deformation history of the target. Accuracy is distinct from precision: a dataset may be precise but inaccurate if a systematic bias is introduced by an unstable reference point, and may be accurate but imprecise if the atmospheric correction leaves large random residuals at individual epochs.

For structural monitoring, the most operationally relevant uncertainty metric is not the point-wise phase noise but the precision of the relative displacement between two PS on the same structure, because structural interpretation rests on displacement gradients rather than on absolute velocities. The uncertainty of a relative displacement estimate between two PS is dominated by the correlated

component of the atmospheric residual at the spatial scale of their separation, which decreases as the separation decreases, and by the uncorrelated phase noise of the two individual PS. The total variance of a relative displacement estimate is therefore approximately the sum of the two uncorrelated phase noise variances plus the residual atmospheric covariance term, which tends to zero as the spatial separation approaches the small-area limit discussed in Section 3.3.4. A formal propagation of these contributions into the uncertainty of derived structural metrics, such as differential settlement and tilt estimates, is developed in Section 4.3.3 of Chapter 4, where the engineering interpretation of displacement gradients is addressed in detail.

Chapter 4

Methodology for satellite-based SHM of structures

The processing workflow described in Chapter 3 produces, for each MT-InSAR analysis, a set of measurement points characterised by geographic coordinates, a residual height estimate, a mean LOS velocity, a displacement time series, and a temporal coherence value. These quantities are interferometric products: they describe the radar phase behaviour of individual scatterers over the observation period, referenced to a chosen spatial and temporal datum, and expressed in the geometric coordinate system of the satellite LOS. They do not, in themselves, constitute structural information. A mean LOS velocity of -3 mm/year on a building carries no engineering meaning until the analyst has established where on the structure the scatterer is located, whether the velocity represents ground motion or structural deformation, whether the seasonal component has been accounted for, and whether the value is consistent with the physical constraints of the structural typology. The present chapter develops the methodology for performing this transformation. Each section addresses a specific interpretive step. Section 4.1 covers the selection and geometric attribution of measurement points on built assets. Section 4.2 covers the decomposition of displacement time series into structurally interpretable components. Section 4.3 covers the derivation of relative motion metrics comparable with engineering damage thresholds. Section 4.4 covers the harmonisation of datasets from different sensors and geometries. Section 4.5 covers the integration of satellite observations with complementary information from in-situ monitoring, numerical models, and non-SAR remote sensing. Section

4.6 covers the definition of the diagnostic levels at which the resulting evidence can support engineering judgement.

4.1. Point selection on structures

The MT-InSAR processing chain of Chapter 3 produces a spatially distributed population of candidate PS that satisfies quantitative thresholds on amplitude stability and temporal coherence. From the perspective of the processing, each candidate is an equally valid measurement point. From the perspective of structural interpretation, however, the candidate population is not homogeneous. Some candidates originate from the structural elements that the monitoring programme seeks to characterise, while others arise from ground-level features, adjacent infrastructure, or transient objects whose phase behaviour carries no information about the structure under examination. Point selection on structures is the operation by which this distinction is drawn, transforming a processing-level population of coherent scatterers into an interpretation-level subset that can enter the diagnostic chain. The operation addresses three related problems: the physical characterisation of scatterers on built assets and the identification of their likely origin within the structural geometry (Section 4.1.1), the treatment of scatterers whose phase stability spans only a sub-interval of the stack and whose lifecycle may itself carry diagnostic meaning (Section 4.1.2), and the attribution of each retained scatterer to a specific structural element through geometric plausibility reasoning (Section 4.1.3).

4.1.1. Persistent scatterers on built assets

A PS on a built asset is a resolution cell whose scattering response is dominated by a single stable physical object over the entire observation period, producing an amplitude and a phase history that satisfy the diagnostic criteria established in Chapter 3. In terms of the indices introduced in Section 3.4.1, a PS candidate on a structure exhibits an ASI close to unity (equivalently, an amplitude dispersion D_A typically below 0.2) and a temporal coherence γ_p above the threshold adopted for the specific analysis, generally in the range 0.6 to 0.8 depending on the stack size and the deformation regime. The combination of these two conditions, stable amplitude and model-consistent phase, identifies cells in which the interferometric measurement can be interpreted as the response of a single physical reflector rather than as the stochastic output of a distributed scattering environment.

The radar backscattering intensity of a structural element, and therefore its likelihood of generating a usable PS, is governed by the interaction of four physical factors: material electromagnetic properties, surface geometry relative to the satellite LOS, surface roughness at the scale of the radar wavelength, and element dimensions relative to the resolution cell.

Material electromagnetic properties exert the strongest first-order control on backscatter intensity (Ulaby et al., 1986). Metallic surfaces, including steel parapets, aluminium cladding, iron railings, copper roofing, and galvanised fixtures, are strong radar reflectors because metals are electrically conductive and produce near-total backscattering of the incident microwave energy. Concrete and stone surfaces produce moderate backscatter that depends on surface roughness and moisture content; dry smooth concrete can behave as a specular reflector that directs energy away from the sensor, while rough or weathered concrete scatters more diffusely and returns a measurable signal. Masonry surfaces, particularly coursed stone and exposed brick, exhibit stable but moderate reflectivity because the joint pattern provides a consistent roughness at the centimetre scale that is well matched to both C-band and X-band wavelengths. Glass, polymer membranes, and thin composite panels are partially or wholly transparent to microwave radiation and may produce negligible backscatter despite being visually prominent in optical imagery. Vegetated surfaces, including green roofs, climbing plants on facades, and trees adjacent to the structure, decorrelate rapidly because the scattering elements (leaves, branches, stems) change position between acquisitions at a rate that destroys phase stability (Woodhouse, 2006).

Surface geometry determines whether the backscattered energy is directed toward the satellite antenna or away from it. Flat horizontal surfaces such as roof slabs and bridge decks act as specular reflectors that redirect energy away from the sensor except at near-normal incidence, which rarely occurs for spaceborne SAR. Vertical surfaces oriented perpendicular to the LOS direction, by contrast, can form dihedral reflectors in combination with the adjacent ground or horizontal surface, producing very strong returns. The classical dihedral geometry, in which a vertical wall meets a horizontal pavement at a right angle, generates some of the brightest PS observed in urban environments. Trihedral geometries, formed by the intersection of three mutually perpendicular surfaces (for example, a corner between two walls and a floor), produce even stronger returns that are stable over a wide range of incidence angles (Richards, 2009). Many of the PS identified on buildings in the case studies of Chapter 5 are physically associated with such corner geometries at the junction of façades, parapets, and roof edges.

Surface roughness interacts with wavelength to determine the scattering regime. A surface is electromagnetically smooth when its root-mean-square height variation is much smaller than the radar wavelength, and rough when the variation is comparable to or larger than the wavelength. At X-band ($\lambda \approx 3.1$ cm), even moderately textured concrete and dressed stone surfaces may satisfy the roughness criterion for diffuse scattering, while at C-band ($\lambda \approx 5.55$ cm) the same surfaces may behave as smoother reflectors. The practical consequence is that X-band sensors such as COSMO-SkyMed tend to produce denser PS populations on concrete and masonry structures than C-band sensors such as Sentinel-1, an effect that is observed in the comparative analyses of Chapter 5.

The dimension of the structural element relative to the resolution cell determines whether the element can produce a dominant scatterer. An element much smaller than the resolution cell contributes to the aggregate backscatter of that cell but may not dominate it; an element that fills the cell or produces a backscatter intensity substantially higher than the background can generate a point-like response that satisfies the PS selection criteria. Natural structural corners formed at the intersection of walls, parapets, and roof edges exploit this principle without any deliberate design, and constitute the main source of PS on existing built assets.

Where the natural density of PS on a structure is insufficient for the monitoring objective, artificial corner reflectors can be installed to guarantee the presence of a phase-stable measurement point at a known and precisely surveyed position. A corner reflector is a metallic device, typically a trihedral or dihedral configuration, designed to produce a strong and temporally stable backscatter response within a single resolution cell. Figure 4. 1 shows an example of a corner reflector installed on the *Sanctuary of Vicoforte*.




Figure 4. 1. Example of an artificial corner reflector mounted on the *Sanctuary of Vicoforte* to enhance the radar visibility of selected structural points and to support the identification of stable measurement targets in the multi-temporal interferometric stack.

The device dimensions are calibrated to the radar wavelength: for X-band applications, a trihedral with an inner edge length of approximately 30 to 50 cm produces a radar cross-section of the order of 10^3 square metres, sufficient to dominate the surrounding clutter at typical urban signal levels (Garthwaite, 2017). Corner reflectors have been deployed on bridges, dams, and heritage structures where natural PS coverage is sparse or geometrically ambiguous, and their use is explicitly foreseen in the monitoring programme of the *Sanctuary of Vicoforte* (Section 5.4.2). The installation transforms the satellite observation from an opportunistic measurement on targets of unknown physical origin into a controlled measurement at a structurally defined location, with known coordinates and a predictable phase centre. Two operational considerations limit indiscriminate deployment. The first is that a reflector installed within a resolution cell that already contains strong natural scatterers may produce only a marginal increase in backscatter intensity, because the existing clutter dominates the cell amplitude; reflector placement is therefore most effective when directed toward cells with otherwise weak or unstable scattering. The second is that corner reflectors require physical access to the structure, a stable mounting surface, protection against vandalism and environmental degradation, and periodic verification of alignment. Their deployment is accordingly justified on structures of high importance where long-term dedicated monitoring is planned and where natural PS density has been assessed and found insufficient through a preliminary InSAR analysis.

A practical distinction of operational relevance concerns PS located at different positions along the structural height. PS on pavements or low-rise elements near the building perimeter tend to reflect the combined movement of the soil-structure

system at foundation level, whereas PS on upper façades or roof-level metallic elements incorporate the same foundation movement plus any additional deformation accumulated over the structural height. The ReLUIS guidelines for the use of interferometric satellite data in structural applications recommend that the selection of PS for building-scale analysis always be accompanied by an explicit assessment of the vertical position of each point, distinguishing between ground-level and elevated PS before computing differential displacements (Prota et al., 2025). When the observation period is sufficiently long and the spatial density of PS is adequate, the vertical stratification of PS on a structure allows the analyst to separate foundation settlement from superstructure deformation, providing a richer diagnostic input than a single aggregated velocity estimate for the building as a whole.

The distribution of PS on a given structure depends on the acquisition geometry as well as on the material and geometric properties described above. A facade oriented towards the satellite generates PSs from its exposed surface, while a facade oriented away from the satellite produces no detectable coherent scatterers regardless of its material composition. For Sentinel-1 in IW mode, the relatively coarse azimuth resolution of approximately 20 m limits the capacity to resolve individual structural elements, and many PSs on buildings integrate the response of multiple objects within a single resolution cell. In contrast, COSMO-SkyMed in Stripmap HIMAGE mode resolves scatterers at the scale of individual façade elements, and PS density on building façades may be sufficient to characterise displacement gradients at the floor level. The engineering consequence of this density difference is that Sentinel-1 data are better suited to characterising displacement patterns at the building or block scale, while X-band high-resolution data offer the spatial selectivity required for element-level structural diagnosis. In both cases, the attribution of a PS to a structural element rather than to adjacent ground or infrastructure requires cross-referencing the geocoded PS position with a geometric model of the structure, as discussed in Section 4.1.3.

4.1.2. Temporary scatterers and target lifecycle

The PS framework described in Section 4.1.1 rests on the assumption that each retained measurement point maintains phase stability over the full observation window. Not all coherent scatterers on a built structure satisfy this assumption. A class of targets generates a stable and reliable phase response only during a sub-interval of the stack, appearing and disappearing at datable epochs that correspond to physical changes in the scattering environment. The present thesis formalises this

class of targets as a systematic diagnostic category, designating them Temporary Scatterers (TS) and developing a framework that generalises the PS concept by admitting time-bounded coherence windows as a legitimate analytical state. The TS framework constitutes one of the methodological contributions of this thesis to the multi-temporal interferometric processing of structures.

Recognition that a radar target may remain phase-coherent over only part of an observation interval did not originate in structural engineering. It arose within the geophysical and geodetic InSAR community. The persistent scatterer formulation retains only targets that remain stable across the entire stack (Ferretti et al., 2001) and discards those coherent over shorter intervals. Subsequent work questioned that exclusion. The SqueeSAR algorithm recovered partially coherent ground through a distributed-scatterer model (Ferretti et al., 2011). Within the same line, temporary coherent scatterers, coherent only inside a finite temporal window, were treated as an exploitable signal rather than as noise (Novali et al., 2011). Hu et al. (2019) later placed the idea on a systematic footing through adaptive temporal subsets, estimating the coherent interval individually for each target before solving its displacement. Across these formulations the application domain remained geophysical: subsidence, landslides, and vegetated terrain, where coherence over the full stack is the exception.

In each of those domains the temporariness of a scatterer was treated as an obstacle to overcome, a way to extend coverage over ground that conventional processing would leave unmeasured. By itself, the temporal limit of coherence carried no interpretive weight. The present thesis reverses this emphasis. On a structure, the epoch at which a target ceases to register can coincide with the physical removal or destruction of the surface that produced it. Conversely, the epoch at which a new target begins to register can mark the emergence of debris or of a temporary construction element. The boundaries of the coherent interval therefore fix discrete structural events in time at the resolution of the satellite revisit cycle, a function the geophysical antecedents were never designed to serve. Read in this way, a time-bounded target is not a measurement to be salvaged but a record of structural change. Earlier exploitation of time-bounded coherence remained within geophysical and geodetic monitoring; its formalisation as a parametrised structural-diagnostic category appears not to have been undertaken systematically before. What makes such a target diagnostic is therefore not the displacement it records but the timing of its coherence: the epochs at which it begins and ceases.

The lifecycle of each scatterer is described by two epoch indices: T_{on} , the first acquisition at which the scatterer exhibits coherent behaviour, and T_{off} , the last acquisition at which coherence is maintained. A PS is the special case in which $T_{on} = 1$ and $T_{off} = N$, where N is the total number of acquisitions in the stack: the scatterer is phase-stable from the first image to the last. A TS is any point for which $T_{on} > 1$ or $T_{off} < N$, or both. The framework therefore unifies the PS and TS concepts within a single temporal description, with the PS as the limit case in which the coherence window coincides with the observation window. In the implementation adopted here, the T_{on} and T_{off} indices are computed through an on-off model of the amplitude time series that identifies the acquisitions bounding the live interval; these two indices are exported as the fields “in” and “fin” associated with each point, enabling the analyst to identify which portion of the temporal axis each displacement time series covers.

Three operationally distinct sub-types arise from this definition. The first sub-type ($T_{on} > 1$, $T_{off} = N$) describes a scatterer that appears during the observation window and remains present until the end: a newly installed metallic fixture, a construction element erected during a building intervention, or post-event debris that creates a new reflecting surface. The second sub-type ($T_{on} = 1$, $T_{off} < N$) describes a scatterer that is present at the start of the observation window but disappears before its end: a structural element removed during demolition, a surface destroyed by collapse, or a reflecting configuration altered by restoration works. The third sub-type ($T_{on} > 1$, $T_{off} < N$) describes a scatterer whose entire lifecycle falls within the observation window, bounded on both sides: a temporary construction element installed and subsequently removed, or a reflecting surface that exists only during an intermediate phase of a structural event sequence. The four lifecycle configurations of the framework are summarised by their amplitude profile across the acquisition sequence. The Persistent Scatterer maintains a high and stable amplitude throughout the stack, while each of the three Temporary Scatterer sub-types exhibits a distinct on-off pattern bounded by the T_{on} and T_{off} epochs of its live interval, as shown in Figure 4. 2.




Figure 4. 2. Schematic representation of the four scatterer lifecycle configurations, expressed through the amplitude profile across the acquisition sequence. Panel (a): Temporary Scatterer of the first sub-type ($T_{on} > 1$, $T_{off} = N$), in which the scatterer appears at a datable epoch within the observation window and remains coherent until the end of the stack. Panel (b): Temporary Scatterer of the second sub-type ($T_{on} = 1$, $T_{off} < N$), in which the scatterer is present from the first acquisition and disappears at a datable epoch before the end of the stack. Panel (c): Temporary Scatterer of the third sub-type ($T_{on} > 1$, $T_{off} < N$), in which the entire coherence window of the scatterer falls within the observation period, bounded by datable T_{on} and T_{off} epochs on both sides. Panel (d): Persistent Scatterer ($T_{on} = 1$, $T_{off} = N$), with high and stable amplitude maintained throughout the entire stack.

In the case studies of Chapter 5, the second sub-type dominates the collapse scenarios investigated in the *Himera Viaduct* and *Albiano-Magra Bridge* analyses, where TS on the failed decks (*Albiano-Magra*) and on the post-event debris adjacent to the structure (*Himera*) exhibit T_{off} values coinciding with the collapse or demolition dates. The first sub-type is prevalent in the post-event phase of the Palisades Fire study, where new scatterers emerge on fire-generated debris and on reconstructed elements, with T_{on} values clustered in the weeks following the fire. The third sub-type is comparatively rare in the case studies, occurring on scatterers

linked to transient interventions such as temporary shoring during restoration works or cleared debris piles.

A preliminary indication of TS behaviour can be obtained before the MT-InSAR inversion by examining the amplitude channel of the SAR stack through multi-channel colour composites. The technique assigns three acquisitions from different epochs to the red, green, and blue channels of a single RGB image. A target present in all three epochs produces a white or grey pixel, indicating consistent backscatter across the observation window. A target present in only one epoch produces a pixel coloured in the corresponding primary: red for the earliest acquisition, green for the intermediate, blue for the latest. A target that disappears between the first and third epoch produces a red pixel at its location, indicating that backscatter was present at the start of the observation window but absent at the end. The complementary case, a target that appears between the first and third epoch, produces a blue pixel. Mixed colours indicate partial presence across the three dates.

When applied to a structural monitoring scenario, the RGB composite provides an immediate visual survey of scattering changes on and around the structure. Construction scaffolding installed between the first and second epoch appears as a green or cyan feature at the scaffold location. A building demolished between the second and third epoch appears as a red or yellow feature at the former building footprint. The technique does not replace the formal TS identification through the sparse-point inversion described below, because it operates on amplitude only and does not exploit phase stability. However, it serves two preparatory functions that improve the efficiency of the subsequent analysis. The first is spatial guidance: the RGB composite indicates where on the structure or in the surrounding area temporal changes have occurred, allowing the analyst to focus the small-area processing on the zones most likely to contain TS. The second is temporal bracketing: the colour of a changed feature constrains the epoch of the change to the interval between the corresponding acquisition dates, providing an initial estimate of the TS lifecycle window that can be refined through the phase-based analysis.

The application of the technique to a representative scene is illustrated in Figure 4. 3, which combines the amplitude maps of three successive acquisitions over a generic port area with the corresponding RGB composite. The sea surface appears uniformly dark across the three epochs because the smooth water surface at the operating incidence angle behaves as a specular reflector and redirects almost all incident energy away from the antenna, producing the black background visible in

the composite. Built structures with stable corner-reflection geometry produce strong and temporally consistent returns and appear as white pixels in the composite. Coloured pixels mark the locations where the backscatter response has changed between the three acquisition epochs, with the colour distribution allowing the temporal bracketing of the change. The composite therefore provides at a glance the spatial distribution of changing scatterers on and around the structures of interest and supports the subsequent identification of Temporary Scatterer candidates through the phase-based analysis described in the remainder of this section.




Figure 4. 3. Multi-channel change detection through RGB amplitude composite over a representative urban port area. Panels (a), (b), and (c): mean amplitude maps computed from three SAR acquisitions at distinct epochs, with brightness proportional to the

backscattered intensity of each scene. Panel (d): RGB composite obtained by assigning the amplitude maps of panels (a), (b), and (c) to the red, green, and blue channels respectively. Targets present in all three epochs appear as white or grey pixels, while targets present in only one or two epochs produce coloured features whose hue identifies the temporal interval of their appearance or disappearance.

A second preliminary visualisation technique exploits the interferometric coherence rather than the amplitude. While the amplitude composite described above detects changes in the backscattered intensity, the coherence between two acquisitions measures the stability of the complex phase within the estimation window. The two quantities carry partially independent information: a structural surface may maintain a constant amplitude because its material and geometry remain unchanged, while losing coherence because the distribution of minor scatterers within the resolution cell has changed due to surface weathering, moisture variation, or microstructural modification. Conversely, a surface may exhibit a change in amplitude because a new element has been added or an existing one removed, while the adjacent pixels remain coherent. The application of the technique to the same port area used in Figure 4. 3 is shown in Figure 4. 4. Panels (a), (b), and (c) report the coherence maps of three interferometric pairs spanning successive intervals of the observation window, and panel (d) presents the corresponding RGB composite. The composite identifies surfaces whose phase coherence has changed between the three intervals, including infrastructure elements with stable amplitude that nonetheless exhibit progressive loss of coherence and that would not be detected by the amplitude-based composite of Figure 4. 3. A coherence change detection composite, constructed by assigning coherence maps computed from interferometric pairs at different epochs to the R, G, B channels, can reveal changes that are invisible in the amplitude channel and that may indicate surface degradation, moisture infiltration, or physical modification at a scale below the amplitude detection threshold (Lu et al., 2022). The coherence composite therefore provides a complementary screening layer that can guide the subsequent phase-based TS analysis toward structural elements undergoing subtle scattering changes.




Figure 4. 4. Coherence change detection composite on the same representative port area of Figure 4. 3. Panels (a), (b), and (c): coherence maps of three interferometric pairs computed on intervals of progressively later epochs within the observation window. Panel (d): RGB composite obtained by assigning panels (a), (b), and (c) to the red, green, and blue channels of a single image. Pixels with high coherence across all three pairs appear as white or light grey in the composite, while coloured pixels identify surfaces whose coherence is preserved in only a subset of the three intervals.

Beyond the visual screening provided by the amplitude composite of Figure 4. 3 and by the coherence composite of Figure 4. 4, the candidate population for the TS analysis can be pre-filtered before the formal two-state fitting through a quantitative criterion based on the ASI, a computationally inexpensive screening criterion. The two visualisation techniques discussed above identify spatial regions of interest within the scene, while the ASI-based filter operates on the sparse-point

candidate population produced by the multi-temporal interferometric processing of Chapter 3, and reduces the candidate set to the points most likely to exhibit TS behaviour. PS occupy the high-ASI tail of the distribution, with values typically above 0.6 to 0.7 for well-selected targets on built assets. Pixels with ASI close to zero exhibit amplitude behaviour indistinguishable from random noise and are not candidate TS. TS occupy a characteristic intermediate band of the distribution: their amplitude is not stable across the full stack (which would produce a high ASI), but it is not randomly variable either (which would produce an ASI near zero). The two-state pattern produces a moderate standard deviation driven by the amplitude contrast between the live and dead intervals, yielding ASI values typically in the range 0.0 to 0.5. Pixels with negative ASI, as discussed in Section 3.4.1, correspond to extreme cases in which the standard deviation exceeds the mean, typically because a strong scatterer is present in a small number of acquisitions and absent in the majority; these are candidate TS with very short live windows relative to the full stack. Selecting candidate points within the intermediate ASI band and the negative ASI tail, while excluding the high-ASI PS population, reduces the candidate set to the points most likely to exhibit TS behaviour before the computationally more demanding two-state fitting is applied.

The spatial distribution of the ASI over the candidate population and the corresponding histogram are reported in Figure 4. 5, which complements the amplitude composite of Figure 4. 3 and the coherence composite of Figure 4. 4 by extending the screening from the full SAR scene to the sparse-point candidate set. The map in panel (a) adopts a colour scale from zero to unity, with negative ASI values, corresponding to pixels for which the temporal standard deviation of the amplitude exceeds the temporal mean, collapsed onto the zero level and distinguished by a separate colour. The histogram in panel (b) quantifies the relative population of each ASI regime and supports the choice of the filtering thresholds discussed above.




Figure 4. 5. ASI of the candidate population over the same representative port area of Figure 4. 3. Panel (a): spatial map of the ASI, with the value at each candidate point reported on a colour scale ranging from zero for highly variable amplitude records to unity for stable persistent scatterers. Negative ASI values, which arise when the temporal standard deviation of the amplitude exceeds the temporal mean, are collapsed onto the zero level for the purpose of the colour mapping, and a distinct colour is assigned to identify their occurrence on the map. Panel (b): histogram of the ASI distribution across the same candidate population.

The formal identification of TS from the multi-temporal amplitude record is performed through a two-state fitting procedure applied to each candidate point independently. The procedure partitions the amplitude time series into two intervals: a live interval in which the mean amplitude is significantly above the background level, and a dead interval in which the mean amplitude is at or near the noise floor. For each candidate partition boundary, corresponding to each possible acquisition epoch within the stack, the algorithm computes the mean amplitude in the two resulting intervals and evaluates the goodness of fit of the two-state model against the observed amplitude sequence. The optimal partition is the one that maximises the contrast between the two states while minimising the residual variance within each state. The fit quality, expressed as a normalised metric, quantifies how well the observed amplitude behaviour conforms to the idealised on-off pattern and serves as the primary criterion for accepting or rejecting the TS classification. A high fit quality indicates a clear transition between two distinct amplitude regimes, consistent with a physical event that created or destroyed the scattering surface. A low fit quality indicates either a gradual amplitude change that does not conform to the two-state assumption, or a noisy amplitude record in which

no clear transition is identifiable. The two-state model extends naturally to a multi-state formulation for scatterers that exhibit more than one lifecycle transition: a three-state model accommodates a scatterer that is absent, appears, and subsequently disappears (third sub-type), by fitting three amplitude regimes separated by two transition boundaries. The computational cost of the multi-state search increases with the number of states, but for typical stack lengths encountered in structural monitoring (50 to 200 acquisitions) the processing remains tractable for candidate populations of several thousand points.

The detection procedure can be illustrated through a representative didactic scenario. Consider a hypothetical multi-temporal SAR stack of 80 acquisitions spanning five years over a bridge undergoing maintenance. During the first three years, a metallic temporary barrier is installed along the deck edge, producing a strong and stable radar return in the corresponding resolution cell. The barrier is removed at the start of the fourth year, and the amplitude at that location drops to the background level for the remaining acquisitions. The ASI for this pixel is approximately 0.3, placing it within the TS candidate zone. The two-state fitting identifies the optimal transition boundary at the acquisition closest to the removal date, producing a high fit quality. The resulting T_{off} provides an interferometric estimate of the removal date with a temporal resolution equal to the satellite revisit interval.

The lifecycle parameters T_{on} and T_{off} identified through the two-state amplitude fitting are not only descriptive metadata: they enter the interferometric estimation as operational weights. When the TS lifecycle has been determined, the periodogram search described in Section 3.1.2 is restricted to the acquisitions within the interval $[T_{\text{on}}, T_{\text{off}}]$, excluding all images in which the scatterer was not physically present. The differential phases corresponding to the dead intervals are omitted from the summation in Equation 3.12, and the temporal coherence is computed over the reduced set of M' observations corresponding to the live interval only, with M' less than the total number of interferograms M . The effect is twofold. First, the temporal coherence increases on average because the incoherent observations from the dead intervals are excluded. These observations would otherwise contribute random phase residuals and depress the coherence. The increase is not deterministic for any individual point, because the residual phases within the live interval may themselves be noisy. The expected value of the coherence under the null hypothesis of random residuals is reduced according to the relationship between expected coherence and stack size discussed in Section 3.4.3. The calibration of the coherence threshold should be adjusted accordingly

when interpreting TS-derived measurements. A point that has entered the candidate population through the ASI-based pre-filtering described above and that was subsequently rejected as incoherent over the full stack may yield a high temporal coherence when the estimation is restricted to its live window, revealing displacement information that would otherwise be lost. Points that have not entered the TS candidate population, on the other hand, remain rejected by the standard PSI workflow and are not recovered. Second, the displacement time series produced for the TS is bounded by T_{on} and T_{off} , so that no displacement values are estimated or reported outside the live interval, ensuring that the output represents only the period during which the phase measurement was physically meaningful.

Within the processing workflow, the amplitude-temporal baseline plot associated with each candidate PS provides a complementary diagnostic for confirming TS behaviour identified through the two-state fitting. A scatterer whose amplitude is stable for short temporal baselines from the master but decays for longer baselines, forming a compact high-amplitude cluster at short separations and a dispersed low-amplitude cloud at long separations, is consistent with a target that remains coherent only during a specific subset of the observation period. The position of this subset along the timeline depends on the lifecycle of the individual scatterer relative to the master acquisition epoch, and may or may not include the master itself; the master is selected to optimise the amplitude stability across the entire stack rather than to fall within the live window of any specific TS, so the diagnostic must be interpreted in light of this distinction. The temporal coherence estimated over sub-periods of the full stack, rather than over the entire observation interval, can be used to refine the lifecycle boundaries initially identified by the amplitude fitting and to quantify the phase quality within the live interval, which determines whether the TS displacement time series is usable for structural interpretation.


The application of the workflow is applied to a representative high-rise building over an urban area. A X-band dataset of seventy-two images is processed in user-driven mode. The building was under construction at the start of the observation period and reached completion during the dataset, providing a case in which the chronology of the structural surfaces is documented a priori and can be compared with the output of the two-state fitting. The three-dimensional view of the building reported in panel (a) of Figure 4. 6 shows the PS and TS identified on the structure, colour-coded by the estimated T_{on} . The PS, characterised by T_{on} coincident with the first acquisition of the stack, are concentrated on the left wing of the building and

on the adjacent pre-existing urban context, while the TS, characterised by T_{on} successive to the first acquisition, are distributed on the right wing of the building.


The diagnostic confirmation of the TS classification through the amplitude-temporal baseline plot is illustrated in panel (b) of Figure 4. 6 for a representative TS selected from the upper portion right wing. The plot exhibits the diagnostic signature anticipated in the preceding paragraphs: the high-amplitude observations cluster in a compact region of the timeline that coincides with the live interval bounded by the estimated T_{on} and T_{off} , while the observations outside that region exhibit amplitudes close to the noise floor. The compact cluster does not coincide with the master acquisition epoch but is offset by approximately one year, consistent with the construction date of the corresponding storey, and confirms that the master selection criterion is independent of the lifecycle of any specific TS.

The displacement time series of the same representative TS is reported in panel (c) of Figure 4. 6. The time series is restricted to the live interval identified by the two-state fitting.


The physical cause of a TS lifecycle transition, whether the appearance of the scatterer at T_{on} or its disappearance at T_{off} , cannot be determined from the amplitude and phase record alone. A scatterer that emerges at a datable epoch may correspond to the construction of a new structural element, the installation of temporary works, or the growth of a stable reflecting surface from environmental processes, while a scatterer that loses coherence at a datable epoch may have been affected by structural failure, routine maintenance, environmental modification, or seasonal vegetation growth. Additional contextual information, such as maintenance records, photographic surveys, inspection reports, or LiDAR point clouds acquired before and after the event, is required to assign a physical cause to the observed lifecycle transition. The retrospective event analysis methodology developed in Section 4.3.4 provides the interpretive framework for this attribution in the specific context of structural failure investigation.



(a)



(b)



(c)

Figure 4. 6. Application of the TS workflow to a high-rise building under construction during the observation period. Panel (a): three-dimensional view on Google Earth of the PSs and TSs identified on the building, colour-coded by the estimated T_{on} , with PSs appearing at the first acquisition shown in dark blue and TSs appearing at successive epochs shown in progressively warmer colours. Panel (b): amplitude-temporal baseline plot of a representative TS selected from the upper portion of the new building, with the high-

amplitude cluster identifying the live interval bounded by the estimated T_{on} and T_{off} . Panel (c): displacement time series of the same representative TS restricted to the live interval. In panel (b), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panel (c), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.

4.1.3. Geometric plausibility: roofs, facades, edges, and metallic elements

The geocoded position of a PS provides a planimetric coordinate pair and a residual height estimate relative to the external DEM. As discussed in Section 3.5.1, however, the layover geometry of side-looking SAR means that a PS detected at a given range-azimuth coordinate may correspond physically to any point along the layover column projected between the satellite and the ground. On a building of ten storeys, the layover column may span the full height of the facade, and a PS nominally located at mid-facade in the geocoded output could originate from ground level, from an intermediate floor, or from the rooftop parapet, depending on the actual three-dimensional position of the dominant scatterer within the resolution cell. Geometric plausibility assessment is the procedure by which this ambiguity is reduced, if not fully resolved, by combining the geocoded PS position with prior knowledge of the structure's geometry and the satellite acquisition parameters.

The assessment follows a three-step logic. First, the incidence angle and heading of the satellite are used to compute the projection of each structural surface onto the LOS direction, identifying which faces, edges, and elements are visible to the sensor and which fall in shadow or layover. Second, the residual height estimate produced by the PS time-series inversion, which represents the difference between the true elevation of the scatterer and the DEM elevation at that planimetric position as established in Section 3.1.2, is compared against the expected height of candidate structural elements such as parapet tops, floor slab edges, or metallic rooftop features. Agreement between the residual height and the expected structural elevation provides corroborating evidence for the attribution. Third, the spatial clustering of multiple PSs at similar residual heights and within the same geometric sector of the building suggests a common physical origin and increases the confidence of the attribution. Where high-resolution optical imagery or LiDAR data are available, the three-dimensional coordinates of the PSs can be compared directly against the point cloud or building model, providing a more objective basis for the attribution than geometric reasoning alone. In the absence of such data, the analyst must document the attribution explicitly as a geometric hypothesis, stating

the assumed physical origin of each PS and the evidence supporting that assumption. This documentation is a prerequisite for the structural interpretation steps described in Section 4.3, where relative displacements between PS attributed to different structural positions are used to derive gradient and tilt estimates.

The geometric plausibility assessment can be illustrated through a representative didactic scenario that applies the three-step logic to a hypothetical building observed in both ascending and descending geometry. Consider a rectangular building with its longer axis oriented north-south, located in a mid-latitude European city. The building has four vertical facades (east, west, north, south) and a flat roof. In descending orbit geometry, a right-looking Sentinel-1 satellite flies from north to south and observes from east to west. The incidence angle at mid-swath is approximately 39 degrees from the vertical. Under these conditions, the east-facing facade is oriented toward the satellite and is illuminated at an oblique angle: the radar signal strikes the facade and the adjacent ground simultaneously, producing a dihedral reflection that generates strong backscatter. The west-facing facade is oriented away from the satellite and falls in the radar shadow of the building itself: no signal reaches this surface, and no PS can originate from it. The flat roof is visible but produces weak backscatter because horizontal surfaces redirect most of the incident energy away from the sensor at the prevailing incidence angle. Metallic fixtures on the roof, such as ventilation units, antenna mounts, or railing posts, may act as isolated strong scatterers whose phase centre is located at the fixture height above the roof surface. The north and south facades are oriented parallel to the satellite flight direction and are nearly perpendicular to the range direction; they contribute minimally to the backscatter because neither dihedral nor trihedral geometries are formed with respect to the LOS.

In ascending orbit geometry, the same satellite flies from south to north and observes from west to east. The illumination geometry is reversed: the west-facing facade is now toward the satellite and produces the dihedral reflection, while the east-facing facade falls in shadow. The roof conditions remain essentially unchanged, and the north-south facades remain geometrically unfavourable.

The practical consequence for structural interpretation is that ascending and descending datasets sample different portions of the same building. PSs identified in the descending analysis are physically located on or near the east-facing facade and on roof-level metallic elements; PSs identified in the ascending analysis originate from the west-facing facade and from the same category of roof fixtures. Combining the two geometries increases the spatial coverage of the building but

does not produce overlapping measurements on the same physical surface. When the analyst interprets a spatial velocity gradient between two PSs on the same building, the gradient must be evaluated in the context of which facade each PS originates from: a velocity difference between an east-facade PSs (descending) and a west-facade PSs (ascending) may reflect differential structural behaviour but may also reflect the geometric projection of a uniform motion vector onto two different LOS directions.

For buildings whose principal facades are not aligned with the cardinal directions, the illumination conditions are intermediate and must be computed from the specific heading angle and incidence angle of the orbit. The general principle remains: the analyst must determine, before interpreting the displacement data, which structural surfaces are geometrically visible in each orbit and where the PS are physically located on the building.

A further source of uncertainty in the physical attribution of a PS to a specific structural element arises from the concept of electromagnetic barycentre. When multiple reflecting surfaces within the same resolution cell contribute to the backscattered signal with comparable intensities, the recorded phase does not correspond to any single physical object but to the amplitude-weighted average of the contributing phase centres. The geocoded position of the resulting PS, derived from the estimated residual height and the LOS geometry, represents this weighted average rather than the position of a discrete structural feature. On a building facade, for example, a window frame, a balcony railing, and a downpipe may all contribute to the backscatter within a single Sentinel-1 resolution cell. The PS geocoded at that location represents the electromagnetic barycentre of these three contributions, which may not coincide spatially with any of the individual objects. For COSMO-SkyMed, where the finer resolution cell of approximately $3\text{ m} \times 3\text{ m}$ reduces the number of contributing elements, the barycentre effect is less pronounced but not absent, particularly on facades with closely spaced metallic fixtures.

The barycentre effect has two practical implications for structural interpretation. First, the physical origin of a PS on a building cannot be determined unambiguously from its geocoded coordinates alone; the analyst must consider all plausible contributing surfaces within the resolution cell and assess which is likely to dominate on the basis of material, geometry, and orientation. Second, the barycentre may shift between acquisitions if the relative contributions of the individual scatterers change, for example due to the installation or removal of

temporary elements such as scaffolding, signage, or construction equipment. Such a shift would manifest as a phase change that could be misinterpreted as structural displacement if the barycentre migration is not recognised. The amplitude time series of the PS, examined alongside the displacement time series, can provide diagnostic evidence of barycentre instability: a sudden change in amplitude without a corresponding change in the displacement trend suggests a modification of the scattering configuration rather than a structural event. The geometric plausibility logic illustrated through the didactic scenario above is operationalised in the case studies of Chapter 5, where the orientation of the building facades, the satellite heading, and the incidence angle of each acquisition are reconstructed from the documentation of the structure and from the orbital parameters of the dataset.

4.2. Time-series analytics for structural interpretation

The displacement time series extracted from a MT-InSAR analysis is not a direct measurement of structural behaviour. It is a sampled record of LOS range change that contains, in superposition, contributions from distinct physical processes operating at different temporal scales. A linear trend, if present, may reflect a sustained geotechnical process such as consolidation settlement, long-term creep, or progressive structural damage. A periodic oscillation at the annual frequency is typically driven by thermal expansion and contraction of structural materials, by seasonal soil moisture cycles, or by both mechanisms acting jointly. Departures from these regular patterns, whether abrupt or gradual, may indicate a change in the physical state of the structure or its foundation, and their detection and characterisation constitute the analytical objective most directly relevant to SHM. Treating the raw time series as a single undifferentiated signal and extracting only a mean velocity, as is done in standard EGMS products, collapses this information into a single scalar and discards the diagnostic content carried by the higher-order temporal structure. A decomposition-based approach is therefore necessary whenever satellite data are intended to support structural assessment rather than simple screening. The three components of the decomposition, their physical interpretation, and their estimation in the structural monitoring context are addressed in Sections 4.2.1 through 4.2.3.

4.2.1. Linear trends

The mean LOS velocity, estimated by fitting the linear component of the parametric model to the displacement time series of each PS over the full observation period, is the primary output of the PSI inversion and the quantity reported in most

published InSAR products, including the EGMS. In the parametric model adopted in this thesis, the linear velocity is estimated jointly with the residual height and the thermal coefficient, and the simultaneous resolution of the three contributions reduces the variance of the velocity estimate compared with the simpler two-parameter case. In the structural monitoring context, a statistically significant negative velocity, meaning a displacement in the range-increasing direction, indicates that the target is moving away from the satellite at a sustained rate and may be consistent with ongoing settlement or differential compaction of the foundation system. A positive velocity suggests uplift or shortening of the LOS path and may reflect swelling, frost heave, or rebound following load removal. The internal precision of the linear velocity estimate improves with the three-halves power of the observation period, as derived from the stochastic model of the parametric inversion (Hanssen, 2001), which is the principal reason why multi-year stacks are preferred for the detection of slow deformation rates of the order of 1 mm per year or less. For a coherent persistent scatterer drawn from a ten-year Sentinel-1 archive and processed with effective atmospheric phase screen mitigation, the internal precision of the linear velocity is of the order of 0.1 to 0.3 mm per year (Crosetto et al., 2016; Ferretti et al., 2001). The absolute accuracy validated against external geodetic networks remains at the level of 1 to 2 mm per year (Cigna et al., 2021; Martins et al., 2024; Vrati et al., 2023). Slow foundation movements are accordingly within the detection capability of the technique, provided that the precision and accuracy figures are not conflated.

However, the linear model is an adequate representation of the time series only when the deformation process is genuinely steady-state over the observation window. Many structural and geotechnical processes are inherently non-linear. Primary consolidation decelerates over time following a hyperbolic or logarithmic law. Damage accumulation may accelerate as a structure approaches a critical state. Load changes introduced by adjacent excavations or phased construction produce step-like responses superimposed on any background trend. Fitting a single linear function to a time series that contains these features yields a mean velocity that is a weighted average over a non-stationary process, and the resulting estimate may be misleading if interpreted as a current rate of movement. The residuals of the linear fit, defined as the difference between the observed displacement at each epoch and the fitted linear model, therefore carry information that the mean velocity alone cannot convey. A systematic pattern in the residuals, such as a concave-upward curvature or a progressive increase in scatter toward the end of the record, may indicate that the linear approximation is inadequate and that a more flexible model is required. The assessment of residual behaviour is accordingly a necessary step

before the linear velocity estimate is used for structural diagnosis, and it provides the analytical motivation for the non-linear and change-point approaches addressed in Section 4.2.3.

4.2.2. Seasonal and thermal components

Periodic oscillations at the annual frequency are present in the displacement time series of virtually all PSs located on built structures in temperate climates, and they typically constitute the largest-amplitude signal component after the linear trend has been removed. Two physically distinct mechanisms contribute to this annual cycle. The first is the direct thermal response of structural materials: steel, reinforced concrete, and masonry all undergo cyclic elongation and contraction proportional to the temperature change and the material thermal expansion coefficient, producing a LOS displacement that follows the annual temperature cycle with a phase lag determined by the thermal inertia of the structure. For a steel bridge deck of approximately 30 m length subjected to a 30-degree Celsius annual temperature range, the linear thermal expansion amounts to approximately 10-11 mm (computed as $\alpha \cdot L \cdot \Delta T$ with $\alpha \approx 12 \times 10^{-6}$ per °C), which projects onto LOS amplitudes of the order of 10 to 15 mm depending on the geometry of the deck relative to the satellite viewing direction (Qin et al., 2018). On masonry heritage buildings, the thermal expansion coefficient is lower, typically between 5×10^{-6} and 10×10^{-6} per degree Celsius, but the resulting annual amplitude remains well above the detection threshold of Sentinel-1 and COSMO-SkyMed for structures of moderate height (Fornaro et al., 2013; Monserrat et al., 2011). The second mechanism is seasonal environmental forcing transmitted through the foundation: variations in soil moisture, groundwater depth, and the volumetric behaviour of near-surface soils generate a seasonal displacement pattern at the structure's base that propagates upward and is recorded by PSs at all heights. The two mechanisms are distinguishable in principle by their phase relationship with the temperature record: a purely thermal structural response is nearly in phase with the temperature cycle, whereas a soil-driven seasonal response may lag by several weeks or months depending on the hydraulic diffusivity of the foundation soil.

The standard parametric approach to estimating the periodic component models the annual oscillation as a sinusoidal function of known frequency, fitting its amplitude and phase to the time series jointly with the linear trend. Where an independent temperature record is available from a nearby meteorological station or from a weather data service integrated within the processing environment, a more physically grounded model can be adopted, regressing the LOS displacement

directly onto the temperature variable and recovering a thermal sensitivity coefficient expressed in millimetres per degree Celsius (Monserrat et al., 2011). The thermal sensitivity coefficient carries structural information in its own right: an anomalous value relative to the expected expansion coefficient of the structural material may indicate a change in boundary conditions, the presence of a crack that modifies the effective free length of a structural element, or a foundation constraint that partially suppresses thermal movement. Two risks attend the removal of the periodic component. The first is over-subtraction: if the fitted sinusoid or thermal model absorbs displacement that is genuinely structural in origin but happens to be correlated with the temperature cycle, the residual time series will understate the true structural movement. The second is under-subtraction: if the seasonal component is not adequately captured by a single sinusoid, for example when higher harmonics are present or when the amplitude is non-stationary, the residuals will retain spurious periodic content that may be misinterpreted as a structural signal. Both risks reinforce the need to interpret the residuals of the seasonal fit with the same rigour applied to the residuals of the linear fit, and to report the explained variance of the model alongside the estimated coefficients. The residuals remaining after removal of both the linear trend and the periodic component define the input to the change-point and non-linear analysis described in the following section.

The estimation of the thermoelastic coefficient requires a temperature record at the acquisition dates of the SAR stack. Three sources of temperature data are available in practice, differing in spatial representativeness and ease of access. The first is the air temperature recorded by local meteorological stations, which provides the most direct proxy for the thermal load experienced by the structure but may not be available at sufficiently close distance for all sites. The second is the temperature record provided by global weather data services, such as the Visual Crossing platform (Visual Crossing, 2024), which interpolates station observations and reanalysis products onto a regular grid and can be queried automatically within the processing environment for any location worldwide. The spatial resolution of such products, typically 10 to 30 km, is adequate for structures in temperate climates where the temperature gradient over the area of interest is small relative to the annual temperature range. The third source, used in the *Vicoforte* case study (Section 5.4.4), is satellite-derived LST from Sentinel-3 SLSTR, which provides a measurement of the radiative surface temperature rather than the air temperature and may capture local thermal effects (urban heat island, topographic shading, surface material) that station-based air temperature records do not resolve. The choice among these sources affects the physical meaning of the estimated coefficient: a coefficient estimated against air temperature represents the structural

response to ambient thermal forcing, while a coefficient estimated against surface temperature incorporates the thermal properties of the building envelope. For engineering interpretation, the distinction matters when the coefficient is used to identify anomalous thermal behaviour that may indicate material degradation or hidden structural defects, as discussed in the synthesis of Section 5.5.

The thermoelastic displacement component can be incorporated directly into the multi-temporal phase model as a third unknown estimated jointly with the residual height Δh_p and the linear velocity Δv_p . The two-parameter model of Equations 3. 10 and 3. 11 in Section 3.1.2 expresses the double-differential phase $\varphi_i^H(p)$ and $\varphi_i^D(p)$ as linear functions of the perpendicular baseline B_{ni} and the temporal baseline B_{ti} respectively. The three-parameter extension adds a third independent variable, the temperature difference ΔT_i between the slave acquisition and the master image, coupled to the thermoelastic coefficient $\Delta \alpha_p$ of the scatterer. The thermal phase contribution follows Equation 4. 1:

$$\varphi_i^{thermal}(p) = \frac{4\pi}{\lambda} \cdot \Delta T_i \cdot \Delta \alpha_p \quad 4. 1$$

where $\Delta \alpha_p$ is the thermoelastic coefficient expressed in millimetres of LOS displacement per degree Celsius, ΔT_i is the temperature difference between slave acquisition i and the master image in degrees Celsius, and the factor $4\pi/\lambda$ follows the two-way phase convention of Equation 3. 3. The total double-differential phase model of Equation 3. 5 is accordingly extended to include $\varphi_i^{thermal}(p)$ as a fourth physically deterministic component, leaving only the atmospheric and noise residuals after the three-parameter fitting. The periodogram search described in Section 3.1.2 now operates in a three-dimensional parameter space $(\Delta h_p, \Delta v_p, \Delta \alpha_p)$, maximising the temporal coherence γ_p over the set of M interferograms. The computational cost of the search increases relative to the two-parameter case, but for typical stack sizes of 50 to 200 images and reasonable parameter ranges, the processing remains tractable within the small area module of the processing environment.

Three conditions must be satisfied for the three-parameter estimation to produce physically meaningful results. The first is the availability of a temperature record at each acquisition date, as discussed above: the series ΔT_i must be known for all M slaves. The second is sufficient seasonal temperature variability within the observation period. If the temperature range across the stack is small, as may occur

in tropical climates or in short observation windows confined to a single season, the leverage of the temperature variable in the periodogram is insufficient to separate the thermal component from the noise, and the estimated $\Delta\alpha_p$ becomes unreliable. A temperature range of at least 15 to 20 degrees Celsius across the stack, typical of multi-year observations in temperate European climates, generally provides adequate leverage. The third condition is that the structure must exhibit a measurable thermoelastic response: elements whose LOS displacement is insensitive to temperature, such as PS on the ground surface or on massive concrete elements with high thermal inertia, produce $\Delta\alpha_p$ values close to zero regardless of the temperature range.

The structural significance of the estimated thermoelastic coefficient extends beyond the correction of the displacement time series. The coefficient provides, for each PS, a quantitative measure of the sensitivity of that specific structural element to thermal forcing. On a steel-framed building, the coefficient integrates the combined effects of the thermal expansion of the structural members, the boundary conditions at column bases and beam connections, and the differential stiffness between the structural frame and the cladding. On a masonry building, the coefficient reflects the thermal expansion of the masonry fabric, the moisture-dependent dielectric properties of the wall surface, and the constraints imposed by the building geometry on the free thermal deformation. An anomalous coefficient, one that deviates from the value predicted by the material properties and the structural geometry, may indicate a modification of the boundary conditions such as a bearing that has seized or a joint that has lost capacity, or a change in the load path that alters the distribution of thermal stresses. The attribution of an anomalous coefficient to a structural cause requires the exclusion of alternative explanations, including changes in the surface properties of the scattering element, differential exposure to solar irradiation along the diurnal cycle, and moisture-driven dielectric variations on hygroscopic materials such as masonry. The thermoelastic coefficient is therefore not merely a nuisance parameter to be removed from the displacement signal but a diagnostic variable in its own right, a point that is developed further in the cross-case-study discussion of Section 5.5 and operationalised in the integrated displacement model of the *Sanctuary of Vicoforte* in Section 5.4.

When the three-parameter estimation is applied to PS distributed along the height of a building, the thermoelastic coefficient $\Delta\alpha_p$ exhibits a characteristic vertical gradient that reflects the cumulative nature of thermal expansion in a multi-storey structure. For a regular frame in steel or reinforced concrete founded on a rigid base, the thermal expansion of each column between two consecutive floor

levels produces an incremental vertical displacement that is proportional to the inter-storey height and to the coefficient of thermal expansion of the structural material. The LOS projection of this incremental displacement, modulated by the satellite incidence angle, accumulates from the base upward: a PS at the n -th floor experiences the sum of the thermal displacements of all n storeys below it. The estimated $\Delta\alpha_p$ therefore increases approximately linearly with the height of the PS above the base, reaching its maximum value at the roof level and approaching zero at the ground, where the thermal displacement of the foundation is negligible. The same principle applies to masonry buildings, although the smaller thermal expansion coefficient of masonry produces gradients of reduced magnitude, and the presence of vaulted or domed elements alters the simple linear cumulation pattern that holds for orthogonal frames.

This vertical gradient constitutes a diagnostic signature of structurally regular thermal behaviour. A statistically significant departure from the expected linear gradient, beyond the noise floor of the coefficient estimation, may indicate a local anomaly in the structural system. A storey at which $\Delta\alpha_p$ is substantially lower than predicted by the linear trend may correspond to a zone of increased structural stiffness, such as a transfer slab, a belt truss, or a strengthened connection, that constrains the thermal expansion locally. Conversely, a storey at which $\Delta\alpha_p$ exceeds the predicted value may indicate a zone of reduced restraint, such as a deteriorated bearing, a cracked connection, or a floor slab that has partially lost composite action with the frame. The attribution of a deviation to a specific structural cause requires the exclusion of alternative explanations, including the variation of solar exposure with floor height, the differential moisture content of cladding materials, and the local geometry of the scattering elements. The diagnostic value of the vertical gradient is greatest when the observed profile can be compared against the predicted distribution from a structural model, which converts the satellite-derived coefficient into a quantitative constraint on the stiffness distribution of the building. In the case studies of Chapter 5, the thermal coefficient is estimated where the temperature data and the dataset configuration permit, and the observed vertical distribution is discussed in the cross-case-study synthesis of Section 5.5.

The methodological gain produced by the inclusion of the thermoelastic coefficient in the parametric estimation can be illustrated through the application of the workflow to a representative high-rise building, over an urban area in temperate climate. A X-band dataset of seventy-two images processed in user-driven mode is used as the empirical setting for the comparison. The processing output is summarised in Figure 4. 7. The reflectivity map of the area, shown in panel (a) and

computed as the temporal mean of the SAR amplitude stack, identifies the building as a coherent reflector against a background of mixed urban scatterers, and provides the geometric reference for the subsequent interpretation of the parametric estimation. The reference point of the small-area processing is selected on a stable ground surface adjacent to the building footprint, in a position that satisfies the three operational criteria of long-term stability, absence of significant thermal expansion, and high amplitude stability. The selection of a ground-level reference point with negligible thermal sensitivity ensures that the thermoelastic coefficients estimated for the PS on the building are interpretable as absolute values rather than as differences with respect to a thermally responsive reference.

The temporal coherence obtained with the two-parameter model, which estimates only the residual height and the linear velocity, is reported in panel b) and exhibits a spatially heterogeneous distribution on the building. The PS at ground level and at the lower storeys reach high coherence values, while the PS at the upper storeys exhibit progressively lower coherence. The pattern is the diagnostic signature of an unmodelled deterministic component that grows with the height of the PS above the base, consistent with the cumulative nature of thermal expansion in a multi-storey structure. The unmodelled component is absorbed into the residual phase of the periodogram and depresses the coherence according to the relationship between expected coherence and modelled signal discussed in Section 3.4.3.

The extension of the model to three parameters, with the addition of the thermoelastic coefficient as a third unknown, produces a substantial improvement in the coherence distribution shown in panel c). The coherence at the upper storeys increases to values comparable with those at the lower storeys, and the spatial heterogeneity of the previous estimation is largely resolved. The comparison between panel b) and panel c) confirms that the residual phase under the two-parameter model contained a deterministic thermal component whose explicit modelling reduces the variance of the residuals and increases the discriminative power of the coherence threshold. The fraction of PS that satisfy the coherence threshold for inclusion in the final dataset increases correspondingly, and the displacement time series of the retained points become interpretable as residuals from a model that accounts for the thermal forcing rather than as composite signals containing an unmodelled deterministic contribution.

The spatial distribution of the estimated thermoelastic coefficient on the building, shown in panel d) and expressed in radians per degree Celsius according to the native output of the processing chain, exhibits the characteristic vertical

gradient predicted by the cumulative nature of thermal expansion. The PS at ground level and at the lowest storeys exhibit coefficients close to zero, consistent with the negligible thermal expansion of the foundation system. The coefficient increases approximately linearly with the height of the PS above the base, reaching values of the order of zero point three to zero point four radians per degree Celsius at the roof level. The empirical observation confirms the analytical derivation of the parameter K developed earlier in this subsection and provides the diagnostic signature against which structural anomalies of the type discussed in Section 5.5 can be evaluated.




Figure 4. 7. Application of the parametric PSI estimation with thermoelastic coefficient on a representative high-rise building. Panel (a): reflectivity map of the area, computed as the temporal mean of the 72-image SAR amplitude stack. Panel (b): temporal coherence obtained with the two-parameter model that estimates residual height and linear velocity. Panel (c): temporal coherence obtained with the three-parameter model that adds

the thermoelastic coefficient as a third unknown. Panel (d): spatial distribution of the estimated thermoelastic coefficient on the building, expressed in radians per degree Celsius, exhibiting the characteristic vertical gradient that increases with the height of the PS above the base.

4.2.3. Change-points and non-linear behaviour

After the linear trend, the periodic component, and where applicable the thermal component have been estimated and removed, the residual time series may contain displacement patterns that none of the deterministic models can account for. These residuals represent the analytically most informative content of the time series for structural diagnosis, because they record departures from the expected background behaviour of the structure that may be attributable to changes in loading, boundary conditions, or structural integrity. Two classes of departure are relevant.

The first class is an abrupt change-point, defined as a statistically significant discontinuity in the mean level or slope of the residual series occurring at a specific epoch. A change-point in the mean level corresponds to a step-like displacement, consistent with a sudden loss of support, a shear failure in a weak structural plane, or the onset of a sliding mechanism. A change-point in the slope corresponds to a change in the rate of displacement, consistent with a modification of the loading regime, the completion of a nearby construction activity, or the progressive mobilisation of a failure mechanism that had previously been dormant. Standard statistical methods for change-point detection, including the cumulative sum control chart, the Pettitt test, and Bayesian online change-point detection (Adams & MacKay, 2007), are applicable to PS residual series and have been used in the InSAR literature to identify the onset of accelerating deformation on slopes and infrastructure (Berti et al., 2013). Their application to structural PS time series follows the same logic, with the additional requirement that the detection threshold accounts for the measurement uncertainty characterised in Section 3.5.3, so that only changes exceeding the noise floor are reported as statistically significant.

The second class of departure is a gradual non-linearity, in which the cumulative displacement curve exhibits a progressive change in curvature without a discrete transition between regimes. An accelerating trend, in which successive displacement increments grow over time, is consistent with a process that amplifies as it develops, such as the progressive opening of a crack under sustained eccentric loading or the rheological softening of a foundation soil under cyclic wetting and drying. A decelerating trend, in which displacement increments diminish over time,

is consistent with a self-limiting process such as primary consolidation approaching its asymptote or the closing of a seasonal crack as moisture conditions return to baseline. Distinguishing between these two cases carries direct diagnostic implications: acceleration warrants escalating monitoring intensity and, potentially, intervention, while deceleration suggests that the process is stabilising. In practice, the identification of gradual non-linearity in PS time series requires a record of sufficient length, generally three years or more for Sentinel-1, to separate a genuine curvature from the noise floor. Piecewise linear models, polynomial fits, or data-driven decompositions such as Empirical Mode Decomposition can be applied to the residual series to characterise the non-linear component, and their output should always be accompanied by an uncertainty estimate that reflects the propagated measurement precision. The combination of linear velocity, seasonal amplitude, thermal sensitivity coefficient, change-point dates, and non-linear curvature constitutes the full analytical description of a PS time series, and this description provides the input to the structural response metrics derived in Section 4.3.

A residual ambiguity affecting the interpretation of non-linear or rapidly varying displacement time series derives from the fundamental wrapping property of the interferometric phase. As established in Section 3.1.1, the observed phase is confined to the interval $[-\pi, +\pi]$, corresponding to a maximum unambiguous LOS displacement of $\lambda/2$ per acquisition interval. The periodogram search described in Section 3.1.2 identifies the velocity, residual height, and thermal coefficient that maximise the temporal coherence, and the resulting phase residuals are implicitly assigned to a single phase cycle. If the true displacement between two consecutive acquisitions exceeds $\lambda/2$, the observed phase wraps into the adjacent cycle, and the periodogram may assign the observation to the wrong cycle without producing a detectable reduction in temporal coherence, particularly when the number of images is modest or when the displacement rate changes abruptly.

The phase replica diagnostic provides a visual test for this condition. For each observation i in the estimated displacement time series, the analyst plots not only the nominal displacement value d_i but also the values $d_i + \lambda/2$ and $d_i - \lambda/2$, corresponding to the adjacent phase cycles above and below the estimated value. The resulting plot displays three parallel copies of each time-series point, separated vertically by $\lambda/2$. If the estimated time series is correctly unwrapped, the nominal values form a smooth trajectory and the replicas remain visually disconnected from it. If one or more observations have been assigned to the wrong phase cycle, the nominal value appears as an outlier that departs from the smooth trend, while the corresponding replica at $+\lambda/2$ or $-\lambda/2$ aligns with the trajectory that the remaining

observations define. The diagnostic is particularly informative when the deformation is non-linear: a time series that follows a linear trend for several years and then accelerates may exceed the $\lambda/2$ threshold at the transition point, causing the periodogram to assign the post-acceleration observations to the wrong cycle and producing an apparent reversal or discontinuity in the displacement record.


For the sensors used in this thesis, the unambiguous displacement interval is approximately 15.6 mm for COSMO-SkyMed (X-band, $\lambda/2 = 1.56$ cm) and 27.8 mm for Sentinel-1 (C-band, $\lambda/2 = 2.78$ cm). At the typical displacement rates encountered in routine structural monitoring of heritage buildings and infrastructure, which range from sub-millimetre seasonal oscillations to a few millimetres per year of long-term trend, the displacement accumulated between two consecutive acquisitions remains well below the unambiguous interval at both wavelengths, and cycle-slip ambiguity is negligible. The risk becomes relevant in two specific scenarios. The first is the analysis of structures undergoing rapid displacement events such as collapses, landslides, or accelerating failure mechanisms, where the displacement between two consecutive acquisitions may exceed $\lambda/2$, particularly for X-band sensors whose shorter wavelength halves the unambiguous interval. The second is the PS-by-PS inversion in the immediate vicinity of the reference point, where the common-mode phase components of the reference propagate into the differential phase of each candidate and may produce apparent step-like patterns in the unwrapped time series that are not physical. The replica diagnostic is therefore applied routinely to all X-band time series of structures exhibiting rapid displacement, and to any time series that exhibits discontinuities or reversals that cannot be explained by physical processes. The collapse scenarios of Section 5.3 and the fire-damaged buildings of Section 5.6 provide the most direct application of this diagnostic in the case studies.

The application of the diagnostic can be illustrated through a representative scatterer affected by non-linear deformation, identified within the small-area processing of a X-band dataset of seventy-two images. The output is summarised in Figure 4. 8. Panel (a) reports the displacement time series estimated by the periodogram inversion conditioned on the linear trend model, with the phase replicas. The time series exhibits a smooth quasi-linear trajectory consistent with the model, and the replicas remain visually disconnected from the nominal values throughout this portion. Beyond a transition epoch corresponding to the onset of an accelerated deformation regime, the nominal displacement values depart from the smooth trend and several observations align with the lower replica at $-\lambda/2$ rather than with the main trajectory. The pattern is the diagnostic signature of phase


unwrapping errors induced by inter-acquisition displacements that exceed the $\lambda/2$ threshold, and confirms that the linear trend model fails to capture the actual kinematic behaviour of the scatterer over the full observation period.

Panel (b) of Figure 4. 8 reports the displacement time series of the same scatterer estimated without the imposition of the linear trend model, with the periodogram inversion restricted to the residual height parameter only. In this configuration, the displacement at each acquisition is recovered from the differential phase without any parametric constraint on its temporal evolution, and the resulting trajectory reflects the actual kinematic behaviour of the scatterer up to the limit imposed by the phase ambiguity itself. The reconstructed time series exhibits a coherent non-linear pattern in which the early quasi-linear portion is preserved and the post-transition acceleration emerges as a continuous departure from the linear trend. The example illustrates the methodological principle that the imposition of a parametric model on the displacement evolution is a working assumption rather than a physical constraint, and that the reconstruction of the time series in free configuration is the appropriate diagnostic tool when the parametric assumption is violated.

The choice between the parametric and the free reconstruction is governed by the analytical objective of the processing. The parametric reconstruction provides a compact summary of the kinematic behaviour through a small set of estimated parameters and is appropriate when the deformation is approximately consistent with the assumed model. The free reconstruction provides the full evolution of the displacement without parametric assumptions and is appropriate when the deformation regime is unknown a priori or when the parametric reconstruction reveals diagnostic signatures of model violation through the phase replica analysis.



(a)



(b)

Figure 4. 8. Phase replica diagnostic and displacements reconstruction for a representative scatterer affected by non-linear deformation. Panel (a): displacement time series estimated by the periodogram inversion conditioned on the linear trend model, with the phase replicas. Panel (b): displacement time series of the same scatterer estimated without the imposition of the linear trend model, with the periodogram inversion restricted to the residual height parameter only.

4.3. From LOS measurements to structural response metrics

The analytical decomposition described in Section 4.2 produces, for each PS, a set of characterised time-series components expressed in LOS displacement units. These components are physically meaningful descriptions of the temporal behaviour of individual scatterers, but they do not yet constitute structural response metrics in the engineering sense. A structural response metric is a quantity that can be related to the mechanical state of the structure, such as a differential settlement between two foundation zones, a tilt angle of a wall, or a curvature increment along a structural element. The translation from LOS components to structural response metrics requires four sequential operations. The first, addressed in Section 4.3.1, is the computation of relative motion between PS pairs, which removes the arbitrary referencing offset and isolates the differential displacement carrying structural information. The second, addressed in Section 4.3.2, is the spatial interpretation of relative displacements across the structure, deriving settlement patterns and gradient estimates that can be compared with engineering damage thresholds. The third, addressed in Section 4.3.3, is the application of sanity checks that verify the physical consistency of the derived metrics before they enter the diagnostic chain. The fourth, addressed in Section 4.3.4, extends the methodology to retrospective event analysis, where satellite data are exploited to reconstruct the displacement history of a failed structure from archival observations.

4.3.1. Relative motion and reference dependence

All LOS displacement time series produced by MT-InSAR are expressed relative to a chosen spatial reference, as established in Sections 3.4.2 and 3.5.2. The absolute value of the displacement at any PS therefore contains a contribution from the true motion of that PS and an additive offset equal to the motion of the reference point, which is assumed to be zero but may not be in practice. For structural interpretation, the diagnostic interest lies not in the absolute displacement of individual PS but in the differential motion between PS pairs attributed to different positions on the same structure. The differential LOS displacement between two PSs eliminates the common reference offset, provided both PS belong to the same interferometric network and are processed with the same reference. The resulting quantity reflects the change in relative distance between the two structural locations in the LOS direction. If one PS is attributed to a ground-level position and the other to an elevated structural element, the differential displacement approximates the LOS component of the relative structural deformation accumulated between the two points over the observation period.

In the absence of decorrelation or geometric ambiguity, the precision of a differential displacement estimate is governed by the uncorrelated phase noise of the two PSs, which propagates in quadrature, and by the spatially correlated atmospheric residual at the scale of their separation, which decreases as the points are closer together. The choice of which PS pairs to use for relative motion computation should follow the geometric plausibility criteria of Section 4.1.3, so that only pairs with a physically interpretable separation in the structural geometry are included in the analysis.

The differential nature of the InSAR measurement has consequences for the interpretation of results on tall buildings that are particularly instructive and that are not always appreciated by analysts accustomed to absolute geodetic measurements. Consider a high-rise building observed with an X-band sensor over several years. The building has PSs distributed along its facade from ground level to the roof, and the surrounding area provides additional PS on the ground and on adjacent low-rise structures. The periodogram estimates, for each PS, a residual height relative to the reference point and a LOS velocity relative to the same reference. Both quantities change sign, magnitude, and distribution depending on which PS is designated as the reference. The dependence is most clearly observable on the residual height, which is the parameter that the periodogram inversion estimates as the difference between the actual elevation of the scatterer and the elevation predicted by the

external DEM used for the topographic phase removal. Consider the case in which the reference point is selected on the ground adjacent to the building and the external DEM is a Digital Terrain Model (DTM) that does not describe the building itself. The residual heights of the facade PSs are positive and increase progressively from the base to the top of the building. The actual elevation of each scatterer exceeds the elevation predicted by the DTM by the height of the corresponding storey above the foundation. If the external DEM is a Digital Surface Model (DSM) that already includes the building, the residual heights of the facade PSs are close to zero, because the actual elevation and the predicted elevation coincide within the accuracy of the DEM. The reference point choice introduces a constant offset on the estimated residual heights of all PSs in the network, equal to the residual height of the reference point itself with sign reversed. The velocity field reveals any differential motion between the building and the reference, including long-term settlement and the thermoelastic component. The temporal coherence is highest for PSs near the reference point, on the ground and at the base of the building, and decreases toward the top, because the differential atmospheric residual, the unmodelled thermal displacement, and the geometric sensitivity all increase with the vertical and horizontal distance from the reference.

If the reference point is moved to the top of the building, the residual height of all other PS in the network is shifted by the residual height of the original reference point relative to the new reference, with sign reversed. In the typical case of a DTM-based processing, the new residual heights are negative for all points below the reference, with the most negative values at the ground level, because the new reference is now located at the highest position of the network. The velocity field is now relative to the top of the building: any settlement of the building as a whole is invisible, because both the reference and the measured points participate in the same rigid-body motion. Only differential displacements within the building, such as relative settlement between the top and the base, or lateral deformation of the facade, are observable. The temporal coherence distribution inverts: it is highest near the top and lowest at the base and on the ground, because the atmospheric and thermal residuals are now largest for the points farthest from the reference. If the reference point is placed at mid-height on the facade, an intermediate condition arises, with residual heights positive above the reference and negative below, and the coherence distribution symmetric about the reference level.

The practical implication for structural monitoring is that the position of the reference point must be chosen deliberately to match the diagnostic question being addressed. If the objective is to detect settlement of the building relative to its

foundation, the reference should be on the ground or on a neighbouring stable structure. If the objective is to detect differential deformation within the building, such as inter-storey drift, facade bulging, or differential thermal expansion between structural elements, the reference should be on the part of the building that is assumed to be most stable. If the objective is to characterise the thermoelastic response of the building envelope, the reference should be at ground level where the thermal displacement is minimal, so that the thermoelastic coefficient estimated for upper-storey PS reflects the full cumulative thermal expansion from the base upward.

The dependence of the residual height distribution on the reference point selection is illustrated empirically through three independent processing runs on the same high-rise building introduced in Sections 4.1.2 and 4.2.2. The processing runs differ exclusively in the position of the reference point, while all other processing parameters, including the candidate selection threshold, the periodogram configuration, and the temporal coherence acceptance criterion, are kept identical across the three runs. The output is summarised in Figure 4. 9.

Panel (a) reports the residual height distribution obtained with the reference point at the base of the building, on a stable ground surface adjacent to the building footprint. Consistent with the use of a DTM as external DEM in the processing chain, the residual heights of the PSs on the building are positive and increase progressively from the base to the upper storeys, reaching the maximum value at the roof level. Panel (b) reports the distribution obtained with the reference point relocated to a position at mid-height on the building. The residual heights of the PSs below the reference are negative, with the most negative values at the base, while the residual heights of the PSs above the reference are positive and reach progressively higher values toward the roof. Panel (c) reports the distribution obtained with the reference point at the top of the building. The residual heights of all PSs below the reference are negative, with the most negative values at the base, while the residual heights at the reference level itself are zero by construction.

The three panels demonstrate that the same physical building, observed by the same SAR stack and processed with identical algorithmic parameters, yields three substantively different apparent distributions of residual height as a function of a single processing choice. The numerical values of the residual heights at any individual PS change between the three panels by an amount equal to the elevation of the corresponding reference point, and the pattern of progressively increasing or decreasing residual height along the building height is preserved in absolute terms

but inverted in sign between the base-reference case and the top-reference case. The empirical observation confirms the methodological principle established in the discursive treatment above: the residual height is a relative quantity whose absolute interpretation requires explicit knowledge of the reference point location, and the analyst is responsible for selecting the reference position that aligns with the diagnostic question being addressed.




Figure 4. 9. Residual height distribution on a high-rise building observed by X-band stack of 72 acquisitions, processed with three different reference point positions. Panel (a): reference point at the base of the building, on a stable ground surface adjacent to the building footprint. Panel (b): reference point at mid-height on the building. Panel (c): reference point at the top of the building.

4.3.2. Settlement patterns, gradients, and tilt-like behaviour

The spatial distribution of relative LOS velocities across a structure, computed with respect to a common PS on a stable reference element, forms a displacement pattern that can be interpreted in terms of the structural behaviour of the building or infrastructure. Three classes of pattern are of primary diagnostic relevance.

The first class is a spatially uniform velocity field, in which all PSs on the structure move at the same rate in the LOS direction. This pattern is consistent with a rigid-body settlement of the entire structure on its foundation, without differential deformation between structural zones, and has limited diagnostic significance beyond confirming that the structure is settling at a measurable rate. The second class is a spatially graded velocity field, in which the LOS velocity varies systematically across the structure, increasing from one side to the other or from lower to upper structural levels. A transverse gradient is consistent with differential settlement between foundation zones, which may be attributable to variable soil stiffness, eccentric loading, or localised foundation deterioration. A vertical gradient, with greater LOS velocity at upper PS than at lower ones, may indicate that the structure is deforming under its own weight, potentially through creep or progressive cracking. The third class is a spatially irregular pattern, in which adjacent PSs exhibit markedly different velocities without a coherent spatial structure. Irregular patterns may indicate genuine structural heterogeneity, such as the presence of structural joints, construction phases with different foundation systems, or localised damage. However, they may also reflect residual geocoding errors, atmospheric artefacts, or PSs attributed to different structural elements that are not kinematically connected. Distinguishing between these two explanations requires the sanity checks described in Section 4.3.3.

The spatial gradient of differential settlement constitutes the primary engineering metric for relating InSAR observations to the established damage classification framework for masonry structures proposed by Burland et al. (1977). The angular distortion β between two points on a structure is defined as the ratio of the differential settlement δ to the horizontal separation L between the two points, as expressed by Equation 4. 2:

$$\beta = \frac{\delta}{L} \quad 4. 2$$

where δ is the differential settlement between two PS attributed to geometrically distinct structural positions, L is the horizontal distance between them, and β is the resulting angular distortion, a dimensionless quantity typically expressed as a fraction such as 1/500 or 1/300. The Burland et al. (1977) classification associates threshold values of β with progressively more severe damage categories for masonry buildings: values of β below 1/500 are associated with negligible damage, values in the range 1/500 to 1/300 with slight damage, values in the range 1/300 to 1/150 with moderate to severe damage, and values exceeding 1/150 with severe to very severe damage potentially affecting structural stability. Subsequent refinements, notably by Boscardin & Cording (1989), introduced additional parameters such as deflection ratio and horizontal strain to capture the interaction between differential settlement and building stiffness, and the Burland classification should be understood as the baseline framework from which these refinements departed. The application of this classification to InSAR-derived differential settlements requires that two preliminary conditions be satisfied. The first is that the pair of PSs used for the computation of β be physically attributable to structurally connected elements belonging to the same foundation or superstructure unit, so that their differential displacement reflects a genuine angular distortion of the structure rather than a relative motion between disconnected zones. The second is that the differential settlement δ exceed twice the propagated uncertainty of the two single-PS velocity estimates over the observation period, as quantified in Section 3.5.3, so that the resulting β is statistically distinguishable from noise.

When both conditions are met, the angular distortion computed from satellite observations provides a quantitative link between the interferometric measurement and the engineering framework routinely used for the structural assessment of existing masonry buildings. The application of Equation 4. 2 to the spatial velocity field of a structure produces a map of angular distortion rates (in units of radians per year or equivalent dimensionless rate per unit time), which when integrated over a representative observation period yields a cumulative distortion estimate directly comparable against the Burland thresholds.

4.3.3. Sanity checks for structural interpretation

Before any structural conclusion is drawn from the derived metrics, a set of physical plausibility checks must be applied to verify that the patterns identified in Section 4.3.2 are consistent with the geometry of the measurement, the precision of the data, and the expected range of structural behaviour. Four checks are recommended.

The first is a magnitude check: the derived differential velocity or displacement must be compared with the measurement uncertainty quantified in Section 3.5.3. Any pattern whose amplitude does not exceed twice the propagated uncertainty at the relevant spatial scale cannot be distinguished from noise and should not be reported as a structural observation. The second is a geometric consistency check: the sign and direction of the differential LOS displacement must be consistent with the geometry of the postulated structural mechanism. The interpretation of the sign depends on the acquisition geometry because the same physical displacement projects with opposite signs onto the ascending and descending LOS directions. A PS that moves westward, for example, produces a LOS-shortening signature in a descending right-looking geometry (where the satellite observes from east to west, so westward motion reduces range) and a LOS-lengthening signature in an ascending right-looking geometry (where the satellite observes from west to east). A differential settlement in which a PS on the western side of a building moves faster in the LOS direction than a PS on the eastern side is geometrically consistent with an ascending acquisition looking from west to east, but would require additional explanation if observed in a descending geometry where the relationship is inverted. Inconsistency between the observed sign and the geometrically expected sign is a strong indicator that the PS attribution is incorrect or that the postulated mechanism does not apply.

The third check is a temporal consistency check: the pattern identified from mean velocities must be confirmed in the cumulative displacement time series, verifying that the differential signal is present throughout the observation period and not driven by a single anomalous epoch. A differential pattern that appears only at one epoch and disappears thereafter is more likely attributable to a residual atmospheric artefact than to a structural process. The fourth check is a cross-dataset consistency check: where both ascending and descending datasets are available, the differential patterns derived from the two geometries must be mutually consistent under the LOS projection of the postulated displacement vector. Agreement between the two geometries substantially increases confidence in the structural interpretation, while disagreement indicates either a geometric inconsistency in the PS attribution or a displacement mechanism that is not adequately described by the assumed model. The outcome of these checks, when documented explicitly, constitutes the evidentiary basis on which any subsequent structural diagnostic conclusion can be defended.

4.3.4. Retrospective event analysis and forensic constraints

The interpretation methodology developed in Sections 4.1 through 4.3.3 is designed primarily for prospective monitoring: the analyst processes an ongoing or recently completed SAR archive to characterise the current displacement behaviour of a structure whose condition is unknown. A distinct but methodologically related application arises when the analysis is conducted retrospectively, after a structural failure has occurred, with the objective of reconstructing the displacement history of the failed element from archival SAR data that were not acquired for monitoring purposes. The forensic application of satellite InSAR exploits the fact that the continuously accumulating radar archive constitutes an unintentional displacement record of every structure within the sensor coverage area.

The methodological steps of retrospective event analysis follow the same sequence described in the preceding sections: point selection, time-series decomposition, metric derivation, and plausibility verification. Three additional constraints, however, apply specifically to the forensic context and must be accounted for before any conclusion can be drawn.

The first constraint is temporal. The analyst must define the event window, the interval over which the structural failure is known or hypothesised to have occurred, and partition the displacement time series into pre-event, peri-event, and post-event segments. For sudden collapses, the event window is bounded by two consecutive satellite acquisitions that straddle the failure date. The temporal resolution of the forensic analysis is therefore limited by the revisit time of the sensor: 6 to 12 days for Sentinel-1 and 4 to 16 days for COSMO-SkyMed, depending on the constellation configuration. TS, as defined in Section 4.1.2, provide the most direct forensic timestamp: the T_{off} value of a TS located on the failed element constrains the date of the physical event to the interval between the last coherent acquisition and the first incoherent one. Where multiple TS on the same structure exhibit consistent T_{off} values, the convergence of independent temporal constraints strengthens the chronological inference.

The second constraint is statistical. Pre-collapse displacement anomalies, the signals most sought after in forensic analysis, are typically small relative to the noise floor of the MT-InSAR technique. Velocity anomalies of 2 to 5 mm per year on a bridge deck or building facade are close to the detection margin for C-band sensors and within the range of atmospheric residuals, orbital ramp errors, and thermal effects that have not been fully removed (Parizzi et al., 2020). The

attribution of such anomalies to structural distress rather than to environmental or instrumental causes requires explicit uncertainty quantification. At a minimum, the analyst must demonstrate that the reported anomaly exceeds the estimated standard deviation of the velocity estimate for the specific measurement point, and that the spatial pattern of the anomaly is consistent with a physically plausible deformation mode for the structural typology under investigation. The exchange between Milillo et al. (2019) and Lanari et al. (2020) on the Morandi Bridge illustrates the consequence of omitting this step: the same dataset, processed with different techniques and different assumptions about the noise model, yielded contradictory conclusions about the presence of pre-collapse deformation.

The third constraint is epistemological. Satellite InSAR measures the quasi-static component of LOS displacement over time. Structural failures result from combinations of progressive material degradation, loss of boundary conditions, dynamic loading, and sudden brittle fracture. The satellite record captures only the quasi-static surface expression of processes that may be occurring within the structural volume at scales and rates that the sensor cannot resolve. A displacement anomaly that is spatially consistent with the failed portion of a structure is a necessary but not sufficient condition for attributing the failure to the process that produced the anomaly. The absence of a unique mapping from InSAR displacement to failure mechanism is a fundamental constraint on the forensic reach of the method, and the analyst must state explicitly which alternative interpretations have been considered and excluded, and on what basis.

These three constraints do not invalidate the forensic use of InSAR. They define the conditions under which forensic conclusions are defensible. The post-event portion of the SAR archive carries information that is less epistemologically fraught than the pre-event portion: the disappearance of scatterers on collapsed elements, the emergence of TS on debris, and the re-establishment of stable PS on reconstructed elements constitute a legible structural chronicle that does not require causal attribution. The forensic value of satellite data is therefore greatest when the analysis combines pre-event kinematic characterisation, subject to the constraints above, with post-event scattering change detection and TS lifecycle analysis, producing a displacement biography of the structure that an engineering investigation can interpret in the context of material evidence, inspection records, and numerical simulation.

4.4. Handling heterogeneous datasets

MT-InSAR analyses of the same structure may draw on datasets from different sensors, processed independently and potentially over different temporal windows. The case most frequently encountered in the structural monitoring literature, and the one most relevant to the case studies of Chapter 5, is the joint use of Sentinel-1 C-band data and COSMO-SkyMed X-band data. The two datasets are complementary: Sentinel-1 provides a long, freely accessible observation record at moderate spatial resolution, while COSMO-SkyMed offers higher resolution, greater phase sensitivity to displacement, and the capacity to resolve individual structural elements at the building or bridge scale. However, their direct comparison is not straightforward. Differences in wavelength, spatial resolution, PS density, temporal sampling, reference frame, and atmospheric correction all introduce discrepancies between the two velocity fields that may be misinterpreted as genuine structural displacement differences if not properly accounted for. The following subsections address these discrepancies in sequence. Section 4.4.1 characterises the measurement differences between the two sensors at the level of the individual PS. Section 4.4.2 describes the harmonisation and co-referencing procedure required before joint interpretation is attempted. Section 4.4.3 discusses the interpretability problems that persist after harmonisation and the conditions under which the two datasets can and cannot be directly compared.

4.4.1. Sentinel-1 versus COSMO-SkyMed

The primary difference between Sentinel-1 and COSMO-SkyMed for structural monitoring purposes lies in their spatial resolution and the consequent density of detectable PS on built assets. Sentinel-1 in IW mode achieves a nominal resolution of 5 m in range and approximately 20 m in azimuth, whereas COSMO-SkyMed in Stripmap HIMAGE mode provides approximately 3 m in both directions, and in Spotlight mode sub-metre resolution is achievable. On a building facade of typical urban dimensions, the higher resolution of COSMO-SkyMed translates directly into greater PS density and into the possibility of attributing individual PS to specific structural elements, a capability that Sentinel-1 generally cannot provide at the same level of geometric discrimination.

The second difference concerns wavelength. The C-band wavelength of 5.55 cm means that each fringe of the Sentinel-1 interferometric phase corresponds to a LOS displacement of approximately 2.78 cm, while the X-band wavelength of 3.1 cm for COSMO-SkyMed gives a fringe spacing of approximately 1.56 cm. The X-

band system is therefore approximately 1.8 times more sensitive to LOS displacement per unit phase change, an advantage for detecting small structural movements, but also a disadvantage in terms of susceptibility to phase aliasing and temporal decorrelation in environments with seasonal surface moisture variations.

The third difference concerns the temporal sampling and observation record. Sentinel-1 offers a six-day repeat cycle over most of Europe following the restoration of the two-satellite configuration and a continuous public archive from 2014 onwards. COSMO-SkyMed acquisitions are typically commissioned on demand or obtained through the ASI Research and University Programme, resulting in temporal gaps that may reduce the diagnostic value of the time series for slow processes. The PS velocity precision of the two sensors under comparable observation period and atmospheric conditions differs in proportion to their wavelength ratio: for equal phase noise, the COSMO-SkyMed velocity estimate carries approximately 1.8 times the displacement sensitivity of Sentinel-1 per unit time, but the shorter wavelength also makes the measurement more susceptible to phase aliasing if rapid displacement events occur between consecutive acquisitions, as discussed in Section 4.2.3. These complementary characteristics have motivated the joint use of the two sensors in infrastructure monitoring studies, and their respective strengths and weaknesses inform the multi-sensor strategy adopted in the case studies of Chapter 5.

The qualitative difference between the two sensors is illustrated in Figure 4. 10, which compares the PS populations identified on a motorway junction in Rome through a medium-resolution C-band Envisat dataset acquired between 2002 and 2010 and a high-resolution X-band COSMO-SkyMed dataset acquired between 2011 and 2014 (Macchiarulo, Giardina, et al., 2022). The denser PS population produced by the X-band dataset and the finer geometric discrimination of structural elements visible in panel b) of the figure provide an empirical reference against which the case studies of Chapter 5 can be interpreted, where Sentinel-1 and COSMO-SkyMed are deployed in a complementary multi-sensor strategy.




Figure 4. 10. PSs identified over a motorway junction in Rome using medium-resolution C-band Envisat data acquired between 2002 and 2010 in panel (a) and high-resolution X-band COSMO-SkyMed data acquired between 2011 and 2014 in panel (b). Each PS is represented by a dot colour-coded by cumulative LOS displacement, with positive and negative values indicating motions towards and away from the satellite respectively. The contrast between the two panels illustrates the substantial gain in PS density achievable with X-band sensors on a built infrastructure asset, and the corresponding improvement in the spatial selectivity of the displacement characterisation at the element level..

4.4.2. Harmonization and co-referencing

Before the velocity or displacement fields from Sentinel-1 and COSMO-SkyMed can be compared or jointly interpreted on the same structure, three harmonisation operations must be performed.

The first operation is temporal alignment. The two datasets rarely share identical observation windows. If one dataset spans 2015 to 2022 and the other spans 2018 to 2020, the mean velocities derived from the two are averaged over different periods and cannot be directly compared without first verifying that the deformation process was stationary over both windows, or without restricting the comparison to the common temporal interval.

The second operation is reference frame alignment. As established in Section 3.5.2, each PS-InSAR dataset is referenced to a locally chosen reference PS, and there is no guarantee that the two reference points have the same absolute velocity. A systematic velocity difference between the two datasets, manifesting as a spatially uniform offset in the difference map, is the expected consequence of different reference choices. The offset can be estimated and removed by identifying a set of PS that are common to both datasets or spatially proximate, computing the mean velocity difference over that set, and subtracting the estimated offset from one of the two velocity fields. Where GNSS permanent stations are available, both datasets can be independently anchored to the same geodetic datum using the procedure described in Section 3.5.2, which removes the reference bias without requiring a common PS set.

The third operation is LOS projection alignment. Sentinel-1 and COSMO-SkyMed are typically processed in ascending and descending geometries that do not coincide, and even within the same orbit direction the incidence angles and headings may differ. A velocity value in Sentinel-1 ascending LOS and a velocity value in COSMO-SkyMed ascending LOS are not directly comparable unless both are projected onto a common physical direction, such as the vertical, using the LOS unit vectors appropriate to each acquisition geometry. The uncertainty introduced by this projection depends on the sensitivity of the assumed decomposition to errors in the PS attribution, as discussed in Section 2.2.2, and should be propagated explicitly into the harmonised product.

4.4.3. Bias removal and interpretability issues

Even after temporal alignment, reference frame correction, and LOS projection, residual discrepancies between Sentinel-1 and COSMO-SkyMed velocity fields on the same structure may persist, and some of these are not removable by any linear correction procedure. Three sources of residual bias are of practical relevance.

The first is the thermal sensitivity difference. As discussed in Section 4.2.2, PS on metallic or reinforced concrete elements exhibit a thermal LOS sensitivity $\Delta\alpha_p$ that is a function of the structural geometry, the material expansion coefficient, and the satellite viewing angle. Because the incidence angles and headings of the two sensors differ, the thermal sensitivity coefficient of the same physical scatterer is different in the two LOS systems. If the thermal component has not been explicitly estimated and removed from both time series before comparing mean velocities, the velocity difference will contain a residual that reflects the difference in thermal sensitivity rather than a genuine displacement discrepancy.

The second source is PS population mismatch. The PS sets of the two datasets are generally not coincident in space, because the higher resolution of COSMO-SkyMed may resolve scatterers on individual structural elements that are blended into a single resolution cell in Sentinel-1, and because layover geometry differs between the two sensors. Comparing velocities at geometrically proximate but physically distinct PS implicitly assumes that the two points share the same displacement history, which may not hold on a geometrically complex structure where adjacent PS belong to different structural zones.

The third source is the difference in temporal coherence thresholds and atmospheric correction quality between the two processing chains. COSMO-SkyMed data processed with a strict temporal coherence threshold may retain only the most stable scatterers, which systematically sample different structural positions than the full PS set obtained from Sentinel-1 under a more permissive threshold. Residual biases of this kind cannot be removed after processing; they can only be mitigated by documenting the selection criteria and processing parameters of each dataset explicitly and by restricting the joint interpretation to structural conclusions that are robust to the population mismatch. In practice, this means using the two datasets for complementary rather than identical diagnostic purposes: Sentinel-1 for long-term trend characterisation at the building scale, and COSMO-SkyMed for high-resolution element-level diagnosis over shorter observation windows.

4.5. Integration with complementary information

Satellite interferometric data provide a spatially distributed and temporally dense record of LOS displacement that no conventional monitoring system can replicate at comparable cost and coverage. However, the inherent constraints of the technique, documented throughout this chapter, mean that satellite data alone cannot resolve all the questions that arise in structural assessment. The LOS measurement integrates contributions from multiple physical processes that must be disentangled analytically. Geometric ambiguities limit the attribution of PS to specific structural elements. Referencing uncertainty affects the absolute accuracy of velocity estimates. These constraints do not prevent satellite data from being useful; they define the conditions under which their use is appropriate and the forms of complementary information that strengthen the conclusions. Section 4.5.1 addresses the cross-validation of satellite-derived displacement metrics against in-situ monitoring data, which provides an independent verification of both the measurement accuracy and the physical interpretation. Section 4.5.2 addresses the integration of satellite observations with numerical structural models, which transforms observed displacement patterns into quantitative estimates of structural condition. Section 4.5.3 addresses the fusion of InSAR with non-SAR remote sensing data at the building scale, with particular reference to the airborne LiDAR integration methodology developed for the post-fire assessment application of Chapter 5.

4.5.1. Cross-validation with in-situ monitoring

Cross-validation between satellite-derived displacements and in-situ monitoring data serves two distinct purposes. The first is empirical accuracy verification: if an independent sensor records a displacement that is geometrically consistent with the LOS projection of a PS at the same structural location, agreement between the two estimates provides evidence that the satellite measurement is free from the systematic errors discussed in Section 3.5. The second purpose is physical interpretation support: in-situ sensors measure quantities, such as absolute tilt, crack width, or piezometric level, that the satellite cannot access directly, and their time series may explain displacement patterns that would otherwise remain ambiguous in the satellite record alone.

The monitoring instruments most commonly used for cross-validation in the structural InSAR literature include GNSS receivers, precision levelling campaigns, terrestrial laser scanners, clinometers, and piezometers. Each provides a different

type of information and is subject to different geometric and temporal sampling constraints. GNSS receivers yield three-dimensional absolute displacement vectors at continuous or daily temporal resolution, making them the most direct comparator for InSAR velocity estimates once the GNSS-derived displacement is projected onto the satellite LOS (Bähr & Hanssen, 2012). Precision levelling provides high-accuracy vertical displacement at campaign intervals and is particularly valuable for verifying differential settlement gradients derived from the spatial distribution of PS velocities. TLS point clouds, acquired at two or more epochs, yield three-dimensional surface deformation fields at millimetre-level spatial resolution and can resolve displacement patterns on building facades that are too spatially complex for PS-InSAR to characterise without ambiguity.

In practice, the available in-situ data rarely coincide exactly in time and space with the satellite PS set. Temporal alignment between the two records requires either interpolation or restriction to a common observation window, and spatial alignment requires that the in-situ sensor location be matched to the nearest geometrically plausible PS, accepting the localisation uncertainty described in Section 3.5.1. Where the agreement between the two records falls within the combined measurement uncertainty, the satellite-derived metrics can be treated as validated. Where systematic discrepancies persist after all known bias sources have been accounted for, the discrepancy itself carries diagnostic information, potentially indicating a physically real difference between the deformation at the instrument location and that at the PS location.

4.5.2. Integration with numerical models

Numerical structural models provide the physical reference frame within which observed displacement patterns can be interpreted in terms of structural condition. A finite element model of a building or bridge, calibrated against the as-built geometry and material properties, predicts the displacement field that the structure would exhibit under a specified set of loads and boundary conditions. Comparing this predicted field with the satellite-observed PS velocity map reveals discrepancies that may indicate deviations between the modelled and actual structural behaviour, attributable to changes in material stiffness, boundary condition alterations, or damage-induced stiffness reductions.

The process of adjusting model parameters so that the model output matches observed data is known as model updating (Mottershead & Friswell, 1993) and has been applied in the SHM literature primarily with vibration-based measurements.

The application of satellite displacement data to model updating is conceptually analogous but operationally distinct, because the observable is a quasi-static long-term displacement rather than a dynamic modal parameter, and the measurement operator that relates model degrees of freedom to LOS displacements must account for the satellite geometry, the PS position within the structural model, and the LOS projection of the displacement vector at that point. Where PS are attributed to specific nodes of the structural model, the measurement operator can be written explicitly, and the model updating problem becomes a constrained optimisation in which the model parameters are adjusted to minimise the difference between predicted and observed LOS displacements across the full PS set.

The utility of this integration depends on the density and spatial distribution of PS relative to the degrees of freedom of the model: a sparse PS set constrains only a small subset of model parameters, while a dense set from high-resolution COSMO-SkyMed data may support a more complete characterisation of the deformed shape. In all cases, the uncertainty of the satellite-derived displacements, propagated as described in Section 3.5.3, must be incorporated into the updating formulation as observational noise, so that the model updating result reflects the actual information content of the satellite data rather than treating it as exact.

4.5.3. Fusion with non-SAR remote sensing data

The preceding sections address the integration of satellite InSAR data with information sources that either share the same measurement principle, as in the case of multiple SAR datasets treated in Section 4.4, or belong to the domain of structural engineering, as in the case of in-situ sensors in Section 4.5.1 and numerical models in Section 4.5.2. A third category of complementary information is provided by remote sensing systems that operate on physical principles different from radar interferometry and that measure quantities other than LOS displacement. Airborne LiDAR, optical satellite imagery, and thermal infrared sensors fall within this category. Among these, airborne LiDAR is the most directly relevant to building-scale SHM because it provides three-dimensional geometric information at a spatial resolution comparable to that of high-resolution SAR and at a level of geometric detail that InSAR cannot achieve.

The methodological challenge of combining InSAR with non-SAR remote sensing data differs qualitatively from the harmonisation problems addressed in Section 4.4. When two SAR datasets are combined, the measured quantity is the same (LOS displacement), the measurement geometry is related through known

orbital parameters, and the harmonisation consists of aligning reference frames, projecting onto common LOS directions, and accounting for differences in wavelength, resolution, and temporal sampling. When InSAR is combined with LiDAR, the measured quantities are fundamentally different. InSAR provides a temporal record of phase change at discrete scatterer locations, from which displacement histories, velocities, and lifecycle indicators are derived. LiDAR provides a dense three-dimensional point cloud at each acquisition epoch, from which surface elevation, building height, volumetric change, and geometric deformation metrics are derived. The two data types cannot be compared point-to-point; they must be fused at a higher level of spatial aggregation.

The appropriate spatial unit for this aggregation depends on the application. For building-scale post-event assessment, the natural aggregation unit is the building footprint polygon, derived from cadastral records, building inventory databases, or automated extraction from optical imagery. Each building footprint serves as a spatial container within which the two data sources are independently summarised into per-building feature vectors. The InSAR contribution to this vector may include the number of PS within the footprint, the pre-to-post-event PS density ratio, the fraction of event-induced TS, and summary statistics of the pre-event displacement time series. The LiDAR contribution may include the mean elevation change between pre-event and post-event normalised DSM, the standard deviation of the elevation change, the volume change proportion, and the footprint coverage fraction. The resulting per-building feature vector is multi-dimensional and multi-modal, combining temporal information from InSAR with geometric information from LiDAR.

The spatial attribution of InSAR scatterers to building footprints requires particular care because the geocoded position of a PS carries an uncertainty that depends on the accuracy of the residual height estimate and on the layover geometry of the urban environment. In dense residential areas, a PS whose geocoded coordinates fall within one building footprint may physically originate from an adjacent structure whose facade is projected into the same slant-range resolution cell. The reverse geocoding procedure, in which building footprint polygons are projected into the SAR slant-range geometry to define attribution zones, reduces this ambiguity by accounting for the known building height and the satellite viewing geometry. The residual attribution uncertainty should be propagated into the per-building feature statistics, for example by computing the PS density ratio and TS fraction as expected values over a probabilistic attribution rather than as deterministic counts.

Once per-building feature vectors have been assembled from both sources, the fusion itself can be conducted through supervised classification, unsupervised clustering, or regression, depending on the monitoring objective. Supervised classification requires a labelled training set in which buildings are assigned to known condition categories (for example destroyed, damaged, intact) on the basis of ground-truth inspection. The classifier learns the joint discriminative structure of the multi-modal feature space and produces per-building condition estimates for the unlabelled population. The choice of classifier should balance predictive performance against interpretability: ensemble methods such as Random Forest provide both high accuracy and feature importance rankings that allow the analyst to assess the relative contribution of each sensor to the classification outcome.

Two conditions govern the diagnostic value of multi-modal fusion relative to single-sensor analysis. The first is complementarity: the fusion adds value only if the two sensors capture different aspects of the structural condition that are jointly more informative than either aspect alone. For post-event assessment, this condition is typically satisfied because LiDAR quantifies geometric change with high spatial fidelity but lacks temporal context, while InSAR provides temporal context through the pre-event displacement record and the scatterer lifecycle but lacks three-dimensional geometric resolution. The second condition is sufficient spatial overlap: the fusion requires that both sensors provide usable data over a common population of buildings. Where one sensor has no coverage, for example because InSAR point density is zero on a specific building type or because LiDAR acquisition does not extend to a given area, the fusion degrades to single-sensor analysis for the affected buildings.

The methodological framework described in this section, which constitutes one of the original contributions of this thesis to multi-modal SHM, is examined empirically in the ongoing research presented in Section 5.6, where PS-InSAR scatterer lifecycle metrics and bi-temporal airborne LiDAR building change metrics are combined at the per-building scale for damage classification following the 2025 Palisades Fire. The preliminary results obtained through the PS-InSAR component of the analysis are reported therein, while the full fusion framework, the classification performance assessment, and the feature importance analysis remain directions for completion of the research.

4.6. Toward decision-ready diagnostics: from displacement to engineering judgement

The preceding sections have defined the technical operations required to transform the raw output of a MT-InSAR processing chain into structural information. These operations comprise the selection of geometrically plausible measurement points, the decomposition of their displacement time series, the derivation of relative motion metrics, the verification of physical consistency, the harmonisation of heterogeneous datasets, and the integration of satellite observations with complementary evidence. Each of these operations reduces the distance between the interferometric measurement and the engineering interpretation. The question that remains is how far, precisely, the satellite evidence can travel along this path, and at what point the responsibility for judgement passes from the data to the engineer.

The term “decision-ready diagnostics”, introduced in this thesis as the organising principle of the four-level diagnostic framework, designates the condition in which the evidence derived from satellite interferometric data has been processed, validated, and contextualised to a degree sufficient to support an actionable engineering judgement about the structure. The term does not imply that the satellite data alone produce a diagnosis. Rather, the term designates an output state in which the satellite contribution has been brought to a level of analytical maturity. At this level, the satellite evidence can enter a decision process alongside other sources of evidence, without requiring the decision-maker to perform additional interferometric processing or to interpret raw LOS displacement values.

The diagnostic capacity of satellite observations is not uniform across all monitoring scenarios. The four-level articulation introduced in the present work organises this non-uniformity into a structured progression, in which each level is characterised by a distinct information content and a distinct set of conditions that must be satisfied before the output can be considered reliable. The four levels constitute a methodological contribution of this thesis, and are summarised in Table 4. 1 for reference throughout the empirical evaluation of Chapter 5.

The first level is screening. At this level, the satellite data serve to identify structures or structural elements that exhibit displacement anomalies relative to their surroundings or relative to a population of comparable assets. The information content is binary: a structure is either flagged or not flagged. The conditions for reliability are relatively modest: a sufficient number of measurement points within

or near the structure, a velocity estimate with an uncertainty smaller than the anomaly of interest, and a reference frame that is consistent across the portfolio being screened. Screening does not require user-driven processing; standardised EGMS products are generally adequate for this purpose.

The second level is displacement characterisation. At this level, the satellite data provide a quantitative description of the displacement behaviour of the structure over the observation period: the mean velocity, the amplitude and phase of the seasonal component, the presence or absence of change-points, and the spatial pattern of relative motion across the structural footprint. The information content is richer than screening, but it remains kinematic: it describes how the structure moves, not why. The conditions for reliability are correspondingly more stringent. Displacement characterisation requires user-driven processing with structure-specific parameter choices, a reference point whose stability can be justified on physical grounds, and a temporal coherence threshold calibrated to the expected deformation regime rather than to a default regional value.

The third level is diagnostic interpretation. At this level, the characterised displacement pattern is attributed to a physical cause, or to a ranked set of plausible causes, through comparison with structural knowledge, geotechnical context, and, where available, complementary monitoring data. The information content is causal: the analysis concludes, for example, that the observed differential settlement pattern is consistent with consolidation of a compressible soil layer beneath the eastern foundation, or that the seasonal oscillation amplitude is compatible with the thermal expansion coefficient of a concrete deck at the observed height. The conditions for reliability include all those of the preceding level, plus cross-validation against at least one independent information source and a demonstrated ability to distinguish the structural signal from environmental and instrumental artefacts.

The fourth level is continuous monitoring with automated anomaly detection. At this level, a calibrated surrogate model, trained on the satellite displacement record and on environmental input variables, generates a normalised monitoring statistic that flags deviations from the expected baseline behaviour. The information content is prospective: the system does not merely describe past displacement but detects departures from the established pattern that may signal a change in structural condition. The conditions for reliability are the most demanding of the four levels. The calibration period must be sufficiently long and temporally dense. The environmental input variables must capture the dominant sources of reversible

displacement. The data selection strategy must be coherence-aware and exclude unreliable input points. A validated statistical framework must define the alert thresholds.

The progression through these levels is not automatic. Each level requires a deliberate analytical investment that is justified only when the monitoring objective demands it and when the data conditions permit it. A screening-level analysis on a portfolio of two thousand bridges is a legitimate and complete diagnostic output; an attempt to extract diagnostic interpretation from a structure on which only three measurement points pass the coherence threshold is not. The methodology developed in this chapter provides the technical toolkit for each level; the decision of which level to pursue for a given structure depends on the risk profile of the asset, the available data, and the resources that the monitoring programme can allocate.

The four-level framework reframes the role of satellite InSAR within structural health monitoring as a progression of analytical maturity rather than as a binary capability. The progression is conditioned by data conditions, by the availability of complementary information sources, and by the explicit quantification of uncertainty at each level. The empirical investigation of Chapter 5 examines whether and to what extent each level becomes achievable on structures of different typology, data availability, and monitoring complexity, and identifies the recurrent factors that govern the transition from one level to the next.

Table 4. 1. Synthesis of the four-level diagnostic framework introduced in Section 4.6, reporting for each level the information content delivered to the engineering user, the data and processing requirements, the conditions for reliability that must be satisfied, and an indicative application scenario. The table is the definitional reference for the framework and is recalled in the cross-case-study assessment of Section 5.5.

Level	Information content	Required data and processing	Conditions for reliability	Indicative application scenario
Level 1 Screening	Binary identification of structures or structural elements that exhibit displacement anomalies relative to	Standardised products (e.g. EGMS Calibrated). User-driven processing not required. Sufficient measurement	Velocity uncertainty smaller than the anomaly of interest. Reference frame consistent across the screened portfolio. No requirement	Portfolio-scale screening of bridge or building stocks for prioritisation of inspection.

Level	Information content	Required data and processing	Conditions for reliability	Indicative application scenario
	surroundings or to a comparable population.	points within or near the structure; portfolio-consistent reference frame.	for structure-specific parameter tuning.	
Level 2 Displacement characterisation	Quantitative kinematic description of structural displacement: mean velocity, seasonal amplitude and phase, change-points, spatial pattern across the footprint.	User-driven processing of raw SLC data with structure-specific parameters. Reference point physically justified. Coherence threshold calibrated to the deformation regime.	Conditions of Level 1, plus: structure-specific reference, application-specific coherence threshold, and verification of phase aliasing absence given the expected displacement rate.	Building-scale or single-asset analysis where the engineering question concerns the kinematic behaviour of the structure.
Level 3 Diagnostic interpretation	Causal attribution of the characterised displacement pattern to a physical mechanism or to a ranked set of plausible mechanisms.	All operations of Level 2, plus integration with structural knowledge, geotechnical context, and at least one independent information source.	Conditions of Level 2, plus: cross-validation against an independent observable, demonstrated separation between structural and environmental contributions, and physical consistency of the proposed cause.	Forensic analysis of structural events or causal investigation of long-term deformation on individual heritage and infrastructure assets.
Level 4 Continuous monitoring	Prospective detection of departures from the	All operations of Level 3, plus a calibrated	Conditions of Level 3, plus: sufficiently long and dense	Continuous monitoring of high-value assets (e.g.

Level	Information content	Required data and processing	Conditions for reliability	Indicative application scenario
	calibrated baseline through a normalised monitoring statistic suitable for automated alerting.	surrogate model trained on satellite displacements and environmental drivers, and a validated statistical framework for alert thresholds.	calibration period, environmental inputs capturing the dominant reversible drivers, and coherence-aware data selection that excludes unreliable input points.	heritage monuments, critical infrastructure) where prospective alerting is the operational requirement.

Chapter 5

Applications and case studies¹

5.1. Screening and prioritization using EGMS products

The translation of satellite-derived displacement data into actionable engineering information requires a structured approach to data consumption. For national or regional infrastructure stocks, the direct application of user-driven processing to every individual asset is neither feasible nor warranted at a first stage. A more practicable strategy consists in exploiting freely available, pre-processed products as a first tier of analysis, reserving user-controlled workflows for those assets that exhibit anomalous or ambiguous behaviour. The EGMS provides exactly this function. Covering the European territory with Sentinel-1 C-band data and delivering three progressively refined product levels, EGMS constitutes the most operationally accessible entry point to satellite interferometric data for structural and infrastructure engineers working at portfolio scale (Ferretti et al., 2023; Kotzerke et al., 2022).

The quality attributes distributed with EGMS products, described in detail in Section 3.2.3, constitute the primary basis for screening. Mean LOS velocity, temporal coherence, and the residual root mean square error of the time series are the three attributes most directly relevant to structural prioritisation. Mean LOS velocity provides a first-order indicator of displacement rate. Assets warrant further

¹ Part of the work described in this chapter was also published in (Coccimiglio et al., 2026; Tarantini, Coccimiglio, et al., 2026; Tarantini et al., 2025; Tarantini, McAvoy, et al., 2026; Tarantini, Miraglia, et al., 2026)

attention when their mean velocities clearly exceed the velocity precision of the EGMS Calibrated product. The velocity precision is specified at 0.7 mm per year as the mean velocity standard deviation across the Basic, Calibrated, and Ortho products of the service (Kotzerke et al., 2022). The specification has been empirically confirmed at the millimetre level against external geodetic networks, with reported root mean square errors in the range of 0.5 to 2 mm per year depending on the validation site and the geodetic reference (Crosetto et al., 2025; Esteves Martins et al., 2024), warrant further attention. Temporal coherence quantifies the stability of the scattering behaviour over the observation period and effectively filters out points whose time series are too noisy to support reliable trend estimation. Points with coherence below 0.5 are excluded from EGMS products during processing, and an analyst-defined coherence threshold, such as the retention of points with coherence equal to or above 0.8, further reduces the set to those whose time series are sufficiently reliable for comparative evaluation. The residual root mean square error provides a complementary measure of point-level precision, capturing the residual variance of the time series after model fitting, and can be used to flag points where systematic behaviour may be masked by high noise levels.

The use of large-scale MT-InSAR products for infrastructure prioritisation has been explored in several recent studies on Italian bridge portfolios. Calò et al. (2023) proposed a framework that combines MT-InSAR data with structural information and environmental records to estimate deformations over time on simply supported concrete girder bridges, formalising the link between interferometric observations and structural condition indicators. The same group developed an early warning system based on MT-InSAR data and ancillary information for the monitoring of single-span and multi-span simply supported concrete girder bridges within a bridge portfolio (Calò et al., 2024). Nettis et al. (2023) proposed a methodology for the portfolio-scale detection of structural deformations of bridges through MT-InSAR data, with a geoprocessing chain that supports the automatic interpretation of bridge-specific deformation scenarios and the definition of monitoring and assessment priority classes. These contributions collectively support the view that MT-InSAR data, whether derived from user-driven processing or from standardised products such as EGMS, are most productive when used as a prioritisation filter rather than as a diagnostic endpoint. The specific choice of data source depends on the temporal coverage required, the spatial resolution needed, and the availability of processing infrastructure.

The interpretative constraints of the screening stage deserve explicit acknowledgement. Several limitations discussed in Chapter 3 apply with particular

force at this level of analysis. The EGMS temporal window currently covers the period from approximately 2018-2019 onwards, which means that structural events occurring before this baseline, including pre-collapse signatures such as those documented for the Morandi Bridge in Genoa (Milillo et al., 2019), fall entirely outside the reach of any EGMS-based assessment. The spatial resolution of Sentinel-1 C-band data, approximately 20 m in range for the IW mode, restricts the density of measurement points on narrow structures such as bridge decks and single-span buildings, where point counts may be too low to support spatial pattern analysis within the EGMS framework. The referencing conventions of EGMS products introduce a further systematic constraint: any velocity anomaly identified through EGMS screening is relative to the internal reference of the processed tile, so that a structure exhibiting near-zero mean velocity may still be undergoing absolute motion if the surrounding tile is itself moving coherently. Absolute velocity referencing through the Calibrated product mitigates but does not fully eliminate this concern, since the GNSS-based harmonisation introduces its own uncertainty, particularly where the density of reference stations is low.

Despite these constraints, the practical value of EGMS as a screening layer is well established in the literature and confirmed by the applications presented in this thesis. The ReLUIIS guidelines for the use of interferometric data in structural engineering (Prota et al., 2025) explicitly position EGMS products as a first-order resource for identifying structures warranting deeper investigation, while cautioning against their use as stand-alone evidence for structural assessment. The present thesis adopts this same framework and applies it consistently across the case studies. In each case, the EGMS layer is used to establish a displacement context, identify anomalies or confirm expected behaviour, and inform the decision to proceed to user-driven processing. The transition from EGMS-based screening to user-driven processing is not governed by a single universal threshold, but responds to a combination of indicators whose formal articulation is deferred to Section 5.5. The case studies that follow illustrate how the screening logic operates in contrasting structural contexts, namely a masonry heritage building in an urban setting in Section 5.2, reinforced concrete infrastructure subjected to catastrophic collapse in Section 5.3, and a monumental structure where the interaction between soil deformation, environmental variables, and structural response demands a coherent integration of multiple data sources in Section 5.4.

5.2. Case study I - Santissima Trinità in Turin: building-scale SHM of a masonry heritage structure

The first case study is designed to test three components of the methodology developed in Chapter 4 under conditions that are representative of a recurrent monitoring scenario in Italian urban heritage: a compact masonry building of complex geometry, embedded in a dense historic centre, with limited prior instrumentation and no dedicated radar reflectors. The first component is the point selection logic of Section 4.1, concerning how many geometrically plausible and phase-stable measurement points can be identified on a building whose radar-facing surfaces are partially occluded by adjacent structures. The second component is the comparison between standardised EGMS products and user-driven processing, which provides an empirical basis for the transition criteria discussed in Section 4.6: under what conditions does a screening-level analysis yield sufficient information, and what does a user-driven analysis add in terms of point density, measurement quality, and structural interpretability? The third component is the time-series analytics of Section 4.2, applied here to assess whether a structure that exhibits no velocity anomaly at the EGMS screening level can nonetheless be characterised more precisely through the decomposition of seasonal and residual displacement components in a higher-coherence user-driven dataset. The case study therefore operates at the boundary between Levels 1 and 2 of the diagnostic framework of Section 4.6, and its principal finding concerns the conditions under which progression from one level to the next is both feasible and informative for heritage building management.


The Church of the *Santissima Trinità* in Turin represents one of the most geometrically distinctive examples of late-Renaissance sacred architecture in Piedmont. The choice of this structure as the first case study is motivated by three convergent considerations. First, the church belongs to the category of monumental masonry heritage buildings whose structural condition requires continuous attention but whose monitoring history is fragmented, making satellite-based retrospective observation particularly valuable. Second, the building has been the subject of a parallel programme of structural and dynamic investigation within the Department of Structural, Geotechnical and Building Engineering (DISEG) research group at Politecnico di Torino (PoliTO), providing an interpretative framework against which satellite-derived observations can be contextualised. Third, the case allows the comparison of two operationally distinct workflows, namely EGMS-based pre-processed data and user-driven processing, on a single

building of bounded footprint embedded in a dense urban environment, thereby isolating the effect of the processing strategy on point density, coherence, and interpretability.


The following subsections develop each component of this case study in sequence. Section 5.2.1 provides the structural and historical context. Section 5.2.2 describes the available satellite datasets and their acquisition characteristics. Sections 5.2.3 and 5.2.4 present the EGMS-based and user-driven analyses respectively. Section 5.2.5 places the two approaches in direct comparison across point density, coherence quality, and interpretability. Section 5.2.6 draws the implications for the SHM of masonry heritage buildings in general.

5.2.1. Structural and historical context

The Church of the *Santissima Trinità* occupies a prominent position in the heart of Turin's historic centre, embedded within one of the most dense and architecturally stratified urban blocks of the city. The building was commissioned in 1598 by the confraternity of the same name and designed by Ascanio Vitozzi, ducal architect to Carlo Emanuele I of Savoy. The central plan is organised around a circle inscribed within two superimposed equilateral triangles, forming a six-pointed star that gives expression to the Trinitarian symbolism of the patron confraternity. Structurally, this geometry produces an alternating sequence of six convex chapel apses and six re-entrant angles at the perimeter, resulting in a planimetric irregularity that concentrates stress in the transition zones between the radiating chapels and the central cylindrical body. The dome and lantern dominate the composition in proportional terms, rising well above the relatively low cylindrical drum walls, and the lateral thrusts generated by the dome mass act on a drum whose geometry does not provide the uniform perimeter continuity that circular-plan domes normally rely upon (De Lucia et al., 2020, 2021). The compact footprint of the church, surrounded on three sides by buildings of comparable or greater height, defines a radar visibility envelope that the satellite analyses of Sections 5.2.3 and 5.2.4 must accommodate, as shown in Figure 5. 1.



(a)



(b)

Figure 5. 1. Panel (a): Aerial view of the *Santissima Trinità* in Turin, embedded within the surrounding historic block. Panel (b): Plan of the urban block with the footprint of the church highlighted.

The exterior of the church presents a compact volumetric profile whose section reveals the full height hierarchy: a lower cylindrical body, an eight-pilaster drum, an oval-plan dome, and a lantern. The plan organisation of the church and the geometry of the dome and lantern, which carry the principal radar-reflecting surfaces analysed in Sections 5.2.3 and 5.2.4, are shown in Figure 5. 2. A distinctive feature of the building, with direct relevance for the interpretation of radar scattering behaviour, is that the lantern shaft has remained an unfinished brick surface since its original seventeenth-century construction. The exposed brickwork of the lantern shaft constitutes a primary radar-reflecting masonry surface of the building, although it is not the only such surface, as the discussion of the four scattering regimes below makes clear.




Figure 5. 2. *Santissima Trinità* in Turin: panel (a), plan; panel (b), longitudinal section of the dome and lantern.

The dome envelope is not homogeneous in its external materials and presents a stratified configuration that the radar interrogates as a sequence of distinct surfaces. The drum, which is the cylindrical body supporting the dome, is clad on its external face with metallic sheet, plausibly copper, by visual identification on close-range aerial imagery and by analogy with the documented Turin Baroque conservation

practice on comparable structures such as the Chiesa di San Filippo Neri. Above the drum, the masonry shell of the dome is concealed beneath an upper tile cover supported by radial masonry false struts that rise from the dome base. The radar-facing surface of the dome at this elevation is therefore the tile cover; whether a metallic layer is interposed between the masonry shell and the tile cover, as suggested by close-range visual inspection in some areas, cannot be confirmed from the documentary record consulted, and is in any case shielded from the radar line of sight by the tile cover above. From the dome cover emerges the lantern shaft, in exposed brick masonry, and the lantern is closed by a small cap (cupolino) clad with metallic sheet.

The dome envelope therefore presents four distinct radar-scattering regimes. The first is the metallic-clad drum, whose backscatter is dominated by specular reflection and edge diffraction at C-band and X-band wavelengths, with bright dihedral returns at the seams between sheets and at the joints with the masonry below. The second is the masonry false struts, which are vertical radial walls of brick that intersect the metallic drum cladding below and the tile cover above, generating dihedral configurations that are the canonical hosts of stable double-bounce PS in heritage monitoring. The third is the tile cover of the dome, whose ceramic surface scatters as a low-loss dielectric of intermediate dielectric constant, with a backscatter coefficient between that of brick and that of metal but with no specular contribution; this is the radar-facing surface at the elevation of the dome, irrespective of any subjacent layer that the tile cover may shield from the LOS. The fourth is the exposed-brick masonry of the lantern shaft, whose dielectric properties vary with moisture content and temperature on timescales ranging from hours to weeks. The metallic cap of the lantern adds a fifth, geometrically minor element. The interpretation of the displacement field over the church requires this stratified envelope to be recognised, because the scattering signature of each PS depends on which of these regimes the geocoded position can be plausibly attributed to. The four scattering regimes of the dome envelope are visible in close-range aerial photography, descending from the metallic “*cupolino*” of the lantern at the top, through the exposed-brick lantern shaft and the tile-covered dome supported by masonry false struts, to the metallic-clad drum at the base, as shown in Figure 5. 3.



(a)



(b)



(c)

Figure 5. 3. Close-range aerial view of the dome system of the *Santissima Trinità*. (a) “Cupolino” of the lantern, clad with metallic sheet. (b) Lantern shaft in exposed brick and tile cover of the dome supported by masonry false struts. (c) Metallic-clad drum at the base of the dome.

The interior, originally articulated in the austere stone vocabulary of late mannerist Piedmonts architecture, was extensively redecorated in the eighteenth century by Filippo Juvarra and his pupil Giovanni Pietro Baroni, who introduced polychrome marble cladding and stucco decoration in the chapels and the presbytery. The choir volume attached to the rear of the presbytery was heavily damaged by Allied bombing in 1943 and subsequently rebuilt with reinforced concrete structural elements in the post-war period. The result is a building that combines seventeenth-century masonry of varying composition and quality, eighteenth-century additions in partially heterogeneous materials, and a post-war reinforced concrete volume whose boundary with the historic fabric introduces a structural discontinuity that is of direct relevance to any attempt at satellite-based settlement pattern analysis. The masonry materials throughout include a mix of irregular rubble in lime mortar for internal walls and more regularly coursed brick for the drum and dome, with the foundation system resting on the alluvial deposits characteristic of the Turin plain, whose long-term consolidation behaviour has not been fully characterised for this specific building.

A programme of conservation interventions was carried out on the church between 2014 and 2021, commissioned by the Arciconfraternita della *Santissima*

Trinità and funded principally by the Compagnia di San Paolo foundation under its Edifici Sacri programme. The programme was organised in three sequential lots. The first lot, executed between March and November 2016, comprised the conservation of the dome, drum, lantern, and upper dome cap, and was carried out by specialists from the Centro Conservazione e Restauro La Venaria Reale using interior scaffolding that reached up to the lantern at approximately 47 m height. The second lot, executed during 2018 and concluded at the beginning of 2019, treated the presbytery and the nave, including the polychrome marble surfaces and the main altar. The church was reopened to the public on 16 March 2019 following the completion of the second lot. A third lot, addressing the lateral altars, was carried out in 2021. No external scaffolding on the dome exterior is documented at any point in this programme; all conservation activities relied on internal scaffolding installed inside the church volume. The exterior surfaces of the dome envelope were accordingly not subjected to any documented direct intervention during the 2014-2021 period. This chronology carries specific implications for the interpretation of satellite time series spanning the 2018-2023 observation window, discussed in detail in Section 5.2.4.

From a structural monitoring perspective, the *Santissima Trinità* presents challenges that are representative of a broad category of masonry heritage buildings in Italian urban contexts. The coexistence of original seventeenth-century masonry, eighteenth-century additions in heterogeneous materials, and post-war concrete reconstruction creates a structural system whose response to long-term loading is not easily predicted from monolithic models. Previous investigations by the DISEG research group have addressed the geometric characterisation of the dome through three-dimensional laser scanner survey integrated with finite element modelling, including linear and nonlinear structural analyses of the dome-drum system that identified the significance of irregularities in the dome intrados and their relationship to observed crack patterns (De Lucia et al., 2020, 2021). More recently, the group applied experimental dynamic data to improve the numerical description of the structural behaviour of a structurally analogous Baroque church (Ceravolo et al., 2024), developing a model-updating procedure that constitutes a methodological reference for the interpretation of the *Santissima Trinità* displacement data in relation to its finite element description. Parallel methodological work at Politecnico di Torino on adaptive hyperparameter selection for modal tracking algorithms in masonry monumental structures (Coccimiglio et al., 2025) provides the analytical framework within which long-term modal monitoring of heritage structures can be coupled with the satellite-derived quasi-static record analysed in the present case study.

Two vibration-based investigation campaigns have been conducted on the *Santissima Trinità* by the DISEG group. The first was carried out in 2010 and established the fundamental dynamic parameters of the structure through ambient vibration testing. The second, more recent campaign was conducted between December 2022 and March 2023, organised in two measurement windows between 22 December 2022 and 9 January 2023 and between 21 February and 17 March 2023. Three seismic accelerometers of the Integrated Circuit Piezoelectric type, model 393B12, were installed on the lantern to capture the diaphragmatic vibrations of this component, with two sensors oriented transversally and one longitudinally. The experimental setup exploited operational modal analysis techniques under ambient excitation, enabling the extraction of natural frequencies and mode shapes without the need for artificial forcing (Lenticchia et al., 2023). The natural frequencies identified in these campaigns constitute the observational target for model updating procedures of the type described by Ceravolo et al. (2024) and by Lenticchia et al. (2025), which calibrate finite element descriptions of the structural behaviour using experimental data as constraints.

The structural assessment conducted by De Lucia et al. (2020, 2021) on the basis of three-dimensional laser scanner survey integrated with finite element modelling provided the first quantitative characterisation of the geometric irregularities of the dome intrados, including depressions and deviations from the theoretical elliptical surface. These irregularities, which reflect the accumulated effect of differential foundation settlements and long-term creep of the masonry, were found to be structurally significant in the context of nonlinear analyses using a Drucker-Prager constitutive model for the dome-drum system. The relationship between foundation settlements and the observed crack pattern in the drum and the lower cylindrical walls constitutes one of the primary structural hypotheses that the satellite displacement data are called upon to test: if differential settlement is ongoing, it should be detectable as a spatially non-uniform velocity field across the building footprint, potentially concentrated in the transition zones between the chapels and the cylindrical body, and potentially correlated with the crack distribution documented in the structural surveys.

The monitoring of masonry heritage buildings through dynamic methods has been extensively discussed in the literature, with reference to both permanent systems and periodic campaigns (Ramos et al., 2010; Saisi & Gentile, 2015). The common finding across these studies is that environmental variables, particularly temperature, humidity, and soil moisture, produce systematic oscillations in the identified natural frequencies that must be separated from damage-related changes

before any structural diagnosis can be drawn. The satellite-based contribution examined in this case study addresses a different but complementary dimension of the same monitoring problem: rather than tracking the dynamic response, it aims to characterise the quasi-static displacement behaviour of the structure over the five-year EGMS observation window, and to assess whether the displacement amplitudes and spatial patterns are consistent with a stable, undamaged masonry body or suggest a more complex long-term deformation mechanism.

5.2.2. Available satellite datasets

Two distinct satellite datasets were employed in this case study, differing in their origin, temporal coverage, processing level, and geometric configuration. Their complementarity is central to the comparative argument developed in the subsequent sections.

The first dataset consists of EGMS Calibrated products downloaded directly from the EGMS data portal for the area of interest centred on the *Santissima Trinità* and its surrounding urban block. The EGMS Calibrated product was selected in preference to the Basic product because its GNSS-referenced absolute displacement values allow comparison between measurement points distributed across the structure and the adjacent buildings without the tile-referencing constraints that affect the Basic product. As described in Section 3.2.3, the Calibrated product provides LOS mean velocity, temporal coherence, RMSE, and full displacement time series for each measurement point, derived from Sentinel-1 IW TOPS acquisitions (Ferretti et al., 2023; Kotzerke et al., 2022). Data are available in both ascending and descending geometries, with temporal coverage spanning from January 2018 to December 2023, corresponding to the 2018-2022 and 2019-2023 EGMS processing releases. The acquisition mode is IW with VV polarisation, providing a nominal ground resolution of approximately 20 m in range and 22 m in azimuth before multi-looking, reduced to approximately 5 m in range for IW SLC products. The revisit interval of the Sentinel-1 constellation over the Turin area is 6 days in the standard dual-satellite configuration, though effective temporal sampling in the EGMS products is typically 12 days per geometry due to single-orbit selection in processing.

The second dataset consists of raw Sentinel-1 SLC images downloaded from the ASF Data Search Vertex portal and processed entirely within the user-driven workflow. The dataset covers the period from 2019 to 2023, comprising 169 scenes acquired in IW mode with descending orbit geometry and VV polarisation. The

restriction to descending geometry in the user-driven processing was motivated by the illumination geometry of the church within its urban block. The two free façades of the *Santissima Trinità* open onto the corner of via Garibaldi to the south and via XX Settembre to the east, while the remaining two flanks of the building are embedded in the surrounding urban fabric. A descending Sentinel-1 pass over Turin illuminates the ground towards the west-northwest, by virtue of the right-looking convention applied to a near-polar orbit with a heading of approximately 190 degrees from north, and therefore illuminates the south-facing and east-facing free façades of the church. The ascending pass, by contrast, illuminates from the east-northeast and reaches predominantly the embedded north-western flanks of the building. The empirical distribution of the persistent scatterers retained in the user-driven solution, which concentrates on the east façade of the church, is consistent with this geometric argument. The master image was selected automatically using an optimisation algorithm that minimises the combined spatial baseline, temporal baseline, and Doppler centroid frequency differences across all acquisitions in the stack, yielding a master date of 22 July 2020. An external high-resolution DEM from the Copernicus DEM programme was integrated to support accurate geocoding and to generate the synthetic amplitude image used for GCP selection. The principal characteristics of the two datasets are summarised in Table 5. 1.

Table 5. 1. Summary of the two satellite datasets used in the case study of the *Santissima Trinità*.

Field	EGMS Calibrated	User-driven Sentinel-1
Sensor	Sentinel-1 (ESA, Copernicus)	Sentinel-1A
Product	EGMS Calibrated, releases 2018-2022 and 2019-2023	SLC, IW TOPS
Temporal coverage	January 2018 - December 2023	2019 - 2023
Number of acquisitions	n/a (pre-processed product)	169
Orbit geometry	ascending and descending	descending
Polarisation	VV	VV
Acquisition mode	IW	IW
Incidence angle (sub-swath IW2)	approximately 39°	approximately 39°

Field	EGMS Calibrated	User-driven Sentinel-1
Wavelength (C-band)	5.55 cm	5.55 cm
Revisit interval	6 days nominal; 12 days effective per geometry	6 days nominal
Master image date	n/a	22 July 2020
Processing	EGMS Calibrated, GNSS-referenced	user-driven, small-area

The urban context of the *Santissima Trinità* introduces several geometric constraints on radar target visibility that must be stated before the analysis is presented. The church is surrounded by continuous building fronts of comparable or greater height, with the main façade opening through a narrow street connection and the rear of the choir abutting directly onto the adjacent urban block. In the Sentinel-1 IW descending geometry, the incidence angle over Turin is approximately $38\text{--}40^\circ$ for sub-swath IW2, which is the sub-swath covering the city centre. At this incidence angle, the dome and lantern of the church, rising to approximately 40 m above ground level, fall within the layover shadow of the taller surrounding buildings when viewed from the descending pass, with the exact extent of the layover column depending on the precise azimuth of the satellite track relative to the building orientation. The practical effect is that some PS candidates nominally geocoded to the dome or drum position may in fact originate from offset ground-level reflecting surfaces whose slant-range position coincides with the dome location in the SAR image plane. The geometric plausibility checks described in Section 4.1.3 are therefore applied with particular rigour in this case study.

The temporal overlap between the two datasets and the conservation chronology summarised in Section 5.2.1 carries specific implications for the interpretation of the radar time series. The EGMS observation window opens in January 2018, approximately fifteen months after the completion of Lot 1, which addressed the dome and lantern interiors between March and November 2016 with interior scaffolding only. The user-driven stack opens in 2019 and overlaps with Lot 2 (presbytery and nave, completed at the beginning of 2019) and with Lot 3 (lateral altars, 2021), both carried out within the church volume and without external scaffolding on the dome envelope. The exterior surfaces of the dome envelope are accordingly not affected by documented conservation activity during the joint EGMS and user-driven observation windows, and the radar time series are expected to reflect the stable long-term scattering configuration of the four envelope regimes identified in Section 5.2.1, without superimposed transient effects from

scaffolding. The scattering implications of the four regimes are discussed in Section 5.2.4.

5.2.3. EGMS-based analysis

As a first analytical step, we downloaded and examined the EGMS Calibrated product for the area of interest centred on the *Santissima Trinità* and its surrounding urban block. The Calibrated product was preferred over the Ortho product for the reasons set out in Section 5.2.2: the 100 m regular grid of the Ortho product would reproject many of the measurement points detected on and around the church onto interpolated positions that no longer correspond to identifiable structural surfaces, whereas the Calibrated product preserves the native point positions derived from the PS selection algorithm with sub-pixel geolocation accuracy relative to the SAR image geometry. The data were downloaded for both ascending and descending geometries and for the two available temporal releases covering 2018-2022 and 2019-2023 respectively.

The spatial distribution of EGMS measurement points across the area of interest reflects the typical density expected in a dense Italian historic urban centre processed with Sentinel-1 C-band data. The urban fabric surrounding the church, composed of continuous residential and institutional building fronts of five to seven storeys, provides a rich stock of stable radar reflectors (roof edges, metallic balcony rails, window frames, and parapets) that generate a dense population of PS candidates throughout the tile. Within the church footprint itself, the number of EGMS points attributable to the structure is more limited, as expected from the geometric constraints discussed in Section 5.2.2. The dome surface of exposed brick produces a spatially heterogeneous backscatter pattern that is detectable in the reflectivity map, but the combination of the incidence angle and the surrounding building heights places a significant portion of the dome in partial layover, redistributing its radar signature in the range direction in a manner that complicates unambiguous attribution to specific structural elements.

We selected four measurement points for detailed examination: two from the EGMS dataset (P20 and P24) and two from the user-driven dataset (P1803 and P1765), the latter pair discussed in Section 5.2.4. The EGMS selection focused on points geocoded in close proximity to the church body, on the metallic-clad drum and on the cylindrical-body masonry below it, with coherence values as high as available and mean velocities within the range expected for a stable urban structure.

The complete quality attributes of all four selected points are reported in Table 5. 2.

Table 5. 2. Quality attributes of the four selected MPs on the *Santissima Trinità* (P20 and P24 from EGMS Calibrated; P1803 and P1765 from user-driven Sentinel-1).

	ID	Lat	Lon	h [m]	Vel [mm/year]	Temporal coherence	Disp _{std} [mm]	Disp _{mean} [mm]
EGMS	P ₂₀	45.07197	7.683567	258.9	0.5	0.86	2.75	1.282
	P ₂₄	45.07196	7.683607	252.9	-0.1	0.66	4.05	0.466
User-driven	P ₁₈₀₃	45.07209	7.683617	264.1	-0.2	0.92	2.25	0.548
	P ₁₇₆₅	45.07195	7.683671	248.3	-0.5	0.86	1.39	-0.252

Point P20 exhibits a temporal coherence of 0.86 and a mean LOS velocity of approximately 0.5 mm/year, consistent with a stable scatterer that undergoes no significant long-term trend over the 2018-2023 observation period. The geocoded position of P20 places it at the level of the drum below the dome, in the zone where the metallic-clad drum cladding meets the masonry false struts that support the tile cover above, producing a dihedral configuration that is the canonical host of stable double-bounce PS in heritage monitoring (Cigna et al., 2014). The high coherence of P20 and the absence of long-term trend in its displacement time series are consistent with this attribution. The geocoded height of P20 (258.9 m) carries a typical single-PS vertical uncertainty of the order of a few metres for Sentinel-1 IW data (Ferretti et al., 2023), and the assignment to this elevation rather than to the cylindrical wall below or to the upper dome surface above rests on the relative ordering of the four selected points and on the local urban geometry rather than on the absolute height value alone. The displacement time series of P20 oscillates within a band whose amplitude does not exceed the range typical of urban C-band scatterers subject to annual thermal cycling. The time series plots, shown in Figure 5. 4, reveals a seasonal component visible as a low-amplitude sinusoidal modulation of period approximately 12 months. The seasonal component at this amplitude is consistent with the thermoelastic response of an urban masonry structure to annual temperature variations, as discussed in Section 4.2.2, although the metallic component of the drum-level reflectors at the P20 position introduces a thermal expansion contribution from the cladding itself that adds to, rather than substitutes for, the masonry response. No systematic drift, change-point, or anomalous excursion is detectable in the P20 time series over the full five-year window,

suggesting that the radar-reflecting fabric attributed to this measurement point has remained stable throughout the observation period.

Point P24, by contrast, presents a temporal coherence of 0.66, considerably lower than P20 and below the threshold of 0.80 that would normally be applied in a structure-specific quality filter. The mean LOS velocity of P24 remains close to zero, at approximately -0.1 mm/year, for the majority of the time series. The displacement time series of P24, shown in Figure 5. 4, panel (b), exhibits a systematically higher oscillation amplitude than P20 throughout the observation window, and the amplitude appears to increase further toward the end of the 2019-2023 period. Two hypotheses can be considered to explain this behaviour. The first is a random degradation of the scattering configuration attributable to weathering or local surface modification; the second is a physical mechanism related to the material characteristics of the reflecting surface itself. The second hypothesis is supported by the position of P24 relative to the structural zones of the church. The geocoded height of P24, at 252.9 m, is the lowest among the four selected points and is approximately 6 m below the elevation of P20; allowing for the typical few-metre vertical uncertainty of single-PS Sentinel-1 heights (Ferretti et al., 2023), this places P24 on a portion of cylindrical-body masonry below the drum, that is, on the exposed-brick fabric of the lower body of the church rather than on the metallic-clad ancillary roofs above. Under this attribution, the lower temporal coherence and larger oscillation amplitude of P24 become physically interpretable as manifestations of the intrinsic thermal and hygric sensitivity of exposed brickwork. Seventeenth-century brick masonry of the type used for the *Santissima Trinità* is porous, retains moisture after rainfall, and exhibits dielectric properties that vary with both temperature and moisture content on timescales ranging from hours to weeks. The backscatter behaviour of such a surface at C-band is accordingly more variable than that of dry, hydrophobic, or metallic surfaces, which is sufficient to produce the observed coherence reduction without invoking any structural or maintenance-related cause.




Figure 5. 4. Time series of selected points: (a) P₂₀, (b) P₂₄, (c) P₁₈₀₃, and (d) P₁₇₆₅

The increase in oscillation amplitude is concentrated in the portion of the time series from early 2022 onward, when the dispersion of the displacement values around the multi-year mean grows visibly with respect to the preceding years. Two physical mechanisms could account for this signature. The first is a moisture-driven decorrelation of the brick scattering surface during the prolonged drought that affected northwestern Italy from December 2021 onward (Toreti et al., 2022), which produced a multi-month departure from the average moisture state of the masonry and therefore a shift of its dielectric response outside the regime within which the temporal coherence assumption holds. The second is a localised modification of the scattering surface at the P24 position, of an undocumented nature, that progressively altered the scattering geometry and degraded the coherence of the corresponding pixel without producing a structurally meaningful displacement signal. The two hypotheses are not mutually exclusive and cannot be discriminated on the basis of the EGMS time series alone; they are presented here as alternative explanations of the same empirical pattern, both consistent with the broader interpretation of P24 as a low-coherence point whose displacement signature reflects properties of the reflecting surface rather than of the structural fabric beneath. A third factor that may have contributed to the post-2021 decorrelation, and that cannot be ruled out from the available data, is the loss of Sentinel-1B on 23 December 2021, which doubled the nominal Sentinel-1 revisit interval over Turin from six to twelve days for the remainder of the observation

window, with a known effect on the temporal coherence of C-band interferometric stacks (see Section 3.5)

The lower temporal coherence and larger oscillation amplitude of P24 may be consistent with, rather than definitively caused by, the intrinsic thermal and hygric sensitivity of the exposed brick surface of the lower cylindrical body. This interpretation remains a working hypothesis that cannot be discriminated conclusively from alternative explanations (weathering, local surface modification, or residual processing artefacts) on the basis of the EGMS record alone, but it is the more parsimonious explanation given the documented physical characteristics of the radar-reflecting masonry fabric of the church. The coherence value of 0.66 falls in the range where EGMS does not exclude the point entirely but where the time series should be treated with caution (Kotzerke et al., 2022; Prota et al., 2025). The practical implication is that the interpretation of P24 as a structural signal would be unsound; its inclusion in a gradient analysis across the building would propagate the dielectric variability of the brick surface into an apparent structural displacement that does not reflect mechanical behaviour.

This observation has a general methodological implication. The EGMS point selection algorithm applies a temporal coherence threshold of approximately 0.5 to each subswath independently, and measurements across different subswaths are not cross-validated before point delivery. A point with coherence 0.66 is therefore included in the downloadable dataset and appears in the spatial distribution map with no visual flag distinguishing it from higher-quality neighbours. The user who applies no additional coherence filter above the platform default retains P24 alongside P20, giving both equal analytical weight in any spatial or temporal pattern analysis. The implication of this for building-scale SHM is that a coherence threshold of at least 0.80 should be imposed as a post-download filter when the analytical objective is structural interpretation rather than broad ground motion mapping. This threshold is consistent with the guidance of the ReLUIS interferometric guidelines (Prota et al., 2025) and with practice reported for heritage building applications in the literature (Cavalagli et al., 2019; Cigna et al., 2014).

The EGMS analysis of the area surrounding the church, extending to the adjacent building block and the facing building fronts, reveals a spatially coherent pattern of near-zero to slightly negative mean LOS velocities. The velocity field across the urban block does not show systematic gradients that would suggest differential foundation settlement at block scale, and the mean velocities of all points with coherence above 0.80 fall within the uncertainty range of the Calibrated

product, estimated at approximately ± 1 mm/year for the Turin area under typical atmospheric conditions. The results therefore appear to indicate that the area surrounding the *Santissima Trinità* is characterised by long-term stability in the satellite observation window, with no significant ongoing subsidence or differential settlement detectable at EGMS resolution and sensitivity.

The EGMS analysis, while informative at the level of the surrounding urban block, yields only a small number of measurement points attributable with reasonable confidence to the church structure itself. The practical effect of the per-swath processing and the 0.5 coherence threshold is that the density within the church footprint appears higher than what would survive a structure-specific quality filter. The enriched but lower-quality point set is one of the distinguishing characteristics of EGMS relative to user-driven processing, and forms the primary axis of the comparison developed in Section 5.2.5. The user-driven analysis presented in Section 5.2.4 approached the same structure with tighter quality controls and a building-scale area-of-interest definition.

5.2.4. User-driven Sentinel-1 analysis

We conducted the user-driven analysis on the 169-image descending-orbit Sentinel-1 SLC stack described in Section 5.2.2, following the multi-temporal PS-InSAR workflow of Chapter 3.

We co-registered all 169 SLC images to the master acquisition of 22 July 2020 and constructed the interferometric network as a star graph centred on the master, with normal baselines ranging from approximately -120 m to +270 m, all well within the critical baseline for C-band Sentinel-1. We integrated an external Copernicus DEM at the preliminary processing stage to remove the bulk topographic phase contribution from each interferogram, and we selected a GCP on a stable, high-amplitude scatterer in the immediate vicinity of the church to anchor the LOS solution and suppress residual long-wavelength phase ramps attributable to unmodelled orbital and atmospheric contributions. The four quality attributes of the user-driven Sentinel-1 dataset, namely the coherence-coloured star graph of the interferometric network, the distribution of the normal baselines, the distribution of the Doppler centroid frequencies, and the local temperature at the acquisition dates, are shown in Figure 5. 5.




Figure 5. 5. Quality attributes of the user-driven Sentinel-1 dataset: (a) star graph of the interferometric network centred on the master image with the perpendicular baseline of each spoke and the colour of each spoke encoding the temporal coherence of the corresponding pair; (b) distribution of normal baselines across the 169 acquisitions; (c) distribution of Doppler centroid frequencies; (d) local temperature at the acquisition dates.

The reflectivity map, computed as the pixel-wise mean amplitude across the full 169-image stack, reveals a spatially heterogeneous backscatter pattern across the urban block. The *Santissima Trinità* presents the stratified envelope described in Section 5.2.1, with reflectivity contributions distributed across four distinct regimes. The metallic-clad drum and the metallic cupolino of the lantern produce high-amplitude specular returns at the elevations corresponding to their respective LOS-facing facets; the masonry false struts at the dome base produce dihedral returns by interaction with the underlying drum cladding; the tile cover of the dome produces a moderate and spatially distributed reflectivity signature consistent with a low-loss dielectric surface; the exposed-brick lantern shaft produces a localised dielectric scattering signature whose amplitude varies with the temporal evolution of moisture content. The cylindrical walls and drum base generate moderate

reflectivity at the expected slant-range positions, though as anticipated from the layover geometry, some of this energy is redistributed towards the reflectivity of the surrounding lower buildings. Sparse point selection was implemented using the local maxima criterion on the reflectivity map, with an amplitude threshold selected to retain only pixels that are simultaneously local maxima and exceed a threshold more stringent than the default EGMS processing criterion. The purpose of this threshold choice is to select structurally attributable targets rather than to maximise point count. The mean amplitude of the user-driven Sentinel-1 stack over the small-area window centred on the church are shown in Figure 5. 6.




Figure 5. 6. Small-area maps for the user-driven Sentinel-1 processing of the *Santissima Trinità*: mean amplitude across the 169-image stack.

We defined a small-area processing window to encompass the church body and the immediately surrounding perimeter extending to the first row of adjacent buildings. We deliberately restricted the spatial extent of the small-area window to confine the residual atmospheric correction to a scale comparable to the dimensions of the target structure, on the basis of the consideration that atmospheric phase contributions decorrelate over distances of several kilometres while structural displacement gradients on a building of the *Santissima Trinità*'s footprint are expected to act over distances of tens of metres at most. A wider processing window would have averaged the atmospheric correction over an area in which the structural signal of interest is spatially negligible, while a narrower window would have produced too few candidate scatterers for a robust local reference selection. As described in Section 3.3.4, the small-area approach allows a denser candidate grid

and a locally adapted reference point to be applied, reducing residual atmospheric contributions across the small spatial extent of the structure. A temporal coherence threshold of 0.6 was applied as the minimum quality criterion for point retention, consistent with the ReLUIS interferometric guidelines for Sentinel-1 datasets of this density and duration (Prota et al., 2025). This threshold is more stringent than the EGMS default of approximately 0.5 and ensures that the retained points have maintained consistent scattering behaviour throughout the 2019-2023 observation window. The final geocoding step exported the retained PSs to the Google Earth reference system using the GCP correction, enabling direct spatial overlay with architectural and structural reference data.

We identified two measurement points, P1803 and P1765, on or immediately adjacent to the church structure with coherence values above the 0.6 threshold and geocoded positions attributable to the upper and lower portions of the radar-facing structural envelope respectively. The geocoded heights of the two points (264.1 m for P1803 and 248.3 m for P1765) span a vertical interval of approximately 16 m across the height hierarchy of the church. The vertical separation is consistent with the assignment of P1803 to the elevation of the dome base, where the masonry false struts support the tile cover, and of P1765 to the cylindrical-body elevation below the drum. As for the EGMS points, the absolute height values carry a typical single-PS uncertainty of the order of a few metres for Sentinel-1 IW data (Ferretti et al., 2023). Their full quality attributes are reported alongside those of the EGMS points in Table 5. 2. Quality attributes of the four selected MPs on the *Santissima Trinità* (P20 and P24 from EGMS Calibrated; P1803 and P1765 from user-driven Sentinel-1). of Section 5.2.3. Point P1803 exhibits a temporal coherence of 0.92, a mean LOS velocity of approximately -0.2 mm/year, and a displacement standard deviation of approximately 2.25 mm over the observation period. Point P1765 exhibits a temporal coherence of 0.86, a mean LOS velocity of approximately -0.5 mm/year, and a displacement standard deviation of approximately 1.39 mm. Both velocity values are within the precision floor of the user-driven Sentinel-1 solution, estimated at approximately ± 1 mm/year for mean velocity in urban environments (Fattahi & Amelung, 2014; Hanssen, 2001), and neither indicates any detectable long-term trend that would warrant a structural alert.

The displacement time series of P1803 and P1765, shown in Figure 5. 4, panels (c) and (d), exhibit a low-amplitude seasonal modulation of period approximately 12 months, consistent with the thermoelastic response of a masonry drum and cylindrical wall of this size under the Turin thermal cycle. The seasonal amplitude does not exceed the values reported for comparable heritage structures in the

literature (Monserrat et al., 2011). No change-point, accelerating trend, or anomalous excursion is detectable in either time series over the full 2019-2023 window. The standard deviation of P1803 is larger than that of P1765, which may reflect a marginally less stable scattering configuration at the P1803 position. The temporal coherence of 0.92 nevertheless indicates that the scattering mechanism remained sufficiently stable throughout the period for the displacement estimate to be reliable.

As established in Section 5.2.1, no documented restoration activity affected the exterior surfaces of the church during the observation window of the user-driven dataset. The 2018 works of Lot 2 addressed the interior of the presbytery and the nave, with scaffolding installed inside the church volume, and were completed before the dataset enters its statistically dense portion. The third lot of 2021 involved conservation of the lateral altars, again within the interior volume of the church. The user-driven time series are therefore expected to reflect the stable long-term scattering configuration of the church exterior rather than transient effects from conservation works. The two retained points P1803 and P1765, attributable respectively to the upper and lower portions of the radar-facing structural envelope, do not exhibit any amplitude anomaly or coherence degradation pattern that would be consistent with an undocumented intervention on their reflecting surfaces. The interpretation of P1803 and P1765 as structurally stable scatterers on the permanent masonry fabric is therefore supported by both the processing diagnostics and the documented conservation history of the building.

5.2.5. Comparison of point density, quality, and interpretability


The comparison between the two datasets rests on three analytically distinct axes: the number of measurement points attributable to the structure and its immediate surroundings, the quality indicators associated with those points, and the degree to which the resulting displacement information supports structural interpretation. The findings on each axis carry different implications for the design of satellite-based SHM workflows.

The spatial distribution of the EGMS Calibrated and user-driven measurement points over the urban block, together with the four points selected for detailed examination on the church, is shown in Figure 5. 7. Visual inspection reveals that the EGMS Calibrated product yields a higher apparent point count within the urban block than the user-driven solution. In the area immediately surrounding the *Santissima Trinità*, EGMS provides several tens of measurement points, of which


a subset can be attributed with varying degrees of confidence to the church structure or its base-level perimeter. The user-driven small-area analysis identifies a more restricted set, with two points (P1803 and P1765) retained after the combined reflectivity and coherence filter as structurally attributable candidates. The numerical difference, however, does not reflect a greater physical coverage by EGMS. The higher EGMS count is a consequence of two processing decisions that inflate apparent density without a corresponding gain in structural information: the application of a temporal coherence threshold of approximately 0.5 rather than 0.6, and the execution of candidate selection independently per subswath rather than across the full burst stack. In the per-subswath approach, a point that reaches the coherence threshold in one subswath is included in the delivered dataset even if it would fail a combined multi-subswath stability test. The user-driven workflow, by contrast, performs the full-stack assessment across all subswaths simultaneously, so that a candidate appearing coherent in one subswath but unstable in another is discarded before the final solution. The practical consequence is that EGMS point density tends to be inflated in areas where the spatial coverage of subswaths partially overlaps the target structure, and the inflation is concentrated precisely at low-to-moderate coherence values, the range most likely to introduce unreliable displacement estimates into a structural analysis (Crosetto et al., 2016).

The numerical values of the quality indicators make the quality difference explicit. The two user-driven points exhibit temporal coherence values of 0.92 (P1803) and 0.86 (P1765), with displacement standard deviations of approximately 2.25 mm and 1.39 mm respectively. The two EGMS points exhibit coherence values of 0.86 (P20) and 0.66 (P24), with standard deviations of approximately 2.8 mm and 4.1 mm. The EGMS point P20, with coherence 0.86, is qualitatively comparable to the user-driven point P1765 and would survive a coherence filter of 0.80 applied post-download. P24, with coherence 0.66, would not survive such a filter, yet appears in the downloaded dataset with no default visual distinction from P20. The displacement standard deviation of P24 at approximately 4.1 mm is approximately three times that of P1765, a factor that is directly relevant when displacement amplitudes of structural significance are expected in the sub-millimetre to low-millimetre range. In such conditions, retaining a point with σ of this magnitude introduces a noise level that masks the signal of interest rather than characterising it. The ReLUIS interferometric guidelines explicitly note that the structural significance of displacement estimates is conditional on a demonstrated signal-to-noise ratio, and that noise characterisation through standard deviation and coherence is a prerequisite for any structural interpretation rather than an optional

post-processing step (Prota et al., 2025). The user-driven workflow enforces this prerequisite by design, whereas the EGMS workflow delegates it to the analyst.



(a)



(b)

Figure 5. 7. Measurement points from the EGMS Calibrated product and from the user-driven Sentinel-1 processing over the *Santissima Trinità* in Turin, displayed on Google Earth in three-dimensional projection: panel (a) reports the area-wide distribution of all retained measurement points across the urban block centred on the church, while panel (b) reports the detail of the four selected points on the church itself, namely P₂₀ and P₂₄ from the EGMS Calibrated product and P₁₈₀₃ and P₁₇₆₅ from the user-driven processing. The four points are the basis for the comparative analysis of Section 5.2.4.

Despite the density and quality differences, the displacement values recorded by the two approaches on the comparable points are numerically consistent and mutually corroborating. Mean LOS velocities range from approximately -0.5 to +0.5 mm/year across all four points, time series oscillate within a band consistent with thermoelastic seasonal modulation in both datasets, and the seasonal signature visible in both EGMS and user-driven records is of comparable amplitude and phase. The agreement between the two independent processing chains on the fundamental characterisation of the structure as stable and non-trending is the most important finding of the comparison: neither dataset raises a displacement signal that would trigger deeper structural investigation under the decision framework of Section 5.5.

The interpretation of the displacement field at building scale, however, favours the user-driven solution in one important respect. Because the user-driven processing applies the small-area module with a locally defined reference point and a tighter coherence filter, the residual atmospheric signal across the limited extent of the church footprint is reduced relative to the larger tile-level residuals that remain in the EGMS Calibrated product. For a building of the *Santissima Trinità*'s dimensions, the structural displacement gradient between any two measurement points on the same structure is expected to be of the order of a fraction of a millimetre over distances of 10 to 30 m. At this level, even modest tile-level atmospheric residuals can obscure intra-structure gradients entirely. The small-area solution mitigates this effect without eliminating it, but the mitigation is sufficient to make the spatial pattern of P1803 and P1765 more directly interpretable in structural terms than the EGMS map with its larger residual baseline. Studies applying PS-InSAR to comparable masonry heritage structures in Italian urban contexts have consistently reported that local processing and reference selection are necessary conditions for reliable gradient-level structural interpretation (Cavalagli et al., 2019; Cigna et al., 2014).

The comparison also underscores a consideration relevant to FE model calibration. The user-driven displacement time series of P1803 and P1765, with coherence values of 0.92 and 0.86 respectively and well-characterised noise levels, could in principle be used as observational constraints in a model-updating procedure of the type described by Ceravolo et al. (2024). The EGMS point P24, with its elevated standard deviation, would introduce excessive uncertainty into such a calibration and would be appropriately excluded. The EGMS point P20, comparable in quality to P1765, could serve as a supplementary constraint provided the temporal baseline difference and referencing offset between the two datasets are

accounted for in the model formulation. The practical implication is that EGMS provides a useful screening layer and a corroborating source, but user-driven processing is the necessary precondition for displacement data of sufficient quality to support FE model updating in heritage structural applications.

5.2.6. Implications for masonry heritage SHM

The primary finding of this case study is also its most structurally reassuring: the *Santissima Trinità* in Turin does not exhibit, within the 2018-2023 observation window, any satellite-detectable displacement trend, spatial gradient, or change-point behaviour that would indicate active settlement, ongoing differential deformation, or progressive structural deterioration. Mean LOS velocities across all retained points from both processing approaches fall within the precision floor of the respective datasets, time series oscillate within a band consistent with thermoelastic seasonal response rather than irreversible displacement, and the two processing chains agree on this characterisation independently. Under the decision framework described in Section 5.5, this outcome does not trigger a recommendation for deeper investigation, and the satellite monitoring contributes a negative result of genuine informational value: a stable baseline against which any future deformation signal can be assessed.

The significance of a negative result should not be underestimated in the context of masonry heritage monitoring. A large proportion of historic buildings in Italy are monitored only intermittently, if at all, due to the cost and logistical difficulty of installing and maintaining dedicated sensor networks. For structures in this category, the EGMS observation window provides the first multi-year deformation record at no cost and with no field intervention, and even a finding of stability is informative because it establishes a reference state against which future change can be measured. This retrospective function of satellite data has been highlighted in the literature on large-scale heritage monitoring and is one of the main reasons why the ReLUIS interferometric guidelines position EGMS as a first-tier resource for the systematic screening of the Italian built heritage stock (Prota et al., 2025). The *Santissima Trinità* case illustrates how this function operates in practice: the EGMS data confirm that the structural stability of the church is consistent with the absence of detectable ground-level or intra-structure displacement over the observation period, providing external corroboration for the structural assessments conducted through the parallel vibration-based campaigns.

The comparison between EGMS and user-driven processing establishes a methodological implication that extends beyond the *Santissima Trinità*. For heritage buildings embedded in dense urban environments, the apparent density advantage of EGMS conceals a coherence cost that is directly relevant to building-scale structural interpretation. The per-subswath processing strategy that generates higher EGMS point counts simultaneously increases the proportion of points with temporal coherence below 0.80, the threshold below which displacement standard deviations grow to levels that exceed the structural signal of interest in a stable masonry building. When the monitoring objective is limited to confirming stability or flagging large-amplitude deviations, this trade-off may be acceptable, because even low-coherence points can detect metre-level displacements. When the objective is to characterise the displacement field at the level of intra-structure gradients relevant to differential settlement analysis, or to provide observations suitable for FE model calibration, the coherence cost of the EGMS strategy becomes prohibitive. In this context, user-driven processing with a tighter coherence threshold and a locally defined reference is a necessary, rather than optional, upgrade from the screening level.

A second methodological implication concerns the interpretation of low-coherence measurement points on heritage structures with stratified envelopes. The reinterpretation of EGMS point P24 developed in Section 5.2.3 suggests that coherence values in the intermediate range (0.5 to 0.8) on masonry heritage buildings may reflect the intrinsic thermal and hygric sensitivity of the reflecting surface rather than structural instability or processing artefacts. Exposed brickwork, unrendered masonry, and other porous surfaces exhibit dielectric variability that manifests as moderate coherence reduction without any structural significance. The *Santissima Trinità* is representative of a broader category of heritage buildings in which different portions of the same envelope present radically different scattering behaviour: exposed brick on the dome shell and on the lantern shaft, metallic cladding on the drum and on the cupolino, polychrome marble or rendered surfaces on lower elevations, and dihedral configurations at the interfaces between adjacent components. Each of these regimes maps onto a distinct coherence and amplitude statistic, and the same numerical threshold therefore samples different physical surfaces with different effective stringency. The analyst who interprets low-coherence points as structural anomalies, or who includes them in spatial gradient analyses without a coherence-weighted uncertainty treatment, risks converting material properties of the building into spurious structural signals. This consideration applies with particular force to masonry heritage structures in which the exterior surfaces combine materials of different finishing and weathering

histories, as is commonly the case for buildings that have undergone multi-century accumulations of conservation interventions.

Several limitations of the present analysis should be acknowledged before the implications are generalised. The user-driven processing was restricted to the descending orbit, so the displacement component captured is the scalar LOS projection specific to that acquisition geometry. Vertical settlement and east-west horizontal motion contribute to the descending LOS measurement in proportions determined primarily by the incidence angle, but a rigorous decomposition into structural displacement components would require ascending-orbit data of comparable quality and a sufficient number of co-located points on the same structural elements in both geometries. At the resolution of Sentinel-1 C-band data, the number of structurally attributable points per element of the *Santissima Trinità* is too small for this decomposition to be reliable, and the single-geometry result must accordingly be interpreted as a scalar constraint on a linear combination of vertical and horizontal motion rather than as a direct measurement of settlement. This limitation is intrinsic to C-band urban monitoring at building scale and motivates the recommendation, discussed further in Section 5.5, that structures for which rigorous decomposition is needed should be targeted with higher-resolution X-band data.

The analysis lacks simultaneous in-situ displacement observations with which to cross-validate the satellite measurements. The vibration-based monitoring campaigns conducted by the DISEG group in 2010 and 2022-2023 characterise the dynamic rather than the static structural response, and their observables (natural frequencies and mode shapes) are not directly comparable with the quasi-static LOS displacement recorded by InSAR. The combination of the two monitoring sources is not without structural value: the modal frequencies identified in the 2022-2023 campaign, interpreted through the thermoelastic model updating framework of Ceravolo et al. (2024) and the long-term mode-tracking strategy proposed by Coccimiglio et al. (2023), constrain the effective stiffness distribution of the structure, which in turn determines the deformation modes that a given foundation settlement pattern would induce. The satellite displacement data establish that no such pattern has exceeded the detection threshold of the EGMS or user-driven solution over the five-year window, which is consistent with the absence of any degradation trend in the modal frequencies identified in the dynamic campaigns. The two monitoring dimensions are therefore mutually corroborating even in the absence of simultaneous measurements.

Looking forward, the most productive development for this case study would be the integration of the user-driven displacement time series into the existing FE model of the dome-drum system as boundary condition data, jointly with the experimental modal frequencies, in a combined static-dynamic updating procedure. This approach, conceptually aligned with the methodology described by Ceravolo et al. (2024) and Lenticchia et al. (2025), would allow the mechanical parameters and foundation compliance to be estimated from both static and dynamic observational evidence simultaneously. The satellite data would constrain the long-term foundation displacement component of the model, while the modal frequencies would constrain the stiffness distribution within the structural body. Such integration would convert the satellite observation from a standalone monitoring product into an active input for structural identification, and would represent a concrete step towards the decision-ready diagnostic framework articulated in Section 4.5.3.

5.3. Case study II - Retrospective forensic analysis of infrastructure collapses: the Himera Viaduct and the Albiano-Magra Bridge

The second case study applies the retrospective event analysis methodology developed in Section 4.3.4 to two Italian infrastructure collapses that occurred within the temporal coverage of available SAR archives. The *Himera Viaduct*, which partially collapsed in April 2015 due to a landslide in central Sicily, is analysed using both post-reconstruction EGMS data and user-driven Sentinel-1 processing spanning the pre- and peri-collapse period. The *Albiano-Magra Bridge*, which collapsed in April 2020 in north-western Tuscany, is analysed using a unified COSMO-SkyMed archive covering 2011-2022 in both ascending and descending geometries. The two cases share four characteristics that make them jointly informative for the purposes of this thesis. Both are reinforced concrete structures in geotechnically hazardous environments. Both produced TS whose lifecycle boundaries constrain the event chronology. Both exhibit pre-collapse kinematic signatures that are spatially correlated with the failed portions but are close to the statistical detection margin of the technique. Both allow the direct comparison of standardised and user-driven processing in a forensic context. The case study tests the three forensic constraints identified in Section 4.3.4 (temporal resolution, statistical detectability, epistemological gap) against empirical evidence and assesses the diagnostic level that retrospective satellite analysis can realistically achieve for infrastructure failure investigation.

The Italian infrastructure stock provides a particularly relevant context for this type of investigation. The national bridge inventory, estimated at approximately 60,000 viaducts and bridges along the primary and secondary road network, is characterised by a high proportion of structures that were designed and built in the 1950s, 1960s, and 1970s to design standards that are no longer current, and that have now reached or exceeded their nominal service life without systematic health monitoring in place. The collapse of the Morandi Bridge in Genoa in August 2018, which caused 43 fatalities, catalysed a national regulatory response that culminated in the ministerial guidelines for the classification and safety management of existing bridges adopted in 2020. Within this regulatory framework, satellite interferometric data have been explicitly recognised as a potential tool for portfolio-scale screening and for the identification of structures warranting priority inspection (Prota et al., 2025). The Morandi case also constitutes the most prominent reference point for the epistemological debate on the retrospective use of InSAR data in infrastructure forensic analysis. Milillo et al. (2019) reported pre-collapse displacement anomalies on the Morandi tower using COSMO-SkyMed data, while Lanari et al. (2020) subsequently disputed the statistical significance of those findings. The exchange, discussed in Section 4.3.4, frames the interpretation of pre-collapse signatures in both *Himera* and *Albiano-Magra*: the spatial correlation between an anomaly and a subsequent failure is a necessary but not sufficient condition for attributing the anomaly to the causal mechanism of the collapse. The two bridges, before and after their respective collapses, are shown in Figure 5. 8..

The structure of the case study is as follows. Section 5.3.1 presents the *Himera Viaduct* analysis, combining EGMS post-reconstruction data with user-driven Sentinel-1 processing of the pre- and post-collapse window. Section 5.3.2 presents the *Albiano-Magra Bridge* analysis, based on COSMO-SkyMed X-band data processed in user-driven mode across both the pre- and post-collapse phases. Section 5.3.3 consolidates the TS observations and the pre- and post-event evidence specific to the two case studies, while the cross-cutting methodological discussion of retrospective inference and the early-warning question is deferred to the synthesis of Section 5.5.




Figure 5. 8. The two infrastructure case studies of Section 5.3: (a) *Himera Viaduct*, collapsed deck (April 2015); (b) *Himera Viaduct*, rebuilt deck; (c) *Albiano-Magra Bridge*, historical view; (d) *Albiano-Magra Bridge*, collapse (April 2020).

5.3.1. Himera Viaduct: EGMS and user-driven Sentinel-1 analysis

The *Himera Viaduct*, located along the A19 motorway in central Sicily, represents one of the most instructive collapses in recent Italian infrastructure history. On 10 April 2015, a section of the original reinforced concrete deck failed following progressive instability of the Rocca di Sciara slope, which had been undergoing slow but measurable displacement for several years prior to the event. The collapse claimed no fatalities owing to the prior closure of the carriageway, yet the structural and geotechnical sequence leading to failure offers an unusually well-documented case for retrospective InSAR analysis. Two questions frame the investigation reported here: whether satellite-derived deformation data, processed at the spatial resolution available from Sentinel-1 C-band imagery, could have detected diagnostic signals before the collapse; and how the post-reconstruction EGMS product compares with a dedicated user-driven analysis in terms of both coverage and interpretability.

The user-driven processing was performed on two independent Sentinel-1 stacks. The ascending stack comprised 45 SLC scenes acquired along path 117 between 13 October 2014 and 29 April 2016, with a master image dated 2 September 2015. The descending stack comprised 39 scenes acquired along path 124 between 14 October 2014 and 30 April 2016, with a master dated 17 July 2015. Both stacks were co-registered using a star graph network tied to the respective master, with orbital corrections applied using precise state vectors. Topographic phase contributions were removed using the SRTM 90 m DEM. PS candidates were identified by applying an amplitude dispersion threshold consistent with the sparse built environment of the viaduct. The temporal coherence threshold for the final PS estimation was set at 0.45, reflecting the short temporal baseline of the stacks and the expectation of non-stationary scattering behaviour associated with the evolving structural condition. The interferometric network and the reflectivity field of the user-driven Sentinel-1 stacks in both orbital geometries are shown in Figure 5. 9.




Figure 5. 9. User-driven Sentinel-1 processing of the *Himera Viaduct* over the period October 2014 to April 2016. Panels (a) and (b) report the star graphs of the interferometric networks in the ascending and descending orbits respectively, with each spoke representing

one interferometric pair connecting a slave acquisition to the master, and with the colour of each spoke encoding the temporal coherence of the corresponding pair. Panels (c) and (d) report the reflectivity maps of the small-area processing window in the ascending and descending orbits respectively, with the position of the GCP selected to anchor the LOS solution highlighted in each map. The dual-orbit configuration provides the geometric diversity required to distinguish horizontal from vertical components of the displacement field..

In the ascending orbit, we identified a limited set of PS on the viaduct piers and abutments, consistent with the concrete and metallic surfaces of the standing portions of the structure. Mean LOS velocities across these scatterers were close to zero, confirming that the surviving structural elements were not undergoing broad settlement during the observation window. The ascending stack also revealed one TS, whose temporal behaviour merits detailed discussion. The TS was located in the area immediately adjacent to the bridge, geographically consistent with the foot of the Rocca di Sciara slope. The point appeared in the displacement time series only after the landslide event of April 2015, and its signal persisted until the onset of the demolition activities, at which point the associated coherence fell below the threshold and the TS ceased. The temporal behaviour of this scatterer is consistent with the hypothesis that it represents a temporary surface modification introduced by the landslide itself, most probably the deposition of displaced material that offered a coherent radar return for a finite period before being subsequently cleared from the site. No event-related TS was identified on the viaduct deck itself, in either the ascending or the descending orbit. This absence is not a null finding to be discounted: it is a substantive geometric observation that reflects the position of the viaduct within the *Himera* river valley. The ascending pass, with right-looking illumination from the east-northeast, suffers a symmetric occlusion from the eastern flank. Under these conditions, the deck surfaces that would otherwise host the high-amplitude scatterers associated with collapse-related TS are not consistently illuminated, and the scatterer population on the deck is correspondingly sparse. A second factor compounds the geometric limitation. The IW acquisition mode of Sentinel-1 has a ground sampling of approximately 5 m in range and 20 m in azimuth, which means that each resolution cell on the deck integrates the radar response of a sizeable extent of the structural surface and of any adjacent ground or vegetation that falls within the same cell. At this resolution, the probability that a single pixel on the bridge deck is dominated by a coherent scatterer suitable for PS or TS extraction is intrinsically low, particularly for a narrow elongated structure of the order of 10 m in transverse dimension. The combination of the valley occlusion and the coarse resolution explains why the user-driven analysis recovers

only one event-related TS and only on the slope-adjacent area where the line of sight is unobstructed and the scattering surface (the post-landslide debris field) is sufficiently large to fill several resolution cells with a coherent return. The amplitude-temporal baseline and the displacement time series plots of the slope-adjacent TS, with the compact high-amplitude cluster opening after the landslide event and closing at the start of the demolition phase, are shown in Figure 5. 10.





Figure 5. 10. Slope-adjacent temporary scatterer identified on the *Himera Viaduct* in the ascending Sentinel-1 stack, interpreted as transient debris subsequently cleared from the site: (a) amplitude-temporal baseline plot; (b) displacement time series. In panel (a), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panel (b), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.

The interpretive significance of these findings must be situated within a broader body of evidence. Macchiarulo, Milillo, et al. (2022) conducted a retrospective analysis of the *Himera Viaduct* using COSMO-SkyMed X-band data acquired between 2011 and 2014, a period preceding the Sentinel-1 archive used here. That study identified a group of anomalous deformation points in the Rocca di Sciara area and detected measurable displacement activity on the slope in the years prior to the collapse, providing direct evidence that the geotechnical instability was

developing well before the structural failure. The pre-collapse COSMO-SkyMed PS velocity field over the *Himera Viaduct* and the adjacent Rocca di Sciarra slope, with the spatial distribution of the landslide system documented in the regional geomorphological inventory, is shown in Figure 5. 11. The displacement time series of the two outliers, one ascending and one descending, are shown in Figure 5. 12, where the dates from which the relative deformation exceeds the two-sigma threshold are marked.



(a)




Figure 5. 11. COSMO-SkyMed PSs over the *Himera Viaduct* and the adjacent Rocca di Sciara slope, processed in user-driven mode for the pre-collapse period from 2011 to 2014, colour-coded by mean LOS velocity. Panel (a) reports the ascending orbit, panel (b) the descending orbit. The dual-orbit comparison enables the qualitative attribution of velocity signatures to vertical or horizontal motion components on the structure and on the surrounding terrain..




Figure 5. 12. Relative displacement time series of two outliers on the *Himera Viaduct* identified by COSMO-SkyMed processing of the 2011-2014 archive: (a) ascending orbit,

relative LOS velocity 6.11 mm/year, two-sigma exceedance from mid-2013; (b) descending orbit, relative LOS velocity 5.68 mm/year, two-sigma exceedance from early 2014. Adapted from (Macchiarulo, Milillo, et al., 2022).

The Sentinel-1 analysis reported here does not extend to that earlier period; the available archive begins in October 2014, only six months before the collapse, which substantially limits the potential for detecting gradual precursory trends. The temporal gap between the two datasets underlines a general constraint of Sentinel-1-based PS analysis applied to infrastructure events that predate the operational phase of the Copernicus programme. The absence of a velocity anomaly in the ascending PS map should therefore not be interpreted as evidence of absence of precursory activity; it reflects, rather, the short observation window and the inherent difficulty of resolving slow deformation signals from a six-month pre-event stack. The observability framework developed by Cernuto et al. (2026) for landslide-induced bridge deformation provides a useful reference for assessing the geometric conditions under which slope-induced displacements would be detectable from a given orbit geometry, and the *Himera* case aligns with the scenario class characterised by predominantly slope-parallel motion with limited LOS projection.

The EGMS Calibrated product provides a temporally distinct and complementary perspective. Covering the period from 2015 to 2021 (extended to 2023 in later releases), we interrogated the EGMS dataset over the reconstructed viaduct, which was opened to traffic in 2020 following a major engineering intervention that replaced the failed span. New PS emerged on the reconstructed deck segments, with coherence values consistently above 0.5 and mean displacement rates in the low millimetre per year range. These values are not necessarily indicative of structural anomaly: moderate rates of this order are expected on newly constructed reinforced concrete elements during the post-construction consolidation phase and may also reflect residual environmental moisture effects on radar backscatter. The presence of interpretable PS on the reconstructed sections demonstrates that the EGMS product is capable of monitoring newly built infrastructure within its standard processing cycle, a capability of direct relevance to post-intervention acceptance monitoring. The spatial distribution of the EGMS Calibrated measurement points on the rebuilt deck and the displacement time series of two representative points are shown in Figure 5. 13. The three InSAR datasets that have been applied to the *Himera Viaduct*, namely the COSMO-SkyMed pre-event analysis of Macchiarulo, Milillo, et al. (2022), the user-driven Sentinel-1 analysis of the peri-collapse window reported in

this section, and the EGMS post-reconstruction record, are summarised in Table 5.3.




Figure 5.13. EGMS Calibrated measurement points on the rebuilt deck of the *Himera Viaduct* (2018-2023): spatial distribution of the points on the reconstructed segments, with a panel in the lower left comparing the displacement time series of two representative measurement points.

Table 5.3. Comparative summary of the three InSAR datasets applied to the *Himera Viaduct*.

Field	COSMO-SkyMed (Macchiarulo et al., 2022)	User-driven Sentinel-1 (this thesis)	EGMS Calibrated (Copernicus)
Sensor and band	COSMO-SkyMed, X-band (~3.1 cm)	Sentinel-1A, C-band (5.55 cm)	Sentinel-1, C-band (5.55 cm)
Source and processing	Author-controlled processing of CSK SLC	Author-controlled processing of S-1 SLC, SARPROZ	EGMS pre-processed product
Temporal coverage	2011-2014 (pre-collapse)	October 2014 - April 2016 (peri-collapse)	January 2018 - December 2023 (post-reconstruction)
Number of acquisitions	n.r. for the <i>Himera</i> segment	45 ascending; 39 descending	n/a (pre-processed product)

Field	COSMO-SkyMed (Macchiarulo et al., 2022)	User-driven Sentinel-1 (this thesis)	EGMS Calibrated (Copernicus)
Orbit geometries	ascending and descending	ascending and descending	ascending and descending
Ground sampling (range $\times$ azimuth)	approximately 3 m $\times$ 3 m	approximately 4 m $\times$ 14 m	n/a (geocoded product)
Master image date	n.r. in source	02 September 2015 (asc); 17 July 2015 (desc)	n/a
Selection criterion	ASI > 0.6 (general) and ASI > 0.2 within infrastructure buffer	amplitude threshold 1.5/3 and coherence > 0.45	EGMS default per-subswath selection
PS on the viaduct	966 ascending; 1346 descending	two TS in ascending; none in descending	new PS on the rebuilt deck
Key findings	pre-collapse PS velocities up to ~ 6 mm/year on the Rocca di Sciarra slope; two outliers on the viaduct (one per orbit) with relative LOS velocity 6.11 and 5.68 mm/year, exceeding the 2σ threshold from mid-2013 (asc) and early 2014 (desc)	a TS in ascending, near the bridge on transient debris cleared after the landslide; no event-related TS in descending orbit	new PS on the rebuilt deck, coherence > 0.5, velocities 0.2 to 5.5 mm/year

Taken together, the Sentinel-1 and EGMS analyses of the *Himera Viaduct* illustrate two distinct phases of a structural lifecycle in which satellite InSAR can contribute to monitoring. In the pre- and peri-collapse phase, the value of a dedicated user-driven analysis lies in the identification of TS that encode the chronology of physical events at and around the structure. The detection of velocity anomalies is precluded by the short archive window. The visibility of such TS is conditioned by two factors: the radar geometry of the structure within its surrounding terrain, and the spatial resolution of the satellite sensor relative to the structural dimensions. The *Himera* case demonstrates this conditioning in a particularly direct way: the valley topography occludes the deck from both ascending and descending Sentinel-1 LOS directions, and the coarse Sentinel-1 IW

resolution further limits the recovery of coherent scatterers on the narrow deck of the viaduct. The *Albiano-Magra* case, analysed with COSMO-SkyMed Stripmap data at approximately 3 m resolution, demonstrates the converse situation, in which the higher-resolution X-band sensor recovers a dense scatterer population on the bridge deck and on the surrounding terrain that the C-band Sentinel-1 archive could not have produced. In the post-reconstruction phase, the EGMS product offers a systematic and low-cost means of verifying the deformation behaviour of new structural elements, where the visibility constraints are partially relaxed by the elevation of the rebuilt deck above the original collapsed segment. The transition between these two monitoring modes highlights a gap that is, at present, difficult to close with either tool alone: the period immediately surrounding a structural failure, when archive data are limited and the structure is being progressively modified, represents the most challenging observational window for satellite InSAR. The implications of this limitation for early-warning applications, and the role of structure-terrain geometric conditioning in determining the achievable diagnostic level, are addressed in the cross-cutting discussion of Section 5.5.

5.3.2. Albiano-Magra Bridge: high-resolution COSMO-SkyMed interpretation

The *Albiano-Magra Bridge*, also known as the *Caprigliola Bridge*, crossed the Magra River between the municipalities of Caprigliola on the east bank and Albiano Magra on the west bank, in north-western Tuscany. Approximately 260 m long and 7 m wide, the structure consisted of five slender reinforced concrete arch spans resting on four river piers and two abutments. Originally completed in 1908 and subsequently reconstructed following damage sustained during the Second World War, the bridge carried a single carriageway along state road SS330 and served as a primary connectivity link between the two banks for the surrounding communities. On 8 April 2020, during the COVID-19 national lockdown, the bridge collapsed abruptly into the Magra River. The entire deck and most of the arch spans failed, while portions of the substructure remained partially standing. The collapse caused no fatalities and only minor injuries, owing to the reduced traffic at the time, yet the structural failure attracted immediate scientific attention as a case in which a multi-year satellite record was available for retrospective analysis.

The Albiano-Magra Bridge has been the subject of a prior in-depth InSAR analysis by Farneti et al. (2023a), which reconstructed two-dimensional deck displacements with explicit uncertainty quantification in a pre-collapse retrospective study. The

present analysis is methodologically complementary rather than competitive: where Farneti et al. (2023a) established quantitative pre-collapse kinematics with propagated error budgets, the TS lifecycle approach developed in Section 4.1.2 contributes a temporal observability axis through the appearance and disappearance epochs T_{on} and T_{off} of intermittent scatterers, encoding the chronology of structural events at the resolution of the satellite revisit cycle. The geotechnical context, in particular the role of slope-foundation interaction at the Caprigliola abutment, is consistent with the landslide-bridge interaction framework recently formalised by Cernuto et al. (2025).

The satellite dataset used for the present analysis consists of COSMO-SkyMed first-generation Stripmap acquisitions in HH polarisation. Two independent stacks were assembled: 173 acquisitions in ascending geometry and 94 acquisitions in descending geometry, together covering the period from 2011 to 2022. The full temporal span of each stack encompasses both the pre-collapse and post-collapse phases within a single, consistent processing chain, which constitutes a methodological choice of consequence. Rather than treating the two phases as separate datasets, we processed the entire archive as a unified stack allows the displacement and amplitude statistics to be estimated across the collapse boundary without introducing artificial discontinuities from separate reference frames or independent velocity inversions. The distinction between pre- and post-collapse behaviour is consequently carried out a posteriori, by interpreting the reconstructed time series against the known failure date.

Master image selection, co-registration, topographic phase removal, and sparse candidate identification followed the workflow of Chapter 3, with master dates of 18 July 2017 for the ascending geometry and 19 July 2018 for the descending geometry, and topographic phase removal based on the Copernicus DEM. A coherence threshold of 0.50 was applied as the minimum quality criterion for retaining measurement points in the final solution. This threshold is consistent with established practice in multi-temporal PS-InSAR analyses over urban and infrastructure environments, where values in the range 0.40-0.60 are commonly adopted to balance measurement density against phase reliability (Crosetto et al., 2016). Geocoding was performed using the Google Earth reference system, and a GCP was selected in an area of presumed stable backscattering in each orbit to anchor the corresponding displacement field. The interferometric networks and the reflectivity fields of the user-driven COSMO-SkyMed stacks of the *Albiano-Magra Bridge* in both orbital geometries are shown in Figure 5. 14.

The use of X-band imagery at Stripmap resolution offers a material advantage over C-band Sentinel-1 data for bridge-scale structural analysis. The shorter wavelength (approximately 3.1 cm for COSMO-SkyMed, compared with 5.55 cm for Sentinel-1) provides approximately twice the displacement sensitivity per unit of phase change, as discussed in Section 4.4.1. The finer spatial resolution of Stripmap mode (approximately 3 m in both range and azimuth, compared with 5 m by 20 m for Sentinel-1 IW) increases the density of coherent scatterers on metallic and concrete surfaces because the smaller resolution cell resolves individual structural elements such as parapets, deck joints, and pier capitals that remain unresolved at C-band. These characteristics, which have been exploited in multiple infrastructure monitoring studies, make X-band Stripmap data the most suitable satellite source currently available for detailed pre-collapse kinematic characterisation of bridge structures.




Figure 5. 14. User-driven COSMO-SkyMed processing of the *Albiano-Magra Bridge* (2011-2022): (a, b) coherence-coloured star graphs of the interferometric networks in ascending (a) and descending (b) orbits; (c, d) reflectivity maps of the Magra River crossing with the Ground Control Point indicated, ascending (c) and descending (d).

The pre-collapse portion of the time series, covering the period from 2011 to April 2020, reveals a pattern of predominantly low mean LOS velocities across the standing structural elements of the bridge, consistent with a structure that was not experiencing broad or uniform deformation throughout the monitoring period. A subset of scatterers located on the spans that subsequently failed exhibits localised positive LOS trends of modest but persistent character. Average velocities for these points remain low, yet the cumulative LOS displacements reach values of the order of 10-15 mm by the time of failure. A second cluster of scatterers exhibiting more pronounced pre-collapse displacement trends is located on the Caprigliola side of the bridge and on the Caprigliola hill above the eastern abutment. On the bridge end facing the hill, a small population of PSs shows mean LOS velocities of approximately +1 to +2 mm/year in the ascending orbit and approximately 0 to -1 mm/year in the descending orbit. Beyond the bridge, on the slope of the Caprigliola hill itself, a more clearly defined cluster of PSs exhibits mean LOS velocities of approximately +3 to +4 mm/year in the ascending orbit and approximately -5 to -6 mm/year in the descending orbit. The opposite-sign behaviour of the LOS velocities of the Caprigliola cluster across the two orbits is diagnostically significant: vertical motion produces same-sign LOS values in both ascending and descending acquisitions, while horizontal motion produces opposite-sign values. The observed signature is therefore consistent with a predominantly horizontal displacement field, with the magnitudes of the two orbital projections suggesting that the field has a substantial east-west component. The qualitative interpretation of the cluster as slope-driven motion of the Caprigliola hill in the years preceding the bridge collapse is supported by the spatial coherence of the cluster and by its persistent presence across the multi-year archive. A quantitative decomposition into vertical and east-west components, which would require a joint inversion of the two LOS time series, is not pursued here; the decomposition is in principle feasible from the data of Figure 5. 15 and could be undertaken in subsequent work. The spatial distribution of the PSs in the bridge area and on the Caprigliola hill, colour-coded by mean LOS velocity for the pre-collapse period, is shown in Figure 5. 15.



(a)



(b)

Figure 5. 15. Pre-collapse PSs over the *Albiano-Magra Bridge* and the surrounding terrain, processed from the COSMO-SkyMed Stripmap archive over the period 2011 to April 2020 in user-driven mode and colour-coded by mean LOS velocity. Panel (a) reports

the ascending orbit, panel (b) the descending orbit. The Caprigliola hill on the eastern side of the bridge is included within the spatial extent of both panels, allowing the joint observation of bridge-scale displacement signatures and slope-scale ground motion. The pre-collapse Caprigliola cluster discussed in Section 5.3.2 is visible in both panels with opposite-sign LOS velocities consistent with a predominantly horizontal displacement field..

The spatial correlation between two distinct anomalous populations, the failed spans on the bridge itself and the slope-bound cluster on the Caprigliola hill, suggests that the pre-collapse kinematic signature reflects a combination of structural changes in specific bridge components and geotechnical activity in the surrounding terrain, the latter potentially loading the eastern abutment through slope-induced foundation movement. Farneti et al. (2023a) reached a broadly consistent interpretation from an independent COSMO-SkyMed analysis of the same bridge, decomposing LOS measurements into vertical and longitudinal components in a bridge-aligned reference frame and detecting displacement anomalies at the left abutment in the final period before failure. Scattarreggia et al. (2022) complemented this observational evidence by integrating pre-collapse deformation patterns with non-linear finite element collapse simulations, exploring alternative failure mechanisms and the sensitivity of the structural response to different boundary condition assumptions. Taken together, these independent lines of evidence reinforce the interpretation that the bridge was exhibiting kinematic symptoms of structural change in the years preceding the collapse, with a possible geotechnical contribution from the surrounding slope, even if the causal relationship between the observed displacements and the failure mechanism cannot be established from InSAR data alone.

The post-collapse portion of the time series reveals a qualitatively different scattering regime. On-deck scatterers that had maintained coherent phase histories throughout the pre-collapse period disappear abruptly from the sparse-point solution following the failure date, as the structural elements hosting them were either removed by the collapse or heavily damaged, eliminating the stable reflecting surfaces that had sustained the interferometric signal. New scatterers emerge on the debris field and on temporary works introduced during the demolition and reconstruction phases. These points exhibit the characteristic amplitude-temporal baseline signature of TS, with compact high-amplitude clusters at short temporal baselines and rapid amplitude decay as the temporal separation increases beyond the lifetime of the temporary reflecting configuration. The contrast between the pre-collapse stability of long-lived PS on the still-standing piers and the bounded-

coherence behaviour of TS on the failed spans is documented in Figure 5. 16 and Figure 5. 17, where the amplitude-temporal baseline plots and the displacement time series of a representative PS on a still-standing pier and a representative TS on a failed span are presented for both orbital geometries.




Figure 5. 16. Amplitude-temporal baseline plots and displacement time series of a representative PS on a still-standing pier and a representative TS on a failed span of the *Albiano-Magra Bridge* in the ascending orbit (COSMO-SkyMed, 2011-2022): (a) PS

amplitude-temporal baseline plot; (b) PS displacement time series; (c) TS amplitude-temporal baseline plot; (d) TS displacement time series. In panels (a) and (c), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panels (b) and (d), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.




Figure 5. 17. Amplitude-temporal baseline plots and displacement time series of a representative PS on a still-standing pier and a representative TS on a failed span of the


Albiano-Magra Bridge in the descending orbit (COSMO-SkyMed, 2011-2022): (a) PS amplitude-temporal baseline plot; (b) PS displacement time series; (c) TS amplitude-temporal baseline plot; (d) TS displacement time series. In panels (a) and (c), black dots represent individual amplitude measurements, while red circles indicate the mean value. In panels (b) and (d), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.

The PS on the still-standing pier exhibits a stable amplitude across the full observation window and a displacement time series with no significant trend over the 2011-2022 period, a pattern that holds in both the ascending orbit shown in Figure 5. 16 and the descending orbit shown in Figure 5. 17. The TS on the failed span, by contrast, exhibits in both orbits a compact high-amplitude cluster confined to the pre-collapse window, with rapid amplitude decay after the April 2020 failure date, and a displacement time series that records a localised pre-collapse trend of the order of 10-15 mm of cumulative LOS displacement before the abrupt loss of phase stability at the collapse. The two-orbit confirmation of the TS lifecycle bounding strengthens the interpretation of the lifecycle as the interferometric signature of the structural event, since a processing artefact would not be expected to produce a coherent two-orbit signature with the same temporal closure date. The comparison between the two scatterer types provides direct visual evidence of the divergence in kinematic and coherence behaviour between the surviving and the failed structural elements.

The temporal distribution of TS death dates across the bridge area reveals a coherent spatial pattern that supports the interpretation of the TS lifecycle as an interferometric signature of the structural event. Figure 5. 18 displays the persistent and temporary scatterers retained in both orbital geometries on a Google Earth substrate, colour-coded by the date at which each scatterer ceased to satisfy the coherence criterion. The TS associated with the collapse cluster on the failed deck spans, with death dates concentrated around 8 April 2020, while the persistent scatterers on the still-standing piers and on the Caprigliola hill retain coherence beyond the collapse date and through the end of the observation window in 2022. This spatial-temporal pattern reinforces the diagnostic value of the TS lifecycle as a structural-event marker and visualises the spatial extent of the structural change in the time domain.



(a)



(b)

Figure 5. 18. Spatial distribution of persistent and temporary scatterers identified by user-driven COSMO-SkyMed processing on the *Albiano-Magra Bridge* and on the *Caprigliola hill* (2011-2022) on a Google Earth substrate, colour-coded by death date in both orbital geometries: (a) ascending orbit; (b) descending orbit.

The identification of TS on the collapsed structure carries methodological significance beyond the *Albiano-Magra* case. TS whose life span is demonstrably bounded by the collapse chronology constitute an interferometric signature of the structural event in the time domain. Their presence in the sparse-point solution is not a processing artefact to be discarded; rather, it encodes information about the physical sequence of events. The death date of a TS on a structural element provides a lower bound on when that element ceased to function as a stable radar reflector, which, in the case of collapse scenarios, is directly related to the timing of structural failure or subsequent demolition. The pre-collapse portion of a TS time series, where it exists with sufficient length, can contribute additional localised kinematic evidence on the deforming element. The post-collapse dispersion of the time series, conversely, reflects the transition from a coherent structural surface to a chaotic scattering environment, and the abruptness of this transition in the amplitude-temporal baseline domain offers a means of approximately dating the event from the satellite record alone.

The interpretation of the pre-collapse anomalies must, however, be situated within the epistemological framework established in Section 4.3.4 and in the introduction to this case study. The observed cumulative displacements of the order of 10-15 mm on the bridge spans, while spatially correlated with the failed portions, remain modest relative to the scatter introduced by atmospheric residuals, reference-frame uncertainties, and the seasonal component of the thermal response of concrete structures. The Caprigliola cluster, by contrast, exhibits LOS velocities of 3 to 6 mm/year, three to four times higher than the rate implied by the on-deck cumulative displacement, and a two-orbit signature consistent with a horizontal motion field, providing a geotechnical signal whose statistical significance is correspondingly higher than the structural one. As Farneti et al. (2023a) noted in their residual service-life analysis of the same bridge, the reliability of deformation-based prognostic estimates depends critically on the length of the available monitoring record and on the assumptions embedded in the deterioration model. The present analysis does not claim to demonstrate that the pre-collapse InSAR signal on the bridge itself constitutes a robust early-warning indicator; such a claim would require prospective validation against structures that did not subsequently collapse, and it would demand explicit uncertainty quantification that is beyond the scope of a retrospective forensic analysis. The observed anomalies are better regarded as kinematic symptoms that could, in an operational SHM context, have prompted targeted inspection of the affected spans and of the surrounding slope, and provided spatial prioritisation for further investigation including geotechnical reconnaissance of the eastern abutment area. The question of whether satellite

InSAR can contribute to operational early warning for bridge infrastructure, and the role of joint structural-geotechnical interpretation in that context, is addressed in the cross-cutting discussion of Section 5.5.

5.3.3. Temporary scatterers, pre-/post-event analysis, and early-warning issues

The user-driven processing of both the *Himera* and *Albiano-Magra* datasets identified measurement points whose phase stability is bounded by a finite temporal window rather than sustained across the full observation period. These TS, introduced in Section 4.1.2 as a diagnostic category, provide two forms of evidence that complement the displacement time series of long-lived PS.

The first form of evidence is chronological. The death date of a TS, defined as the acquisition epoch at which coherence falls irreversibly below the acceptance threshold, can be associated with the physical disappearance or modification of the scattering surface when combined with knowledge of the event timeline. In the *Himera* case, the user-driven analysis identified one event-related TS in the ascending orbit, located at the foot of the *Rocca di Sciara* slope and interpreted as transient debris cleared from the site after the landslide. The TS death date coincides with the onset of the demolition phase, providing an interferometric timestamp for the post-event clearance operations. No event-related TS was identified on the viaduct deck itself, a finding attributable to the valley geometry of the site that occludes the deck from the line of sight of both ascending and descending Sentinel-1 acquisitions. In the *Albiano-Magra* case, TS located on the arch deck in both ascending and descending orbits exhibited death dates aligned with the April 2020 collapse, providing an interferometric timestamp for the structural event with a temporal resolution determined by the satellite revisit cycle.

The second form of evidence is kinematic. Where the coherence window of a TS is sufficiently long, the pre-event portion of its displacement time series contributes spatially localised information on the behaviour of the structural element before failure. In the *Albiano-Magra* analysis, TS on deck spans that subsequently collapsed exhibited cumulative LOS displacements of the order of 10-15 mm over the pre-collapse period, spatially correlated with the failed portion of the structure. When read alongside the time series of stable PS on adjacent piers and abutments, these transient records permit a more spatially disaggregated interpretation of the pre-failure deformation field than would be available from long-lived PS alone.

The identification of TS in standardised products such as EGMS is substantially more limited than in user-driven workflows. Standardised processing chains apply uniform coherence thresholds calibrated for the detection of long-lived, stable scatterers on slowly deforming terrain. TS with bounded coherence windows frequently fail these filters and are discarded before the output is generated. The forensic and diagnostic information encoded in collapse-related TS is therefore inaccessible from standardised products and can be recovered only through dedicated user-driven analysis. The principal pre-collapse and peri-collapse InSAR findings of the two case studies, together with the visibility limitations and the assessment of statistical robustness, are summarised in Table 5. 4.

The pre-collapse displacement anomalies observed in the Albiano-Magra case study, and the pre-collapse slope kinematics observed at *Himera* by Macchiarulo, Milillo, et al. (2022), raise questions about the statistical robustness of retrospective attribution and about the conditions under which satellite InSAR can contribute to operational early warning for infrastructure. These questions extend beyond the specifics of *Himera* and *Albiano-Magra* and are addressed, together with the evidence from the other case studies, in the cross-cutting synthesis of Section 5.5.

Table 5. 4. Cross-case summary of the pre-collapse and peri-collapse InSAR evidence for the two infrastructure case studies of Section 5.3.

Field	Himera Viaduct	Albiano-Magra Bridge
Bridge and location	<i>Himera Viaduct</i> , A19 motorway, central Sicily	<i>Albiano-Magra Bridge</i> , SS330, north-western Tuscany
Collapse date	10 April 2015 (partial collapse)	8 April 2020 (full collapse)
Trigger	landslide on the Rocca di Sciara slope	structural failure with possible geotechnical contribution
Satellite sensor and band	COSMO-SkyMed X-band (Macchiarulo et al., 2022); Sentinel-1 C-band (this thesis); EGMS C-band	COSMO-SkyMed X-band (this thesis)
Observation period	2011-2014 (CSK); October 2014 - April 2016 (S-1); 2018-2023 (EGMS)	2011-2022 (single unified archive)
Orbit geometries	ascending and descending	ascending and descending

Field	Himera Viaduct	Albiano-Magra Bridge
Pre-collapse anomaly on the structure	two CSK outliers on the deck (one per orbit), relative LOS velocities of 6.11 mm/year (asc) and 5.68 mm/year (desc), 2σ exceedance from mid-2013 (asc) and early 2014 (desc); no Sentinel-1 anomaly detectable due to short archive	PS subset on subsequently failed spans with cumulative LOS displacement of 10-15 mm by April 2020; localised positive LOS trends of modest character
Pre-collapse anomaly in the surrounding terrain	Rocca di Sciara slope: pre-collapse PS velocities up to ~6 mm/year (CSK 2011-2014); slope-adjacent TS appearing after April 2015 landslide and lasting until demolition (S-1)	Caprigliola hill: PS cluster with LOS velocities of +3 to +4 mm/year (asc) and -5 to -6 mm/year (desc); opposite-sign signature consistent with predominantly horizontal motion field
TS death-date alignment	S-1 ascending TS near the slope: death date coinciding with the onset of the demolition phase	TS on failed spans in both orbits: death dates aligned with the April 2020 collapse
Visibility limitations	valley topography occludes the deck from both ascending and descending Sentinel-1 LOS directions; no event-related TS recoverable on the deck itself	no equivalent terrain occlusion; CSK X-band density and ascending-descending coverage allow joint structural and geotechnical reading
Statistical robustness	structural anomaly close to detection margin (1-2 mm/year for C-band); slope anomaly more robust due to higher velocities and persistent multi-year signal	structural anomaly close to detection margin (modest cumulative displacement); Caprigliola cluster well above margin (3-6 mm/year in LOS magnitude, opposite signs across the two orbits consistent with a predominantly horizontal motion field, two-orbit signature statistically robust)

5.4. Case study III - The Sanctuary of Vicoforte: coherence-aware satellite monitoring and soil-structure interaction

The *Sanctuary of Vicoforte*, officially known as the Basilica Regina Montis Regalis, occupies a singular position in both the cultural heritage of north-western Italy and the history of structural monitoring research. Located in the municipality of *Vicoforte* in the Province of Cuneo, the building is best known for its oval masonry dome, whose internal axes of 37.15 m and 24.50 m make it the largest elliptical masonry dome in the world. The structure has been the subject of systematic structural investigation for more than four decades, and since 2004 it has hosted a permanent multi-sensor monitoring system that constitutes one of the most comprehensive long-term SHM datasets available for a masonry heritage building in Europe. The case study reported here draws on this exceptional monitoring context to investigate the contribution of satellite InSAR to an integrated sensing framework, and to develop a coherence-aware data-driven model that links satellite-derived displacement measurements to the environmental variables that govern soil-structure interaction at the site. An aerial view of the Sanctuary in its surroundings and an internal view of the elliptical dome, together with a plan and a section of the structure, are shown in Figure 5. 19 and Figure 5. 20, respectively.



Figure 5. 19. Geometric and structural overview of the *Sanctuary of Vicoforte*: (a) aerial view of the Sanctuary and its immediate surroundings; (b) Internal view of the oval masonry dome.




Figure 5. 20. Geometric and structural overview of the *Sanctuary of Vicoforte*: (a) plan of the structure, highlighting the oval dome footprint; (b) longitudinal section through the dome and the drum.

The selection of *Vicoforte* as a case study for satellite-assisted SHM reflects a deliberate methodological progression from the preceding case studies. Case study I demonstrated that satellite InSAR can provide interpretable displacement measurements on a compact urban masonry building in the absence of any in-situ reference, and assessed the relative merits of standardised EGMS products versus user-driven processing. Case study II demonstrated that satellite InSAR, when applied to infrastructure collapse scenarios, can contribute forensic evidence through TS analysis and can support post-reconstruction stability monitoring. Both cases involved deformation regimes in which the primary analytical challenge was one of data quality and geometric interpretation. *Vicoforte* presents a qualitatively different challenge. The structure is essentially stable in the sense that no gross settlement or progressive failure is anticipated, yet the displacement time series extracted from InSAR observations exhibit systematic non-stationarities that are not attributable to structural damage. These non-stationarities arise from the interaction between the building foundations and a heterogeneous subsoil whose mechanical response is modulated by seasonal variations in temperature and moisture content. Their correct separation from any genuine structural signal is a prerequisite for a reliable SHM interpretation, and this separation problem defines the methodological core of the present case study.

The soil-structure interaction context at *Vicoforte*, which underlies the analytical choices of the present case study, has been characterised through a long history of geotechnical and structural investigation. The interplay between the differential stiffness of the foundation soil and the environmental modulation of its mechanical response defines the displacement regime that the satellite observations are called upon to characterise. Section 5.4.1 develops the geotechnical and structural context in detail; the present paragraph anticipates only the two features that bear directly on the monitoring strategy. The first is the heterogeneous bearing conditions of the foundation, which produce a systematic asymmetry in the long-term settlement behaviour of the structure. The second is the seasonal modulation of the clayey-silt deposit by temperature and moisture variations, which generates a cyclic displacement component that must be separated from any deterioration trend before structural conclusions can be drawn (Ceravolo et al., 2021; Coccimiglio, Coletta, et al., 2022; Coletta et al., 2019).

The present case study is structured around four interconnected questions. First, what is the quality of the PS-InSAR signal obtainable from satellite data over a structure with the geometric and material characteristics of the *Vicoforte* basement, and how does temporal coherence vary with processing parameters and the physical nature of the reflecting surfaces? Second, how do environmental variables derived from satellite geophysical products or from low-cost in-situ sensors relate to the displacement behaviour of basement PS? Third, can a data-driven model be developed that uses environmental inputs to reproduce and predict the displacement behaviour of individual PS, enabling a continuous monitoring framework whose operation does not depend on the availability of dense in-situ sensor networks? Fourth, what are the implications of this integrated framework for the SHM of monumental buildings in which the deformation signal is dominated by environmental loading rather than by structural degradation?

The structure of the case study follows these four questions across Sections 5.4.1 to 5.4.6. Section 5.4.1 provides the structural and geotechnical context in detail. Section 5.4.2 describes the available in-situ and satellite datasets. Section 5.4.3 presents the PS-InSAR processing and the coherence-based quality assessment. Section 5.4.4 examines the relationship between environmental variables and the structural monitoring record. Section 5.4.5 develops the coherence-aware data-driven displacement model through a SBL framework. Section 5.4.6 draws the implications for integrated SHM practice.

5.4.1. Structural and geotechnical context

The *Sanctuary of Vicoforte* was commissioned at the close of the sixteenth century by Duke Carlo Emanuele I of Savoy, who intended the building to serve as the dynastic mausoleum of the House of Savoy. The project was assigned to the court architect Ascanio Vittozzi, who refined an earlier oval-plan design by Ercole Negri di Sanfront. Construction began in 1596 but was interrupted almost immediately by foundation problems attributable to the weak subsoil conditions, as well as by the successive deaths of the architect in 1615 and the duke shortly afterwards. Financial and political instability prolonged the abandonment for nearly a century, leaving the structure incomplete at the level of the dome impost. Work resumed in the early eighteenth century under the architect-engineer Francesco Gallo, who completed the high drum and the dome between approximately 1728 and 1732, with technical support from the royal architect Filippo Juvarra. To counteract the horizontal thrust generated by the dome geometry, Gallo installed three iron rings with a combined cross-sectional area of approximately 140 cm² within the masonry. The building was designated a National Monument in 1880 and has been the subject of scientific investigation for more than four decades.

The structural vulnerabilities of the Sanctuary were documented as early as the mid-twentieth century by engineer Martino Garro, who identified extensive cracking concentrated at the drum openings and attributed the damage to differential settlements of the eight foundation buttresses, with recorded displacements reaching up to 30 cm between adjacent pillars. Geotechnical investigations conducted between 1976 and 2008 established the principal features of the subsoil stratigraphy. The stratigraphic succession reconstructed through continuous core drilling and in-situ tests reveals lateral homogeneity with limited variations in thickness across the site. A superficial horizon of fill material and topsoil, generally less than 1.0 to 1.5 m thick, rests above a level of clayey silts and silty clays of grey-brown colour, locally with fine sandy intercalations. These soils display medium-to-stiff consistency, medium plasticity, and water content close to or slightly above the plastic limit, with moderate undrained cohesion and low effective friction angles indicative of predominantly cohesive behaviour. Laboratory tests highlight medium-to-high compressibility, with oedometer moduli compatible with weakly to normally consolidated fine soils, and low permeability that influences the dissipation times of excess pore pressures under rapidly applied loads. At greater depths, more compact layers of marls and marly clays (sometimes weakly cemented) constitute the geotechnical reference substrate, characterised by

increased shear strength and reduced deformability, with modest effective cohesion values but higher friction angles than the overlying layers (Chiorino et al., 2008).

The more stable marl layer dips in a south-westerly direction, so that the overlying clayey-silt deposit is approximately 5 m deep on the south-western side and reduces to a few centimetres on the eastern side. As a consequence, five of the eight foundation buttresses rest on the compressible clayey-silt deposit, while the remaining three rest directly on the marlstone. The differential stiffness between these two foundation conditions generates a systematic asymmetry in the settlement behaviour of the structure, and the consequent shear stresses in the masonry are responsible for the crack pattern that has been observed and measured since at least the nineteenth century (Garro, 1962). The variable thickness of the clayey-silt deposit and the differential bearing conditions of the eight foundation buttresses are shown in Figure 5. 21, panel (a), with the structural plan of the Sanctuary at foundation level and the buttress numbering shown in panel (b). A survey of the cracks on the longitudinal section technical drawing and a detail of the cracks on the openings are shown in Figure 5. 22.




Figure 5. 21. Schematic cross-section of the subsoil stratigraphy beneath the *Sanctuary of Vicoforte*, showing the variable thickness of the clayey-silt deposit (approximately 5 m on the south-western side, reducing to centimetres on the eastern side), the marlstone substratum, and the differential bearing conditions of the eight foundation buttresses (a). Structural plan of the *Sanctuary of Vicoforte* at foundation level, the eight foundation buttresses are labeled (b).




Figure 5.22. Survey of the cracks on the longitudinal section (a) and detail of the cracks on the openings (b).

The structural response to this geotechnical heterogeneity is not static. The clayey-silt deposit is susceptible to seasonal variations in moisture content and temperature, which modulate its compressibility and produce cyclic displacement components at the foundation level. These environmental effects propagate upward through the masonry and are detectable in both the static monitoring record and, as shown in the subsequent sections, in the PS-InSAR displacement time series. The post-tensioned tie-rod system installed in 1987 and re-tensioned in 1997 was designed to arrest the long-term outward displacement of the drum and limit the opening of the main cracks. The system consists of 56 Dywidag steel tie rods of diameter 32 mm and nominal yield strength approximately 1080 MPa, embedded in the masonry at the drum level and anchored at 14 perimeter points through steel frames. Load cells on each tie rod allow continuous monitoring of the tensile force, and the time series of these cells constitute one of the primary static monitoring outputs available for integration with satellite data (Figure 5.23).




Figure 5. 23. Section of the dome: original iron circular ring and modern post-tensioning ring on the drum (left) and detail of the post-tensioning ring system (right).

The structural and geotechnical history of the Sanctuary has been analysed in depth through a series of finite element models, from the non-linear anisotropic analysis of the dome-drum system by Calderini et al. (2006) to the thermoelastic FE model updating approach of Ceravolo et al. (2020), which demonstrated that model parameters calibrated against environmental inputs can reproduce the observed frequency variations with good fidelity. The three-dimensional geometry of the structure has been characterised through laser scanning, providing the spatial reference for both the FE models and the structural monitoring layouts (Aoki et al., 2003).

The dominant failure mode of concern at Vicoforte is the progressive widening of cracks at the drum and the slow outward displacement of the buttresses under the combined action of dome thrust and differential foundation settlement, rather than catastrophic collapse. Long-term monitoring has confirmed that this process is ongoing at very low rates, with annual displacement increments of the order of tenths of a millimetre, and that its seasonal component is strongly modulated by soil temperature and moisture. The separation of this environmental component from any potential structural degradation trend is the central diagnostic challenge of the monitoring programme, and it defines the problem that the satellite data integration approach of this case study is designed to address.

5.4.2. Available in-situ and satellite data

The permanent monitoring system of the *Sanctuary of Vicoforte* is managed by the Earthquake Engineering and Dynamics Laboratory of Politecnico di Torino under the scientific supervision of Professor Rosario Ceravolo. Since its installation in 2004, the system has undergone substantial evolution and now comprises 130 sensors of various types distributed throughout the structure, its foundations, and the surrounding terrain. The monitoring infrastructure is organised into three functional layers: a low-rate static and environmental system, a high-rate dynamic system, and, since mid-2024, a specialised monitoring system for the north-western bell tower.

The static and environmental system was originally installed in 2004 and remained operational until 2014. Following a renovation completed in 2022-2023, the current configuration includes 56 load cells on the post-tensioned tie rods, 12 crack meters monitoring the principal drum cracks, 2 wire and 4 laser distance meters for displacement and inclination measurement, 34 thermometers distributed within and around the masonry, 4 thermo-hygrometers, 1 barometer, 1 pyranometer on the lantern, 3 piezometers for groundwater monitoring, and 1 soil thermo-hygrometer. The replacement of cable-based distance meters with laser instruments in the 2022-2023 renovation improved the measurement resolution and eliminated errors introduced by cable sag and thermal elongation of the measurement wire. The dynamic monitoring system, installed in 2015, comprises 12 accelerometers that record structural vibrations either at fixed intervals (20 minutes per hour) or upon threshold triggering at a horizontal ground acceleration of 0.042 g. Natural frequencies of the structure have been tracked continuously since 2016 through operational modal analysis, providing a time series of modal parameters that reflects the combined influence of thermal loading, humidity, and structural condition (Pecorelli et al., 2020).

The soil monitoring sub-system deserves specific attention because it provides the physical variables that drive the data-driven displacement model developed in Section 5.4.5. Three piezometers installed in the south-western sector of the Sanctuary, where the clayey silt layer is thicker and more sensitive to hydro-mechanical variations, monitor the water table level. A thermo-hygrometer installed in the same sector records surface soil humidity and temperature at hourly sampling. The time series of these low-cost ground sensors exhibit variations consistent with seasonal cycles and significant meteorological events: maximum temperature points reached in the summer season correspond systematically to lower humidity

values, reflecting the standard inverse correlation between surface soil moisture and air temperature under quasi-stationary conditions. Although the piezometer record is direct and reliable, its high installation and management costs limit the replicability of groundwater monitoring in widespread SHM systems. The temporal behaviour of the groundwater level closely follows that of surface soil moisture measured by the thermo-hygrometer, a correspondence that supports the use of surface moisture alone (together with soil temperature) as the environmental input to the data-driven model, as further discussed in Sections 5.4.4 and 5.4.5. The temporal correspondence between the water-table level measured by the piezometers and the surface soil moisture measured by the soil thermo-hygrometer is shown in Figure 5. 24. The layout of the soil monitoring sub-system, with the position of the three piezometers and the soil thermo-hygrometer, is shown in Figure 5. 25. The native time series of surface soil moisture and temperature recorded by the soil thermo-hygrometer are shown in Figure 5. 26.




Figure 5. 24. Temporal correspondence between the water-table level measured by the piezometers and the surface soil moisture measured by the soil thermo-hygrometer at the *Sanctuary of Vicoforte*. The close agreement between the two signals supports the use of surface moisture as the hydrological input to the data-driven model.



Figure 5. 25. Layout of the soil monitoring sub-system at the *Sanctuary of Vicoforte*. The three piezometers PZ01, PZ02 and PZ03 monitor the groundwater level in the south-western sector of the Sanctuary, where the clayey-silt deposit reaches its maximum thickness. The soil thermo-hygrometer measures surface soil moisture and temperature at hourly sampling.




Figure 5. 26. Native time series of surface soil moisture and temperature recorded by the in-situ soil thermo-hygrometer at the *Sanctuary of Vicoforte* at one-hour sampling. The seasonal cycles and the inverse correlation between temperature peaks and moisture troughs are visible throughout the observation period.

The satellite datasets used in this case study are drawn from two sources. The first is the EGMS Calibrated product, which provides multi-geometry Sentinel-1 C-band displacement time series for the Sanctuary site at no cost and without dedicated processing. Both the 2018-2022 and 2019-2023 temporal windows are available, in descending and two ascending orbit geometries, yielding a combined total of approximately 350 Calibrated measurement points distributed over and around the structure. The Ortho product was considered but rejected for structural analysis: owing to its 100 m grid resampling, Ortho points consistently fall outside the building footprint, rendering them unsuitable for direct monitoring of the dome or drum. The Calibrated points, by contrast, are spatially concentrated on the dome surface and the bell towers, which are the structural elements of primary monitoring interest. The numerical comparison of the measurement point density between the two EGMS products is reported in Table 5. 5. The spatial distribution of the EGMS measurement points on and around the Sanctuary, in both Ortho and Calibrated, is shown in Figure 5. 27.




Figure 5. 27. Spatial distribution of EGMS measurement points on and around the *Sanctuary of Vicoforte*: (a) Ortho dataset; (b) Calibrated dataset. The labels D1st and D2nd identify the descending acquisitions of the 2018-2022 and 2019-2023 windows respectively; A11st and A12nd identify the first ascending track for the same windows; A21st and A22nd identify the second ascending track.

Table 5. 5. Numerical comparison of the EGMS Calibrated and Ortho measurement points on the *Sanctuary of Vicoforte*.

EGMS product	Number of MPs
Calibrated (total, all geometries and windows combined)	≈ 350
Ortho	3

The second satellite source is a stack of COSMO-SkyMed Second Generation SLC images in HH polarisation, processed in user-driven mode. The stack comprises 62 acquisitions (stack identifier STR-005) collected between 29 December 2021 and 14 November 2025 over a scene of approximately $3 \text{ km} \times 3 \text{ km}$ centred on the monument. Acquisitions were collected in Stripmap mode, ascending right-looking geometry, with a scene-centre incidence angle of 29.4° . The processed data have a ground pixel spacing of 2.18 m in range and 2.25 m in azimuth, and the effective revisit time over the site is approximately 23.2 days. The X-band wavelength of 0.0312 m provides approximately twice the displacement sensitivity per unit of phase change compared to Sentinel-1 C-band. All SLCs are

co-registered to a master acquisition dated 13 February 2024. The complementarity between these two satellite sources, EGMS for dome-level structural monitoring and COSMO-SkyMed for basement-level soil displacement estimation, defines the multi-source integration architecture of this case study.

The environmental satellite data used for preliminary correlation analysis in Section 5.4.4 are drawn from two Copernicus geophysical products. LST is derived from Sentinel-3 SLSTR thermal infrared acquisitions, interpolated onto a regular $11 \text{ km} \times 11 \text{ km}$ grid of 121 points centred on the Sanctuary, and serves as a proxy for soil temperature at the surface. The SWI is provided by the Copernicus Global Land Monitoring Service at 1 km resolution and is available for four characteristic time-length parameters representing moisture memory at progressively greater depths. The selection of appropriate parameters for the *Vicoforte* site, and the role of these satellite geophysical products relative to the in-situ measurements in the final data-driven model, are discussed in Section 5.4.4. The principal characteristics of the four satellite datasets used in this case study are summarised in Table 5. 6.

Table 5. 6. Summary of the four satellite datasets used in the *Vicoforte* case study.

Field	EGMS Calibrated	CSK SG SLC (this thesis)	Sentinel-3 SLSTR LST	CGLS SWI
Sensor and band	Sentinel-1, C-band	COSMO-SkyMed Second Generation, X-band	Sentinel-3 SLSTR, thermal infrared	ASCAT and SMOS, C- and L-band
Source / processing	EGMS Calibrated (pre-processed product)	Author-controlled processing of CSK-SG SLC, SARPROZ	Sentinel-3 SLSTR LST product (pre-processed)	Copernicus Global Land SWI (pre-processed)
Wavelength	5.55 cm	3.12 cm	thermal IR (10.85 μm)	5.7 cm and 21 cm
Temporal coverage	2018-2022 and 2019-2023	29 Dec 2021 - 14 Nov 2025	concurrent with SAR archives	concurrent with SAR archives
Number of acquisitions	n/a (pre-processed)	62	n/a	n/a

Field	EGMS Calibrated	CSK SG SLC (this thesis)	Sentinel-3 SLSTR LST	CGLS SWI
Orbit geometries	ascending ($\times 2$) and descending	ascending only	n/a	n/a
Spatial resolution	5×20 m (S-1 IW)	2.18×2.25 m (SG Stripmap)	11×11 km grid (121 points)	1 km global grid
Role in the analysis	dome and bell-tower displacement monitoring	soil-level displacement at the basement	soil temperature proxy (LST)	soil moisture proxy (SWI, four time-length parameters)

5.4.3. PS-InSAR processing and temporal coherence assessment

The PS-InSAR processing for this case study departs from the straightforward implementation applied in Case Studies I and II in one important respect: temporal coherence is treated not solely as a quality indicator for measurement point selection, but as an explicit variable within the modelling framework. The motivation for this methodological choice lies in a property of temporal coherence that is widely acknowledged but rarely exploited in operational practice. Temporal coherence is not a pure noise metric. It correlates with the physical nature of the scatterer and with the deformation process itself, so that coherence-based selection generates a spatially and physically biased sampling of the deformation field. High-coherence points concentrate on man-made structures, rock outcrops, and rigid surfaces, while softer soils, vegetated areas, and surfaces affected by seasonal moisture variations are systematically under-represented. At *Vicoforte*, the target of interest is not the structural deformation of the dome or the drum under mechanical loading, but the quasi-static displacement of the soil-structure interface, which is driven by seasonal environmental inputs and is therefore associated with a non-stationary deformation field at the very margin of the coherence-based sampling bias. Low-coherence PS may, in principle, carry useful information about rapid or anomalous deformation episodes, but their inclusion in the model calibration dataset introduces phase noise that biases the posterior estimates of soil displacement. The framework adopted here therefore propagates coherence values explicitly into the likelihood function of the data-driven model, allowing the model to weight high-coherence and low-coherence observations differentially and to


identify an application-specific coherence threshold that minimises posterior uncertainty.

The EGMS Calibrated product analysis over the Sanctuary site revealed a marked difference between the two temporal windows available. In the 2018-2022 descending geometry, a subset of measurement points on the dome surface exhibits coherence values consistently above 0.6, while the spatial distribution of high-coherence points concentrates on the masonry surfaces of the dome ribs and the bell tower facades, where the stone surfaces and metallic fixtures provide stable radar targets. The 2019-2023 window shows a broadly similar pattern, with a slight increase both in the number of retained measurement points and in the percentage of points above the coherence threshold, consistent with the improved multi-temporal statistics afforded by the longer observation window. The EGMS coherence threshold of approximately 0.5 retains a substantial majority of the points on the dome, while a smaller fraction of points on the vegetated slopes surrounding the building falls below threshold and is discarded. The temporal coherence maps for the EGMS descending and ascending orbits over the Sanctuary, with the 0.6 threshold indicated, are shown in Figure 5. 28 and Figure 5. 29 respectively. The corresponding quantitative summary, comprising the number of available measurement points and the percentages above coherence threshold for each combination of orbital track and observation period, is reported in Table 5. 7.


Table 5. 7. Quality characterisation of the EGMS Calibrated measurement points available on the *Sanctuary of Vicoforte* and its immediate surroundings, organised by orbital geometry and observation period. Three orbital tracks intersect the area, namely one descending and two ascending (Asc.1 and Asc.2), and two temporal releases of the product are considered, covering the periods 2018-2022 and 2019-2023 respectively. For each combination, the table reports the total number of available MPs, the percentage of MPs whose temporal coherence exceeds 0.6, and the percentage classified as type 0 (high-quality permanent scatterers) according to the EGMS quality flag scheme. The full set of MPs satisfies the type 0 condition, and over 90 percent of the MPs satisfies the 0.6 coherence threshold across all six combinations, indicating that the EGMS Calibrated dataset on the Sanctuary is dominated by stable, persistent scatterers throughout the available observation window.

Orbit	2018-2022			2019-2023		
	Available MP	TC > 0.6 (%)	MPtype = 0 (%)	Available MP	TC > 0.6 (%)	MPtype = 0 (%)
Descending	57	91.23	100	59	96.61	100

Orbit	2018-2022			2019-2023		
	Available MP	TC > 0.6 (%)	MPtype = 0 (%)	Available MP	TC > 0.6 (%)	MPtype = 0 (%)
Ascending (1)	88	93.18	100	84	97.62	100
Ascending (2)	73	93.15	100	72	98.61	100




(a)




(b)


Figure 5. 28. Temporal coherence analysis of the EGMS Calibrated measurement points over the *Sanctuary of Vicoforte* in the descending orbit, with the 0.6 coherence threshold indicated by the red line: (a) histogram of the temporal coherence values; (b) spatial distribution of the points retained above and below threshold.



(a)



(b)



(c)




Figure 5. 29. Temporal coherence analysis of the EGMS Calibrated measurement points over the *Sanctuary of Vicoforte* in the two ascending orbital tracks across the 2018-2022 and 2019-2023 windows, with the 0.6 coherence threshold indicated by the red line: (a) and (b) ascending track 1, with histogram and spatial distribution; (c) and (d) ascending track 2, with histogram and spatial distribution.

For the user-driven COSMO-SkyMed analysis at Vicoforte we employed the small-area mode of the multi-temporal workflow described in Chapter 3. We co-registered the 62-image Second Generation SLC stack to the master acquisition of 13 February 2024 and connected the acquisitions through a star-shaped interferometric graph centred on the master. We identified candidate scatterers from the mean-amplitude image using a local-maxima mask, retaining pixels with sufficiently stable amplitude as potential PS. A Copernicus DEM was used to remove the bulk topographic phase, and a GCP located on a reflective and stable surface at coordinates 44.361887° N, 7.862178° E anchored the displacement solution in a geodetic reference frame. Temporal coherence was estimated on the full stack using a Minimum Spanning Tree (MST) graph. MST is a graph-theory algorithm used to optimize interferometric network construction. It connects SAR images (nodes) with the highest possible coherence (edges) while minimizing total spatial/temporal baselines, resulting in a network that minimizes decorrelation, improves phase unwrapping accuracy, and reduces computation time by up to 90%. For each candidate PS, a preliminary height-velocity model was fitted and removed, and the residual phases were used to compute the complex temporal coherence as the average of unit phasors over the interferogram network. The interferometric configuration of the user-driven COSMO-SkyMed Second Generation stack, with the perpendicular baseline distribution, the star-shaped graph centred on the master, and the regional Copernicus DEM, is shown in Figure 5. 30.




Figure 5. 30. Interferometric configuration of the user-driven COSMO-SkyMed Second Generation stack over the *Sanctuary of Vicoforte*: (a) histogram of the distribution of normal baselines across the 62 acquisitions; (b) star-shaped interferometric graph centred on the master, with edges coloured according to the average spatial coherence; (c) Copernicus DEM of the regional setting, with the area of interest indicated by the black crosses.

The modulus of the complex coherence defines the temporal coherence, while the residual phase dispersion provides a complementary noise measure. These quantities were exported as standard PS attributes and used both to select high-quality PS within a small area centred on the structural basement and to define coherence-dependent weights in the subsequent modelling. The mean reflectivity and the spatial coherence of the full COSMO-SkyMed scene over the *Vicoforte* area, including the location of the GCP, are shown in Figure 5. 31.




Figure 5. 31. Mean reflectivity and spatial coherence of the user-driven COSMO-SkyMed Second Generation stack over the *Vicoforte* area: (a) mean amplitude image with the Ground Control Point indicated in red; (b) spatial coherence map.

A central methodological decision concerns the spatial extent of PS selection. For the soil displacement estimation objective, we restricted the PS candidate set to scatterers located at the base of the structural basement, operationally defined as points with geocoded heights between -1 m and +5 m relative to the local ground reference. The justification for this restriction is structural and finds quantitative support in the dynamic characterisation of the Sanctuary. Finite element analyses of the Sanctuary confirm that the basement behaves essentially as a rigid raft for the deformation modes of interest, so that displacement measurements at basement level can be attributed to soil motion rather than to structural deformation of the overlying masonry (Ceravolo et al., 2020). The dominant structural modes of the dome-drum system lie in a relatively low-frequency band, approximately 1 to 7 Hz, while simplified one-dimensional site-response analyses and in-situ tests indicate soil fundamental frequencies at higher values, approximately 9 to 16 Hz. The separation between these frequency ranges confirms that the structural basement and the underlying soil are mechanically coupled in the quasi-static regime relevant to InSAR observations, and that the basement can be treated as a rigid kinematic

interface between soil displacement and structural response. The finite element model of the Sanctuary, with the meshed geometry and the first three modes of vibration on which the rigid-basement assumption rests, is shown in Figure 5. 32.





Figure 5. 32. Finite element model of the *Sanctuary of Vicoforte*: (a) meshed geometry; (b) first mode of vibration; (c) second mode of vibration; (d) third mode of vibration. The model is implemented in Ansys®.

The basement itself is a massive structural element composed of stone and brick masonry, characterised by high stiffness and limited deformability. This behaviour is particularly evident below the elevation of 5 m, where the measured displacements can be reasonably attributed to soil deformation rather than to structural deformation of the building itself. The lower bound of -1 m accounts for residual localisation errors and height biases in the geocoding procedure (imperfect state vectors, DEM-related biases, reference-point-dependent height offsets),

ensuring that true basement-level scatterers slightly below the DEM-derived local ground reference are not erroneously discarded. The dome and drum, which undergo measurable thermoelastic deformation, are accordingly excluded from the soil-monitoring PS set.

For each selected basement PS, the inversion solves a parametric model comprising residual height, a linear LOS velocity term, and a thermoelastic component driven by air temperature. The thermoelastic term explicitly represents the seasonal and temperature-dependent displacement component, so that the residual time series after model subtraction reflects soil behaviour rather than structural thermal response. Temporal coherence and residual phase dispersion are computed for each PS and exported as standard attributes of the processing chain. These quantities are subsequently used in two ways: first, to define an initial quality filter that retains only PSs above a preliminary coherence threshold; second, as inputs to the coherence-aware likelihood function of the SBL surrogate model described in Section 5.4.5. The small-area outputs of the user-driven processing centred on the Sanctuary, including the local reflectivity, the temporal coherence, the mean LOS velocity, and the cumulative LOS displacement of the candidate PSs with their empirical distributions, are shown in Figure 5. 33.






Figure 5.33. User-driven small-area processing outputs for the *Sanctuary of Vicoforte*: (a) reflectivity map of the small area; (b) temporal coherence of the candidate PS, with empirical distribution; (c) mean LOS velocity, with empirical distribution; (d) cumulative

LOS displacement over the observation period, with empirical distribution. The wide spread of the velocity and displacement distributions visible in panels (c) and (d), before coherence-based filtering, motivates the threshold selection procedure of Section 5.4.5.

The identification of an application-specific coherence threshold for the *Vicoforte* basement PS is developed through the SBL uncertainty analysis of Section 5.4.5. The empirical threshold that emerges from this analysis is 0.56, a value determined by the abrupt change in the estimated model uncertainty as a function of coherence. Below this threshold, the uncertainty increases sharply, indicating that predictions become unreliable; above the threshold, predictions are stable and quantitatively meaningful. The threshold therefore emerges from the data themselves through a principled uncertainty-based procedure, rather than from a conventional value adopted in the literature. Its value falls below the conventional threshold of 0.6 commonly applied in user-driven PS-InSAR practice, reflecting the specific noise structure of the *Vicoforte* basement PS set and the regression-based nature of the model that consumes the data. The estimated SBL uncertainty as a function of temporal coherence, with the 0.56 threshold corresponding to the abrupt change point, is shown in Figure 5. 34.




Figure 5. 34. Estimated SBL predictive uncertainty as a function of the minimum temporal coherence threshold applied to the training subset for the *Vicoforte* basement PSs. The uncertainty is computed as the temporal mean of the predictive standard deviation σ normalised by its maximum value, and decreases monotonically with the threshold up to a breakpoint beyond which further exclusion of low-coherence points no longer reduces the predictive uncertainty. The vertical red line indicates the application-specific threshold of 0.56 derived from the breakpoint, which is adopted as the operational coherence threshold for the calibration of the SBL surrogate model in Section 5.4.6..

The PS-InSAR processing of the 62-image COSMO-SkyMed Second Generation stack produced a sparse-point solution over the Sanctuary site with displacement values that span a broad range. The raw LOS displacement time series, before any coherence-based filtering, exhibit cumulative excursions reaching several hundred millimetres over the observation period for the most variable points. The full set of unfiltered LOS displacement time series, with the broad spread reaching cumulative values of approximately ± 400 mm by the end of the observation period, is shown in Figure 5. 35. The distribution of mean LOS velocities across the full PS population extends from large negative to large positive values, with a central mode at near-zero values corresponding to stable urban scatterers and long tails that reflect the inclusion of low-coherence points affected by phase unwrapping errors, atmospheric residuals, and dielectric variability of vegetated and moist surfaces. The cumulative displacement distribution follows a similar pattern, with the central bulk at sub-centimetre values and long tails extending to several hundred millimetres. These raw distributions are not themselves the quantity of interest for soil monitoring. They provide the empirical justification for the coherence-aware selection procedure developed in Section 5.4.5: without a principled threshold, the PS population includes points whose apparent displacement behaviour reflects processing and environmental artefacts rather than the physical process that the model is designed to characterise. The 0.56 threshold identified through the SBL uncertainty analysis reduces the PS population to a small subset of high-reliability basement points whose displacement histories fall within a sub-millimetre to millimetre range, as discussed in Section 5.4.5.






Figure 5. 35. Unfiltered LOS displacement time series for all candidate PS in the small area centred on the *Sanctuary of Vicoforte*. The wide spread of cumulative displacements, reaching approximately ± 400 mm by the end of the observation period, reflects the

inclusion of low-coherence points affected by phase unwrapping errors, atmospheric residuals, and dielectric variability of the surrounding terrain, and motivates the coherence-based selection procedure developed in Section 5.4.5.

The spatial distribution of the selected PSs, rendered in two dimensions in panel (a) and in three dimensions in panel (b), together with a representative amplitude and LOS displacement time series in panel (c), is shown in Figure 5. 36.



(a)



(b)




Figure 5. 36. Spatial distribution and representative time series of the user-driven small-area PSs at the *Sanctuary of Vicoforte*: (a) two-dimensional Google Earth view of the spatial distribution; (b) three-dimensional Google Earth view; (c) example amplitude time series (upper panel) and LOS displacement time series (lower panel) for a representative basement PS, illustrating the high temporal coherence achieved by the user-driven processing and the seasonal thermoelastic component captured by the parametric inversion. In the upper panel of (c), black dots represent individual amplitude measurements, while red circles indicate the mean value. In the lower panel of (c), red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.

5.4.4. Integration with soil temperature and moisture variables

The physical basis for integrating environmental variables with the PS-InSAR displacement record lies in the mechanical behaviour of the clayey-silt deposit beneath the Sanctuary foundations. Seasonal variations in soil temperature modify the volumetric behaviour of the clay mineral matrix through thermoelastic expansion and contraction, producing cyclic vertical displacements at the foundation level that are coherent across the building footprint and detectable in both the in-situ and satellite displacement records. Seasonal variations in soil moisture content modify the effective stress state and therefore the compressibility of the deposit, generating additional displacement components whose phase and amplitude depend on the depth of the moisture signal and on the soil hydraulic

conductivity. Because both processes operate at seasonal frequencies and are spatially correlated over the scale of the building, their signatures in the PS-InSAR time series overlap and must be jointly characterised before any structural trend can be extracted.

The in-situ monitoring system provides direct measurements of both variables: the 34 thermometers distributed within and around the masonry record temperature at multiple depths, and the piezometers and soil thermo-hygrometer monitor groundwater level and near-surface moisture conditions in the surrounding terrain. Long-term analysis of this record by Ceravolo et al. (2021) demonstrated statistically significant correlations between the environmental time series and both the load cell readings on the post-tensioned tie rods and the natural frequencies identified from the accelerometer data. The correlation with natural frequencies is particularly informative, as it reflects the global stiffness of the soil-structure system and is therefore sensitive to changes in foundation stiffness induced by moisture and temperature variations. Coletta et al. (2019) extended this analysis to a cointegration framework, demonstrating that the modal parameters of the Sanctuary and selected environmental variables share a common stochastic trend, which provides a rigorous statistical basis for environmental effect removal in long-term damage detection.

For the specific purpose of driving a data-driven displacement model, two approaches to environmental input definition were considered. The first uses satellite geophysical products (LST and SWI) as spatially continuous proxies for the environmental field over the site. The second uses the local in-situ thermo-hygrometer as a low-cost, high-frequency source of surface temperature and humidity data. Each approach has distinct advantages. Satellite geophysical products provide spatially distributed coverage at no cost and without on-site instrumentation, but at coarser spatial resolution and with systematic biases between surface radiation temperature (LST) and sub-surface soil temperature. In-situ sensors provide direct, well-calibrated measurements of the variables driving the physical process, but their point-wise representativeness of the broader soil volume beneath the foundations requires verification. The principal satellite soil moisture and land surface temperature products that informed the satellite-versus-in-situ comparison are summarised in Table 5. 8 and Table 5. 9 respectively.

Table 5. 8. Satellite soil moisture products relevant to the integration of environmental variables with InSAR displacement analysis for SHM. For each product, the table reports the source mission, the operating agency, the sensing principle (passive microwave, active C-band, radiometric), the native spatial resolution, the effective temporal resolution, the characterised soil depth (surface or root-zone), the access mode, and a concise indication of the typical SHM application regime. The table supports the data selection argument developed in Section 5.4.4, in which the Copernicus Global Land SWI is adopted as the environmental proxy for the *Vicoforte* case study.

Product	Programs/ Satellites	Agency/ Organization	Sensor/ Freq. Band	Spatial resolution	Temporal resolution
SSM	SMOS	ESA	MIRAS L-band	35-50 km	1-3 days
SSM	SMAP	NASA	Radar L-band	3 km	2-3 days
SSM	NISAR	NASA-ISRO	SAR L-band	5-10 m	6-12 days
SSM	GCOM-W1	JAXA	AMSR2 C-band	10 km	1 day
SSM	MetOp	EUMETSAT	ASCAT C-band	12.5 km	1 day
SSM	Copernicus S1	ESA	CSAR C-band	1 km	1.5-4 days
SWI	MetOp + Copernicus S1	EUMETSAT + ESA	ASCAT C-band + CSAR C-band	1 km	1 day

Table 5. 9. Satellite LST products relevant to the integration of thermal environmental variables with InSAR displacement analysis for SHM. For each product, the table reports the source mission, the operating agency, the thermal sensor and spectral domain, the native spatial resolution, the effective temporal resolution, the access mode, and a concise indication of the typical SHM application regime. The table complements the Soil Moisture products tabulation and supports the environmental variable selection argument developed in Section 5.4.4, in which the Sentinel-3 SLSTR LST is adopted as the thermal proxy for the *Vicoforte* case study.

Product	Programs/ Satellites	Agency/ Organization	Sensor/ Freq. Band	Spatial resolution	Temporal resolution
Landsat	NASA	ETM+, TIRS	LST	100 m	16 days

Product	Programs/ Satellites	Agency/ Organization	Sensor/ Freq. Band	Spatial resolution	Temporal resolution
Terra & Aqua	NASA	ASTER, MODIS	LST	1 km	1 day
GCOM-C	JAXA	SGLI	LST	1/12°, 1/24°	1 day
Copernicus S-3	ESA	SLSTR	LST	1 km	hours

The in-situ sensor deployed within the study area is not intended to provide a point-wise physical representation of local soil behaviour. Rather, it is used as a variable that exhibits a strong correlation with the satellite-derived displacement measurements. Under quasi-stationary conditions, the soil surface temperature tends to equilibrate with air temperature, which generally exhibits limited spatial variability within a small range, except for highly localised temporary effects such as shadow. As a consequence, the correlation between temperature and displacement remains valid over spatially distributed zones even when captured by a single in-situ point. The availability of additional sensors would further improve the robustness of the approach, as the corresponding data would enrich the dictionary of basis functions adopted within the SBL framework. The present case study adopts the in-situ thermo-hygrometer as the primary environmental input to the data-driven model, with satellite geophysical products retained as independent validation sources.

The temporal alignment of the in-situ measurements with the satellite observations required a resampling step. The in-situ thermo-hygrometer records temperature and moisture at one-hour sampling, while the COSMO-SkyMed acquisitions have an effective revisit time of approximately 23 days. The surface soil temperature and moisture records were resampled by temporal interpolation onto the satellite acquisition dates, limiting the analysis to the temporal window common to both monitoring systems. Although the resampled series is under-sampled compared to the original hourly signal, it retains the main informational content associated with medium- and long-term variations, particularly seasonal and inter-seasonal components. The loss of information concerns mainly the fast variations (daily and sub-daily), which are not directly observable at the temporal step of satellite acquisitions. The dominant information content relevant to the quasi-static displacement field, namely the inverse correlation between temperature and soil moisture at seasonal and inter-annual timescales, is preserved by the resampling. The surface soil moisture and temperature time series, resampled to the temporal support of the satellite acquisitions, are shown in Figure 5. 37.




Figure 5. 37. Surface soil moisture and surface temperature recorded by the in-situ thermo-hygrometer at the *Sanctuary of Vicoforte*, resampled to the temporal support of the satellite acquisitions (approximately 23 days). The seasonal cycles and the inverse correlation between temperature and moisture are preserved by the resampling.

5.4.5. Data-driven modelling of displacement behaviour

The data-driven modelling framework adopted for this case study is based on SBL, a probabilistic regression approach that constructs a flexible non-linear surrogate model from a dictionary of basis functions while simultaneously inferring the uncertainty on the model predictions. The SBL approach is suited to this problem for two reasons. First, the relationship between the environmental inputs (soil surface temperature and moisture) and the PS-InSAR displacement output is non-linear and not known a priori from first principles, so a parametric mechanical model would require assumptions about soil constitutive behaviour that are poorly constrained by the available data. Second, the coherence values associated with individual PS measurements carry information about the reliability of those measurements, and the SBL framework can propagate this information into the model likelihood function, so that the posterior estimates of displacement reflect the quality of the input observations (Huang & Beck, 2015; Tipping, 2001).

The SBL model takes the displacement along the satellite LOS as the output variable $y(t)$ and the environmental input vector $x(t)$ (Equations 5. 1) (soil temperature and surface moisture) as predictors.

$$x(t) = [x_1(t), x_2(t)]^T \quad 5.1$$

The regression model is shown in Equation 5. 2:

$$y(t) = \Gamma(x(t))^T w + \varepsilon(t) \quad 5.2$$

where $\Gamma(\cdot)$ is a dictionary of non-linear basis functions, w is the weight vector, and $\varepsilon(t)$ is a Gaussian noise term with variance σ^2 .

The zero-mean Gaussian prior distribution on w is (Equation 5. 3):

$$p(w \mid \alpha) = N(0, A^{-1}) \quad 5.3$$

with $A = \text{diag}(\alpha_1, \dots, \alpha_K) \in K \times K$; each hyperparameter α_k controls the relevance of the corresponding basis function. During learning, many α_k tend to large values, driving the associated weights toward zero and producing a sparse representation of the input-output relationship.

The conditional posterior on the weights given the optimal hyperparameters is shown in Equation 5. 4:

$$p(w \mid y) = N(\mu_w, \Sigma_w) \quad 5.4$$

with covariance (Equation 5. 5):

$$\Sigma_w = (A + \sigma^{-2} \Gamma^{*T} \Gamma^*)^{-1} \quad 5.5$$

and mean (Equation 5. 6):

$$\mu_w = \sigma^{-2} \Sigma_w \Gamma^{*T} y \quad 5.6$$

where $\Gamma^* \in N \times K$ is the design matrix built by evaluating the basis functions on the observed inputs N . For a new observation x^* , the Gaussian predictive distribution is shown in Equation 5. 7:

$$p(y^* \mid x^*, y) = N(\mu^*, \sigma^{*2}) \quad 5.7$$

with mean (Equation 5. 8):

$$\mu^* = \Gamma(x^*)^T \mu_w \quad 5.8$$

and variance (Equation 5. 9):

$$\sigma^{*2} = \sigma^2 + \Gamma(x^*)^T \Sigma_w \Gamma(x^*) \quad 5.9$$

The predictive variance σ^{*2} provides an explicit, time-resolved estimate of the uncertainty associated with each prediction, which is the quantity that drives the coherence-threshold selection procedure described below.

The dictionary of basis functions is defined as the combination of two families. The first is a Fourier series that captures periodic behaviour in the input-output relationship. The second is a polynomial series that represents non-oscillatory trends and long-term behaviour. For both families, the series order is determined through optimisation, without a priori imposition of the order value. A multi-start, gradient-based procedure solves the inverse problem from different initial conditions to reduce the risk of convergence to local minima. Starting from a main initial parameter vector, multiple perturbed initial configurations are generated by introducing variations proportional to the parameter values themselves. Each initial configuration is then subjected to a constrained local optimisation using a trust-region quasi-Newton algorithm that exploits gradient information to converge toward a local minimum. The local solutions obtained from all initial points are compared, and the one that minimises the objective function is selected as the final solution. The cost function is minimised using a 50% weight on the training error and a 50% weight on a penalty term defined as the absolute difference between training and validation errors, which limits model complexity and reduces the risk of overfitting. The optimisation is limited to a maximum of 100 iterative cycles, with an early stopping criterion that terminates the algorithm when the normalised difference between two successive weight estimates falls below 1×10^{-6} . We split the dataset into three subsets, allocating 35% to training, 35% to validation, and 30% to testing, with the testing subset excluded from the training phase and used exclusively for the final generalisation assessment.

The coherence-aware extension of the standard SBL framework enters at two points. First, the empirical coherence threshold is identified by examining how the predictive uncertainty, estimated as the temporal mean of the predictive standard

deviation σ^* normalised by its maximum value, varies as a function of the minimum coherence of the PS included in the training dataset. As the coherence threshold is progressively raised, low-coherence PS are excluded and the training dataset becomes more reliable; the estimated uncertainty decreases until a breakpoint is reached, beyond which further exclusion no longer reduces uncertainty. The breakpoint corresponds to the threshold value at which low-coherence, high-noise PS are separated from high-coherence, low-noise ones. For the *Vicoforte* basement PS dataset, the empirical threshold identified through this procedure is 0.56. Below this value, the uncertainty associated with the predictions increases sharply, rendering any regression model unreliable; above it, predictions are stable and quantitatively meaningful. We estimated the threshold by dividing the coherence-ordered PS set into two subsets and identifying the division point that maximises the difference between the root mean square uncertainty of the two groups.

The SBL model calibrated on the selected high-coherence basement PSs, using the in-situ soil temperature and moisture as inputs, achieved an optimal dictionary with polynomial order 3 and trigonometric order 1 for each input variable. By combining the 0.56 coherence threshold with the elevation restriction (-1 m to +5 m), we reduced the full PSs population to five retained basement scatterers that satisfy both quality conditions and are attributable to the soil-structure interface. We chose the elevation interval of -1 m to +5 m to isolate scatterers physically attributable to the basement region of the Sanctuary, where the soil-structure interaction signal is expected to dominate the displacement record. The lower bound of -1 m allows for residual height uncertainty in the geocoded positions of basement-level scatterers, while the upper bound of +5 m excludes scatterers on elevated structural elements whose displacement record would be dominated by the thermoelastic response of the masonry rather than by the soil compressibility cycle that the SBL model targets. The small size of the retained population (five PSs) constitutes a statistical limitation on the generalisability of the SBL calibration; however, the combined use of all five PSs within a hierarchical modelling framework partially compensates for the limited per-PS observation count by exploiting inter-PS coherence constraints on the shared environmental drivers. The dictionary of basis functions is constructed as the union of two modular sub-dictionaries, one for each input variable, with each sub-dictionary comprising a constant term, polynomial terms up to the optimised order, and trigonometric terms up to the optimised order. The modular construction preserves the symmetry between the input variables and facilitates the extension of the model to additional environmental drivers without restructuring the shared component of the dictionary. The constant terms of the two sub-dictionaries are formally collinear and act jointly

as a single intercept of the model: under the Gaussian prior assumed by the Relevance Vector Machine (RVM) algorithm, the weight of the intercept is distributed between the two redundant coefficients, and the effective intercept of the calibrated model is the sum $w_1 + w_7$. The resulting basis function set is shown in Equation 5. 10:

$$\begin{aligned} \mu_j^*(x_1, x_2|c) = & w_1 + w_2 x_1 + w_3 x_1^2 + w_4 x_1^3 + \\ & + w_5 \sin(x_1) + w_6 \cos(x_1) + w_7 + w_8 x_2 + \\ & w_9 x_2^2 + w_{10} x_2^3 + w_{11} \sin(x_2) + w_{12} \cos(x_2) \end{aligned} \quad 5. 10$$

where x_1 and x_2 denote surface temperature and surface humidity respectively, μ_j^* is the j -th predicted LOS displacement, and w_k are the mean estimated weights identified through the RVM algorithm. The two constant-term coefficients w_1 and w_7 exhibit nearly identical estimated values across all five retained PSs, as reported in Table 5. 10, confirming the expected distribution of the intercept weight between the two collinear basis functions under the Gaussian prior of the RVM algorithm. The effective intercept of each calibrated model is therefore the sum of the two coefficients, and its magnitude reflects the mean LOS displacement of the corresponding basement PS over the observation period. The weights are estimated independently for each of the five retained PSs, and the sparsity pattern revealed by the estimated weights confirms that the model identifies a small subset of dominant basis functions per PS while suppressing the remaining terms toward zero. The polynomial terms of low order (constant and linear) and the trigonometric terms capture the seasonal signature, while higher-order polynomial terms remain negligible, consistent with the expected smoothness of the soil-displacement response to environmental forcing. Validation against the withheld test subset confirms that the model generalises well, with normalised error metrics consistent with the posterior uncertainty estimates from the training phase. The Z-score of the model residuals remains below 2 across training, validation, and test subsets, corresponding to normalised error metrics, specifically the Normalised Root Mean Square Error (NRMSE) and the Normalised Mean Absolute Error (NMAE), that fall well within the two-sigma confidence interval implied by the observational uncertainty of the satellite displacement measurements. The LoS displacement time series predicted by the calibrated SBL model for the five retained basement PSs, together with the training, validation, and testing data points and the SBL uncertainty bands, are shown in Figure 5. 38. The average error metrics across the training, validation, testing, and prediction phases are shown in Figure 5. 39. The

estimated mean weight values of the basis functions for the five retained PSs are reported in Table 5. 10.




Figure 5. 38. LoS displacement time series predicted by the calibrated SBL model for the five retained basement PSs at the *Sanctuary of Vicoforte*: (a) point 1; (b) point 2; (c) point 3; (d) point 4; (e) point 5. The blue dots indicate the training subset, the magenta dots the validation subset, the cyan diamonds the testing subset; the solid red line is the SBL prediction; the dashed grey lines are the SBL uncertainty bands; the solid black line interpolates the satellite observations.




Figure 5. 39. Average error metrics of the SBL displacement estimates at the *Sanctuary of Vicoforte*: Z-score (top), Normalised Root Mean Square Error (middle), and Normalised Mean Absolute Error (bottom), reported separately for the training, validation, testing, and total prediction subsets.

Table 5. 10. Estimated mean weight values of the basis functions of Equation 5. 10 for the five retained PSs at the *Sanctuary of Vicoforte*.

μ_{sj}	w_1	w_2	w_3	w_4	w_5	w_6
1	2.1×10^{-1}	-3.7×10^{-9}	-1.9×10^{-3}	3.9×10^{-10}	9.1×10^{-9}	3.5×10^{-2}
2	6.3×10^{-3}	-2.5×10^{-9}	-1.0×10^{-4}	-1.8×10^{-10}	8.9×10^{-9}	2.7×10^{-3}
3	1.2×10^{-1}	-7.0×10^{-9}	-1.1×10^{-3}	3.1×10^{-9}	6.7×10^{-9}	2.1×10^{-2}
4	1.9×10^{-1}	-4.7×10^{-9}	-1.7×10^{-3}	1.3×10^{-9}	9.3×10^{-9}	2.7×10^{-2}
5	-3.0×10^{-1}	-3.4×10^{-10}	2.8×10^{-3}	2.5×10^{-9}	-1.5×10^{-8}	-4.2×10^{-2}
μ_{sj}	w_7	w_8	w_9	w_{10}	w_{11}	w_{12}
1	2.1×10^{-1}	7.1×10^{-9}	3.4×10^{-9}	3.0×10^{-9}	-1.4×10^{-10}	5.0×10^{-9}
2	1.6×10^{-2}	5.1×10^{-9}	1.8×10^{-9}	-9.2×10^{-11}	-6.1×10^{-9}	1.2×10^{-9}
3	1.2×10^{-1}	5.5×10^{-9}	1.5×10^{-9}	1.7×10^{-9}	1.6×10^{-9}	8.7×10^{-9}
4	1.9×10^{-1}	6.3×10^{-9}	2.6×10^{-9}	2.3×10^{-9}	5.8×10^{-10}	7.2×10^{-9}
5	-3.0×10^{-1}	-1.1×10^{-8}	-6.5×10^{-9}	-6.1×10^{-9}	4.3×10^{-9}	-3.5×10^{-9}

The monitoring capability of the trained model is operationalised through a stochastic oscillator, adapted from its original application in financial time-series

analysis as an indicator of extreme-value excursions. The stochastic oscillator $s_o(t)$ is computed from the Z-score time series of the model residuals, where the Z-score normalises the prediction error by the model-estimated variance. Values of the stochastic oscillator close to unity define alert periods during which monitoring should be intensified. A sustained divergence or flattening of the oscillator time series indicates a substantial deviation of the current soil conditions from their historical behaviour, activating a second-level alert that suggests the need for more detailed investigations. The stochastic oscillator is defined in Equation 5. 11:

$$s_o(t) = \frac{Z(t) - Z_{min}}{Z_{max} - Z_{min}} \quad 5. 11$$

where t denotes time, while Z_{min} and Z_{max} represent respectively the minimum and maximum Z-score values calculated over the observation period, so that the indicator takes normalised values between 0 and 1. For the *Vicoforte* dataset, the stochastic oscillator exhibited predominantly low values (approximately 0.25 throughout the observation period) with peaks approaching unity during the warm months of summer. The sustained low values confirm the overall stability of the soil-structure interface during the analysed period: the displacement behaviour is dominated by the seasonal environmental component and shows no evidence of persistent anomalous trends. The summer peaks correspond to periods of maximum temperature during which the non-linearity of the soil thermo-hydrological response generates larger residuals relative to the model predictions, and should not be interpreted as structural warning signals. The stochastic oscillator time series for the five retained PSs, with the two alert levels indicated, are shown in Figure 5. 40. The oversampled SBL predictions, computed at the one-hour sampling of the in-situ environmental inputs and extending beyond the last available satellite observation, are shown in Figure 5. 41.





Figure 5. 40. Stochastic oscillator time series for the five retained PSs at the *Sanctuary of Vicoforte*, computed from the Z-score of the SBL model residuals according to Equation 5. 11. Values close to unity correspond to a transient alert (Level 1, isolated exceedance), while persistent values close to unity correspond to a sustained alert (Level 2, prolonged exceedance), as defined in Section 5.4.6. The oscillator remains predominantly below 0.5 throughout the observation period, with occasional peaks during the warm summer months consistent with thermo-hydrological forcing not fully captured by the calibrated model.






Figure 5. 41. Oversampled LoS displacement time series predicted by the calibrated SBL model for the five retained basement PSs at the *Sanctuary of Vicoforte*: (a) point 1; (b) point 2; (c) point 3; (d) point 4; (e) point 5. The prediction is computed at the one-hour sampling of the in-situ environmental inputs and extends beyond the last available satellite observation, with the SBL uncertainty bands in dashed grey indicating the time-varying confidence interval.

Once calibrated, the SBL model enables two monitoring capabilities that extend beyond the temporal resolution of the satellite record itself. The first is temporal densification: by evaluating the calibrated model on the continuous hourly time series of in-situ environmental measurements while keeping the trained weights fixed, the model produces displacement predictions at hourly sampling rate rather than at the 23-day COSMO-SkyMed revisit interval. The predictions are accompanied by a time-varying uncertainty envelope derived from the predictive variance of the SBL framework. The second capability is blind prediction beyond the last available satellite observation. For the *Vicoforte* dataset, satellite acquisitions are available up to 14 November 2025; beyond this date, the calibrated model generates displacement predictions driven exclusively by the in-situ environmental record, extending the reconstruction through December 2025 and onward as long as the environmental input remains available. This blind prediction window closes the temporal gap between satellite acquisitions and the moment at which a structural diagnosis is required. Because the input environmental variables are available almost in real time from the low-cost thermo-hygrometer, while satellite data have intrinsic acquisition and processing delays, the model anticipates the displacement behaviour that the next satellite acquisition is expected to record,

within known uncertainty bounds. When new satellite data become available, they can be immediately compared with the prediction, allowing for posterior validation and early identification of anomalous deviations from expected behaviour. This capability converts the satellite-derived measurement into an operational continuous monitoring stream supported by a low-cost ground sensor, with a well-defined uncertainty structure that provides a principled basis for alert-level decisions.

5.4.6. Implications for integrated SHM of monumental buildings

The *Vicoforte* case study advances the conceptual framework developed in Case Studies I and II in a direction that is specifically relevant to monumental heritage structures under permanent monitoring. The methodological contribution is threefold: the identification of temporal coherence as a design variable rather than a passive filter, the integration of environmental variables (whether from in-situ sensors or from satellite geophysical products) as drivers of a data-driven displacement model, and the construction of a surrogate model that enables continuous monitoring between satellite acquisitions through a quantitative uncertainty framework. The integrated case-study workflow, encompassing the user-driven PS-InSAR processing chain, the coherence-aware SBL surrogate model, and the operational monitoring layer, is summarised in Figure 5. 42.

The coherence-aware perspective reframes a widely recognised limitation of PS-InSAR as an analytical resource. Temporal coherence has conventionally been treated as a quality criterion for discarding unreliable observations, and uniform thresholds are applied in both standardised products and many user-driven workflows without explicit consideration of how the threshold value affects the inferred displacement field. At *Vicoforte*, the empirical coherence threshold of 0.56 that separates reliable from unreliable PS in the surrogate model framework cannot be determined from geometric or amplitude criteria alone, nor from conventional threshold values adopted in the literature. The SBL-based uncertainty analysis provides a principled method for identifying this threshold from the data themselves, and the resulting model is demonstrably more stable and better calibrated than a model trained on the full PS set without coherence weighting. The approach is transferable to any structure where the scatterer population is heterogeneous, and the monitoring objective requires a data-driven characterisation of displacement behaviour.




Figure 5. 42. Integrated case-study workflow at the *Sanctuary of Vicoforte*: (a) PS-InSAR data processing chain, from data preparation to small-area processing and PS export; (b) SBL surrogate model and monitoring layer, from coherence threshold derivation to stochastic oscillator alert levels. Solid arrows indicate the sequence of tasks within each macro-process; dashed arrows indicate transitions between the two macro-processes.

The integration of low-cost in-situ environmental sensors as model inputs addresses a practical constraint endemic to heritage site monitoring. The density and coverage of in-situ instrumentation are limited by conservation requirements, structural access constraints, and cost. A single thermo-hygrometer measuring surface soil temperature and humidity at hourly sampling, combined with a satellite PS-InSAR archive and the SBL surrogate framework, is sufficient to support a continuous displacement monitoring stream at a fraction of the cost of dedicated displacement sensor networks. The demonstration that satellite-derived environmental products (LST, SWI) are also suitable proxies for the same environmental variables, as discussed in Section 5.4.4, extends the accessibility of the framework to sites where no in-situ sensors are available at all. The methodological implication is that satellite-based SHM of monumental buildings, previously constrained by the revisit-limited temporal resolution of the satellite record, can be operationalised as a continuous monitoring service when complemented by low-cost environmental sensors or by satellite geophysical products.

The framework developed at *Vicoforte* contributes the satellite displacement dimension to an integrated SHM logic that has been extensively discussed in the context of modal parameter tracking and environmental effect removal. The monitoring programme at the Sanctuary now encompasses both the dynamic response of the dome and drum, tracked through accelerometers and operational modal analysis, and the quasi-static displacement of the soil-structure interface, tracked through PS-InSAR and characterised through the SBL model. The two monitoring streams are physically complementary. The dynamic system is sensitive to changes in global structural stiffness but insensitive to slow quasi-static settlement processes; the PS-InSAR system captures quasi-static deformation but is blind to vibration-frequency phenomena. Together, they provide a more complete picture of the structural state than either technique alone, and they can be combined in a joint model-updating procedure that exploits both modal parameters and satellite displacements as observational constraints.

The stochastic oscillator monitoring statistic operationalises the surrogate model output as a decision-support variable. The two alert levels, a transient exceedance (Level 1) and a sustained exceedance (Level 2), provide a graduated response protocol that avoids the false-positive burden of single-threshold alarm systems while ensuring that genuine persistent anomalies are escalated. The validation of this protocol at *Vicoforte* over a multi-year dataset without the occurrence of structural anomalies demonstrates that the system operates stably in the null-hypothesis condition. Prospective validation, meaning the testing of the protocol against a structure that subsequently develops a genuine anomaly, remains a future research objective.

Several limitations of the approach should be acknowledged. The application at *Vicoforte* benefits from a permanent in-situ monitoring system and a well-characterised geotechnical context that are not available at the majority of monitored structures. The transferability of the coherence-aware SBL framework to structures with less favourable data conditions remains an open research question. The basis function dictionary adopted (polynomial order 3 and trigonometric order 1 per input variable) was optimal for the specific combination of soil type, climatic regime, and structural configuration at *Vicoforte*; different sites may require different dictionary configurations, and the sensitivity of the results to this choice has not been systematically explored. The SBL framework assumes Gaussian noise on the displacement observations, an assumption that is reasonable for high-coherence PS but may be violated at low coherence. Finally, the blind prediction capability relies on the continued validity of the calibrated

model, which may degrade over time if the underlying physical process changes due to long-term structural evolution or to environmental shifts beyond the range covered by the training data.

5.5. Synthesis for engineering practice

The three case studies presented in this chapter span a range of structural typologies, monitoring objectives, satellite data sources, and analytical depths. The *Santissima Trinità* represents a compact masonry heritage building in a dense urban environment monitored through Sentinel-1 data at two levels of processing. The *Himera Viaduct* and the *Albiano-Magra Bridge* address infrastructure collapse scenarios using Sentinel-1 and COSMO-SkyMed data respectively, combining pre- and post-event analysis with TS diagnostics. The *Sanctuary of Vicoforte* demonstrates the most complete integration architecture, combining EGMS screening, user-driven COSMO-SkyMed processing, satellite-derived environmental variables, in-situ monitoring, and data-driven surrogate modelling. A fourth application, concerning the post-event assessment of the 2025 Palisades Fire in Los Angeles through the fusion of PS-InSAR and airborne LiDAR, is developed in Section 5.6 as an ongoing research direction that extends the monitoring framework into a disaster-response regime not previously addressed within the architecture of the present thesis.

The diversity of these applications is deliberate. Each case study was selected not only for its intrinsic structural interest but also because it tests a specific combination of conditions from the methodological framework of Chapter 4 and reaches a specific diagnostic level as defined in Section 4.6. The present synthesis draws together the cross-cutting findings that emerge when the three completed case studies are examined comparatively, and organises them around a governing question: under what conditions does satellite InSAR reach each diagnostic level, and what are the recurrent factors that enable or constrain progression from one level to the next?

The synthesis opens with a formalised statement of the tiered monitoring framework that the case studies collectively support (Section 5.5.1). The subsequent sections examine the principal diagnostic bottlenecks identified across the case studies: measurement point density and spatial coverage (Section 5.5.2), the role of temporal coherence as a diagnostic variable (Section 5.5.3), thermal and environmental effects on displacement interpretation (Section 5.5.4), and the diagnostic value of TS together with the limits of retrospective inference and the

early-warning question (Section 5.5.5). The synthesis closes with an evaluation of the diagnostic level achieved by each case study against the framework of Section 4.6, identifying the conditions under which satellite data become decision-ready and the priority developments needed to extend that boundary (Section 5.5.6).

5.5.1. A tiered framework for satellite-assisted SHM: from portfolio screening to structure-specific diagnostics

The progression from portfolio screening to structure-specific diagnostics is not a progression in measurement accuracy alone. It is also a progression in the depth at which the analyst must understand and control the processing chain from which the displacement measurement is derived. Portfolio screening is compatible with standardised products such as those delivered through the EGMS platform, because the diagnostic question at this level, namely whether a given structure is anomalous with respect to its reference population, is answered at a level of abstraction that tolerates fixed processing choices. The uniformity of the processing pipeline across a wide geographic coverage is in fact a methodological advantage at this level, because it ensures the comparability of the velocity and time-series observables across the screened population. Structure-specific diagnostics, by contrast, require the analyst to govern every parametric decision of the workflow. At this level, each processing choice interacts with the structural question being asked, and the assumptions embedded in a standardised pipeline may no longer be consistent with the observational requirements of the specific case. The identification and temporal characterisation of TS, the adoption of application-specific coherence thresholds, the integration of environmental covariates into the phase model, and the spatial restriction of the processing to a small-area window around the target structure are all examples of operations that fall outside the scope of standardised products. A user-driven processing environment is therefore not a methodological preference at the highest diagnostic level but a diagnostic requirement, and the transition from one tier of the framework to the next is accompanied by a corresponding transition in the degree of control that the analyst exercises over the processing chain.

The combined experience of the three case studies supports an operational articulation of the four-level diagnostic framework introduced in Section 4.6. The articulation is organised around three analytical stages that specify how each diagnostic level is reached in practice. Stage one uses the EGMS Calibrated product and corresponds to Level 1 (screening). Stage two introduces user-driven processing and supports Level 2 (displacement characterisation). Stage three combines satellite displacement data with complementary information and enables

Levels 3 and 4 (diagnostic interpretation and continuous monitoring with automated anomaly detection). Each stage may constitute a terminal point if the information obtained satisfies the monitoring objective, or a trigger for escalation to the next stage if the evidence reveals anomalies, ambiguities, or diagnostic questions that the current stage cannot resolve. The mapping between the operational stages of the framework and the diagnostic levels of Section 4.6, with the case studies that exemplify each combination, is summarised in Table 5. 11.

Table 5. 11. Mapping between the operational stages of the satellite-assisted SHM framework (Section 5.5.1) and the diagnostic levels of the analytical framework (Section 4.6). For each stage, the table indicates the typical operational action, the data source, the diagnostic levels accessed, and the case studies that exemplify the stage.

Stage	Operational action	Typical data source	Diagnostic levels accessed	Case studies
Stage 1: Portfolio screening	Identification of structures with kinematic anomalies within a regional inventory; baseline confirmation of stable behaviour	EGMS Calibrated product (or equivalent standardised national-scale displacement service)	L1 (Screening) fully accessible; L2 (Displacement characterisation) partially accessible at the building-block scale	<i>Santissima Trinità</i> (negative screening); <i>Himera Viaduct</i> (post-reconstruction stability)
Stage 2: Structure-targeted analysis	Quantitative characterisation of the displacement field on a specific structure; identification of TS encoding event chronology	User-driven processing of high-resolution SAR archives (Sentinel-1, COSMO-SkyMed)	L2 (Displacement characterisation) fully accessible; L3 (Diagnostic interpretation) partially accessible when complementary information supports the causal interpretation	<i>Santissima Trinità</i> (Section 5.2); <i>Himera Viaduct</i> user-driven analysis (Section 5.3.1); <i>Albiano-Magra Bridge</i> (Section 5.3.2)
Stage 3: Continuous integrated monitoring	Coherence-aware data-driven modelling of displacement behaviour; environmentally-	Multi-source integration: user-driven SAR archive, in-situ environmental	L3 (Diagnostic interpretation) fully accessible; L4 (Prediction) accessible when continuous	<i>Sanctuary of Vicoforte</i> (Section 5.4)

Stage	Operational action	Typical data source	Diagnostic levels accessed	Case studies
	driven prediction with uncertainty quantification; alert-level decision support	sensors, satellite geophysical products (LST, SWI), FE structural model	environmental input and a calibrated surrogate model are available	

The first stage employs the EGMS Calibrated product as a low-cost, immediately accessible instrument for surveillance at the portfolio or territorial scale. The analyst extracts measurement points within and around the structure of interest, examines the mean velocity map for spatial gradients suggestive of differential settlement or ground motion, and inspects individual time series for non-linear trends, seasonal anomalies, or discontinuities. The investment required is minimal: EGMS data are free, pre-processed, and downloadable without specialist SAR expertise. Three distinct functions can be served at this stage. The first is screening, which is the systematic identification, within a portfolio of structures, of those assets exhibiting velocity anomalies above the noise floor of the product. The second is baseline confirmation, which is the documentation of a null result for structures that are genuinely stable, as illustrated by the *Santissima Trinità* where EGMS measurement points above the structure-specific coherence threshold of 0.80 showed near-zero mean velocities and no temporal anomalies, with the only exception of one low-coherence point on the cylindrical body whose post-2022 decorrelation is attributable to the environmental forcing of the masonry rather than to a structural cause. The third is retrospective surveillance, which is the periodic interrogation of EGMS time series over structures that are not permanently instrumented, exploiting the annual updates and the European coverage of the product to query the displacement history of any structure without a dedicated monitoring campaign. At *Vicoforte*, the EGMS Calibrated product provided multi-geometry displacement time series over the dome surface, serving as an independent contextual layer alongside the dedicated COSMO-SkyMed processing.

The limitations of the first stage are well defined by the case studies. The EGMS processing chain applies uniform coherence thresholds calibrated for regional ground motion monitoring, not for building-scale structural analysis. Measurement point density on compact urban structures may be insufficient to resolve displacement gradients between structurally distinct elements. TS, which carry some of the most diagnostically specific information in collapse scenarios,

are systematically discarded by the processing thresholds of standardised products. The reference frame and processing parameters are not adjustable by the user. These constraints define the boundary conditions for escalation to the second stage.

The second stage is activated when the EGMS interrogation reveals anomalies that require spatial or temporal localisation beyond what the standardised product provides, or when the monitoring objective demands control over processing parameters that the EGMS pipeline does not expose. User-driven processing recovers capabilities that are absent from EGMS: the identification and temporal characterisation of TS, the application of structure-specific amplitude thresholds and reference point selection, the adjustment of the coherence criterion to the expected deformation regime, and the ability to process multiple orbit geometries within a unified analytical framework. The case studies demonstrate the added value of the second stage under different conditions. At the *Santissima Trinità*, user-driven processing of 169 Sentinel-1 SLC images produced two measurement points on the church with temporal coherence values of 0.92 and 0.86, substantially higher than the EGMS points on the same structure, enabling a more reliable characterisation of the seasonal displacement signature. At *Himera*, user-driven processing of the peri-collapse Sentinel-1 archive recovered a slope-adjacent TS whose lifecycle boundaries constrain the chronology of the post-event clearance operations, information that is entirely absent from the EGMS post-reconstruction product. The case also exemplifies a geometric limitation specific to valley-confined infrastructure, where the deck itself is occluded from the radar line of sight in both ascending and descending acquisitions. At *Albiano-Magra*, COSMO-SkyMed processing at Stripmap resolution provided scatterer densities on the bridge deck and piers that C-band Sentinel-1 data cannot achieve, and the unified processing of the full 2011-2022 archive across the collapse boundary yielded a continuous structural record spanning pre-failure kinematics, the event itself, and the post-demolition phase. The same processing also revealed a pre-collapse displacement signal on the Caprigliola hill above the eastern abutment, with LOS velocities of opposite sign in the two orbits indicative of horizontal slope motion, an observation that adds a geotechnical dimension to the structural reading of the case and parallels the role of the Rocca di Sciara slope in the *Himera* case.

User-driven processing carries a cost in expertise, computational resources, and turnaround time. The decision to escalate from the first stage to the second should therefore be governed by observable signals in the data rather than by precautionary principles alone. Four practical triggers justify this escalation, identified empirically across the case studies. The first is a spatial velocity gradient across the structure

exceeding 2-3 mm/year over the scale of the structural element. The second is a statistically significant change-point in an individual time series. The third is a systematic decline in temporal coherence over successive observation windows on a specific structural element. The fourth, where the first stage is conducted with EGMS, is an insufficient number of measurement points on the structure to resolve the displacement pattern of interest.

The third stage combines satellite displacement data with complementary information sources to support a structural diagnosis that exceeds what the satellite record alone can provide. Complementary sources may include in-situ monitoring data, geotechnical characterisation, numerical structural models, inspection records, and satellite-derived environmental variables. Integration may take the form of cross-validation, in which satellite and in-situ displacement trends are compared over their overlapping temporal window, or of model calibration, in which the satellite data are used as boundary conditions or calibration targets for a finite element or surrogate model. The *Vicoforte* case study demonstrates the most complete realisation of the third stage in this thesis. The SBL surrogate model, trained on COSMO-SkyMed basement PS displacements and on in-situ environmental inputs, generates a continuous monitoring output (the stochastic oscillator) that normalises the structural condition into a decision-support variable expressed in terms of statistical deviation from a calibrated baseline rather than in raw LOS displacement units. The approach removes the need for the end-user to interpret satellite measurement geometry and presents the monitoring output in a form that is directly compatible with standard SHM decision frameworks.

The third stage is not achievable for all structures. It requires a sufficiently long and temporally dense satellite archive, environmental input variables that capture the dominant sources of reversible displacement, a coherence-aware data selection strategy, and, in most cases, an independent dataset for cross-validation. The *Vicoforte* analysis benefited from a permanent in-situ monitoring system and a well-characterised geotechnical context that are not available at the majority of monitored structures. The transferability of the third stage to structures with less favourable data conditions remains an open question, discussed further in Section 5.5.6.

The three stages are not a rigid procedural sequence, and the appropriate entry point depends on the available data and the monitoring objective. For structures on which a COSMO-SkyMed archive already exists, stage two may be the natural starting point and Level 2 the initial diagnostic target. For structures where in-situ

monitoring data are available, stage three may proceed in parallel with stage one, supporting Level 3 or Level 4 without an intermediate Level 2 pass. For portfolio-scale screening applications, stage one may be the permanent analytical level, with escalation triggered only for the subset of structures that exhibit anomalies. The operational articulation is best understood as a decision tree whose branches are activated by the information content of the data at each step, by the risk profile of the asset, and by the resources that the monitoring programme can allocate, and whose terminal nodes correspond to the four diagnostic levels of Section 4.6.

5.5.2. Point density and spatial coverage on structures

The number of usable measurement points on a structure, and their spatial distribution relative to the elements of engineering interest, emerged across all three case studies as the single most recurrent factor determining whether the satellite analysis could progress beyond the screening level. The comparative evidence reveals that point density is governed by the interaction of four variables: the sensor and its spatial resolution, the processing strategy and its selection criteria, the structural geometry and its relationship to the satellite viewing angle, and the urban or environmental context in which the structure is embedded.

The contrast between Sentinel-1 and COSMO-SkyMed is the most direct illustration of the sensor effect. At *Albiano-Magra*, COSMO-SkyMed Stripmap processing at approximately 3 m ground resolution produced a dense scatterer population on the bridge deck, piers, and abutments, sufficient to characterise the spatial distribution of pre-collapse LOS velocities across individual structural components. An equivalent Sentinel-1 analysis at the same site would yield substantially fewer points on the structure, because the coarser resolution cell (approximately 5 m $\times$ 20 m in IW mode) mixes the contribution of the structural element with the surrounding terrain, reducing the probability that a single dominant scatterer within the cell satisfies the phase stability criterion. The *Vicoforte* case study confirms this pattern from a different angle: at the basement level, where the monitoring objective requires PS on the soil-structure interface, the COSMO-SkyMed X-band data provide a spatially selective measurement set that can be restricted to a narrow elevation band, while the EGMS Sentinel-1 product concentrates its points on the dome and bell towers, where the larger masonry surfaces generate stronger and more stable radar returns at C-band.

The processing strategy exerts a comparably strong influence. At the *Santissima Trinità*, the EGMS product appears to offer a higher point count within

the urban block than the user-driven solution. The apparent density advantage, however, is misleading. The EGMS processing chain analyses each Sentinel-1 subswath independently and applies a temporal coherence threshold of approximately 0.5, retaining points that appear in one subswath but not in others. The user-driven solution performs a combined analysis across subswaths and applies a stricter coherence threshold of 0.6. The result is that the user-driven points, although fewer, exhibit substantially higher temporal coherence (0.86-0.92 versus 0.66-0.86 for the closest EGMS points) and are more reliably attributable to the church structure rather than to adjacent buildings within the same resolution cell. The lesson is general: raw point count on a structure is an unreliable proxy for diagnostic capacity. What matters is the number of points that pass a coherence threshold appropriate to the monitoring objective and that can be attributed, through geometric plausibility analysis, to a specific structural element.

Structural geometry introduces a further constraint that is independent of the sensor and the processing chain. The *Santissima Trinità*, with its compact circular plan, low cylindrical walls, and dome rising above a dense urban block, presents a limited radar-facing surface area and is partially occluded by adjacent buildings on multiple sides. The number of geometrically plausible PS on the church body is inherently small regardless of the processing approach. By contrast, the *Himera Viaduct*, with its linear alignment extending over several hundred metres and its exposed concrete and metallic deck surfaces, offers a much larger radar-illuminated footprint per unit of structural length. The *Vicoforte* Sanctuary occupies an intermediate position: the oval dome and the bell towers generate strong C-band returns at the EGMS level, while the basement, which is the target of the user-driven COSMO-SkyMed analysis, requires higher resolution and careful elevation filtering to isolate PS on the soil-structure interface from those on the surrounding masonry.

The practical implication for the tiered framework of Section 5.5.1 is that point density determines the feasibility of the transition from the first stage to the second. When the EGMS product delivers fewer than approximately three to five measurement points attributable to the structure with coherence above 0.6, the displacement characterisation cannot resolve spatial gradients, and any interpretation beyond a gross velocity check is unreliable. Under these conditions, user-driven processing is warranted because a tailored processing strategy, with structure-specific amplitude thresholds and a locally optimised reference point, may recover additional points or improve the coherence of the existing ones. The escalation does not imply that the EGMS analysis was incorrect. Where user-driven

processing also fails to produce a sufficient measurement set, the analyst must acknowledge that the structure is not adequately observable by satellite InSAR under the available data conditions, and that complementary monitoring approaches are required.

5.5.3. Coherence as a diagnostic variable

Temporal coherence is conventionally treated as a binary quality criterion: measurement points above a chosen threshold are retained, and those below it are discarded. The case studies presented in this chapter suggest that this conventional treatment discards information that is diagnostically relevant and that coherence, when examined as a continuous variable distributed across a structure and evolving over successive observation windows, can contribute to the structural interpretation rather than merely filtering it.

The most developed treatment of coherence as a modelling variable appears in the *Vicoforte* case study. The empirical coherence threshold of 0.56 identified through the SBL uncertainty analysis cannot be determined from geometric or amplitude criteria alone, nor from conventional threshold values adopted in the literature. The SBL-based uncertainty analysis provides a principled method for identifying this threshold from the data themselves, propagating coherence values explicitly into the likelihood function of the surrogate model and weighting high-coherence and low-coherence observations differentially. The resulting coherence-aware modelling framework identifies an application-specific threshold that minimises posterior model uncertainty, providing a principled alternative to the arbitrary thresholds of 0.5 or 0.6 commonly adopted in user-driven practice. The broader methodological point is that temporal coherence is not a pure noise metric: it correlates with the physical nature of the scatterer and with the deformation process itself, so that coherence-based selection generates a spatially and physically biased sampling of the deformation field.

The *Santissima Trinità* case study contributes a complementary perspective on the same theme. The EGMS point P24, geocoded to a portion of cylindrical-body masonry below the drum with a temporal coherence of 0.66, exhibits an increase in time-series oscillation amplitude. The interpretation developed in Section 5.2.3 attributes the increase to the intrinsic thermal and hygric sensitivity of the exposed brick fabric in that elevation range, with a marked decorrelation episode from early 2022 that temporally coincides with the onset of the 2022 northern Italian meteorological drought. The drought consolidated from December 2021, when the

last significant precipitation in Piedmont occurred on 8 December 2021, and produced the driest calendar year recorded in Turin in over two centuries (Avanzi et al., 2024; Montanari et al., 2023; Toreti et al., 2022). Causal attribution of the observed decorrelation episode to drought-induced moisture forcing of the masonry, however, is not established by the present analysis, and three alternative or co-acting factors should be acknowledged. The first is the loss of Sentinel-1B on 23 December 2021, which doubled the nominal Sentinel-1 revisit interval over Turin from six to twelve days for the remainder of the observation window, with a known effect on the temporal coherence of C-band interferometric stacks. The second is the contribution of winter snow events, freeze-thaw cycles, and unmodelled atmospheric phase delays in the Po Valley winter inversion regime. The third is the absence of peer-reviewed precedent for drought-induced façade-moisture decorrelation on masonry heritage buildings, which leaves the proposed mechanism without the validation that would be required for an established attribution.

Under this reading, the low coherence of P24 may carry information about the materiality of the reflecting surface, beyond its role as a quality indicator of the displacement estimate at that specific point. The methodological implication, advanced here as a hypothesis warranting further investigation, is that coherence values in the intermediate range, between 0.5 and 0.8, on masonry heritage buildings may benefit from interpretation that takes into account the materiality of the reflecting surface, and that heterogeneous envelopes such as the Santissima Trinità might warrant differentiated coherence thresholds for surfaces of different physical character. The proposed working hypothesis can be tested in future analyses through the comparison of P24 against control persistent scatterers on stable, non-heritage structures within the same SAR scene, and through the analysis of coherence as a function of the perpendicular and temporal baseline distribution to isolate the dataset effect from any environmental signal.

The infrastructure case studies extend the diagnostic reach of coherence into the temporal domain. A TS is, by definition, a measurement point whose coherence is temporally bounded. The death date of a TS, which is the epoch at which coherence drops irreversibly below threshold, encodes the timing of a physical event: demolition, collapse, material change, or environmental modification of the reflecting surface. In the *Himera* and *Albiano-Magra* analyses, TS death dates aligned with documented structural and post-event events (the post-landslide debris clearance at *Himera*, the deck collapse at *Albiano-Magra*), providing interferometric timestamps at the resolution of the satellite revisit cycle. The

coherence boundary of a TS therefore functions as a diagnostic datum in its own right, with informational content distinct from a simple quality failure.

The fourth manifestation of coherence as a diagnostic variable is the degradation of coherence over successive observation windows on a specific structural element. If the EGMS temporal coherence of measurement points on a building declines systematically between the 2018-2022 and 2019-2023 products, the pattern may indicate progressive surface change, moisture infiltration, or material deterioration. Coherence degradation is a less specific signal than a velocity anomaly, but it is an early and low-cost indicator that can be observed at the first stage of the monitoring framework without dedicated processing.

5.5.4. Thermal and environmental effects on displacement interpretation

Thermally and environmentally driven displacement components are present in the time series of virtually all measurement points located on built structures in temperate climates, and they typically constitute the largest source of periodic variability in the LOS record. The case studies confirm that these components must be explicitly accounted for before any structural conclusion can be drawn, and they reveal a progression from passive recognition of the thermal signature toward its active use as a diagnostic and modelling input.

At the *Santissima Trinità*, the displacement time series of the two user-driven points P1803 and P1765 exhibit a low-amplitude seasonal modulation with a period of approximately twelve months, consistent with the thermoelastic response of the masonry dome and drum to annual temperature variations. The amplitude of this oscillation is small relative to the measurement noise, and the analysis does not attempt to estimate a thermoelastic expansion coefficient. The thermal component is recognised but not exploited: it is treated as a known source of periodic variability that must be distinguished from any structural trend, and its magnitude is consistent with a stable structure undergoing normal environmental cycling.

The infrastructure case studies present a more complex thermal landscape. On concrete bridge decks, the thermal expansion coefficient is larger than on masonry, and the elongated geometry of a bridge span amplifies the LOS projection of the thermal displacement. In the *Albiano-Magra* case study, the COSMO-SkyMed processing included a thermoelastic component in the parametric model (residual height, linear velocity, and thermal coefficient), following the three-parameter

inversion approach described in Section 4.2.2. The inclusion of the thermal parameter improved the coherence of on-deck measurement points and sharpened the separation between the linear velocity field and the seasonal displacement component.

At *Vicoforte*, the environmental dimension extends beyond thermal expansion to include soil moisture effects. The dominant displacement mechanism at the basement level is not the thermoelastic response of the masonry but the seasonal compressibility cycle of the clayey-silt deposit beneath the foundations, driven by coupled temperature and moisture variations. The SBL surrogate model uses in-situ soil temperature and surface moisture as input variables, and the trained model achieves an optimal dictionary with polynomial order 3 and trigonometric order 1, indicating that the relationship between environmental inputs and basement displacement is non-linear and includes both seasonal and inter-annual components. The correlation analysis between satellite-derived environmental variables (LST, SWI) and the in-situ measurements confirms that satellite geophysical products are suitable proxies for sites where no in-situ instrumentation is available.

The progression across the case studies, from passive recognition of the thermal signature at the *Santissima Trinità*, through active estimation of the thermoelastic coefficient at *Albiano-Magra*, to the incorporation of environmental variables into a predictive surrogate model at *Vicoforte*, illustrates a trajectory of increasing analytical sophistication that corresponds to the progression through the diagnostic levels of Section 4.6. At Level 2, the thermal component is characterised but not modelled. At Level 3, it is estimated and attributed to a physical mechanism. At Level 4, it is absorbed into a calibrated model that separates the environmental contribution from the structural residual, enabling prospective anomaly detection. The comparative treatment of thermal and environmental effects across the three case studies, with the corresponding diagnostic level of the treatment, is summarised in Table 5. 12.

Table 5. 12. Comparative treatment of thermal and environmental effects across the three case studies of Chapter 5. The table maps the progression from passive recognition to surrogate modelling onto the diagnostic levels of Section 4.6.

Case study	Dominant environmental driver	Estimation method	Data source for environmental input	Diagnostic level of the treatment
Santissima Trinità	Annual thermoelastic response of masonry dome and drum	Passive recognition; low-amplitude seasonal modulation distinguished from structural trend without quantitative estimation	None (no environmental input used)	L2 (characterisation): thermal component recognised but not modelled
Albiano-Magra Bridge	Annual thermoelastic response of concrete deck spans	Three-parameter inversion: residual height, linear LOS velocity, thermoelastic coefficient	Air temperature record (local meteorological station / global weather service)	L3 (attribution): thermoelastic coefficient estimated and attributed to a physical mechanism
Sanctuary of Vicoforte	Seasonal compressibility cycle of clayey-silt deposit driven by coupled temperature and moisture variations	SBL surrogate model with polynomial and trigonometric basis functions; coherence-aware likelihood; multi-start optimisation	In-situ thermo-hygrometer (primary); Sentinel-3 SLSTR LST and CGLS SWI (validation proxies)	L4 (prediction): environmental contribution absorbed into a calibrated surrogate model that separates environmental signal from structural residual

5.5.5. Temporary scatterers, retrospective inference, and the early-warning question

The case studies of Section 5.3 produced TS whose temporal behaviour was demonstrably bounded by a structural event. In the infrastructure collapse scenarios, TS death dates aligned with documented failure chronologies and provided interferometric timestamps at the resolution of the satellite revisit cycle. The observations support the recognition of TS as a systematic diagnostic category within the PS-InSAR framework, rather than as anomalous points to be filtered out. The present section draws together the cross-cutting methodological implications of these findings, addressing three questions that arise consistently in the literature on InSAR for infrastructure failure analysis.

Two operationally distinct TS types emerged from the infrastructure case studies. The first type is bounded by the physical removal of the scattering surface through demolition or collapse: the reflecting element ceases to exist, and coherence drops irreversibly. The second type is bounded by post-event surface modification, such as debris deposition, which alters the scattering configuration without destroying the underlying geometry. Each type encodes a different physical narrative, and the ability to distinguish among them depends on the availability of complementary information (particularly high-resolution optical imagery) that constrains the physical interpretation of the coherence boundary. The recovery of TS from a SAR archive requires processing choices that differ from standard PS-InSAR practice. Standardised processing chains, including the EGMS pipeline, apply uniform coherence thresholds calibrated for the detection of long-lived scatterers on slowly deforming terrain, and TS with bounded coherence windows frequently fail these filters. User-driven workflows that preserve the full sparse-point population and apply coherence criteria adapted to the specific event timeline are therefore necessary for TS-based diagnostics. The forensic information encoded in collapse-related TS is systematically absent from standardised products and can be accessed only through dedicated processing.

The retrospective analysis of collapse scenarios using archival SAR data has received considerable attention since the 2018 failure of the Morandi Bridge in Genoa, as discussed in Section 4.3.4. The epistemological difficulty of attributing pre-collapse deformation signals to structural causes in a retrospective context has two dimensions. The first is statistical: the noise floor of MT-InSAR imposes a velocity precision of the order of 1 millimetre per year for C-band PSI products under typical observation conditions (Crosetto et al., 2016), with a corresponding

lower bound on the detectable velocity anomaly that scales with the dispersion of the velocity estimates and depends on the time-series length, the acquisition density, and the coherence of the target. Pre-collapse anomalies of 2-5 mm/year, such as those observed on the *Albiano-Magra Bridge* and reported by Farneti et al. (2023a), are close to this detection margin. Their attribution to structural rather than atmospheric or environmental causes requires explicit uncertainty quantification that is rarely performed in the published retrospective literature. The Caprigliola cluster identified in Section 5.3.2 illustrates the converse situation, in which the LOS velocities of the order of 5-6 mm/year and the two-orbit consistency of the signal place the displacement well above the detection margin and within the regime in which retrospective attribution to a geotechnical cause becomes statistically robust. The second dimension is physical: even where a pre-collapse deformation signal is statistically unambiguous, the causal chain between the observed quasi-static LOS displacement and the ultimate failure mechanism is seldom straightforward. Structures fail through combinations of progressive material degradation, loss of boundary conditions, dynamic loading, and sudden brittle fracture. Satellite InSAR is sensitive only to the quasi-static component of this chain and is blind to the sub-revisit-interval dynamics that frequently govern the final failure sequence. Scattarreggia et al. (2022) illustrated this complexity for the *Albiano-Magra* case by showing that non-linear finite element simulations could reproduce the observed collapse through several different failure mechanisms, not all of which predicted the same pre-collapse deformation pattern.

These considerations apply with equal force to the pre-collapse anomalies identified in Case study II. The cumulative displacements of 10-15 mm observed on the *Albiano-Magra* deck spans, while spatially correlated with the failed portion, are at the margin of statistical detectability and are compatible with multiple structural interpretations. The retrospective identification of a pattern consistent with the known failure mechanism does not, on its own, constitute a prospectively actionable diagnostic. The post-event portion of the SAR archive, by contrast, carries information that is less epistemologically fraught. The disappearance of scatterers on collapsed elements, the emergence of TSs on debris and temporary works, and the subsequent re-establishment of stable PSs on reconstructed elements collectively constitute a legible structural chronicle that does not require causal attribution to be diagnostically useful.

Whether satellite InSAR can contribute to operational early warning for bridge infrastructure is a question whose answer depends entirely on the class of failure mechanism under consideration. Selvakumaran et al. (2018) demonstrated, using

TerraSAR-X SBAS analysis of the Tadcaster Bridge in England, that measurable displacement precursors to scour-induced collapse were detectable approximately one month before the event, suggesting that near-real-time monitoring with high-revisit X-band sensors could provide actionable warning. That result rests on a failure mechanism (progressive scour) that produces a relatively gradual and spatially coherent deformation field well suited to InSAR detection. The collapse scenarios investigated in Case study II are structurally different. Both the *Himera* and *Albiano-Magra* failures involved geotechnical or degradation processes whose surface kinematic signature was weak relative to the noise floor, and whose transition from gradual deformation to sudden failure was not preceded by a detectable acceleration in the satellite-observable displacement field.

For collapses driven by progressive, quasi-static processes of sufficient amplitude and spatial coherence, InSAR time series can plausibly serve as a monitoring input that triggers escalated inspection. The relevant processes include foundation settlement, slow landslide-induced pier displacement, and sustained differential movement. Farneti et al. (2023b) formalised this scope by developing a residual service-life prediction framework for bridges undergoing slow deformation, explicitly conditioned on the assumption that the deformation trend is representative of the dominant damage process. Cernuto et al. (2026) extended the analysis to the geometric conditions under which landslide-induced bridge deformation would be detectable from a given orbit geometry. The work provides a pre-assessment tool for determining whether an InSAR-based monitoring scheme is geometrically viable for a particular site. For collapses driven by sudden brittle failure, hidden internal degradation, or dynamic loading events, InSAR provides no warning, and the expectation that it should reflects a misunderstanding of the physical scope of the measurement. The contribution of satellite InSAR to infrastructure safety management is therefore best characterised not as early warning in the strict operational sense, but as risk-informed screening, which is the ability to identify, within a population of structures, those exhibiting kinematic anomalies that warrant priority inspection, while simultaneously documenting the structural lifecycle through TS analysis and supporting post-event stability monitoring through standardised product integration.

5.5.6. Implications for decision-ready diagnostics

The present subsection assesses, against the diagnostic framework of Section 4.6, the level achieved by each of the three completed case studies. The assessment is organised as a per-level discussion of the enabling conditions and the limiting

factors observed empirically, and it feeds into the overall discussion of Chapter 6 without duplicating its analytical scope.

Level 1, which is screening, is confirmed by all three case studies as operationally mature for structures within European territory. The EGMS Calibrated product delivers standardised, geographically consistent, and annually updated displacement data at no marginal cost. Its mean velocity map can be queried against a structure inventory to identify outliers warranting priority inspection. The *Santissima Trinità* case study illustrated the value of a negative screening result: the absence of velocity anomalies above the noise floor of the EGMS product provides a useful baseline confirmation that can inform prioritisation decisions. For the *Himera Viaduct*, the EGMS Calibrated product supported post-reconstruction stability assessment over the 2019-2023 window. The conditions for Level 1 reliability, as defined in Section 4.6, are satisfied by the EGMS product for the majority of urban structures: point density is generally sufficient for identifying anomalies at the building-block scale, velocity uncertainty is well below the thresholds of engineering concern for settlement or subsidence screening, and the GNSS-calibrated reference frame ensures cross-site consistency.

Level 2, which is displacement characterisation, requires user-driven processing and was achieved in all case studies where dedicated user-driven analysis was performed. The *Santissima Trinità* case study demonstrated that user-driven processing recovers measurement points with substantially higher temporal coherence than the corresponding EGMS product on the same structure, enabling a more reliable characterisation of the displacement behaviour at the sub-building scale. The infrastructure case studies demonstrated displacement characterisation under more challenging conditions. For the *Himera Viaduct*, the user-driven ascending and descending orbit analyses characterised the pre-collapse velocity field on the viaduct deck and identified TS whose lifecycle boundaries constrain the chronology of the structural event. For the *Albiano-Magra Bridge*, COSMO-SkyMed processing at Stripmap resolution characterised the spatial pattern of pre-collapse LOS velocities across the structural components, distinguishing between stable piers and spans exhibiting anomalous cumulative displacement. The conditions that limit Level 2 are primarily related to measurement point density and spatial distribution, as discussed in Section 5.5.2.

Level 3, which is diagnostic interpretation, was approached in two case studies but fully achieved in neither, for instructive reasons. At *Vicoforte*, the integration of satellite-derived displacement data with the geotechnical characterisation of the

subsoil, the permanent in-situ monitoring record, and the satellite-derived environmental variables enabled a causal interpretation: the observed LOS displacement at the basement PS is consistent with a seasonal compressibility cycle in the clayey-silt deposit, driven by temperature and moisture fluctuations, and superimposed on a quasi-linear consolidation trend. The coherence-aware SBL model provides a quantitative link between the environmental inputs and the displacement output, and the correlation between satellite-derived and in-situ environmental variables validates the physical plausibility of the model. Diagnostic interpretation was therefore achieved for the soil displacement component of the structural response, although the separation of soil and structural contributions remains an open modelling challenge. For the infrastructure collapse case studies, diagnostic interpretation was pursued retrospectively. The pre-collapse velocity anomalies on the *Albiano-Magra Bridge*, while spatially correlated with the failed spans, are at or below the detection margin of the technique when evaluated prospectively. The retrospective identification of a pattern consistent with the known failure mechanism does not, on its own, constitute a prospectively actionable diagnostic. The epistemological asymmetry between retrospective pattern recognition and prospective anomaly detection is a structural limitation of the satellite evidence in collapse-related scenarios, and it bounds the diagnostic level that can be claimed for these case studies at the lower end of Level 3.

The conditions that limit Level 3 are fundamentally different from those of Level 2. The constraint is not the quality or density of the satellite data, but the availability of complementary information and the complexity of the causal inference. A displacement pattern can be characterised with satellite data alone; attributing that pattern to a physical cause requires structural knowledge, geotechnical context, and, in most cases, an independent measurement against which the satellite interpretation can be tested.

Level 4, which is continuous monitoring with automated anomaly detection, was demonstrated in prototype form exclusively in the *Vicoforte* case study. The SBL surrogate model trained on the selected high-coherence basement PS, using soil temperature and moisture as environmental inputs, produces a continuous displacement prediction against which the observed displacement record is evaluated through the stochastic oscillator statistic. The oscillator normalises the satellite-derived displacement into a dimensionless monitoring variable that flags deviations from the calibrated baseline without requiring the end-user to interpret raw LOS displacement values or geometric projections. The two alert levels defined by the oscillator provide a decision-support framework that is directly compatible

with standard SHM practice. The achievement of Level 4 at *Vicoforte* was enabled by a combination of conditions that are not easily reproduced at other sites: a long and dense COSMO-SkyMed archive, a permanent in-situ monitoring system providing cross-validation data, a well-characterised geotechnical context, and a structure whose dominant displacement mechanism is environmentally driven and therefore amenable to surrogate modelling. Whether the same approach can be transferred to structures where these conditions are only partially met remains an open question, and the substitution of satellite-derived environmental variables for in-situ sensors represents one promising pathway toward broader applicability whose performance under different climatic, geotechnical, and structural conditions has not yet been systematically validated.

Taken together, the assessment across the three case studies establishes that Levels 1 and 2 have been consistently achieved, Level 3 has been reached only in the *Vicoforte* application and, with retrospective qualifications, in the infrastructure collapse analyses, and Level 4 has been reached only in *Vicoforte* under the specific data conditions of that case. The transferability of these findings to broader classes of structures, and the research directions that would extend the diagnostic reach of the approach, are addressed in Chapter 6. The case studies completed in this chapter are now examined collectively against the four-level diagnostic framework introduced in Chapter 4, with the objective of identifying the decision-ready diagnostic outputs that each case has produced and the factors that have governed the achievable diagnostic level. The synthesis of the diagnostic levels achieved by the three case studies against the four-level analytical framework of Section 4.6 is reported in Table 5. 13.

Table 5. 13. Synthesis of the diagnostic levels achieved by the three case studies of Chapter 5 against the four-level analytical framework of Section 4.6. For each level, the table indicates the case studies that achieved it, the enabling conditions, and the empirical limiting factors observed in the case study analysis.

Diagnostic level	Definition	Achieved by	Enabling conditions	Limiting factors
L1 - Screening	Identification of measurement points exhibiting velocity anomalies above the	All three case studies: <i>Santissima Trinità</i> , <i>Himera</i> , <i>Albiano-Magra</i> , <i>Vicoforte</i>	EGMS Calibrated product available for European territory; sufficient point density at building-block scale	Limited to relative kinematic identification; no causal information

Diagnostic level	Definition	Achieved by	Enabling conditions	Limiting factors
	noise floor of the satellite product			
L2 - Displacement characterisation	Quantitative description of the displacement field on a structure, including spatial pattern and temporal behaviour	All three case studies through user-driven processing	User-driven processing; high-resolution X-band data (CSK, CSK-SG); adapted coherence and amplitude criteria	Point density and spatial coverage on the structure; geometric occlusion in valley-confined sites (<i>Himera</i>)
L3 - Diagnostic interpretation	Causal interpretation of the displacement pattern in terms of a physical mechanism	<i>Vicoforte</i> (soil compressibility cycle attribution); <i>Albiano-Magra</i> (retrospective qualifications); <i>Himera</i> (slope-bound TS attributed to landslide)	Complementary information: geotechnical context, in-situ monitoring, FE model, structural inspection records	Epistemological asymmetry between retrospective pattern recognition and prospective anomaly detection; statistical robustness near detection margin
L4 - Prediction	Continuous monitoring with automated anomaly detection and uncertainty quantification, supporting alert-level decisions	<i>Vicoforte</i> (prototype demonstration through coherence-aware SBL surrogate and stochastic oscillator)	Long and dense COSMO-SkyMed archive; permanent in-situ monitoring system; well-characterised geotechnical context; environmentally-driven displacement mechanism	Transferability to sites with less favourable data conditions; basis function dictionary site-specificity; assumptions on noise structure at low coherence

5.6. Ongoing research: towards the fusion of PS-InSAR and airborne LiDAR for building-scale post-event assessment: the 2025 Palisades Fire (Los Angeles, California)

The three case studies presented in the preceding sections addressed structural scenarios within the established scope of satellite InSAR for SHM: long-term monitoring of stable or slowly deforming structures, retrospective forensic analysis of infrastructure collapses, and integrated soil-structure interaction characterisation through surrogate modelling. A fourth application, concerning rapid post-event building-scale assessment following a large-scale wildland-urban interface fire, extends the monitoring framework into a regime that the preceding architecture does not address. The research reported in this section was initiated during a visiting research period at the Qualcomm Institute, University of California San Diego, in collaboration with the Cultural Heritage Engineering Initiative team. The work is ongoing at the time of writing: the PS-InSAR component has produced preliminary results that are reported here, while the airborne LiDAR processing component, conducted by the UCSD collaborators, is at a more advanced stage of development but has not yet been integrated with the satellite radar analysis into a completed fusion framework. The present section therefore differs in character from the preceding Case Studies. It describes a research line whose methodological framework is developed in full and whose preliminary satellite-based results support the feasibility of the approach, while the final integration, classification performance assessment, and operational validation remain directions for completion. The regional context of the Palisades Fire within the January 2025 Los Angeles fire complex, with the relative timing and extent of the concurrent fires, is shown in Figure 5. 43.




Figure 5. 43. Overview of the January 2025 Los Angeles fire complex, showing the perimeters of the Palisades Fire and of the concurrent Eaton, Hurst, Lidia, and other fires, with the regional context of the burn area in relation to the Greater Los Angeles area. Adapted from (NASA Scientific Visualization Studio, 2025).

The decision to report the research at this preliminary stage, rather than to defer it to future work, reflects two considerations. The first is that the methodological framework developed for the Palisades application constitutes a concrete instantiation of the multi-modal fusion architecture introduced in Section 4.5.3 of Chapter 4, and its inclusion in the thesis consolidates the argumentative arc that connects the methodological development in Chapter 4 with the empirical validation in Chapter 5. The second is that the preliminary PS-InSAR results already demonstrate the detectability of the scatterer lifecycle signatures associated with the fire event, which is a necessary precondition for the success of the proposed fusion framework. The section is structured accordingly. Section 5.6.1 establishes the event context. Section 5.6.2 describes the study area and the datasets, including the attribution of the LiDAR processing to the UCSD team. Section 5.6.3 presents the PS-InSAR processing and the scatterer lifecycle characterisation that constitute the work completed within the thesis. Section 5.6.4 describes the conceptual framework for the building-footprint-scale fusion that the ongoing research aims to implement. Section 5.6.5 draws the preliminary implications and identifies the directions for completion.

5.6.1. Event context and motivation for multi-sensor assessment

The structural assessment challenge posed by large-scale wildland-urban interface fires differs qualitatively from the monitoring scenarios addressed in the preceding case studies. In the previous applications, the target structures were either stable or undergoing gradual, quasi-static deformation over observation windows of months to years. In a wildland-urban interface fire event, structural change is abrupt, spatially heterogeneous at the parcel scale, and conditioned by a physical mechanism (direct thermal exposure, ember transport, and secondary combustion) that operates over hours and produces outcomes ranging from complete destruction to no visible damage in adjacent buildings. The assessment problem is therefore primarily one of rapid spatial triage rather than longitudinal monitoring: given a burned area containing thousands of structures, which buildings have sustained significant structural change, and how can that determination be made from space-borne and airborne observations without requiring field access that may be restricted for weeks?

The Palisades Fire ignited at approximately 10:30 a.m. PST on 7 January 2025, in the Pacific Palisades area of Los Angeles, within the Santa Monica Mountains. Extreme Santa Ana wind conditions drove the fire from its origin to over 3,000 acres within twelve hours. Full containment was not declared until 31 January 2025, after 24 days of active suppression operations. The final burned area of approximately 23,000 acres encompasses one of the most densely developed wildland-urban interface zones in California. The CAL FIRE Damage Inspection System recorded several thousand destroyed structures and additional damaged structures from a field inspection campaign within the fire perimeter. The combination of asset density, topographic complexity, and the prevailing wind-driven ember transport mechanism produced a damage field that is spatially heterogeneous at the parcel scale: adjacent buildings on the same street frequently recorded different damage outcomes depending on construction vintage, roofing material, and the presence of combustible materials within the immediately surrounding zone.

The wildland-urban interface context creates specific challenges for post-event damage assessment that single-sensor approaches handle poorly. Optical imagery, whether satellite or aerial, is effective at detecting destroyed building footprints where structural loss is total but provides limited information about partial damage, substructural change, or the deformation state of apparently intact buildings. Thermal infrared imagery can identify hotspots and active damage zones during

and shortly after the event but provides a time-limited observation window and does not yield structural displacement information. Coherence change detection using single-pair DInSAR captures thermal deformation and surface disruption in a two-image comparison but does not exploit the multi-temporal depth of the Sentinel-1 archive and provides no information about the persistence or trajectory of post-event deformation.

Airborne LiDAR and multi-temporal PS-InSAR address complementary aspects of this assessment problem. Airborne LiDAR provides direct, three-dimensional quantification of geometric change between a pre-fire and a post-fire point cloud. A destroyed building leaves a characteristic signature in the bi-temporal normalised DSM difference: a rectangular patch of substantial negative elevation change, typically between 3 and 10 m in magnitude, corresponding to the removal of the roof structure and upper walls. Multi-temporal PS-InSAR, by contrast, provides a continuous record of LOS displacement before and after the event and encodes structural change in the lifecycle of coherent scatterers. A building that maintained stable PS throughout the pre-fire period and then loses them abruptly after the ignition date has, in effect, been identified as having undergone a radical change in its radar reflecting configuration, consistent with either destruction or major structural alteration. The fusion of these two evidence streams at the level of the individual building footprint is the methodological objective that the research reported in this section pursues.

5.6.2. Study area, satellite datasets, and complementary LiDAR material

The study area encompasses the Pacific Palisades community of Los Angeles, characterised by steep east-facing slopes and shallow canyon systems typical of the Santa Monica Mountains, with residential development occupying ridge tops, canyon floors, and the transition slopes between them. Building stock ranges from post-World War II single-storey wood-frame construction to contemporary multi-storey structures, with a substantial proportion predating California's 2008 wildland-urban interface building code amendments.

For the present analysis, we processed the Sentinel-1 constellation archive over the Los Angeles area, spanning from January 2024 to the months following the January 2025 fire. We used both ascending and descending orbit acquisitions in VV polarisation, extending the analysis beyond the descending-only geometry of an earlier preliminary version: the integration of both geometries provides the


displacement decomposition capability that the fusion framework requires. We used a Copernicus DEM at 30 m resolution for topographic phase removal. We processed the stacks in user-driven mode under the POLIMI PSI formulation in star-graph configuration, with a master image selected by optimising the amplitude stability, temporal baseline distribution, and Doppler centroid separation across the full archive. The principal characteristics of the satellite radar and airborne LiDAR datasets used in the Palisades Fire post-event assessment work are summarised in Table 5. 14.

Table 5. 14. Summary of the satellite radar and airborne LiDAR datasets used in the Palisades Fire post-event assessment work of Section 5.6.


Field	Sentinel-1 ASC	Sentinel-1 DESC	LiDAR pre- fire	LiDAR post- fire
Sensor and band	Sentinel-1A, C-band	Sentinel-1A, C-band	Airborne LiDAR (visible-NIR)	Airborne LiDAR (visible-NIR)
Source / processing	Author- controlled processing of SLC archive	Author- controlled processing of SLC archive	LARIAC baseline; processing by UCSD CHEI team	USGS 3DEP emergency response (NV5 Geospatial); processing by UCSD CHEI team
Wavelength	5.55 cm	5.55 cm	n/a (active LiDAR)	n/a (active LiDAR)
Acquisition geometry	Ascending, IW mode, VV polarisation, scene-centre incidence angle 39.12°	Descending, IW mode, VV polarisation, scene-centre incidence angle 39.10°	Airborne, single-epoch survey	Airborne, single-epoch survey
Temporal coverage	Pre- and post- fire archive (master scene 2024-2025)	Pre- and post- fire archive (master scene 2024-2025)	2023 (pre-fire baseline)	2025 (post-fire response)
Number of acquisitions	70	73	1 epoch	1 epoch

Field	Sentinel-1 ASC	Sentinel-1 DESC	LiDAR pre- fire	LiDAR post- fire
Spatial resolution / sampling	3.69 m × 13.94 m (range × azimuth)	3.69 m × 13.94 m (range × azimuth)	high-density point cloud (sub-metre)	high-density point cloud (sub-metre)
Role in the analysis	scatterer lifecycle and density-ratio derivation; FTS and PTS identification	scatterer lifecycle and density-ratio derivation; consistency check across orbits	pre-event geometric reference for bi-temporal differencing	post-event geometry for bi-temporal differencing and per- building geometric change metric

The airborne LiDAR component of the research is conducted by collaborators at the Qualcomm Institute, University of California San Diego, and specifically by Scott McAvoy and Professor Falko Kuester of the Cultural Heritage Engineering Initiative (CHEI) and ALERTCalifornia. The LiDAR processing draws on the pre-fire airborne baseline acquired by NV5 Geospatial in 2023 within the Los Angeles Regional Imagery Acquisition Consortium programme, and on the post-fire airborne survey acquired by NV5 Geospatial between 21 and 22 January 2025 with a Riegl VQ-1560 II sensor at a nominal pulse density of 16 points per square metre, sponsored by the ALERTCalifornia programme of UC San Diego and released to the USGS 3D Elevation Program for emergency response use. The 0.5 m resolution DEMs and DSMs derived from both surveys are publicly available through the USGS 3DEP data portal and OpenTopography, while the full post-fire point cloud has been provided to the research team through the UCSD-NV5 partnership. The LiDAR processing includes point cloud co-registration using the iterative closest point algorithm on stable road surface transects outside the fire perimeter, bi-temporal normalised DSM differencing, and the derivation of per-building geometric change metrics. The thesis does not reproduce the LiDAR processing details, which are the intellectual contribution of the UCSD team and will be reported in the forthcoming joint publication. The present section attributes these elements to the collaborators without further technical development. An illustrative comparison of the pre-fire and post-fire LiDAR point clouds for a representative destroyed building footprint, showing the geometric change detected by the bi-temporal differencing, is shown in Figure 5. 44.



(a)



(b)

Figure 5. 44. Pre-fire and post-fire airborne LiDAR point clouds the Pacific Palisades study area: (a) pre-fire point cloud, with the buildings geometry intact; (b) post-fire point cloud, with the substantial negative elevation change consistent with structural and

vegetation removal. Processing and visualisation by the Cultural Heritage Engineering Initiative and ALERTCalifornia team at the Qualcomm Institute, University of California San Diego.

Building footprints for the study area are derived from publicly available inventories, intersected to produce a building-level polygon layer for the area of analysis. The ground-truth damage classification is provided by the CAL FIRE Damage Inspection System, which records for each building a category (destroyed, damaged, or unaffected) from field inspection. The ground-truth inventory constitutes the validation reference against which the preliminary satellite-derived indicators developed in Section 5.6.3 are compared, and against which the future fusion classifier described in Section 5.6.4 will be evaluated. The per-parcel damage inventory for the Pacific Palisades community, with the destroyed and damaged building categories, is shown in Figure 5. 45.




Figure 5. 45. Per-parcel damage inventory for the Pacific Palisades community after the Palisades Fire of January 2025, showing the spatial distribution of buildings classified as destroyed and damaged by field inspection. Adapted from (Los Angeles County GIS, 2025).

5.6.3. Preliminary PS-InSAR processing and scatterer lifecycle characterisation

The PS-InSAR processing developed within the thesis follows the standard small-area approach introduced in Chapter 3 and applied consistently across the preceding case studies, with two adaptations specific to the post-event context. First, the full archive, encompassing both the pre-fire and post-fire temporal windows, is

processed as a single unified stack rather than as two separate datasets. The unified processing ensures that the displacement history and amplitude statistics of each scatterer candidate are estimated across the event boundary within a consistent reference frame. Second, particular attention is directed to the identification of TS whose coherence window is bounded by the fire date, following the TS characterisation methodology developed in Case study II.

We identified candidate scatterers from the mean amplitude image using a local-maxima approach, with an amplitude dispersion threshold that we calibrated to the urban residential context of Pacific Palisades, where the spatial density of stable artificial reflectors such as metallic fixtures, concrete surfaces, chimney stacks, and window frames is moderate to high. After multi-image PS-InSAR inversion, we applied a temporal coherence threshold of 0.50 to the full-stack solution to define the accepted PS population. We adopted a coherence threshold of 0.50, less stringent than the 0.60 threshold used in the structural case studies of Sections 5.2 to 5.4, on the basis of the analytical objective of the present application. The Palisades analysis seeks to characterise scatterer lifecycle patterns, including the abrupt loss of coherence associated with fire-induced destruction, and the retention of moderate-coherence scatterers in the pre-fire window allows the framework to detect the lifecycle transition with the necessary statistical resolution. A higher threshold would have systematically excluded the candidate population whose lifecycle behaviour the framework is designed to study. The threshold is consistent with practice for urban C-band monitoring and retains measurements with acceptable phase stability while excluding agricultural and vegetated surfaces that exhibit coherence values systematically below this level.

A physical mechanism specific to the fire event governs the lifecycle of the scatterers attributed to destroyed buildings and justifies the analytical emphasis on TS identification rather than continuous displacement tracking. When a building sustains abrupt structural change (collapse of the roof structure, loss of upper wall assemblies, or complete removal of the above-ground envelope by the fire), the reflecting surface hosted by that building ceases to exist as a coherent scatterer between two consecutive satellite acquisitions. If displacement during the event exceeds half the radar wavelength ($\lambda/2$, corresponding to approximately 2.78 cm for Sentinel-1 C-band), the phase change between the two acquisitions bracketing the event exceeds 2π and the interferometric phase becomes ambiguous through cycle slip. In the case of fire-induced destruction, the change is not a displacement in the physical sense but a complete loss of the scattering surface; the result is a discontinuous phase trajectory and an abrupt drop of coherence below the

acceptance threshold, manifesting as the termination of the scatterer time series. This physical mechanism explains why FTS identification is the natural analytical strategy for fire event detection: continuous displacement tracking is mathematically unsound across the event boundary, while TS lifecycle analysis exploits the discontinuity as diagnostic information.

The identification of scatterer lifecycle anomalies proceeds through an examination of the temporal coherence and amplitude time series of each candidate within the footprint of buildings classified as destroyed by CAL FIRE. Three distinct lifecycle patterns emerge from preliminary analysis of the dataset. The first pattern is the abrupt and permanent loss of coherence in the post-fire acquisitions for candidates that exhibited stable PS behaviour throughout the pre-fire period. These scatterers, designated here as Fire-induced Temporary Scatterers (FTS), show a characteristic amplitude-temporal baseline signature in which the pre-fire acquisition cluster is compact and high-amplitude while the post-fire acquisitions exhibit rapid amplitude decay and phase dispersion consistent with the loss of the coherent reflecting surface. The FTS death date aligns with the January 2025 acquisition window immediately following the fire, providing an interferometric timestamp for the structural event that is independent of the optical or thermal evidence. The second pattern is the emergence of new scatterers in the post-fire period on footprints of buildings classified as destroyed; these are attributed to debris fields, partially standing masonry walls, or temporary structures introduced during demolition operations, and they are classified as Post-event Temporary Scatterers (PTS) following the taxonomy developed in Section 5.3.3 of Case study II. The third pattern is unchanged scatterer behaviour across the fire date, observed on the footprints of intact buildings, where the pre-fire PS population is preserved in full and the displacement time series shows no significant discontinuity. Amplitude-temporal baseline scatter plots and LOS displacement time series for four representative scatterers, illustrating the contrast between the FTS and PTS lifecycle patterns, are shown in Figure 5. 46.

Two scalar indicators can be derived from the PS lifecycle analysis for each building footprint, as the building-level inputs to the fusion framework described in Section 5.6.4. The first is the pre-to-post-fire PS density ratio, defined as the count of accepted PSs in the post-fire window divided by the count in the pre-fire window, normalised per unit footprint area. A value close to 1 indicates structural preservation; a value approaching 0 indicates near-complete scatterer loss consistent with destruction. The second indicator is the FTS fraction, defined as the proportion of pre-fire PS candidates on the footprint that exhibit FTS behaviour,

namely coherence values above threshold in the pre-fire window and below threshold in the post-fire window. Taken together, these two indicators provide a two-dimensional representation of the scatterer lifecycle response at building scale that is sensitive to both the magnitude and the abruptness of the structural change encoded in the radar record. The preliminary analysis of the dataset indicates that buildings classified as destroyed by CAL FIRE exhibit density ratio values clustered towards 0 and FTS fraction values clustered towards 1, while intact buildings exhibit the inverse pattern, providing empirical support for the discriminative power of these indicators before their formal integration with the LiDAR evidence.




Figure 5. 46. Amplitude-temporal baseline plots and LOS displacement time series for representative scatterers in the Pacific Palisades study area: Fire-induced Temporary

Scatterer on a destroyed building in ascending (panel a) and descending (panel b) geometry; Post-event Temporary Scatterer appearing in the debris field after the fire in ascending (panel c) and descending (panel d) geometry. In each amplitude-temporal baseline plot, black dots represent individual amplitude measurements, while red circles indicate the mean value. In each displacement time series, red triangles represent individual LOS displacement values, and the solid red line interpolates the displacement trend.

5.6.4. Conceptual framework for multi-modal fusion at the building-footprint scale

The completion of the research requires the integration of the PS-InSAR lifecycle indicators developed in Section 5.6.3 with the LiDAR-derived building change metrics produced by the UCSD collaborators. The fusion is conceived at the level of the individual building footprint, which constitutes the natural analytical unit for post-event damage assessment and resolves the spatial resolution mismatch between the two sensor types. A single Sentinel-1 resolution cell may encompass parts of two or three buildings on the residential lots typical of Pacific Palisades, and a single LiDAR-derived change map pixel at 0.5 m resolution is too fine for direct comparison with a scatterer point. Aggregating both sensor streams to the building footprint level resolves the spatial resolution mismatch and produces a per-building feature space that is consistent and interpretable in engineering terms.

The spatial attribution of PS to building footprints will follow a reverse geocoding procedure in which the building footprint polygons are projected into the SAR slant-range geometry to define the expected spatial extent of scatterers associated with each building (Schneider & Soergel, 2021). The projection accounts for the layover effect, which displaces the apparent SAR position of scatterer candidates on elevated surfaces towards the sensor relative to their true ground position, and ensures that scatterers on roof surfaces and upper facades are correctly attributed to the underlying building rather than to an adjacent footprint. Geolocation refinement using the co-registered LiDAR DTM as a supplementary height reference is expected to reduce the mean residual PS offset to sub-metre level, consistent with the accuracy reported for C-band sensors in the literature (Dheenathayalan et al., 2016). All PS candidates whose corrected geocoded position falls within the buffered building boundary are retained as attributable to that footprint.

The per-building feature vector that the fusion will assemble comprises the two PS-InSAR lifecycle indicators introduced in Section 5.6.3 (pre-to-post-fire PS

density ratio and FTS fraction) and the LiDAR-derived geometric change metrics produced by the UCSD team. The latter comprise the mean elevation change over the footprint area, the standard deviation of elevation change capturing the spatial variability of change within the footprint, the volume change proxy estimating the total volume of structure removed, and the footprint coverage fraction indicating the areal extent of significant elevation loss relative to the building footprint area. The six-dimensional feature space is designed to capture the complementary physical mechanisms of the two sensors (geometry change in the LiDAR dimensions and electromagnetic lifecycle change in the PS-InSAR dimensions) while remaining sparse enough to avoid overfitting given the size of the labelled building inventory.

The fusion itself is conceived as a supervised classification problem trained on the subset of the building inventory with CAL FIRE ground-truth damage labels. The classifier architecture selected for the framework is a Random Forest formulation, motivated by three considerations: its robustness to feature scale differences, which avoids the need for normalisation of the heterogeneous LiDAR and PS-InSAR metrics; its native estimation of feature importance through out-of-bag permutation, which supports interpretable physical analysis of the relative contribution of each sensor to the classification outcome; and its established performance in analogous multi-source damage classification tasks in the earthquake domain (Hajeb et al., 2020; Macchiarulo et al., 2025). The classifier will be evaluated against single-sensor baselines (LiDAR-only and PS-InSAR-only) to quantify the discriminative gain of the fusion. The discriminative gain is expected to concentrate in the damaged class, where single-sensor classifiers exhibit the highest confusion rates and where the complementary evidence from the two sensors is most diagnostically informative.

5.6.5. Preliminary implications and directions for completion

The preliminary work reported in this section supports three conclusions that can be stated with reasonable confidence at the present stage of the research. The first is that the fire event produces a distinctive scatterer lifecycle signature that is detectable in the Sentinel-1 archive and that aligns spatially with the CAL FIRE ground-truth damage inventory. The signature manifests primarily as fire-induced TS on destroyed buildings, whose pre-fire coherent history terminates abruptly at the event date, and as post-event TS on debris fields and temporary works. The signature is diagnostically meaningful at the per-building scale, providing empirical support for the feasibility of a building-level analysis framework. The second

conclusion is that the physical mechanism that governs the fire-induced scatterer response, which is the loss of the coherent reflecting surface rather than a continuous displacement exceeding the phase ambiguity threshold, justifies the analytical emphasis on TS lifecycle analysis rather than on continuous displacement tracking. This mechanism is consistent across the three TS types identified in Case study II and in the Palisades analysis. The third conclusion is that the two scalar indicators derived from the PS-InSAR lifecycle analysis (pre-to-post-fire PS density ratio and FTS fraction) exhibit preliminary discriminative patterns between destroyed and intact buildings that support their inclusion in a multi-modal fusion framework.

Three directions for completion frame the remaining research. The first is the finalisation of the LiDAR processing by the UCSD collaborators and the formal integration of the LiDAR-derived metrics with the PS-InSAR indicators into the per-building feature vector described in Section 5.6.4. The second is the training and evaluation of the Random Forest fusion classifier against the CAL FIRE ground-truth inventory, together with the single-sensor baselines for comparison. The third is the feature importance analysis, which is expected to quantify the relative contribution of the two sensor modalities to the classification outcome and to inform the design of operational workflows that balance data acquisition cost against diagnostic performance.

Several methodological challenges will constrain the direct transfer of the approach to other wildland-urban interface fire events or to contexts where the data conditions differ from those of the Palisades case. The PS density constraint associated with C-band Sentinel-1 on single-storey wood-frame residential construction is a fundamental physical limitation: buildings that host fewer than two or three pre-fire scatterers contribute limited lifecycle information to the density ratio and FTS fraction metrics. In areas with older SAR archives where the pre-fire time series is shorter than three to four years, the statistical confidence of the scatterer lifecycle characterisation is reduced. The level of detection of the bi-temporal LiDAR analysis is adequate for the detection of full or near-full structural destruction but may be insufficient for the reliable detection of minor damage categories. The supervised classification approach requires labelled ground-truth data for training, and the CAL FIRE inventory, while comprehensive, uses a coarser classification resolution than structural engineers would require for detailed post-event assessment.

The operational value of the completed framework, once these directions are addressed, will consist in the capability to produce per-building damage classifications within the time window consistent with emergency management requirements. The post-fire LiDAR acquisition was completed within two weeks of containment, and the Sentinel-1 acquisitions bracketing the fire are available as standard Copernicus open-access products within 24 hours of overpass. The processing chain, from co-registered point clouds and calibrated SLC stacks to a per-building damage classification, is reproducible and scalable. The main operational bottleneck will be the LiDAR acquisition itself, which depends on airborne platform availability and airspace access in the post-event period. Where pre-existing LiDAR baselines are available (as is increasingly the case in California through the LARIAC and 3DEP programmes), the barrier reduces to post-fire survey logistics.

At a broader level, the Palisades application extends the monitoring architecture developed in the preceding case studies into a post-event regime. The standard SHM use case for satellite InSAR, as demonstrated in the three completed case studies, is longitudinal monitoring of structures under operational loads, with the objective of detecting gradual deterioration or anomalous deformation trends against a stable pre-event baseline. The Palisades case inverts this logic: the pre-event archive characterises the normal scattering behaviour of the building stock, and the post-event observations are interpreted as departures from that baseline caused by a sudden, large-scale structural event. The TS lifecycle framework, developed in the context of individual infrastructure failures in Case study II, proves equally applicable at the scale of a neighbourhood-wide disaster event, with the distinction that the building inventory replaces the individual structure as the unit of analysis. The completion of the fusion framework proposed in this section would operationalise this extension and provide a concrete instantiation of the multi-modal fusion architecture introduced in Section 4.5.3.

Chapter 6

Discussion and conclusions

The work concludes with a chapter that draws together the methodological developments of Chapters 3 and 4 and the empirical findings of Chapter 5. The level of abstraction is distinct from that of Section 5.5. Section 5.5 examines how each completed case study performs against the four-level diagnostic framework and articulates the conditions under which each level becomes achievable. The present chapter operates one step removed from the per-case empirical assessment. Methodological contributions, disciplinary implications, scope conditions, and future directions are treated here as objects in their own right.

Six sections compose the chapter. Section 6.1 restates the four methodological contributions and qualifies the level of confidence with which each is offered to the discipline. Table 6. 1 synthesises the contributions and serves as the navigation reference for the rest of the chapter. Section 6.2 closes the research-programme loop opened in Chapter 1 by responding to each of the six research questions formulated in Section 1.4. Section 6.3 articulates the limitations of the work and the threats to validity that bound the generalisability of the findings. Section 6.4 derives audience-specific recommendations for engineering practice. Section 6.5 identifies the future research directions that the work opens. Section 6.6 closes the thesis with a brief restatement of the decision-ready diagnostics commitment that gives the work its title.

Table 6. 1. Synthesis of the four methodological contributions developed in this thesis. For each contribution, the table reports the section in which it is introduced, the case study or case studies that test it empirically, the diagnostic level of the four-level framework to

which it primarily relates, and the level of confidence with which the contribution is offered to the discipline.

Contribution	Description	Empirical testing	Primary level	Confidence
C1 Formalisation of the Temporary Scatterer (TS) as a systematic diagnostic category	Generalisation of the PS framework to admit time-bounded coherence windows as a legitimate analytical state, with explicit lifecycle indices T_{on} and T_{off} and three operational sub-types (Section 4.1.2).	Case study II (<i>Himera Viaduct, Albiano-Magra Bridge</i>); Section 5.6 Palisades application (Fire-induced TS, Post-event TS).	L2 to L3	Established methodological extension. Cross-case empirical evidence on three independent applications.
C2 Coherence-aware Sparse Bayesian Learning (SBL) surrogate modelling framework	Treatment of temporal coherence as a design variable rather than a quality filter, integrated within an SBL surrogate that links satellite displacement to environmental drivers and identifies an application-specific coherence threshold from predictive uncertainty (Section 4.5 for general principles; Section 5.4 for the methodological development and the case-specific implementation).	Case study III (<i>Sanctuary of Vicoforte</i>): empirical coherence threshold of 0.56 from predictive uncertainty; stochastic oscillator monitoring statistic over a multi-year COSMO-SkyMed Second Generation archive. Case study I (<i>Santissima Trinità</i>): supporting evidence through the reinterpretation of P24 advanced in Section 5.5.	L3 to L4	Demonstrated on a single case under favourable data conditions (Vicoforte). Transferability across sites and structural typologies remains to be validated.
C3 Four-level diagnostic framework: from screening to continuous monitoring	Articulation of the diagnostic value of satellite InSAR observations across four progressively more demanding levels (screening, displacement characterisation, diagnostic interpretation, continuous monitoring), with	All three completed case studies (<i>Santissima Trinità, Himera Viaduct and Albiano-Magra Bridge, Sanctuary of Vicoforte</i>) and the ongoing Palisades application. Cross-case mapping in Table 5. 13 (Section 5.5.6).	L1 to L4	Established structural framework. Empirically tested on a heterogeneous case-study sample covering masonry heritage, infrastructure collapse, and monumental buildings.

Contribution	Description	Empirical testing	Primary level	Confidence
	explicit conditions for the achievability of each level (Section 4.6).			
C4 Multi-modal fusion architecture for post-event building-scale damage assessment	Architecture for the integration of PS-InSAR scatterer lifecycle indicators with airborne LiDAR bi-temporal geometric change metrics into per-building feature vectors for post-event damage classification (Section 4.5.3).	Section 5.6: 2025 Palisades Fire ongoing application. Methodological framework developed and preliminary indicators tested; classifier training and validation pending.	L3	Methodological proposal. Empirical validation in progress through the ongoing collaboration with the Qualcomm Institute, UC San Diego.

6.1. The four methodological contributions

The empirical and methodological work of this thesis converges on four contributions offered to the structural engineering and remote sensing communities as transferable instruments. The four contributions are not equally established by the empirical evidence of Chapter 5. Each is qualified below at the level of confidence that the evidence supports. Differentiation is itself a methodological choice, intended to match the strength of each claim to the strength of the empirical base on which it rests.

Contribution C1 is the formalisation of the Temporary Scatterer (TS) as a systematic diagnostic category within the multi-temporal interferometric processing chain. Generalisation of the Persistent Scatterer (PS) framework admits time-bounded coherence windows as a legitimate analytical state. Operationalisation of the category proceeds through the lifecycle indices T_{on} and T_{off} that bound the live interval of each scatterer within the observation period. Three operational sub-types correspond to scatterers that emerge during the observation window, scatterers that disappear before its end, and scatterers whose entire lifecycle falls within the window. Empirical evidence for the diagnostic value of the category comes from Case study II of Section 5.3. Lifecycle boundaries of scatterers identified on the *Himera Viaduct* and on the *Albiano-Magra Bridge* encode the chronology of structural events at the resolution of the satellite revisit cycle. The Palisades Fire application of Section 5.6 extends the diagnostic range of

the category into the Wildland-Urban Interface (WUI) domain through the Fire-induced (FTS) and Post-event (PTS) sub-types. Contribution C1 is offered as an established methodological extension of the persistent scatterer framework, with cross-case empirical evidence on three independent applications.

Contribution C2 is the coherence-aware Sparse Bayesian Learning (SBL) surrogate modelling framework. General principles of the framework are articulated in Section 4.5. The methodological development and the case-specific implementation are concentrated in Section 5.4. Treatment of temporal coherence as a design variable, rather than as a passive quality filter, characterises the framework. Integration of the coherence-aware perspective within a surrogate model links satellite line-of-sight (LOS) displacement to environmental driver variables. Identification of an application-specific coherence threshold proceeds through the predictive uncertainty of the surrogate, providing a principled alternative to the application of conventional thresholds adopted from the regional-scale literature. At *Vicoforte*, the framework identified an empirical threshold of 0.56 from the predictive uncertainty of the calibrated surrogate. The same case operationalised the satellite displacement record as a Stochastic Oscillator (SO) monitoring statistic with continuous prediction capability over the multi-year COSMO-SkyMed Second Generation (CSG) archive. The *Santissima Trinità* case study contributes a complementary perspective on the same theme through the reinterpretation of the European Ground Motion Service (EGMS) point P24, advanced as a working hypothesis in Section 5.5. The hypothesis suggests that coherence values in the intermediate range may carry information about the materiality of the reflecting surface, beyond their role as quality indicators of the displacement estimate at that specific point. Contribution C2 is offered as a framework empirically demonstrated on a single case under favourable data conditions, with supporting evidence from a second case. Transferability to sites with less favourable data conditions, with different soil-structure interaction regimes, or with different basis function dictionaries has not yet been validated.

Contribution C3 is the four-level diagnostic framework introduced in Section 4.6 under the heading “Toward decision-ready diagnostics: from displacement to engineering judgement”. The articulation covers four progressively more demanding levels: screening, displacement characterisation, diagnostic interpretation, and continuous monitoring with automated anomaly detection. Each level is associated with explicit conditions for its achievability. The conditions concern the data, the processing, and the complementary information requirements that govern progression between levels. Empirical validation rests on the cross-case

mapping reported in Table 5. 13 of Section 5.5.6. Levels 1 and 2 are consistently achieved across the three completed case studies, namely the *Santissima Trinità*, the *Himera Viaduct* and *Albiano-Magra Bridge*, and the *Sanctuary of Vicoforte*. Level 3 is reached only at *Vicoforte* and, with retrospective qualifications, in the infrastructure collapse analyses. Level 4 is demonstrated in prototype form exclusively in the *Vicoforte* case. Contribution C3 is offered as an established structural framework for the discipline, empirically tested on a heterogeneous case-study sample that covers masonry heritage, infrastructure collapse, and monumental buildings.

Contribution C4 is the multi-modal fusion architecture for post-event building-scale damage assessment. Methodological development takes place in Section 4.5.3 under the heading “Fusion with non-SAR remote sensing data”. Preliminary empirical testing is concentrated in the ongoing Palisades Fire application of Section 5.6. The architecture integrates Persistent Scatterer Interferometry (PSI) scatterer lifecycle indicators with airborne Light Detection and Ranging (LiDAR) bi-temporal geometric change metrics. Per-building feature vectors are produced by the integration and feed a damage classification stage. The methodological framework is fully developed. Preliminary indicators have been computed on the Sentinel-1 archive of the Pacific Palisades study area. The airborne LiDAR processing component is in advanced development through the collaboration with the Cultural Heritage Engineering Initiative (CHEI) and ALERTCalifornia at the Qualcomm Institute of UC San Diego. Contribution C4 is offered as a methodological proposal whose empirical validation is in progress. Completion of the LiDAR processing, formal integration of the LiDAR-derived metrics with the persistent scatterer indicators, training of the Random Forest fusion classifier against the CAL FIRE ground-truth inventory, and feature importance analysis are the operations that will consolidate the contribution into a complete empirical validation.

Differentiation in confidence levels among the four contributions is the structural feature of the present section. The reader, the reviewer, and the practitioner who encounters the work in the future are entitled to know which methodological claims are supported by cross-case empirical evidence, which are demonstrated on a single case under specific conditions with supporting evidence from additional cases, which are established as structural frameworks supported by the entire empirical corpus, and which are advanced as methodological proposals in progress. Differentiation allows the chapter to articulate ambitious claims for the work without overstating the empirical base on which each claim rests.

6.2. Response to the research questions

Section 1.4 of the present thesis articulated six research questions that defined the scope and the analytical focus of the work. Closure of the research programme requires a precise statement of the answer that the empirical evidence supports for each question. The present section provides the closure.

Research question RQ1 asked under what conditions and within what boundaries satellite Interferometric Synthetic Aperture Radar (InSAR) observations can yield structurally meaningful information at the building scale. Building-scale structural interpretation becomes feasible when four conditions are jointly satisfied. The first is the availability of a sufficient number of Measurement Points (MPs) geocoded to identifiable structural elements. Point density typically requires high-resolution X-band data at the scale of single structural components. C-band data prove sufficient at the scale of urban blocks or compact heritage buildings. The second is the achievability of temporal coherence values appropriate to the deformation regime under analysis. The achievability depends on the materiality of the reflecting surfaces, on the urban context, and on the stability of the interferometric geometry. The third is the resolution of the geometric attribution ambiguities that arise from the layover effect, from the barycentre uncertainty of the geocoded position, and from the residual height estimate of each persistent scatterer. The fourth is the availability of complementary information sources sufficient to support the causal interpretation of the displacement signal at the level of diagnostic ambition pursued by the analysis. Where any condition fails, the analysis remains valid but is bounded to a lower diagnostic level than the one originally targeted.

Research question RQ2 asked how standardised products such as the European Ground Motion Service compare quantitatively against user-driven processing at the building scale, in terms of point density, measurement quality, and structural interpretability. The comparison is structure-specific, and no general superiority can be established between the two strategies. At the *Santissima Trinità* in Turin, the EGMS Calibrated product retained four MPs on the structure with coherence values between 0.66 and 0.86. The user-driven Sentinel-1 processing retained two MPs with coherence values of 0.86 and 0.92. The apparent point-density advantage of the standardised product conceals a coherence cost that the user-driven processing systematically recovers. At the *Sanctuary of Vicoforte*, the EGMS product provided approximately 350 points distributed across the dome and the bell towers. The user-driven COSMO-SkyMed Second Generation processing focused on the basement-

level scatterers with elevation filtering between minus one metre and plus five metres. The user-driven approach provided the spatially selective measurement set required for the soil displacement objective that the standardised product could not have supported in isolation. The two strategies are therefore complementary within a single decision framework. The choice between them is governed by the monitoring objective and by the structural characteristics of the target.

Research question RQ3 asked whether temporal coherence can be treated as a diagnostic variable rather than as a binary quality filter. The affirmative answer is supported through two distinct demonstrations. The first is the *Vicoforte* case study, where temporal coherence is incorporated into the Sparse Bayesian Learning surrogate as a design variable. An application-specific coherence threshold of 0.56 is identified empirically through the predictive uncertainty of the calibrated model. The threshold is therefore data-driven rather than adopted from the regional-scale literature. The second demonstration is the reinterpretation of the EGMS point P24 at the *Santissima Trinità*, advanced as a working hypothesis in Section 5.5. The reinterpretation suggests that coherence values in the intermediate range may carry information about the materiality of the reflecting surface, beyond their role as quality indicators of the displacement estimate at that specific point. The two demonstrations together support the methodological claim that coherence carries diagnostic content at multiple levels of the interpretive chain. Binary treatment of coherence as a quality gate represents an avoidable loss of information for structurally informed analyses.

Research question RQ4 asked whether satellite-derived environmental variables can substitute for in-situ sensors in data-driven displacement models. The answer is partially affirmative. The *Vicoforte* case study has demonstrated that the integration of satellite-derived environmental variables, specifically Land Surface Temperature and the Soil Water Index, within the Sparse Bayesian Learning surrogate produces a calibrated model with predictive performance comparable to that achievable with in-situ environmental sensors, under the specific conditions of the *Vicoforte* site. The substitution potential is therefore demonstrated empirically on a single case. Transferability to sites where the soil and structure characteristics differ from those of *Vicoforte*, where the climatic regime is different, or where the dominant displacement mechanism is not environmentally driven has not been validated within the scope of the present work and remains an open research direction.

Research question RQ5 asked whether Temporary Scatterers can be used to reconstruct the chronology and the spatial extent of structural failures from archival satellite data. The affirmative answer is supported through the empirical demonstrations of Case study II. At the *Albiano-Magra Bridge*, the disappearance of scatterers on the failed deck spans, identified as Temporary Scatterers with T_{off} coinciding with the collapse epoch, encodes the spatial extent of the failure at the resolution of the satellite revisit cycle. At the *Himera Viaduct*, the lifecycle boundaries of scatterers on the slope-adjacent area identified the post-landslide debris field as a coherent target during the live interval following the slope failure. The Palisades Fire application of Section 5.6 extends the demonstration to the wildland-urban interface domain. Fire-induced Temporary Scatterers and Post-event Temporary Scatterers carry diagnostic information about the per-building damage outcomes. The Temporary Scatterer category therefore functions as a diagnostic instrument for the retrospective chronology of structural events. The qualification is that the temporal resolution of the chronology is bounded by the satellite revisit interval, and that the visibility of relevant scatterers depends on the radar geometry of the structure within its surrounding terrain.

Research question RQ6 asked whether persistent scatterer interferometry and airborne LiDAR can be fused at the building scale for post-event damage classification. The methodologically affirmative answer is qualified by the empirical incompleteness of the demonstration. The fusion architecture introduced in Section 4.5.3 is fully developed. Preliminary indicators have been computed on the Pacific Palisades study area. The airborne LiDAR processing component is in advanced development through the collaboration with the Qualcomm Institute of UC San Diego. Empirical validation of the fusion against the CAL FIRE ground-truth inventory, training of the Random Forest classifier, and feature importance analysis are the operations that will consolidate the affirmative answer into a complete demonstration. Methodological viability of the fusion has been established. Preliminary empirical evidence supports its further development. Full validation belongs to the trajectory of the ongoing research that the work opens.

6.3. Limitations and threats to validity

The conclusions articulated in the preceding sections are bounded by six classes of limitations that must be acknowledged before the findings are extended to contexts beyond those examined here. The acknowledgement is structured around the principle that limitations are scope conditions of the contributions rather than defects that weaken them. Explicit identification of the conditions under which each

contribution holds is the appropriate operational response to the question of generalisability.

The first class concerns the intrinsic physical limitations of the satellite measurement. Satellite InSAR measures the projection of the target displacement onto the LOS direction determined by the satellite orbit and the incidence angle. The LOS measurement is a scalar constraint on a three-dimensional displacement vector rather than a direct measurement of any single displacement component. Decomposition into vertical and east-west horizontal components requires the availability of ascending and descending acquisitions with sufficient MPs on the same structural elements, a condition that is not always satisfied at the building scale. The north-south horizontal component remains inaccessible from the near-polar orbital geometry of current SAR missions. Temporal sampling is bounded by the satellite revisit cycle, which has been variable over the observation window of the present work owing to the loss of Sentinel-1B on 23 December 2021 and the subsequent recovery through the launch of Sentinel-1C in December 2024. These constraints bound the kinematic resolvability of structural displacement processes whose characteristic timescales fall below the revisit interval.

The second class concerns the data conditions required for the full achievement of the diagnostic framework. The *Vicoforte* case study demonstrated the most complete integration of the framework and reached the Level 4 diagnostic capability. Several factors made this achievement possible. The first was a permanent in-situ monitoring system. The second was a well-characterised geotechnical context. The third was a long and dense COSMO-SkyMed archive. The fourth was a structure whose dominant displacement mechanism is environmentally driven and therefore amenable to surrogate modelling. These conditions are not easily reproduced at other sites. Transferability of the coherence-aware SBL framework to structures where these conditions are only partially met remains an open research direction.

The third class concerns the epistemological limitations of retrospective analysis. The class is most directly relevant to the infrastructure collapse case studies of Section 5.3. Identification of pre-collapse kinematic anomalies in an archival SAR record, even when the anomalies are spatially correlated with the failed portion of the structure, does not on its own establish that the anomalies would have been detectable prospectively. Nor does it establish that their detection would have been actionable within an operational monitoring programme. Prospective validation of the proposed anomaly detection frameworks against

structures that subsequently develop genuine anomalies is a methodologically necessary step before any operational claim of early-warning capability can be made. The validation is difficult because the base rate of structural failures in monitored infrastructure populations is low, and the required observation windows are long.

The fourth class concerns the limitations of the ongoing research direction reported in Section 5.6. The multi-modal fusion framework for post-event building-scale damage assessment is developed methodologically but not yet fully implemented or validated. Conclusions drawn from the Palisades Fire application at the current stage of the research are preliminary. Integration with the airborne LiDAR processing component, training of the fusion classifier against the ground-truth damage inventory, and feature importance analysis remain to be completed. Preliminary empirical results reported in the chapter support the methodological viability of the fusion but do not establish its quantitative performance.

The fifth class concerns the case-study sample size. The work is supported by three completed case studies and one ongoing application. The sample covers masonry heritage, infrastructure collapse, monumental buildings, and post-event damage assessment. The sample is heterogeneous and covers a substantial portion of the range of structural typologies for which satellite-assisted SHM is currently relevant. The sample is finite, however, and the generalisations supported by the empirical base must be qualified accordingly. The TS category (C1) is supported by two completed case studies and one ongoing application. The four-level diagnostic framework (C3) is supported by all four cases. The coherence-aware SBL framework (C2) is supported primarily by *Vicoforte*, with complementary evidence from the *Santissima Trinità*. The multi-modal fusion architecture (C4) is in progress. The qualification of the confidence level in Section 6.2 reflects this distribution.

The sixth class concerns the user-driven processing assumption that underpins the methodological framework. The contributions have been developed and tested under conditions of full user-driven control over the satellite processing chain. User-driven control includes selection of the master image, definition of the interferometric network, choice of the reference point, and calibration of the coherence and amplitude thresholds. Performance of the framework under the more constrained conditions of standardised products such as the EGMS has been characterised through cross-validation in Case studies I and II. The framework has not been adapted, however, to operate natively on standardised products without

user-driven supplementation. Extension of the framework to the standardised-product regime is a methodological development that the present work does not address and that future work may pursue.

6.4. Recommendations for engineering practice

The empirical evidence assembled in the present work supports a set of recommendations addressed to three audiences within the engineering practice community. The first audience consists of analysts and researchers who conduct the satellite data processing and interpretation. The second consists of structural engineers, asset managers, and decision-makers who consume the output of the satellite analysis and integrate it into their monitoring and maintenance workflows. The third consists of institutional actors, including national maintenance authorities, civil protection agencies, and regulatory bodies, who govern the integration of satellite-based monitoring into the broader infrastructure management ecosystem. The recommendations distinguish between those advanced as original contributions of the work (specifically Recommendation A2 on coherence-aware threshold identification and Recommendation X1 on structured reporting) and those that consolidate consensus already present in the literature with the new empirical evidence assembled in the case studies (Recommendations A1, A3, S0, S1, S2, S3, I1).

Recommendations for the analyst audience.

Recommendation A1, governance of processing strategy. The choice between standardised products and user-driven processing should be governed by the monitoring objective and by the structural characteristics of the target, not by default availability considerations. Standardised products such as the EGMS are appropriate for portfolio-scale screening and for structures where the monitoring objective is limited to confirming stability or to flagging large-amplitude velocity anomalies. User-driven processing is necessary for building-scale structural interpretation, for the identification of TS, and for applications where the coherence threshold or the reference frame must be adapted to the specific deformation regime.

Recommendation A2, treatment of temporal coherence. Coherence should not be treated as a binary quality criterion with a conventional threshold of 0.5 or 0.6 applied uniformly across all applications. The appropriate threshold depends on the monitoring objective, on the physical nature of the reflecting surfaces, and on the

statistical framework within which the displacement data are interpreted. Where a data-driven modelling framework is employed, as in the *Vicoforte* case study, the coherence threshold can be identified empirically through the predictive uncertainty of the model, providing a principled alternative to convention.

Recommendation A3, attribution of MPs to structural elements. The barycentre effect, the layover geometry, and the residual height uncertainty of geocoded PSs contribute to attribution ambiguities that must be resolved before any displacement interpretation can be drawn. Reverse geocoding, in which building footprint polygons or structural element geometries are projected into the SAR slant-range reference frame, is the appropriate operational response to this class of ambiguity.

Recommendations for the structural engineer audience.

Recommendation S0, decision to commission a satellite InSAR analysis. Satellite InSAR is most appropriately commissioned for structures and contexts where four operational conditions are jointly satisfied. The first is the relevance of slow, quasi-static displacement processes to the structural concern at hand, including foundation settlement, soil-structure interaction effects, slope-induced motion, and post-event geometric change. The second is the limitation of physical access for conventional in-situ instrumentation, characteristic of heritage buildings, decommissioned infrastructure, and post-event assessment scenarios. The third is the availability of a sufficient archival depth of satellite acquisitions for the structural typology and geographical region of interest. The fourth is the structural geometry compatibility with the satellite LOS, evaluated through the geometric observability framework discussed in Chapter 4. Integration of satellite InSAR into the standard repertoire of Structural Health Monitoring (SHM) techniques is appropriate when these conditions converge.

Recommendation S1, proportionality between diagnostic level and engineering consequence. The output of a satellite InSAR analysis should be interpreted in the light of the diagnostic level that the analysis achieves. Engineering consequences drawn from the output should be proportionate to that level. A screening-level analysis at Level 1 establishes whether a structure warrants further attention; it does not establish a structural diagnosis, and it should not be treated as such. A Level 2 analysis characterises the displacement behaviour of the structure; the interpretation of that behaviour in causal structural terms requires the integration of complementary information at Level 3.

Recommendation S2, active integration with existing monitoring infrastructure. Integration of satellite data with existing monitoring infrastructure should be pursued actively rather than treated as an optional enhancement. Combination of satellite displacement data with in-situ monitoring, numerical structural models, geotechnical characterisation, or satellite-derived environmental variables multiplies the diagnostic reach of either source alone. The combination enables the achievement of Level 3 and Level 4 under conditions that neither source can support in isolation.

Recommendation S3, exploitation of the temporal dimension as retrospective baseline. The temporal dimension of the satellite record should be exploited to establish reference states against which future monitoring observations can be compared. Even when a current satellite analysis yields no actionable finding, the multi-year displacement record it produces constitutes a retrospective baseline whose value grows as the monitoring history accumulates. Archival value of satellite data is one of the principal distinguishing characteristics of the technique relative to conventional monitoring approaches.

Cross-cutting recommendation.

Recommendation X1, structured reporting of satellite InSAR results for structural monitoring. Reporting of satellite InSAR results for structural monitoring applications should adopt a structured format. The format must explicitly state the diagnostic level achieved, the conditions that enabled that level, the limitations that bound the interpretation, and the uncertainties associated with the displacement estimates. The structured reporting format implemented in the case study chapters supports the communication of satellite InSAR results to non-specialist engineering audiences. The format also reduces the risk of over-interpretation or under-utilisation of the output.

Recommendation for the institutional audience.

Recommendation I1, integration of satellite-based monitoring into national infrastructure management. National maintenance authorities and civil protection agencies should consider the integration of standardised satellite InSAR products into the operational workflows of infrastructure portfolio management. Bridge stocks, dam systems, and heritage building inventories all benefit from such integration. Three engineering and institutional developments would support the integration. The first is the development of standardised interface protocols between

satellite processing platforms and SHM management software. The second is the definition of structured reporting templates that communicate the diagnostic level and its enabling conditions. The third is the establishment of training pathways that bridge the radar-specific expertise required for user-driven processing and the structural engineering expertise required for the interpretation of the output. The Italian ReLUIS guidelines on the use of interferometric satellite data for the interpretation of structural behaviour represent a precedent on which national-scale standardisation may build.

6.5. Future research directions

The methodological framework developed in the present work and the empirical evidence assembled through the case studies open four research directions. Each direction is paired with the methodological contribution to which it primarily relates. Pairing follows the architecture of Section 6.2 and reinforces the argumentative coherence of the chapter.

The first direction concerns the completion of the multi-modal fusion architecture (C4) through the Palisades Fire application of Section 5.6. Several operations will consolidate the contribution into a complete empirical validation. They include the finalisation of the airborne LiDAR processing by the collaborators at the Qualcomm Institute of the University of California San Diego, the formal integration of the LiDAR-derived metrics with the persistent scatterer interferometry indicators into a per-building feature vector, the training and evaluation of the Random Forest fusion classifier against the CAL FIRE ground-truth inventory, and the feature importance analysis that quantifies the relative contribution of the two sensor modalities to the classification outcome. The trajectory of the ongoing collaboration with the Cultural Heritage Engineering Initiative (CHEI) and ALERTCalifornia supports the timely completion of the validation in the period immediately following the present work.

The second direction concerns the prospective validation of the coherence-aware SBL surrogate (C2) developed in the *Vicoforte* case study. The stochastic oscillator-based alert framework has been validated at *Vicoforte* over a multi-year dataset without the occurrence of structural anomalies. The validation demonstrated that the system operates stably in the null-hypothesis condition. Testing of the protocol against a structure that subsequently develops a genuine anomaly, whether through natural structural evolution or through a controlled experimental campaign,

is the essential next step for the demonstration of the prospective monitoring capability of the framework.

The third direction concerns the extension of the TS category (C1) and of the four-level diagnostic framework (C3) to multi-constellation and multi-geometry data. The current case studies employed Sentinel-1 C-band, COSMO-SkyMed X-band first generation (CSK), or COSMO-SkyMed X-band Second Generation data, generally in a single orbit geometry. Combined exploitation of multiple constellations in multiple orbit geometries would enable three-dimensional displacement reconstruction at the structural element scale. The reconstruction would resolve the scalar LOS limitation that currently bounds the structural interpretation. Expansion of the Copernicus constellation through Sentinel-1C and Sentinel-1D, continued accumulation of the COSMO-SkyMed Second Generation archive, and the emergence of additional X-band and L-band platforms in the international satellite community will make multi-sensor fusion increasingly feasible in the coming years.

The fourth direction concerns the development of standardised interface protocols between satellite processing platforms and SHM management software, supporting the implementation of the structured reporting recommended in Recommendation X1. The expertise barrier that currently limits the uptake of user-driven InSAR workflows in structural engineering practice is partly a consequence of the lack of such protocols. Their absence forces the structural engineer to interact directly with raw or semi-processed satellite products whose interpretation requires radar-specific expertise. Standardisation of output formats, definition of structured reporting templates that communicate the diagnostic level and its enabling conditions, and development of training pathways that bridge the two communities are engineering and institutional developments. Pursuit of these developments through coordinated action involving satellite data providers, academic research groups, and engineering practitioners represents a substantial opportunity for the structural monitoring community.

6.6. Closing of the thesis

The transition of satellite InSAR from a research instrument into an operational tool for the structural health monitoring of existing structures is a trajectory rather than a destination. Continuation of the trajectory depends on three convergent developments. The first is the empirical validation of the diagnostic frameworks against broader classes of structures and events. The second is the engineering

development of the interfaces and protocols that bring the capability to the practising engineer. The third is the institutional integration of satellite-based monitoring into national infrastructure management. The work presented in the preceding chapters offers four methodological contributions to that trajectory. The contributions are summarised in Table 6. 1 and articulated through the case studies of Chapter 5: the formalisation of the TS as a systematic diagnostic category, the coherence-aware SBL surrogate modelling framework, the four-level diagnostic articulation that maps satellite observables onto decision-support categories, and the multi-modal fusion architecture for post-event building-scale damage assessment.

The decision-ready diagnostic capability that gives the work its title is not a property of the satellite data alone, nor of any specific processing chain, nor of any single interpretive framework. It is the property of an integrated practice. The practice combines the processing, the interpretation, the cross-validation, and the structured reporting required to bring satellite evidence to the level of analytical maturity at which it can support an actionable engineering judgement. The present work offers a contribution to the construction of that integrated practice and leaves the further development to the trajectory of the discipline.

Bibliography

- Adams, R. P., & MacKay, D. J. C. (2007). Bayesian Online Changepoint Detection. *arXiv preprint*. <https://arxiv.org/abs/0710.3742>
- Airbus built forest monitoring satellite Biomass successfully launched. (2025, aprile 29). <https://www.airbus.com/en/newsroom/press-releases/2025-04-airbus-built-forest-monitoring-satellite-biomass-successfully>
- Airbus Defence and Space. (2022). *Copernicus Digital Elevation Model Product Handbook*. Copernicus.
- Alaska Satellite Facility. (2026a). *Introduction to SAR - HyP3*. https://hyp3-docs.asf.alaska.edu/guides/introduction_to_sar/
- Alaska Satellite Facility. (2026b). *Sentinel-1 InSAR Product Guide*. https://hyp3-docs.asf.alaska.edu/guides/insar_product_guide/
- Alba, M., Barazzetti, L., Scaioni, M., Rosina, E., & Previtali, M. (2011). Mapping Infrared Data on Terrestrial Laser Scanning 3D Models of Buildings. *Remote Sensing*, 3(9), 1847–1870. <https://doi.org/10.3390/rs3091847>
- Aoki, T., Chiorino, M. A., & Roccati, R. (2003). Structural Characteristics of the Elliptical Masonry Dome of the Sanctuary of Vicoforte. *Proceedings of the First International Congress on Construction History*, 203–212.
- Avanzi, F., Munerol, F., Milelli, M., Gabellani, S., Massari, C., Giroto, M., Cremonese, E., Galvagno, M., Bruno, G., Morra di Cella, U., Rossi, L., Altamura, M., & Ferraris, L. (2024). Winter snow deficit was a harbinger of summer 2022 socio-hydrologic drought in the Po Basin, Italy. *Communications Earth & Environment*, 5(1), 64. <https://doi.org/10.1038/s43247-024-01222-z>
- Bado, M. F., & Casas, J. R. (2021). A Review of Recent Distributed Optical Fiber Sensors Applications for Civil Engineering Structural Health Monitoring. *Sensors*, 21(5), 1818. <https://doi.org/10.3390/s21051818>
- Bähr, H., & Hanssen, R. F. (2012). Reliable estimation of orbit errors in spaceborne SAR interferometry. *Journal of Geodesy*, 86(12), 1147–1164. <https://doi.org/10.1007/s00190-012-0571-6>

- Balageas, D., Fritzen, C.-P., & Güemes, A. (2010). *Structural Health Monitoring*. John Wiley & Sons.
- Balaji, D. P. (2024, maggio 16). SAR Series Part 7: SAR Interferometry (InSAR) Quick Reference Guide. *Medium*. <https://medium.com/@preet.balaji20/sar-series-part-7-sar-interferometry-insar-quick-reference-guide-7689681cd3d3>
- Bamler, R., & Hartl, P. (1998). Synthetic aperture radar interferometry. *Inverse Problems*, 14(4), R1. <https://doi.org/10.1088/0266-5611/14/4/001>
- Barrias, A., Casas, J. R., & Villalba, S. (2016). A Review of Distributed Optical Fiber Sensors for Civil Engineering Applications. *Sensors*, 16(5), 748. <https://doi.org/10.3390/s16050748>
- Berardino, P., Fornaro, G., Lanari, R., & Sansosti, E. (2002). A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms. *IEEE Transactions on Geoscience and Remote Sensing*, 40(11), 2375–2383. <https://doi.org/10.1109/TGRS.2002.803792>
- Berti, M., Corsini, A., Franceschini, S., & Iannaccone, J. P. (2013). Automated classification of Persistent Scatterers Interferometry time series. *Natural Hazards and Earth System Sciences*, 13(8), 1945–1958. <https://doi.org/10.5194/nhess-13-1945-2013>
- Biondi, F., Addabbo, P., Ullo, S. L., Clemente, C., & Orlando, D. (2020). Perspectives on the Structural Health Monitoring of Bridges by Synthetic Aperture Radar. *Remote Sensing*, 12(23). <https://doi.org/10.3390/rs12233852>
- Biondini, F., & Frangopol, D. M. (2016). Life-Cycle Performance of Deteriorating Structural Systems under Uncertainty: Review. *Journal of Structural Engineering*, 142(9), F4016001. [https://doi.org/10.1061/\(ASCE\)ST.1943-541X.0001544](https://doi.org/10.1061/(ASCE)ST.1943-541X.0001544)
- Bonaldo, G., Caprino, A., Lorenzoni, F., & Porto, F. da. (2023). Monitoring Displacements and Damage Detection through Satellite MT-InSAR Techniques: A New Methodology and Application to a Case Study in Rome (Italy). *Remote Sensing*, 15(5). <https://doi.org/10.3390/rs15051177>
- Boncori, M., & Peter, J. (2019). Measuring Coseismic Deformation With Spaceborne Synthetic Aperture Radar: A Review. *Frontiers in Earth Science*, 7. <https://doi.org/10.3389/feart.2019.00016>

- Boscardin, M. D., & Cording, E. J. (1989). Building Response to Excavation-Induced Settlement. *Journal of Geotechnical Engineering*, 115(1), 1–21. [https://doi.org/10.1061/\(ASCE\)0733-9410\(1989\)115:1\(1\)](https://doi.org/10.1061/(ASCE)0733-9410(1989)115:1(1))
- Braun, A. (2019). *Radar satellite imagery for humanitarian response. Bridging the gap between technology and application*. <https://doi.org/10.15496/publikation-32698>
- Brownjohn, J. M. W. (2007). Structural health monitoring of civil infrastructure. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, 365(1851), 589–622. <https://doi.org/10.1098/rsta.2006.1925>
- Bürgmann, R., Rosen, P. A., & Fielding, E. J. (2000). Synthetic Aperture Radar Interferometry to Measure Earth's Surface Topography and Its Deformation. *Annual Review of Earth and Planetary Sciences*, 28, 169–209. <https://doi.org/10.1146/annurev.earth.28.1.169>
- Burland, J., Broms, B., & Mello, V. (1977). *Behaviour of foundations and structures*.
- Calderini, C., Chiorino, M. A., Lagomarsino, S., & Spadafora, A. (2006). Non-linear Modelling of the Elliptical Dome of Vicoforte. *V International Conference on Structural Analysis of Historical Constructions (SAHC 2006)*, 2, 1177–1186.
- Calò, M., Ruggieri, S., Nettis, A., & Uva, G. (2023). Multi source Interferometry Synthetic Aperture Radar for monitoring existing bridges: A case study. *Ce/Papers*, 6(5), 787–793. <https://doi.org/10.1002/cepa.2058>
- Calò, M., Ruggieri, S., Nettis, A., & Uva, G. (2024). *BMS-INSAR: GIS PLUGIN FOR BRIDGE PORTFOLIO SERVICE DISPLACEMENT ANALYSIS THROUGH MTINSAR DATA*.
- Capella. (2024, agosto 20). *Capella Space Reveals First Light Imagery from Capella-13 “A Sky Full of SARs” Mission | Capella*. <https://www.capellaspace.com/resources/capella-space-reveals-first-light-imagery-from-capella-13-a-sky-full-of-sars-mission>
- Capes, R., Teeuw, R., & Pilz, T. (2023). European Ground Motion Service calibrated product validation and accuracy assessment. *Remote Sensing*, 15(14), 3588. <https://doi.org/10.3390/rs15143588>
- Caspani, V. F., Tonelli, D., Bado, M. F., Rocca, A., Perissin, D., & Zonta, D. (2026). Bridge monitoring with satellite InSAR technology and clustering

- p algorithm.
- Measurement*
- , 258, 118939.
-
- <https://doi.org/10.1016/j.measurement.2025.118939>
- Cavalagli, N., Kita, A., Falco, S., Trillo, F., Costantini, M., & Ubertini, F. (2019). Satellite radar interferometry and in-situ measurements for static monitoring of historical monuments: The case of Gubbio, Italy. *Remote Sensing of Environment*, 235, 111453. <https://doi.org/10.1016/j.rse.2019.111453>
- Ceravolo, R., Coletta, G., Miraglia, G., & Palma, F. (2021). Statistical correlation between environmental time series and data from long-term monitoring of buildings. *Mechanical Systems and Signal Processing*, 152, 107460. <https://doi.org/10.1016/j.ymssp.2020.107460>
- Ceravolo, R., De Lucia, G., Miraglia, G., & Pecorelli, M. L. (2020). *Thermoelastic finite element model updating with application to monumental buildings*. <https://doi.org/10.1111/mice.12516>
- Ceravolo, R., Lenticchia, E., Miraglia, G., & Scussolini, L. (2024). Improving the dynamic behaviour of historic buildings using experimental data: Application to a Baroque church. *Journal of Civil Structural Health Monitoring*. <https://doi.org/10.1007/s13349-024-00804-x>
- Cernuto, E., Salciarini, D., Ubertini, F., & Giardina, G. (2025). Landslide-Bridge Interaction: A combined approach based on InSAR data and numerical modelling. *International Journal of Disaster Risk Reduction*, 126, 105568. <https://doi.org/10.1016/j.ijdr.2025.105568>
- Cernuto, E., Salciarini, D., Ubertini, F., & Giardina, G. (2026). Assessing InSAR observability of landslides interfering with bridges. *Scientific Reports*. <https://doi.org/10.1038/s41598-026-41011-6>
- Chen, C. W., & Zebker, H. A. (2000). Network approaches to two-dimensional phase unwrapping: Intractability and two new algorithms. *JOSA A*, 17(3), 401–414. <https://doi.org/10.1364/JOSAA.17.000401>
- Chiorino, M. A., Spadafora, A., Calderini, C., & Lagomarsino, S. (2008). Modeling Strategies for the World's Largest Elliptical Dome at Vicoforte. *International Journal of Architectural Heritage*, 2(3), 274–303. <https://doi.org/10.1080/15583050802063618>
- Cigna, F., Esquivel Ramírez, R., & Tapete, D. (2021). Accuracy of Sentinel-1 PSI and SBAS InSAR Displacement Velocities against GNSS and Geodetic Leveling Monitoring Data. *Remote Sensing*, 13(23), 4800. <https://doi.org/10.3390/rs13234800>

- Cigna, F., Lasaponara, R., Masini, N., Milillo, P., & Tapete, D. (2014). Persistent scatterer interferometry processing of COSMO-SkyMed StripMap HIMAGE time series to depict deformation of the historic centre of Rome, Italy. *Remote Sensing*, 6(12), 12593–12618. <https://doi.org/10.3390/rs61212593>
- Coccimiglio, S. (2025). *Integration of Satellite Data into health monitoring protocols for full-scale structures* [Italy]. <https://doi.org/10.13121/polito/porto/3001911>
- Coccimiglio, S., Coletta, G., Dabdoub, M., Miraglia, G., Lenticchia, E., & Ceravolo, R. (2022). Use of Copernicus Satellite Data to Investigate the Soil-Structure Interaction and Its Contribution to the Dynamics of A Monitored Monumental Building. In C. Pellegrino, F. Faleschini, M. A. Zanini, J. C. Matos, J. R. Casas, & A. Strauss (A c. Di), *Proceedings of the 1st Conference of the European Association on Quality Control of Bridges and Structures* (pp. 1171–1179). Springer International Publishing. https://doi.org/10.1007/978-3-030-91877-4_133
- Coccimiglio, S., Miraglia, G., & Ceravolo, R. (2025). Adaptive Hyperparameters Selection for Modal Tracking Algorithm in Structural Health Monitoring of Masonry Monumental Structures. *18th International Brick and Block Masonry Conference (IB2MaC 2024), Lecture Notes in Civil Engineering*. https://doi.org/10.1007/978-3-031-73314-7_83
- Coccimiglio, S., Miraglia, G., & Coletta, G. (2022). Combining satellite geophysical data with continuous on-site measurements for monitoring the dynamic parameters of civil structures. *Scientific Reports*. <https://doi.org/10.1038/s41598-022-06284-7>
- Coccimiglio, S., Miraglia, G., Coletta, G., Epicoco, R., & Ceravolo, R. (2023). Balanced Definition of Thresholds for Mode Tracking in a Long-Term Seismic Monitoring System. *Geosciences*, 13(12), 365. <https://doi.org/10.3390/geosciences13120365>
- Coccimiglio, S., Scussolini, L., Matteini, I., Ceravolo, R., & Ferro, G. A. (2024). Interferometric Satellite Data For The Structural Health Monitoring Of Infrastructures. *Procedia Structural Integrity, II Fabre Conference – Existing bridges, viaducts and tunnels: research, innovation and applications (FABRE24)*, 62, 840–847. <https://doi.org/10.1016/j.prostr.2024.09.113>

- Coccimiglio, S., Tarantini, R., Miraglia, G., Matteini, I., Scussolini, L., Ceravolo, R., & Ferro, G. A. (2026). Analysis of multi-source satellite-derived displacement data for structural monitoring of masonry heritage buildings. *Procedia Structural Integrity, XX ANIDIS Conference*, 78, 1032–1039. <https://doi.org/10.1016/j.prostr.2025.12.132>
- Coletta, G., Miraglia, G., Pecorelli, M., Ceravolo, R., Cross, E., Surace, C., & Worden, K. (2019). Use of the cointegration strategies to remove environmental effects from data acquired on historical buildings. *Engineering Structures*, 183, 1014–1026. <https://doi.org/10.1016/j.engstruct.2019.01.085>
- Colin, E. (2022, novembre 18). *What is the difference between GRD and GRD + Terrain correction?* Medium. <https://elisecolin.medium.com/what-is-the-difference-between-grd-and-grd-terrain-rectification-a2ca4fd20ce>
- Costantini, M., Minati, F., Trillo, F., Ferretti, A., Novali, F., Passera, E., Dehls, J., Larsen, Y., Marinkovic, P., Eineder, M., Brcic, R., Siegmund, R., Kotzerke, P., Probeck, M., Kenyeres, A., Proietti, S., Solari, L., & Andersen, H. S. (2021). European Ground Motion Service (EGMS). *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 3293–3296. <https://doi.org/10.1109/IGARSS47720.2021.9553562>
- Covello, F., Battazza, F., Coletta, A., Lopinto, E., Fiorentino, C., Pietranera, L., Valentini, G., & Zoffoli, S. (2010). COSMO-SkyMed an existing opportunity for observing the Earth. *Journal of Geodynamics*, 49(3–4), 171–180. <https://doi.org/10.1016/j.jog.2009.12.001>
- Crosetto, M., Crippa, B., Mróz, M., Cuevas-González, M., & Shahbazi, S. (2025). Applications based on EGMS products: A review. *Remote Sensing Applications: Society and Environment*, 37, 101452. <https://doi.org/10.1016/j.rsase.2025.101452>
- Crosetto, M., Monserrat, O., Ferro, A., & Serra, M. (2016). Persistent Scatterer Interferometry: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 115, 78–89. <https://doi.org/10.1016/j.isprsjprs.2015.10.011>
- Cusson, D., Trischuk, K., Hébert, D., Hewus, G., Gara, M., & Ghuman, P. (2018). Satellite-Based InSAR Monitoring of Highway Bridges: Validation Case Study on the North Channel Bridge in Ontario, Canada. *Transportation Research Record*, 2672(45), 76–86. <https://doi.org/10.1177/0361198118795013>

- Dabiri, H., Marini, R., Clementi, J., Mazzanti, P., Mugnozza, G. S., Bozzano, F., & Bompa, D. (2025). Monitoring buildings performance using FEA and ML based on the data acquired by InSAR; A case study of Vittoriano building, Rome. *Structures*, 74, 108643. <https://doi.org/10.1016/j.istruc.2025.108643>
- De Lucia, G., Lenticchia, E., & Ceravolo, R. (2020). La geometria della cupola della Ss. Trinità: Dal dato metrico alle valutazioni strutturali. In M. Ruffino (A c. Di), *Una chiesa per il Ducato. La Ss. Trinità di Torino* (pp. 117–142). C.L.U.T. Scrl Editrice. <https://iris.polito.it/handle/11583/2782952>
- De Lucia, G., Lenticchia, E., & Ceravolo, R. (2021). Geometric and Structural Information for the Analysis of Historical Domes: The Case of the SS. Trinità Church in Torino. *Proceedings of the International Conference on Structural Analysis of Historical Constructions*. <https://doi.org/10.23967/sahc.2021.121>
- Delo, G., Civera, M., Lenticchia, E., Miraglia, G., Surace, C., Ceravolo, R., Delo, G., Civera, M., Lenticchia, E., Miraglia, G., Surace, C., & Ceravolo, R. (2022). Interferometric Satellite Data in Structural Health Monitoring: An Application to the Effects of the Construction of a Subway Line in the Urban Area of Rome. *Applied Sciences*, 12(3). <https://doi.org/10.3390/app12031658>
- Deslauriers, B., Del Conte, S., & Fiaschi, S. (2025, agosto 30). *Leveraging InSAR Technology to Optimize Lidar Data Collection*. LIDAR Magazine. <https://lidarmag.com/2025/08/30/leveraging-insar-technology-to-optimize-lidar-data-collection/>
- Dheenathayalan, P., Small, D., Schubert, A., & Hanssen, R. F. (2016). High-precision positioning of radar scatterers. *Journal of Geodesy*, 90(5), 403–422. <https://doi.org/10.1007/s00190-015-0883-4>
- Di Carlo, F., Miano, A., Giannetti, I., Mele, A., Bonano, M., Lanari, R., Meda, A., & Prota, A. (2021). On the integration of multi-temporal synthetic aperture radar interferometry products and historical surveys data for buildings structural monitoring. *Journal of Civil Structural Health Monitoring*, 11(5), 1429–1447. <https://doi.org/10.1007/s13349-021-00518-4>
- Di Mucci, V. M., Cardellicchio, A., Ruggieri, S., Nettis, A., Reno, V., & Uva, G. (2024). *Artificial intelligence in structural health management of existing bridges*. <https://doi.org/10.1016/j.autcon.2024.105719>

- Doebling, S. W., Farrar, C., Prime, M., & Shevitz, D. W. (1996). Damage Identification and Health Monitoring of Structural and Mechanical Systems From Changes in their Vibration Characteristics: A Literature Review. *Technical Report No. LA-13070-MS*, 30. <https://doi.org/10.2172/249299>
- Dong, C.-Z., & Catbas, F. N. (2021). A review of computer vision–based structural health monitoring at local and global levels. *Structural Health Monitoring*, 20(2), 692–743. <https://doi.org/10.1177/1475921720935585>
- Elsisi, A., Zamrawi, A., & Emad, S. (2025). A Comprehensive Review of Structural Health Monitoring for Steel Bridges: Technologies, Data Analytics, and Future Directions. *Applied Sciences*, 15(22), 12090. <https://doi.org/10.3390/app152212090>
- EO59 Ltd. (2026). *SarProZ*© (Versione 24-Nov-2025 14:33:45) [Software]. <https://www.sarproz.com/>
- Esteves Martins, J., Caro Cuenca, M., Davids, B., Zhai, G., González-Alonso, E., Boncori, J. P. M., Kristersen, L., Sala, J., & Solari, L. (2024). Validation of the European Ground Motion Service with Global Navigation Satellite Systems and Corner Reflectors over Deforming Areas. *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*. <https://doi.org/10.1109/IGARSS53475.2024.10641305>
- European Space Agency. (2016). *Sentinel-1 Product Definition*. ESA, European Space Agency, Paris, France.
- Even, M., & Schulz, K. (2020). InSAR Deformation Analysis with Distributed Scatterers: A Review Complemented by New Advances. *Remote Sensing*, 12(5), 744. <https://doi.org/10.3390/rs12050744>
- Farneti, E., Cavalagli, N., Costantini, M., Trillo, F., Minati, F., Venanzi, I., & Ubertini, F. (2023a). A method for structural monitoring of multispan bridges using satellite InSAR data with uncertainty quantification and its pre-collapse application to the Albiano-Magra Bridge in Italy. *Structural Health Monitoring*, 22(1), 353–371. <https://doi.org/10.1177/14759217221083609>
- Farneti, E., Cavalagli, N., Costantini, M., Trillo, F., Minati, F., Venanzi, I., & Ubertini, F. (2023b). Remote monitoring to predict bridge scour failure using Interferometric Synthetic Aperture Radar (InSAR) stacking techniques. *International Journal of Applied Earth Observation and Geoinformation*, 118, 103279. <https://doi.org/10.1016/j.jag.2023.103279>

- Farrar, C. R., Dervilis, N., & Worden, K. (2025). The Past, Present and Future of Structural Health Monitoring: An Overview of Three Ages. *Strain*, 61(1), e12495. <https://doi.org/10.1111/str.12495>
- Farrar, C. R., & Worden, K. (2007). An Introduction to Structural Health Monitoring. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 365(1851), 303–315. <https://doi.org/10.1098/rsta.2006.1928>
- Farrar, C., & Worden, K. (2013). Structural Health Monitoring A Machine Learning Perspective. In *Structural Health Monitoring: A Machine Learning Perspective*. <https://doi.org/10.1002/9781118443118>
- Fattahi, H., & Amelung, F. (2014). InSAR uncertainty due to orbital errors. *Geophysical Journal International*, 199, 549–560. <https://doi.org/10.1093/gji/ggu276>
- Feroz, S., & Abu Dabous, S. (2021). UAV-Based Remote Sensing Applications for Bridge Condition Assessment. *Remote Sensing*, 13(9), 1809. <https://doi.org/10.3390/rs13091809>
- Ferretti, A., Fumagalli, A., Novali, F., Prati, C., Rocca, F., & Rucci, A. (2011). A new algorithm for processing interferometric data-stacks: SqueeSAR. *IEEE Transactions on Geoscience and Remote Sensing*, 49(9), 3460–3470. <https://doi.org/10.1109/TGRS.2010.2087026>
- Ferretti, A., Monti-Guarnieri, A., Prati, C., Rocca, F., & Massonnet, D. (2007). *InSAR Principles: Guidelines for SAR Interferometry Processing and Interpretation* (TM-19). European Space Agency. https://www.esa.int/esapub/tm/tm19/TM-19_ptA.pdf
- Ferretti, A., Passera, E., & Capes, R. (2023). *End-to-end implementation and operation of the European Ground Motion Service. Algorithm Theoretical Basis Document*. <https://land.copernicus.eu/en/products/european-ground-motion-service>
- Ferretti, A., Prati, C., & Rocca, F. (2001). Permanent scatterers in SAR interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 39(1), 8–20. <https://doi.org/10.1109/36.898661>
- Fiaschi, S., & Wdowinski, S. (2020). Local land subsidence in Miami Beach (FL) and Norfolk (VA) and its contribution to flooding hazard in coastal communities along the U.S. Atlantic coast. *Ocean & Coastal Management*, 187, 105078. <https://doi.org/10.1016/j.ocecoaman.2019.105078>

- Fornaro, G., Reale, D., & Verde, S. (2013). Bridge Thermal Dilation Monitoring With Millimeter Sensitivity via Multidimensional SAR Imaging. *IEEE Geoscience and Remote Sensing Letters*, 10(4), 677–681. <https://doi.org/10.1109/LGRS.2012.2218214>
- Franceschetti, G., & Lanari, R. (2018). *Synthetic Aperture Radar Processing*. CRC Press. <https://doi.org/10.1201/9780203737484>
- Garro, M. (1962). *Santuario Basilica di Mondovì presso Vicoforte, Cuneo—Opere di consolidamento e restauro—Technical Report [In Italian]*.
- Garthwaite, M. C. (2017). On the Design of Radar Corner Reflectors for Deformation Monitoring in Multi-Frequency InSAR. *Remote Sensing*, 9(7), 648. <https://doi.org/10.3390/rs9070648>
- Giardina, G., Milillo, P., DeJong, M. J., Perissin, D., & Milillo, G. (2019). Evaluation of InSAR monitoring data for post-tunnelling settlement damage assessment. *Structural Control and Health Monitoring*, 26(2), e2285. <https://doi.org/10.1002/stc.2285>
- Giordano, P. F., Kwapisz, M., Miano, A., Liuzzo, R., Vorwagner, A., Limongelli, M. P., Prota, A., & Ralbovsky, M. (2025). Monitoring of a multi-span prestressed concrete bridge using satellite interferometric data and comparison with on-site sensor results. *Structural Concrete*. <https://doi.org/10.1002/suco.202400881>
- Goldstein, R. M., Zebker, H. A., & Werner, C. L. (1988). Satellite radar interferometry: Two-dimensional phase unwrapping. *Radio Science*, 23(4), 713–720. <https://doi.org/10.1029/RS023i004p00713>
- Grandin, R. (2015, marzo). Interferometric Processing of SLC Sentinel-1 TOPS Data. *FRINGE'15: Advances in the Science and Applications of SAR Interferometry and Sentinel-1 InSAR Workshop, Frascati, Italy, 23-27 March 2015*. <https://doi.org/10.5270/Fringe2015.pp116>
- Grebby, S., Sowter, A., Gluyas, J., Toll, D., Gee, D., Athab, A., & Girindran, R. (2021). Advanced analysis of satellite data reveals ground deformation precursors to the Brumadinho Tailings Dam collapse. *Communications Earth & Environment*, 2(1), 2. <https://doi.org/10.1038/s43247-020-00079-2>
- Guzmán-Acevedo, G. A., Vázquez-Becerra, G. E., Quintana, J. R., Carrion Viramontes, F. J., & Anaya-Díaz, M. (2024). Risk-based management of bridge portfolios integrating satellite radar products. *Structure and Infrastructure Engineering*.

- Hajeb, M., Karimzadeh, S., & Fallahi, A. (2020). Seismic damage assessment in Sarpole-Zahab town (Iran) using synthetic aperture radar (SAR) images and texture analysis. *Natural Hazards*, 103(1), 347–366. <https://doi.org/10.1007/s11069-020-03991-0>
- Hanssen, R. F. (2001). *Radar Interferometry: Data Interpretation and Error Analysis* (Vol. 2). Springer Netherlands. <https://doi.org/10.1007/0-306-47633-9>
- Hooper, A., Zebker, H., Segall, P., & Kampes, B. (2004). A new method for measuring deformation on volcanoes and other natural terrains using InSAR persistent scatterers. *Geophysical Research Letters*, L23611. <https://doi.org/10.1029/2004GL021737>
- Hu, F., Wu, J., Chang, L., & Hanssen, R. F. (2019). Incorporating Temporary Coherent Scatterers in Multi-Temporal InSAR Using Adaptive Temporal Subsets. *IEEE Transactions on Geoscience and Remote Sensing*, 57(10), 7658–7670. <https://doi.org/10.1109/TGRS.2019.2915658>
- Huang, Y., & Beck, J. L. (2015). Hierarchical Sparse Bayesian Learning for Structural Health Monitoring with Incomplete Modal Data. *International Journal for Uncertainty Quantification*, 5(2), 139–169. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2015011808>
- Introduction to SAR—ArcGIS Pro | Documentation.* (2026). <https://pro.arcgis.com/en/pro-app/latest/help/analysis/image-analyst/introduction-to-synthetic-aperture-radar.htm>
- Jónsson, S. (2002). Fault Slip Distribution of the 1999 Mw 7.1 Hector Mine, California, Earthquake, Estimated from Satellite Radar and GPS Measurements. *Bulletin of The Seismological Society of America - BULL SEISMOL SOC AMER*, 92, 1377–1389. <https://doi.org/10.1785/0120000922>
- Just, D., & Bamler, R. (1994). Phase statistics of interferograms with applications to synthetic aperture radar. *Applied Optics*, 33(20), 4361–4368. <https://doi.org/10.1364/AO.33.004361>
- Kotzerke, R., Siegmund, R., & Langenwaller, J. (2022). *End-to-end implementation and operation of the European Ground Motion Service. Product User Manual.* <https://land.copernicus.eu/en/products/european-ground-motion-service>

- Lanari, R., Reale, D., Bonano, M., Verde, S., Muhammad, Y., Fornaro, G., Casu, F., & Manunta, M. (2020). Comment on “Pre-Collapse Space Geodetic Observations of Critical Infrastructure: The Morandi Bridge, Genoa, Italy” by Milillo et al. (2019). *Remote Sensing*, 12(24), 4011. <https://doi.org/10.3390/rs12244011>
- Lazecky, M., Perissin, D., Bakon, M., de Sousa, J. M., Hlavacova, I., & Real, N. (2015). Potential of satellite InSAR techniques for monitoring of bridge deformations. *2015 Joint Urban Remote Sensing Event (JURSE)*, 1–4. <https://doi.org/10.1109/JURSE.2015.7120497>
- Lenticchia, E., Miraglia, G., & Ceravolo, R. (2021). *Exploring Problems and Prospective of Satellite Interferometric Data for the Seismic Structural Health Monitoring of Existing Buildings and Architectural Heritage*.
- Lenticchia, E., Miraglia, G., Cusicanqui Lopez, J. A., & Ceravolo, R. (2025). Sensitivity driven model updating: A multi-step procedure for structural assessment. *Journal of Building Engineering*, 109, 1–23. <https://doi.org/10.1016/j.jobbe.2025.112941>
- Lenticchia, E., Miraglia, G., De Lucia, G., & Ceravolo, R. (2023). Towards the Seismic Monitoring of a Monumental Structure in Mixed Masonry-RC. *19th ANIDIS Conference on Seismic Engineering in Italy*, 1514–1521.
- Lorenz, R., Petryna, Y., Lubitz, C., Lang, O., & Wegener, V. (2024). Thermal deformation monitoring of a highway bridge: Combined analysis of geodetic and satellite-based InSAR measurements with structural simulations. *Journal of Civil Structural Health Monitoring*, 14(5), 1237–1255. <https://doi.org/10.1007/s13349-024-00779-9>
- Los Angeles County GIS. (2025). *2025 LA Fires Damage Assessment Map*. <https://lacounty.maps.arcgis.com/apps/instant/sidebar/index.html?appid=49859fb251d6422cbe8750501f6655b1>
- Lu, P., Casagli, N., Catani, F., & Tofani, V. (2022). Persistent Scatterers Interferometry Hotspot and Cluster Analysis (PSI-HCA) for detection of extremely slow-moving landslides. *International Journal of Remote Sensing*, 35(2), 466–485. <https://doi.org/10.1080/01431161.2013.871084>
- Lynch, J. P. (2006). An overview of wireless structural health monitoring for civil structures. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851), 345–372. <https://doi.org/10.1098/rsta.2006.1932>

- Macchiarulo, V., Giardina, G., DeJong, M., Milillo, P., & others. (2022). Multi-Temporal InSAR for Transport Infrastructure Monitoring: Recent Trends and Challenges. *Proceedings of the Institution of Civil Engineers – Bridge Engineering*, 176(2), 92–104. <https://doi.org/10.1680/jbren.21.00039>
- Macchiarulo, V., Giardina, G., Milillo, P., Aktas, Y. D., & Whitworth, M. R. Z. (2025). Integrating post-event very high resolution SAR imagery and machine learning for building-level earthquake damage assessment. *Bulletin of Earthquake Engineering*, 23, 5021–5047. <https://doi.org/10.1007/s10518-024-01877-1>
- Macchiarulo, V., Milillo, P., Blenkinsopp, C., & Giardina, G. (2022). Monitoring deformations of infrastructure networks: A fully automated GIS integration and analysis of InSAR time-series. *Structural Health Monitoring*, 21(4), 1849–1878. <https://doi.org/10.1177/14759217211045912>
- Mardanshahi, A., Sreekumar, A., Yang, X., Barman, S. K., & Chronopoulos, D. (2025). Sensing Techniques for Structural Health Monitoring: A State-of-the-Art Review on Performance Criteria and New-Generation Technologies. *Sensors*, 25(5), 1424. <https://doi.org/10.3390/s25051424>
- Martins, J. E., Caro Cuenca, M., Davids, B., Zhai, G., González-Alonso, E., Boncori, J. P. M., Kristersen, L., Sala, J., & Solari, L. (2024). Validation of the European Ground Motion Service with Global Navigation Satellite Systems and Corner Reflectors over Deforming Areas. *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*. <https://doi.org/10.1109/IGARSS53475.2024.10641305>
- Massonnet, D., & Feigl, K. L. (1998). Radar interferometry and its application to changes in the Earth's surface. *Reviews of Geophysics*, 36(4), 441–500. <https://doi.org/10.1029/97RG03139>
- Meyer, F. (2019). Spaceborne synthetic aperture radar: Principles, data access, and basic processing techniques. *Synthetic aperture radar (SAR) handbook: Comprehensive methodologies for forest monitoring and biomass estimation*, 21–64.
- Milillo, P., Giardina, G., DeJong, M. J., Perissin, D., & Milillo, G. (2018). Multi-Temporal InSAR Structural Damage Assessment: The London Crossrail Case Study. *Remote Sensing*, 10(2), 287. <https://doi.org/10.3390/rs10020287>

- Milillo, P., Giardina, G., Perissin, D., Milillo, G., Coletta, A., & Terranova, C. (2019). Pre-Collapse Space Geodetic Observations of Critical Infrastructure: The Morandi Bridge, Genoa, Italy. *Remote Sensing*, 11(12). <https://doi.org/10.3390/rs11121403>
- Monserat, O., Crosetto, M., Cuevas-González, M., & Crippa, B. (2011). The Thermal Expansion Component of Persistent Scatterer Interferometry Observations. *Geoscience and Remote Sensing Letters, IEEE*, 8, 864–868. <https://doi.org/10.1109/LGRS.2011.2119463>
- Monserat, O., Crosetto, M., & Luzi, G. (2014). A review of ground-based SAR interferometry for deformation measurement. *ISPRS Journal of Photogrammetry and Remote Sensing*, 93, 40–48. <https://doi.org/10.1016/j.isprsjprs.2014.04.001>
- Montanari, A., Nguyen, H., Rubineti, S., Ceola, S., Galelli, S., Rubino, A., & Zanchettin, D. (2023). Why the 2022 Po River drought is the worst in the past two centuries. *Science Advances*, 9(32), eadg8304. <https://doi.org/10.1126/sciadv.adg8304>
- Moreira, A., Prats-Iraola, P., Younis, M., Krieger, G., Hajnsek, I., & Papathanassiou, K. (2013). *A Tutorial on Synthetic Aperture Radar* [Dataset].
- Mottershead, J. E., & Friswell, M. I. (1993). Model Updating In Structural Dynamics: A Survey. *Journal of Sound and Vibration*, 167(2), 347–375. <https://doi.org/10.1006/jsvi.1993.1340>
- NASA Scientific Visualization Studio. (2025). *Overview Map of the January 2025 Los Angeles Fires*. <https://svs.gsfc.nasa.gov/5568/>
- Nettis, A., Massimi, V., Nutricato, R., Nitti, D. O., & others. (2023). On the Use of Satellite-Based Interferometry for Structural Monitoring of Bridges. *Remote Sensing Applications: Society and Environment*, 29, 100887.
- Ng, A. H.-M. (2010). *Advanced satellite radar interferometry for ground surface displacement detection* [Thesis, UNSW Sydney]. <https://doi.org/10.26190/unsworks/23699>
- Novali, F., Fumagalli, A., Rocca, F., Prati, C., Ferretti, A., & Rucci, A. (2011). Exploitation of temporary coherent scatterers in SqueeSAR analyses. *Proceedings of the Fringe 2011 Workshop*.

- Parizzi, A., Cong, X., & Eineder, M. (2020). First results from multifrequency interferometry: A comparison of different decorrelation time constants at L, C, and X band. *ESA ENVISAT Symposium*.
- Parizzi, A., Perissin, D., Prati, C., Rocca, F., & Ferretti, A. (2006). *Accurate DEM Reconstruction from Permanent Scatterers and Multi-baseline Interferometry* (p. 160). <https://doi.org/10.1109/IGARSS.2006.45>
- Pecorelli, M., Ceravolo, R., & Epicoco, R. (2020). An automatic modal identification procedure for the permanent dynamic monitoring of the Sanctuary of Vicoforte. *International Journal of Architectural Heritage*, 630–644. <https://doi.org/10.1080/15583058.2018.1554725>
- Perissin, D., & Rocca, F. (2006). High accuracy urban DEM using Permanent Scatterers. *IEEE Transactions on Geoscience and Remote Sensing*, 44(11), 3338–3347. <https://doi.org/10.1109/TGRS.2006.877754>
- Perissin, D., & Wang, T. (2012). Repeat-Pass SAR Interferometry with Partially Coherent Targets. *IEEE Transactions on Geoscience and Remote Sensing*, 50(1), 271–280. <https://doi.org/10.1109/TGRS.2011.2160644>
- Perissin, D., Wang, Z., & Wang, T. (2011). The SARPROZ InSAR Tool for Urban Subsidence/Manmade Structure Stability Monitoring in China. *Proceedings of the 34th International Symposium on Remote Sensing of Environment (ISRSE 2011)*.
- Pieraccini, M., & Miccinesi, L. (2019). Ground-Based Radar Interferometry: A Bibliographic Review. *Remote Sensing*, 11(9), 1029. <https://doi.org/10.3390/rs11091029>
- Ponzo, F. C., Iacovino, C., Ditommaso, R., Bonano, M., Lanari, R., Soldovieri, F., Cuomo, V., Bozzano, F., Ciampi, P., & Rompato, M. (2021). Transport Infrastructure SHM Using Integrated SAR Data and On-Site Vibrational Acquisitions: “Ponte Della Musica–Armando Trovajoli” Case Study. *Applied Sciences*, 11(14), 6504. <https://doi.org/10.3390/app11146504>
- Prota, A., Savoia, M., & Ponzo, F. C. (A c. Di). (2025). *Linee guida per l'utilizzo dei dati interferometrici satellitari ai fini dell'interpretazione del comportamento strutturale delle costruzioni*. Doppiavoce. <https://doi.org/10.57580/RELUIS.LGDS>
- Qin, X., Zhang, L., Yang, M., Luo, H., Liao, M., & Ding, X. (2018). Mapping surface deformation and thermal dilation of arch bridges by structure-driven

- multi-temporal DInSAR analysis. *Remote Sensing of Environment*, 216, 71–90. <https://doi.org/10.1016/j.rse.2018.06.032>
- Ramos, L., Marques, L., Lourenco, P., De Roeck, G., Costa, A., & Roque, J. (2010). Monitoring historical masonry structures with operational modal analysis: Two case studies. *Mechanical Systems and Signal Processing*, 24, 1291–1305. <https://doi.org/10.1016/j.ymssp.2010.01.011>
- Richards, J. A. (2009). *Remote Sensing with Imaging Radar*. Springer. <https://doi.org/10.1007/978-3-642-02020-9>
- Riveiro, B., González-Jorge, H., Varela, M., & Jáuregui, D. V. (2013). Validation of terrestrial laser scanning and photogrammetry techniques for the measurement of vertical under clearance and beam geometry in structural inspection of bridges. *Measurement*, 46(1), 784–794. <https://doi.org/10.1016/j.measurement.2012.09.018>
- Rodriguez, E., & Martin, J. (1992). Theory and design of interferometric synthetic aperture radars. *IEEE Proceedings F - Radar and Signal Processing*, 139(2), 147–159. <https://doi.org/10.1049/ip-f-2.1992.0018>
- Rosen, P. A., Hensley, S., Joughin, I. R., Li, F. K., Madsen, S. N., Rodriguez, E., & Goldstein, R. M. (2000). Synthetic aperture radar interferometry. *Proceedings of the IEEE*, 88(3), 333–382. <https://doi.org/10.1109/5.838084>
- Saisi, A., & Gentile, C. (2015). Post-earthquake diagnostic investigation of a historic masonry tower. *Journal of Cultural Heritage*, 16(4), 602–609. <https://doi.org/10.1016/j.culher.2014.09.002>
- Sandwell, D. T., Xu, X., Chen, J., Mellors, R. J., Wei, M., Tong, X., DeSanto, J. B., & Ou, Q. (2025). *Satellite Radar Interferometry: Theory and Practice*. Cambridge University Press. <https://doi.org/10.1017/9781009606226>
- Scattarreggia, N., Salomone, R., Moratti, M., Malomo, D., Pinho, R., & Calvi, G. M. (2022). Collapse analysis of the multi-span reinforced concrete arch bridge of Caprigliola, Italy. *Engineering Structures*, 251, 113375. <https://doi.org/10.1016/j.engstruct.2021.113375>
- Schlögl, M., Widhalm, B., & Avian, M. (2021). Comprehensive time-series analysis of bridge deformation using differential satellite radar interferometry based on Sentinel-1. *ISPRS Journal of Photogrammetry and Remote Sensing*, 172, 132–146. <https://doi.org/10.1016/j.isprsjprs.2020.12.001>

- Schneider, P. J., & Soergel, U. (2021). Segmentation of Buildings Based on High Resolution Persistent Scatterer Point Clouds. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-3–2021, 65–71. <https://doi.org/10.5194/isprs-annals-V-3-2021-65-2021>
- Selvakumaran, S., Plank, S., Geiß, C., Rossi, C., & Middleton, C. (2018). Remote monitoring to predict bridge scour failure using Interferometric Synthetic Aperture Radar (InSAR) stacking techniques. *International Journal of Applied Earth Observation and Geoinformation*, 73, 463–470. <https://doi.org/10.1016/j.jag.2018.07.004>
- Sohn, H. (2006). Effects of environmental and operational variability on structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851), 539–560. <https://doi.org/10.1098/rsta.2006.1935>
- Sohn, H., Worden, K., & Farrar, C. R. (2002). Statistical damage classification under changing environmental and operational conditions. *Journal of Intelligent Material Systems and Structures*, 13(9), 561–574. <https://doi.org/10.1106/104538902028492>
- Talledo, D. A., Miano, A., Bonano, M., Di Carlo, F., Lanari, R., Manunta, M., Meda, A., Mele, A., Prota, A., Satta, A., & Stella, A. (2022). Satellite radar interferometry: Potential and limitations for structural assessment and monitoring. *Journal of Building Engineering*, 46, 103756. <https://doi.org/10.1016/j.jobe.2021.103756>
- Tang, P., Chen, F., Zhu, X., Zhou, W., Tang, P., Chen, F., Zhu, X., & Zhou, W. (2016). Monitoring Cultural Heritage Sites with Advanced Multi-Temporal InSAR Technique: The Case Study of the Summer Palace. *Remote Sensing*, 8(5). <https://doi.org/10.3390/rs8050432>
- Tarantini, R., Coccimiglio, S., Miraglia, G., Grassi, M., Ceravolo, R., & Ferro, G. A. (2025). EGMS and user-driven Sentinel-1 InSAR data for pre- and post-collapse analysis of the Himera viaduct. *ARTISTE 2025 Conference Proceedings*. First International Conference on Artificial Intelligence in Structural Engineering (In press).
- Tarantini, R., Coccimiglio, S., Miraglia, G., Matteini, I., Grassi, M., Ceravolo, R., & Ferro, G. A. (2026). Pre- and post-collapse analysis of infrastructures using COSMO-SkyMed satellite data. In *Procedia Structural Integrity*, 84, 401–408. - III FABRE Conference: Existing bridges, viaducts and tunnels:

research, innovation and applications.
<https://doi.org/10.1016/j.prostr.2026.06.052>

- Tarantini, R., McAvoy, S., Ferro, G. A., & Kuester, F. (2026). *Building-scale wildfire damage assessment by fusing PS-InSAR time series and bi-temporal airborne LiDAR: the 2025 Palisades fire (Los Angeles, California)*. *Remote Sensing (in preparation)*.
- Tarantini, R., Miraglia, G., Coccimiglio, S., Ceravolo, R., & Ferro, G. A. (2026). Soil Displacement Estimation from Integrated Sensing Technologies in Data-Driven Models Biased by Temporal Coherence of PS-InSAR. *Land*, 15(2), 296. <https://doi.org/10.3390/land15020296>
- Tavakkoliestahbanati, A., Milillo, P., Kuai, H., & Giardina, G. (2024). Pre-collapse spaceborne deformation monitoring of the Kakhovka dam, Ukraine, from 2017 to 2023. *Communications Earth & Environment*, 5(1), 145. <https://doi.org/10.1038/s43247-024-01284-z>
- Thenkabail, P. S., Lyon, J. G., & Huete, A. (2018). Advances in Hyperspectral Remote Sensing of Vegetation and Agricultural Crops. In *Fundamentals, Sensor Systems, Spectral Libraries, and Data Mining for Vegetation* (2^a ed.). CRC Press.
- Thöns, S. (2018). On the Value of Monitoring Information for the Structural Integrity and Risk Management. *Computer-Aided Civil and Infrastructure Engineering*, 33(1), 79–94. <https://doi.org/10.1111/mice.12332>
- Tipping, M. E. (2001). Sparse Bayesian Learning and the Relevance Vector Machine. *Journal of Machine Learning Research*, 211–244. <https://doi.org/10.1162/15324430152748236>
- Tonelli, D., Caspani, V. F., Valentini, A., Rocca, A., Torboli, R., Vitti, A., Perissin, D., Zonta, D., Tonelli, D., Caspani, V. F., Valentini, A., Rocca, A., Torboli, R., Vitti, A., Perissin, D., & Zonta, D. (2023). Interpretation of Bridge Health Monitoring Data from Satellite InSAR Technology. *Remote Sensing*, 15(21). <https://doi.org/10.3390/rs15215242>
- Tonelli, D., Valentini, A., Rocca, A., Zorzi, S., Lotti, A., & Zonta, D. (2023). Uncertainty quantification of satellite InSAR-monitoring of bridges: A case study. *Ce/Papers*, 6(5), 900–906. <https://doi.org/10.1002/cepa.1991>
- Toreti, A., Bavera, D., Acosta Navarro, J., Cammalleri, C., De Jager, A., Di Ciollo, C., Hrast Essenfelder, A., Maetens, W., Magni, D., Masante, D., Mazzeschi, M., Niemeyer, S., & Spinoni, J. (2022). *Drought in Europe – August 2022*

- (GDO Analytical Report, JRC Technical Report EUR 31192 EN, JRC130493). Publications Office of the European Union. <https://doi.org/10.2760/264241>
- Touzi, R., Lopes, A., Bruniquel, J., & Vachon, P. W. (1999). Coherence estimation for SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 37(1), 135–149. <https://doi.org/10.1109/36.739146>
- Types of orbits. (2020, marzo 30). https://www.esa.int/Enabling_Support/Space_Transportation/Types_of_orbits#LEO
- Ulaby, F. T., Moore, R. K., & Fung, A. K. (1986). *Microwave Remote Sensing: Active and Passive, Volume III: From Theory to Applications*. Artech House.
- Visual Crossing. (2024). *Weather Data & Weather API Services*. <https://www.visualcrossing.com/>
- Vradi, A., Sala, J., Solari, L., & Balasis-Levinsen, J. (2023). Validating the European Ground Motion Service: An Assessment of Measurement Point Density. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-4/W7-2023, 247–253. <https://doi.org/10.5194/isprs-archives-XLVIII-4-W7-2023-247-2023>
- Weissgerber, F., Colin-Koeniguer, E., Nicolas, J.-M., & Trouvé, N. (2017). 3D Monitoring of Buildings Using TerraSAR-X InSAR, DInSAR and PolSAR Capacities. *Remote Sensing*, 9(10), 1010. <https://doi.org/10.3390/rs9101010>
- Werninghaus, R., & Buckreuss, S. (2010). The TerraSAR-X Mission and System Design. *IEEE Transactions on Geoscience and Remote Sensing*, 48(2), 606–614. <https://doi.org/10.1109/TGRS.2009.2031062>
- Woodhouse, I. H. (2006). *Introduction to Microwave Remote Sensing*. CRC Press. <https://doi.org/10.1201/9781315272573>
- Worden, K., Cross, E. J., Dervilis, N., Papatheou, E., & Antoniadou, I. (2015). Structural Health Monitoring: From Structures to Systems-of-Systems★. *IFAC-PapersOnLine, 9th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes SAFEPROCESS 2015*, 48(21), 1–17. <https://doi.org/10.1016/j.ifacol.2015.09.497>
- Yang, W., He, Y., Zhu, Q., Zhang, L., & Jin, L. (2024). Unwrap-Net: A deep neural network-based InSAR phase unwrapping method assisted by airborne

- LiDAR data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 218, 510–529. <https://doi.org/10.1016/j.isprsjprs.2024.11.009>
- Yang, Y., Lacidogna, G., Elchalakani, M., & Hancock, C. M. (2024). Special Issue Editorial: “Remote Sensing in Structural Health Monitoring”. *Remote Sensing*, 16(24), 4748. <https://doi.org/10.3390/rs16244748>
- Zebker, H. A., & Villasenor, J. (1992). Decorrelation in interferometric radar echoes. *IEEE Transactions on Geoscience and Remote Sensing*, 30(5), 950–959. <https://doi.org/10.1109/36.175330>