Bridge Failure Risk Prediction Using Geospatial Data Processing via Multi-Head Attention Deep Learning Model

IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 24 November 2025, accepted 20 December 2025, date of publication 26 December 2025,
date of current version 2 January 2026.

Digital Object Identifier 10.1109/ACCESS.2025.3648853

RESEARCH ARTICLE

Bridge Failure Risk Prediction Using Geospatial Data Processing via Multi-Head Attention Deep Learning Model

FRANCESCO DELLA SANTA[1,2], FEDERICO SISCI[1], FARBOD KHOSRO ANJOM[3], FLAVIO PINO[4], AND MARCO CIVERA[5]

[1]Department of Mathematical Sciences, Politecnico di Torino, 10129 Turin, Italy
[2]Gruppo Nazionale per il Calcolo Scientifico INdAM, 00185 Rome, Italy
[3]Department of Environment, Land, and Infrastructure Engineering, Politecnico di Torino, 10129 Turin, Italy
[4]Department of Management and Production Engineering, Politecnico di Torino, 10129 Turin, Italy
[5]Department of Structural, Geotechnical, and Building Engineering, Politecnico di Torino, 10129 Turin, Italy

Corresponding author: Flavio Pino (flavio.pino@polito.it)

The work of Francesco Della Santa was supported by the Future Artificial Intelligence Research (FAIR) from the European Union Next-GenerationEU (Piano Nazionale di Ripresa e Resilienza (PNRR)–Missione 4, Componente 2, Investimento 1.3) under Grant D.D. 1555 11/10/2022, PE00000013. The work of Farbod Khosro Anjom was supported by Multi-Risk Science for Resilient Communities under a Changing Climate (RETURN) Extended Partnership from the European Union Next-GenerationEU (National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3) under Grant D.D. 1243 2/8/2022, PE0000005–Vertical Spoke 2. The work of Francesco Della Santa, Farbod Khosro Anjom, Flavio Pino, and Marco Civera was supported by the Bridge Risk Identification Using Data-Driven Geospatial Evaluation (BRIDGE) Research from Politecnico di Torino as one of the three winner teams of Poli Hack Days 2024.

ABSTRACT Bridge failures are a significant threat to infrastructure safety and public security, which demand cost-effective and scalable monitoring systems. This work proposes a novel data-driven framework for early warning of bridge collapse risk, leveraging Interferometric Synthetic Aperture Radar (InSAR) displacement time series and Deep Learning. A mathematical formulation of a bridge-collapse risk index is introduced, allowing quantitative estimation of the probability of failure from displacement data. To overcome data scarcity, a synthetic bridge dataset is generated through a combination of geometrical transformations and stochastic perturbations applied to real InSAR observations. Then, a Multi-Head Attention-based Neural Network is trained to predict the risk increment over time windows, using both on-bridge and surrounding geospatial points as input. The model is validated on the historical collapse of the Tadcaster bridge and stress tested on the Cantiano bridge, which failed during the 2022 Marche flood in Italy. The results show that the proposed approach effectively captures the temporal evolution of structural instability, with a conservative (risk-overestimating) but consistent prediction trend. These findings demonstrate the potential of combining InSAR data and attention-based Deep Learning models for scalable, non-invasive bridge health monitoring.

INDEX TERMS Bridge collapse prediction, InSAR, deep learning, geospatial data processing, multi-head attention, structural health monitoring, synthetic data generation.

I. INTRODUCTION

Bridge failures represent a critical threat to public safety and infrastructure reliability, which underscores the urgent need for better monitoring and failure detection strategies; for example, there have been almost 250 total and partial collapses in less than 25 years in Italy alone [7]. Traditional inspection protocols, predominantly based on periodic visual assessments, are often insufficient to detect latent structural vulnerabilities, especially in the absence of surface-level damage. Alternatively, monitoring systems based on physical sensors can be applied to bridges, but they are complicated

The associate editor coordinating the review of this manuscript and approving it for publication was Yiqi Liu.

and expensive solutions, limited in their applicability and scalability, as each bridge would need its own set of sensors (e.g., [1], [17]). However, geospatial data from Interferometric Synthetic Aperture Radar (InSAR) [11] with Persistent Scatterers (PS) [16] are abundant and cheaper than physical sensors or visual assessments.

In recent years, the use of InSAR data has significantly widened its scope, moving from classic subsidence studies to a broader set of real monitoring applications. Beyond ground deformation [9] and railway anomaly detection [28], PS-InSAR is now used for landslide evolution tracking, especially in slow-moving slopes where long time series help reveal seasonal and environmentally driven deformation trends (e.g., [6], [15]). It has also become a standard tool in urban stability assessment—monitoring building settlement, tunneling effects, and metro-line construction impacts—thanks to its millimetric sensitivity in dense built environments [8], [13]. Applications in volcanic areas and geothermal fields are also growing, where PS-InSAR helps capture subtle inflation–deflation cycles linked to subsurface processes [14].

As of now, InSAR data are typically used retrospectively for bridges, investigating the causes of mechanical failures and thus collapses rather than anticipating them; see, for example, [26], [27]. Their correlation with environmental data is also very novel and has barely been investigated, except for recent advanced applications (e.g., see [21]).

In this work, we propose a prototype of a data-driven early warning system for bridge risk assessment, which leverages space-borne InSAR data to predict the likelihood of structural failure. Specifically, given the time series of vertical or Line-of-Sight (LoS) displacements of geospatial points near/on the partially collapsed Tadcaster bridge (see [25], and Figures 1), we formally define formulas to quantify the concept of bridge-collapse risk; then, we generate a synthetic dataset of similar collapsed bridges to train a Deep Learning (DL) model that is able to estimate the current bridge-collapse risk looking at past displacement time series. After the training on these synthetic data, the DL model is validated on the Tadcaster bridge historical data, and a stress test is performed on a completely different bridge (Cantiano bridge), belonging to an Italian region hit by a flood in September 2022 that caused many bridges to (partially or totally) collapse, included the latter one [20].

This work is particularly innovative for the topic of structural health monitoring. Indeed, to the best of our knowledge, it is the first attempt in the literature of training a DL model able to estimate the risk of collapse for a bridge looking at InSAR-derived displacement time series. The proposed model is designed to enable the predictive use of InSAR data, providing actionable insights to optimize bridge maintenance schedules and prioritize structures that require early intervention.

The field of Structural Health Monitoring (SHM) is quite extensive for civil structures and infrastructures; the vast majority of field case studies, however, rely on physically-attached sensors. Only recently, non-contact methodologies (both proximal and remote sensing) have seen an increase in the use of this application, for example, resulting in the recently written Italian guidelines for satellite-based SHM [22]. However, satellite-based analysis is often performed qualitatively by domain experts, for an a posteriori understanding of at what point in time the bridge started behaving anomalously, in correlation with environmental phenomena, interpreting it as a collapse precursor (e.g., see [32]).

This paper also provides contributions in terms of mathematical formulations and modeling. In particular, this work presents: (*i*) a novel, data-driven, and mathematically well-defined quantification of the bridge-collapse risk (to the best of the authors' knowledge, such kind of risk index is still missing in the literature); (*ii*) the definition of a learning task that aims to predict the *increment* of the risk in a time window by observing the time series behaviors, for making the DL model applicable to any generic bridge; (*iii*) a formal procedure for generating a dataset of InSAR signals of synthetic bridges, given few data available from real-world measurements; (*iv*) a PCA-based re-orientation strategy for pre-processing and creating uniform representation of the data as uniform input to DL model; (*v*) a novel Neural Network (NN) architecture for the DL model, based on Multi-Head Attention (MHA) layers [33] and on a custom NN layer that is able to reorder the inputs according to the geospatial coordinates.

Our experiment results suggest that InSAR data can be processed through modern DL techniques to contribute meaningfully to early warning system development. The proposed approach improves scientific understanding of the precursors of structural failures and provides concrete practical value to civil authorities and infrastructure managers looking for cost-effective risk mitigation tools. This value can be further increased in the future by increasing data availability for training the DL model.

The paper is organized as follows: in Section II, we describe the real-world problem of bridge monitoring, and introduce its mathematical modeling and the corresponding learning task for the DL model training; in Section III, we describe the architecture of the NN we used in the work; in Section IV, we report the training details of the NN model, its performance on the test set, and its validation on the Tadcaster bridge data and the Cantiano bridge data; we end with some conclusions drawn in Section VI.

II. PROBLEM DESCRIPTION AND FORMULATION

The availability of geospatial data has recently opened new opportunities for remote bridge monitoring. In particular, InSAR technology has been proven to be effective in detecting millimeter-scale ground and infrastructure displacements [10], [26], but it is still mainly used in retrospective analyzes of structural failures performed by domain experts, rather than in automated pre-warning frameworks. However, these data are still more abundant and cheaper than physical

sensors or visual assessments. For this reason, we propose a novel strategy for remote bridge monitoring based on DL and InSAR data.

In this work, our objective is to develop a predictive DL model that estimates a bridge collapse risk index by analyzing displacement time series derived from InSAR data. Such a model would enable the anticipation of structural instabilities, thereby supporting timely and preventive maintenance actions. The proposed framework is expected to provide both economic and safety-related advantages. Indeed, by leveraging remotely acquired InSAR measurements, this framework could: significantly reduce the need for expensive instrumentation; support effective targeted maintenance with additional cost savings; allow large-scale monitoring, extending surveillance to a higher number of structures, if compared to traditional methods.

To pursue these objectives, we start with a mathematical formulation of the problem and of the task that the DL model has to learn. Additionally, we describe a criterion for generating synthetic InSAR data for the DL model training, due to the scarce data availability and for a meaningful validation when applying the trained model on real data.

A. RISK MODELIZATION

The InSAR data of $K \in \mathbb{N}$ geospatial points consists in K time series describing the vertical/LoS displacements in millimeters (mm) of each geospatial point, respectively. For each time series $\{d_t^{(k)}\}_{t\in\mathbb{N}} \subseteq \mathbb{R}$, $k = 1, \ldots, K$, the value at the instant t defines the difference between the position of the point at this instant and at the starting instant $t = 0$; then, $d_0^{(k)} = 0$, for each k. Typically, these time series are analyzed a posteriori to understand at what point in time the bridge started to show anomalous displacements, which can be interpreted as precursors of collapse, possibly in correlation with critical environmental phenomena in its region in the analyzed time period [32]. Nonetheless, this displacement analysis is often performed qualitatively by domain experts, and a quantitative measure of the "risk-of-collapse" is still missing to the best of the authors' knowledge. In particular, a risk-of-collapse measure is crucial for properly defining a learning task and for training a Neural Network (NN) model able to predict such quantity for real-time monitoring of bridges' structural health.

The core idea for defining the *risk index* quantity for a collapsed bridge, is to define an index that increases from 0 (i.e., 0% risk of collapse) to 1 (i.e., 100% risk of collapse) in the time range $[0, T]$, according to the behavior of the displacement time series. A naive approach could be that of defining an arbitrary function that, in our opinion, reflects the possible increase of risk (e.g., see Figures 1-3), but such an approach can introduce users' biases in the task that the NN model will learn.

For this reason, a different approach to define the quantity of the risk index can be to infer the risk totally from the displacement data; in particular, we emphasize the fact that

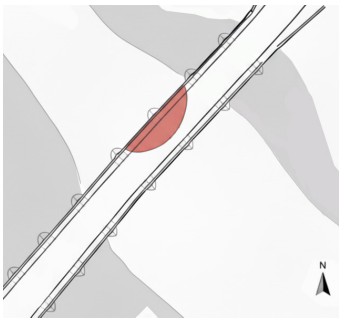

FIGURE 1. Scheme of the partial collapse of the Tadcaster bridge. Image realized by the authors, based on the map available in [25].

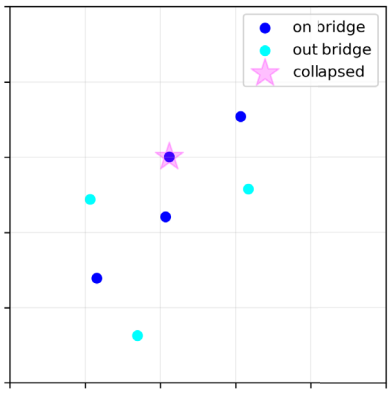

FIGURE 2. Geospatial points for the Tadcaster bridge.

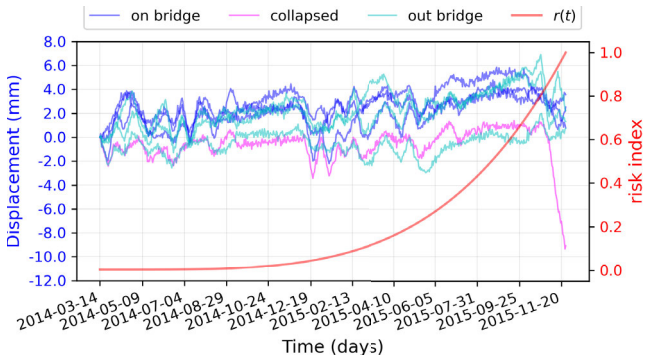

FIGURE 3. Displacements (blue, cyan, and magenta curves) and heuristic polynomial risk index (red curve) for the Tadcaster Bridge. Left-hand side y-axis for the displacements, right-hand side y-axis for the risk index.

this data-driven index is useful not only for defining a learning task, but also for a posteriori analyzes of bridge collapses. The data-driven *risk index* can be defined as follows.

Definition 1 (Risk Index): Let $K \in \mathbb{N}$ be the number of available geospatial points $\boldsymbol{p}^{(1)}, \ldots \boldsymbol{p}^{(K)} \in \mathbb{R}^2$ located on a collapsed bridge, and let $\{d_t^{(1)}\}_{t=0}^{T}, \ldots, \{d_t^{(K)}\}_{t=0}^{T}$ be their corresponding displacement time series, respectively. The last instant $t = T$ of the time series denotes the last measurement before the instant $t = t^*$ of the bridge collapse (i.e., $T \lesssim t^*$).

Let $\mu^{(k)}$ and $\sigma^{(k)}$ denote the sample mean and the sample standard deviation of $\{d_t^{(k)}\}$, for each k, and let

$\tau^{(k)} := \mu^{(k)} - \sigma^{(k)}$. Then, we define *risk increment* at the instant t the quantity

$$\rho_t := \frac{1}{K}\sum_{k=1}^{K} \max\left\{\frac{\tau^{(k)} - d_t^{(k)}}{|\tau^{(k)}|}\ 0\right\}; \tag{1}$$

summarizing, the value ρ_t is the average relative displacement below the threshold $\tau^{(k)}$ at instant t, for all the K points. Then, we define the *absolute risk* and the *relative risk* at the instant t as

$$R(t) := \sum_{i=0}^{t} \rho_i \quad \text{and} \quad r(t) := \frac{R(t)}{\sum_{i=0}^{T} \rho_i} = \frac{R(t)}{R(T)}, \tag{2}$$

respectively.

In Figure 4 we illustrate the displacements of the Tadcaster bridge and the corresponding relative risk index $r(t)$ obtained from these data.

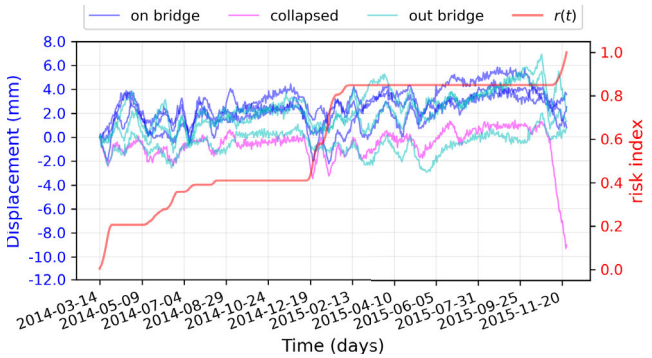

FIGURE 4. Displacements (blue, cyan, and magenta curves) and data-driven relative risk index (red curve) for the Tadcaster Bridge. Left-hand side y-axis for the displacements, right-hand side y-axis for the risk index.

An alternative approach to defining the risk index can be a trade-off between the data-driven approach of Defition 1 and a heuristic approach. The motivation behind the introduction of trade-off risk indices is to avoid the unrealistic case of periods with zero risk increment, which instead characterizes the data-driven risk; for this reason, in this work, we will train our NN model on the latter kind of risk index and not on the data-driven one.

We introduce a trade-off risk by perturbing the null data-driven risk increments with heuristically determined noise, i.e.:

$$r'(t) := \frac{\sum_{i=0}^{t} \rho'_i}{\sum_{i=0}^{T} \rho'_i}, \tag{3}$$

where

$$\rho'_t := \begin{cases} \rho_t, & \text{if } \rho_t > 0 \\ \varepsilon, & \text{otherwise} \end{cases}, \tag{4}$$

ρ_t is the risk increment in (1), and $\varepsilon \sim \mathcal{U}(2e\text{-}4, 1.4e\text{-}3) + \mathcal{U}(0, 5e\text{-}5)$. In Figure 5, we illustrate the different behaviors of risk $r(t)$ and risk $r'(t)$.

Both these definitions of the risk index assume that at the instant $t = 0$ the bridge is in perfect structural health,

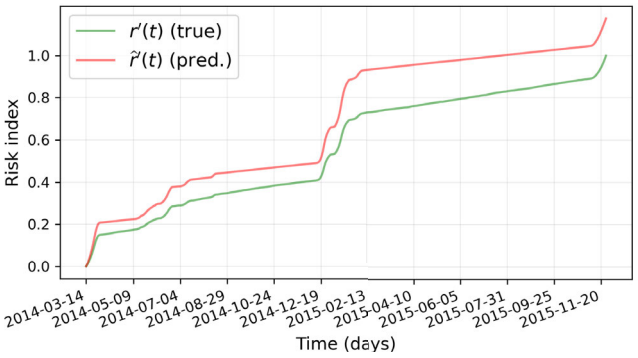

FIGURE 5. Comparison of (relative) risk indices for the Tadcaster Bridge.

with null risk, and that the data-driven thresholds $\tau^{(k)}$ are all negative (even if (1) is well defined also for $\tau^{(k)}$ positive). Obviously, the assumption of perfect structural health cannot always be true, depending on the history of the bridge preceding the start of the InSAR monitoring. We postpone to future work the definition of risk indices that involve also the previous bridge history; nonetheless, as illustrated in the numerical results, this assumption does not ruin the efficacy of the NN predictions.

B. LEARNING TASK FORMULATION

Given the definitions of risk index, we can define a learning task for training a NN model that is able to evaluate the risk-of-collapse of a bridge by looking at InSAR data of $M \in \mathbb{N}$ geospatial points in its surroundings (i.e., both on the bridge and its neighborhood); in particular, we consider only the risk index $r'(t)$ defined in equations (3) and (4). Concerning the motivation behind the choice of using also geospatial points nearby the bridge, it comes from the opinion of domain experts, who often emphasize that displacements of the terrain around the bridge may have an indirect influence on the bridge's structural health.

Since the prediction of the risk index at instant t given all the past InSAR data can be a complicated and resource-expensive task, we define a simpler but equivalently efficient task. We consider the learning task of predicting the increment of risk index for a time window of $W \in \mathbb{N}$ consecutive instants $(\bar{t}\ \bar{t}+1, \ldots, \bar{t}+(W-1))$, given the displacement subsequences $\{d_t^{(1)}\}_{t=\bar{t}}^{\bar{t}+W-1}, \ldots, \{d_t^{(M)}\}_{t=\bar{t}}^{\bar{t}+W-1}$ of the $M \in \mathbb{N}$ geospatial points considered. This latter approach is versatile and cheap, suitable to be generalized to bridges characterized by behaviors of the risk index different from the risk behavior of the bridges available for the training.

Indeed, a model trained to predict the risk increments $(\rho_{\bar{t}}, \ldots, \rho_{\bar{t}+(W-1)})$, for each given time window, is not aware about the previous health conditions of the bridge, and has no information about the previous displacements signals. In this sense, the model learns a mapping between the kinematics of deformation and the slope of the risk index, rather than the direct value, or a binary classification of "normal" versus "abnormal" behavior. This predictive structure is particularly suitable for extending predictions

also to bridges with different displacement behaviors and, therefore, different dation rates and collapse times (see Section IV-E).

Following the notation in Defition 1, equation (1), for the ease of notation, we denote by ρ_t the risk increment at the instant t; indeed, we want to define the learning task independently on the type of risk considered (even if we focus on $r'(t)$). According to this, the learning task we define can be summarized in the problem of approximating the function

$$\begin{aligned} \boldsymbol{F} : \mathbb{R}^{M\times W} \times \mathbb{R}^{M\times 2} \times \{0,1\}^M &\to \mathbb{R}^W \\ (D, X, \boldsymbol{b}) &\mapsto \rho\,, \end{aligned} \tag{5}$$

where $D \in \mathbb{R}^{M\times W}$ is the matrix made by the displacement subseries of the M geospatial points for a given time window of W instants, $X \in \mathbb{R}^{M\times 2}$ is the matrix of the geospatial points' coordinates, $\boldsymbol{b} \in \{-1,1\}^M$ is the vector denoting if a point is on the bridge (one) or not (zero), and $\rho \in \mathbb{R}^W$ is the vector of corresponding risk increments, instant by instant.

For example, if we assume without loss of generality that the first K points, of the M total ones, are all the points on the bridge (the remaining points are nearby the bridge), we have that:

$$D = \begin{bmatrix} d_{\bar{t}}^{(1)} & \cdots & d_{\bar{t}+(W-1)}^{(1)} \\ \vdots & & \vdots \\ d_{\bar{t}}^{(M)} & \cdots & d_{\bar{t}+(W-1)}^{(M)} \end{bmatrix}, \tag{6}$$

$$X = \begin{bmatrix} x^{(1)} & y^{(1)} \\ \vdots & \vdots \\ x^{(K)} & y^{(K)} \\ x^{(K+1)} & y^{(K+1)} \\ \vdots & \vdots \\ x^{(M)} & y^{(M)} \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} b^{(1)} \\ \vdots \\ b^{(K)} \\ b^{(K+1)} \\ \vdots \\ b^{(M)} \end{bmatrix} = \begin{bmatrix} 1 \\ \vdots \\ 1 \\ -1 \\ \vdots \\ -1 \end{bmatrix}, \tag{7}$$

and

$$\boldsymbol{F}(D, X, \boldsymbol{b}) = \rho = (\rho_{\bar{t}}, \ldots, \rho_{\bar{t}+(W-1)})\,. \tag{8}$$

C. METHODOLOGY FOR SYNTHETIC DATA GENERATION

Given the current scarcity of open source InSAR displacement data, synthetic data generation is necessary here to train a reliable DL model for bridge failure risk assessment, while keeping real data of a separated bridge collapse only as a test to validate the trained model.

A critical aspect is that the generated synthetic bridges' data must be realistic and representative of failing systems. An analytical validation of this property, or a data-driven validation based on a comparison with data of other collapsed bridges would be ideal; nonetheless, there are no methods in the literature that are able to predict an imminent collapse of the bridge just by looking at the displacement time series (see Section I and Section II-A) and, as mentioned above, there are no comprehensive, publicly accessible datasets containing multiple correlated bridge failures with InSAR displacement data.

For these reasons, we imposed the reliability of the synthetic bridges indirectly, by carefully selecting the construction operations for their generation; such reliability is furthermore validated by the experiment results observed for the trained DL model with respect to the real bridges of Tadcaster and Cantiano (see Sections IV-D IV-E).

The synthetic data generation procedure we propose is based on geometric operations and displacement perturbations that transform the coordinates and displacements of given M points in the surroundings of a bridge, respectively, but preserving the structural failing system of the bridge. In particular, after rescaling the coordinates of the points to fit the $[0,1]^2$ square (but maintaining bridge proportions, e.g., see Figure 2), we apply a combination of one or more of these operations:

1) reflection with respect to the vertical axis, the horizontal axis, and/or one of the two diagonals;
2) perturbation of the coordinates of one or more points by applying a white noise $\epsilon_x, \epsilon_y \sim \mathcal{N}(0, 0.05/3)$;
3) perturbation of the displacement time series of one or more points by element-wise multiplication with a white noise vector $\delta^{\text{mul}} = (\delta_0^{\text{mul}}, \ldots, \delta_T^{\text{mul}})$, such that $\delta_i^{\text{mul}} \sim \mathcal{N}(1, 0.1/3)$, and/or by summation of a noise vector $\delta^{\text{add}} = (\delta_0^{\text{add}}, \ldots, \delta_T^{\text{add}})$, such that $\delta_i^{\text{add}} \sim \mathcal{N}(0, 1/3)$;
4) creation of a new synthetic geospatial point on the bridge and the corresponding synthetic displacement time series;

The transformations related to items (1)-(3), are almost straightforward, whereas the creation of a new synthetic geospatial point with its own synthetic displacement time series (i.e., item (4)) deserves more details, which are illustrated in the following.

Without loss of generality, assume that the first K points, of the M total ones, are all points on the bridge. The synthetic new geospatial point $\boldsymbol{p}^{(M+1)}$ on the bridge is generated randomly as a point belonging to the convex hull of the previous bridge-points and such that it is sufficiently far from them (according to an arbitrary threshold distance $h \in \mathbb{R}_+$); i.e.:

$$\boldsymbol{p}^{(M+1)} \sim \text{ConvHull}\left\{\boldsymbol{p}^{(1)}, \ldots, \boldsymbol{p}^{(K)}\right\} \tag{9}$$

and

$$\text{dist}\left(\boldsymbol{p}^{(M+1)}, \boldsymbol{p}^{(k)}\right) \geqslant h\,, \quad \forall\, k = 1, \ldots, K\,. \tag{10}$$

After generating $\boldsymbol{p}^{(M+1)}$, we generate its corresponding displacement time series as a weighted average of all the other bridge-points' time series, where the weight is inversely proportional to the distance between $\boldsymbol{p}^{(M+1)}$ and the other points. Specifically, $\{d_t^{(M+1)}\}_{t=0}^T$ is such that

$$d_t^{(M+1)} = \frac{1}{\text{dist}_{\text{tot}}} \sum_{k=1}^{K} \frac{\text{dist}_{\min}}{\text{dist}\left(\boldsymbol{p}^{(M+1)}, \boldsymbol{p}^{(k)}\right)} d_t^{(k)}\,, \tag{11}$$

for each $t = 0, \ldots, T$, where

$$\text{dist}_{\min} := \min_{k=0,\ldots,K} \left\{ \text{dist}\left(\boldsymbol{p}^{(M+1)}, \boldsymbol{p}^{(k)} \right) \right\} . \tag{12}$$

and

$$\text{dist}_{\text{tot}} := \sum_{k=1}^{K} \text{dist}\left(\boldsymbol{p}^{(M+1)}, \boldsymbol{p}^{(k)} \right) . \tag{13}$$

Before concluding this subsection, we analyze why the transformations (1)-(4) applied to the Tadcaster's data generate a set of new synthetic bridges that are a physically-reliable representation of collapsed bridges.

Transformations (1) are reflections with respect to an axis, and are thus equivalent to a change of the coordinate system. Transformations (2)-(3) are designed to slightly move the coordinates of the geospatial points from their original location, and to apply small perturbations to the displacements; therefore, the shape and structural characteristics of the failing system are not altered. In the end, transformation (4) is similar to typical InSAR data preprocessing methods, based on interpolation (e.g., see [26]); therefore, any new generated point is coherent with the structural failing system described by the other points used for generating it.

The result of these transformations applied to the Tadcaster bridge is a set of synthetic bridges that are geometrically characterized by a structure that is equivalent to the Tadcaster's one, and with displacements that are constrained to have a general behavior analogous to the Tadcaster's displacements; therefore, each synthetic bridge exhibits a monotonic risk progression leading to collapse at time T.

In summary, the synthetic bridges remain structurally consistent approximations of realistic collapse scenarios, suitable for training a generalized learning framework until larger real datasets become available.

III. NEURAL NETWORK ARCHITECTURE ARCHETYPE

In this section, we describe the NN architecture archetype that we adopt to build a DL model suitable for the learning task illustrated in Section II-B.

The main idea behind the architecture we propose for learning the risk increment prediction task is to exploit the Multi-Head Attention (MHA) layers and the encoding structure of Transformer models, i.e. architectures that have been originally developed for Language Models (see [33]). Indeed, we noticed some analogies between the interaction of word tokens in a sentence and the structural interaction of geospatial points on a bridge and its surroundings. Nevertheless, the applications of such architectures are different from Language Models; specifically, the NN architecture is not structured for predicting the next element of a series (e.g., the next token of a sentence) given the previous ones, but has the target of extrapolating implicit, nonlinear, relationships among multiple time series (i.e., the displacement time series) for inferring the corresponding risk increment resulting from these interactions. In brief, the target is to approximate the function $\boldsymbol{F}$ defined in (5) by exploiting the attention mechanism of the MHA layers.

In the following subsections, we describe: (*i*) a custom layer specifically developed for this architecture, having the role of reordering the inputs according to the geospatial coordinates; (*ii*) the general scheme of the NN architecture.

A. REORDERING LAYER FOR INPUTS

As the word tokens given as input to Language Model Transformers follow the logic rules of languages, similarly the inputs related to the M geospatial points can be reordered according to proper geospatial criteria.

Actually, such reordering procedure is a pre-processing operation of the inputs and can be performed outside the model. However, we desired to embed such procedure into the NN architecture in order to ease the global pre-processing operations, focusing them on the data values rather than on the order of the elements. Indeed, thanks to this embedded pre-processing, the NN can receive unordered inputs because it automatically applies the proper reordering.

In this work, the reordering criterion we embed in the NN through a proper custom layer follows the "*left-to-right, top-to-bottom*" order. Mathematically, given a generic input $(D, X, \boldsymbol{b}) \in \mathbb{R}^{M\times K} \times \mathbb{R}^{M\times 2} \times \{0,1\}^M$ (see equations (6) and (7)), the *Reordering* (RO) Layer transforms it into the matrix $(\widetilde{D}, \widetilde{X}, \widetilde{\boldsymbol{b}})$ such that

$$\widetilde{D} = \begin{bmatrix} d_{\tilde{t}}^{(\ell_1)} & \cdots & d_{\tilde{t}+(W-1)}^{(\ell_1)} \\ \vdots & & \vdots \\ d_{\tilde{t}}^{(\ell_M)} & \cdots & d_{\tilde{t}+(W-1)}^{(\ell_M)} \end{bmatrix} , \tag{14}$$

$$\widetilde{X} = \begin{bmatrix} x^{(\ell_1)} & y^{(\ell_1)} \\ \vdots & \vdots \\ x^{(\ell_M)} & y^{(\ell_M)} \end{bmatrix} , \ \widetilde{\boldsymbol{b}} = \begin{bmatrix} b^{(\ell_1)} \\ \vdots \\ b^{(\ell_M)} \end{bmatrix} , \tag{15}$$

where $\{\ell_1, \ldots, \ell_M\} = \{1, \ldots, M\}$ and for each $i, j \in \{1, \ldots, M\}$, $i < j$, it holds $x^{(\ell_i)} < x^{(\ell_j)}$ or $y^{(\ell_i)} \leqslant y^{(\ell_j)}$, if $x^{(\ell_i)} = x^{(\ell_j)}$.

B. MULTI-HEAD ATTENTION-BASED ARCHITECTURE

The architecture we define for our DL model is mainly based on MHA layers, Normalization Layers, and Residual connections (see [2], [12], [33], respectively); also standard Fully-Connected (FC) layers are included into the architecture, similarly to the structure of the Transformer's encoder illustrated in [33].

Before the Transformer-like blocks of layers, we apply the RO layer and an element-wise multiplication of the reordered input $\widetilde{D}$ with an M-by-W matrix that is returned by an FC layer applied to the matrix $(\widetilde{X} \,|\, \widetilde{\boldsymbol{b}})$. This latter operation has the role of encoding the reordered coordinates and, then, of encoding the displacements according to the transformed coordinates. Therefore, the main body of our architecture

receives and transforms $\mathbb{R}^{M\times W}$ matrices, returning matrices of the same shape.

For obtaining an output suitable for the learning task (i.e., vectors in $\mathbb{R}^W$), at the end of the NN's main body, we attach a custom layer that sums up all the rows of the M-by-W matrix, obtaining a vector; then, we conclude with a FC layer with W units and *softplus* activation function. The motivation of using such activation for the output layer is that the risk increments we consider in our experiments are always positive, by definition of $\rho'(t)$ in equation (4).

IV. NUMERICAL EXPERIMENTS

In this section, we report the training procedure and the validation results for a DL model characterized by an architecture of the type describe in the previous section.

Before the training procedure, efforts have been spent in preparing a proper synthetic dataset. This operation consisted in collecting the Tadcaster displacement data from the literature (the details are reported in subsection Section IV-A) and generating a dataset of synthetic bridges following the steps illustrated in Section II-C.

Concerning the training, after a preliminary hyper-parameter selection procedure, a NN architecture of the type described in Section III is trained and evaluated on the synthetic data, with respect to the risk index $r'(t)$. To ease the learning task, particular attention has been dedicated to develop a PCA-based re-orientation pre-processing strategy of the bridge points, in order to automatically rotate the bridge in horizontal position.

Then, the trained model is validated on the original data from the Tadcaster Bridge. The results show good predictive performance, although the model tends to overestimate the risk (which is preferable to underestimation in risk assessment); in particular, these results are promising for potential real world applications, because the Tadcaster real data are new, unseen data for the models (trained on the synthetic ones).

Despite the good results for the Tadcaster bridge, we conclude the section by illustrating a further stress test for the trained model. In this test we illustrate the prediction performances of the model for a completely different bridge, the Cantiano bridge, belonging to an Italian region hit by a flood on 15 September 2022 that caused many bridges to collapse (including the latter one). We classify this second validation step as a stress test for the trained DL model because it is applied to longer displacement time series, characterized by different statistics, and because the structure of the bridge was partially different from the Tadcaster one (Cantiano bridge was for pedestrians). Nonetheless, the risk predictions returned by the model for the risks $r'(t)$ prove to be interesting and promising, highlighting the concrete potential to extend and enhance the framework presented in this work in order to develop a real-time bridge monitoring system based on InSAR data.

A. REAL DATA GATHERING AND SYNTHETIC DATA GENERATION

The displacement data of 7 geospatial points of the Tadcaster bridge has been collected from [25]. This case study is well known within the InSAR structural health monitoring community. In particular, the authors of [25] identified anomalous LoS deformation on a key geospatial point located on the north/central span, on the side that eventually failed (see Figures 1 and 2). The last two displacement measurements at this point exceeded the background variability (≈ 3.88 mm) and were interpreted as precursory deformation signals prior to collapse. During this period, the Wharfe River (crossed by the Tadcaster bridge) experienced exceptionally high flows, consistent with scour processes that likely contributed to the bridge failure; as a consequence, the displacements' behavior are characterized by this phenomenon and, with their shapes, they describe the indirect effect that high flows and scour processes had on the bridge structure. Nonetheless, in future work we plan to include environmental data as input features for improving the modelization of the learning task (see Section VI). For more details about the InSAR data of the Tadcaster bridge, we refer to Appendix A and the references therein.

After being recovered, the Tadcaster's displacement time series have been processed through a data augmentation procedure, in order to have daily displacement data. In particular, the data augmentation procedure consists of the combination of a linear interpolation of the collected data and the application of perturbations like δ^{mul} in item (3), Section II-C.

Given these new Tadcaster's displacement histories, a dataset of 159 synthetic bridges, with at most $M = 9$ points, is generated according to the content of Section II-C. These bridges are all similar but different from the Tadcaster one and they are used for training the risk increment prediction model. The number of synthetic bridges is such that we have 160 bridges in total if considering also the true Tadcaster bridge.

B. DATA PRE-PROCESSING AND PCA-BASED BRIDGE ORIENTATION STRATEGY

In DL, data pre-processing is often a crucial step to improve and easen the training of NNs. In this work we apply one regression target pre-processing operation and one input pre-processing operation.

The regression target pre-processing is relatively simple. The target vectors ρ associated to the inputs $(D, X, \boldsymbol{b})$ are normalized with respect to the value $\rho_{\max}\cdot 0.97$, where $\rho_{\max}$ is the maximum value of risk increment in the training set. We do not normalize with respect to $\rho_{\max}$ in order to prepare the model to see target values also greater than one, in case in the validation or test set the maximum risk increment is greater than $\rho_{\max}$.

The input pre-processing is not focused on the displacement values, but on the coordinates of the bridge points. The

idea is to rotate and rescale the geospatial points in order to feed the DL model always with a bridge in horizontal position, centered in the origin, and with the leftmost and the rightmost points on it that have x-coordinate equal to -1 and 1, respectively (nearby points can lie beyond this values). While the rescaling operation is trivial, given the bridge in horizontal position, the re-orientation procedure cannot be performed by rotations given a priori. Therefore, we propose a re-orientation technique based on performing Principal Component Analysis (PCA) with respect to the bridge's points. In this way, we obtain a roto-reflection matrix that transforms the geospatial points, re-orienteering the bridge in horizontal position.

Additionally, the input pre-processing also includes the generation of fake, synthetic, geospatial points in case the bridge has a number of points $M' < M$, where M is the number of points that the NN model expect to see (in the architecture considered, $M = 9$). If $M' < M$, the pre-processing operation, after re-orientation and scaling of the points, adds $(M - M')$ fake, not-on-bridge, points in the coordinate $(x_{\min} - 1, 0)$, with null displacement time series. It is important to observe that this pre-processing operation permits to have a DL model that can process bridges with a variable number of geospatial points, up to a quantity equal to M.

For an example of re-orientation preprocessing, see Figure 6 and compare it with Figure 2.

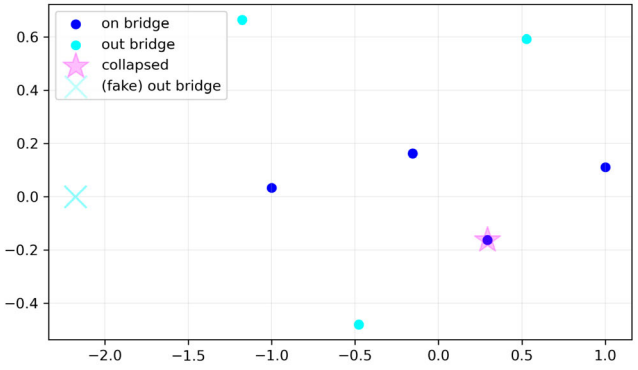

FIGURE 6. PCA-based pre-processing applied to the geospatial points for the Tadcaster bridge (see. Figure 2).

C. TRAINING AND PERFORMANCES ON SYNTHETIC DATA

Given the dataset of 159 synthetic bridges generated from the Tadcaster bridge data (see Section IV-A), with at most $M = 9$ points, we randomly extract $N = 500$ time windows of length $W = 14$ days for creating a dataset of input-target pairs $\{((D_i, X_i, \boldsymbol{b}_i)\ \rho_i)\}_{i=1}^{N}$. This dataset is split into a training set, validation set, and test set according to the percentages 40%, 10%, and 50%, respectively. The early training attempts showed that using about 40% of the data was enough to train a well-fitting model. As a result, we kept a large portion of the data (50%) for testing on unseen samples to make sure the model could generalize well to different data. The data are all pre-processed as described in the previous subsection.

A preliminary hyper-parameter random search strategy [4] is adopted with respect to the hyper-parameters illustrated in Table 1. This random search procedure returns the set of optimal hyper-parameters reported in the rightmost column of Table 1.

TABLE 1. Hyper-parameter values (second column) and random search optimal selection (third column).

Hyper-parameters	Values	Optimal
# Transf.-like blocks	$1, \ldots, 8$	6
# Heads (MHA)	4, 6, 8, 10, 12	12
# Keys (MHA)	$8, \ldots, 16$	13
Drop-out prob. (MHA, FC)	$0.00, 0.05, \ldots, 0.30$	0.00
Act.Func. (FC)	*elu, tanh, leaky relu, swish*	*leaky relu*
# Units (1^{st} FC, each block)	$2W, 3W, 4W$	$4W = 56$
# Units (FC, after blocks)	$W/2, W/2 + 4, \ldots$ $\ldots, W/2 + 20$	$W/2 + 12 = 19$
Mini-batch size	32, 64	32

Then, we train the NN architecture, based on the set of optimal hyper-parameters, for learning the risk index $r'(t)$. The training options are the following:

- Mean Squared Error (MSE) loss function;
- Adam optimizer [18];
- Maximum number of epochs: 1500;
- Early stopping [30] (patience 200 epochs);
- Reduce learning rate on plateaus [31] (factor 0.5, patience 75 epochs).

At the end of the training, the NN is evaluated on the test set, with respect to the loss function value and the Mean Absolute Error (MAE); in particular, the MAE is measured after the inversion of the pre-processing operation for the targets (see Section IV-B). The results obtained show very good prediction performances (see Table 2).

TABLE 2. Test set performances of the three NNs trained with respect to the risk increments of the risk indices $r^0(t)$, $r'(t)$, and $r''(t)$.

Error Measure	$r'(t)$
MSE* (loss)	5e-3
MAE**	1.1e-4

* without pre-proc. inversion on targets/preds.
** with pre-proc. inversion on targets/preds.

D. VALIDATION ON THE TADCASTER BRIDGE

The good prediction results measured on the test set are promising (see Table 2), but they are not enough to understand if the model is good in predicting the risk index with respect to a sliding time window of $W = 14$ days containing the displacement values of all the $M \leqslant 9$ points of a bridge. For this reason, we perform a comparison of the risk index reconstructed from the NN predictions for the Tadcaster bridge and the ground-truth indices.

We recall that the Tadcaster bridge's data are not included in the dataset used for the training set, the validation set, and the test set; therefore, the model has never seen these data.

However, we cannot ignore that the Tadcaster data necessarily have strong similarities with the synthetic data concerning the displacement general behaviors, since the synthetic data have been generated by transformations and perturbations of the Tadcaster's original ones.

To perform a comparison between the NN predicted risk index and the corresponding ground truth, we first need to define a procedure to reconstruct the predicted risk given the risk increment predictions returned by the DL model; indeed, we recall that the model predicts risk increments for time windows of length W, but it is completely blind to the passage of time and to the previous value of the risk index.

Let us denote by $\widehat{\rho}'^{(t)} \in \mathbb{R}^W$ the risk increment prediction returned by the NN, obtained for the Tadcaster's input data $(D^{(t)}, X^{(t)}, \boldsymbol{b}^{(t)})$ that correspond to the time window $(t, t+1, \ldots, t+(W-1))$. Then, the risk index predicted by the NN is

$$\widehat{r}'(t) := \frac{\sum_{i=0}^{t} \mu(\widehat{\rho}'_i)}{\sum_{i=0}^{T} \mu(\widehat{\rho}'_i)}, \tag{16}$$

where, for each $i = 0, \ldots, T$, we have that $\mu(\widehat{\rho}'_i)$ is the average of all the predictions made by the NN model for the instant i, while the time window is sliding; i.e.:

$$\mu(\widehat{\rho}'_t) := \frac{1}{W} \sum_{w=0}^{W-1} \widehat{\rho}'^{(t-w)}_{w+1}, \quad \forall\, t = 0, \ldots, T, \tag{17}$$

assuming by convention that the predictions for instants $t < 0$ are equal to the predictions for the instant $t = 0$, and that the predictions for instants $t > T$ are equal to the predictions for the instant $t = T$. The general idea behind this risk reconstruction procedure is to update the risk increment prediction as the time window slides, taking the average of the predictions at each instant in order to obtain a more precise risk prediction.

Given the risk prediction formula (16), we visualize the risk index curve predicted by the NN, and we compare it with the ground-truth curve. What we observe is a good behavior prediction for risk r', with a generic tendency to overestimate the risk (see Figure 7 and Table 3); indeed, the shape of the predicted risk curve mimics the shape of the ground truth curve, but increasing faster with respect to time. For a more accurate and quantitative analysis, we measure the days t on which the bridge actually reached specific risk levels, as well as the days t on which these risk levels were predicted by the model. In addition to the value 1, representing imminent collapse, the considered risk thresholds include $0.50, 0.55, \ldots, 0.75$ (see Table 3); these values define an operationally meaningful range in which maintenance actions can be scheduled, avoiding both the danger of an impending failure and the risk of performing unnecessary maintenance. Looking at the "early warning" values in Table 3, we notice that for almost all the risk values we have an overestimate of the risk that is of approximately 1-3 weeks ($\sim$ 2.5% of the monitoring interval of $T = 623$ days); the only exceptions are the risk value 0.75, that is reached by the model approximately 2 months earlier, and the value 1, that is reached by the model approximately 4 months earlier.

Although the results shown in Figure 7 and Table 3 highlight the presence of margins for improvement for reducing the risk overestimation, the model nonetheless demonstrates promising capabilities. Indeed, it is important to highlight that, in the context of structural health monitoring, an overestimation of the risk is preferable to an underestimation. Indeed, predicting a risk value higher than the ground truth is less problematic than predicting a lower one; e.g., like in medical applications, where false positives are more acceptable than false negatives.

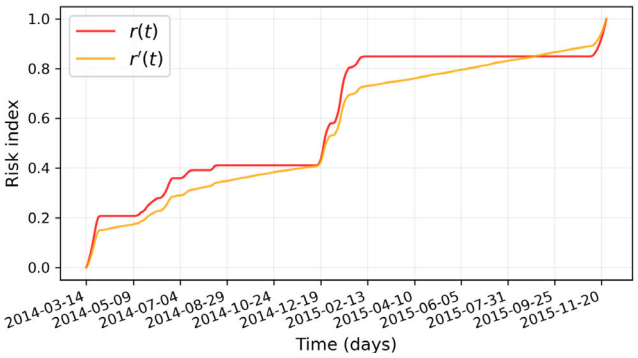

FIGURE 7. $r'(t)$ risk prediction and ground truth for the Tadcaster bridge.

TABLE 3. Quantification of the risk overestimate of the model for the Tadcaster bridge.

risk value	day t		early warning	
	pred.	true	days	% (over T)
0.50	279	286	-7	-1.12%
0.55	283	300	-17	-2.73%
0.60	285	304	-19	-3.05%
0.65	290	307	-17	-2.73%
0.70	301	320	-19	-3.05%
0.75	304	373	-69	-11.08%
1.00	496	623	-127	-20.39%

In general, looking at the results obtained on the Tadcaster bridge data, we can state that our model is characterized by promising generalization abilities. For this reason, in the next section we perform a stress test for the trained model, in order to better understand how much general this model can be.

E. VALIDATION ON THE CANTIANO BRIDGE

The Cantiano bridge was situated in Marche, an Italian region hit by a flood on 15 September 2022, which caused this bridge to totally collapse (at least nine other bridges in the area were similarly affected, with partial or total collapses. The structure of the Cantiano bridge

is different from the Tadcaster one, as the latter was a road bridge, while Cantiano was a footbridge with a much shorter and narrower span. However, they suffered a very similar failure mechanism, as both collapsed during a flood (see Figure 8).

FIGURE 8. Pictures of the Cantiano bridge before (top) and after (bottom) the flood.

We extracted the Cantiano bridge data from a collection of displacement time series measured for the Italian region it belongs to; all these data have been purchased by the authors through research funding for developing this work and future improvements. For more details about the InSAR data of the Cantiano bridge, we refer to Appendix B and the references therein.

The displacement time series of the Cantiano bridge cover a much longer time period than the Tadcaster time series (approximately 11 years and a half for Cantiano, and 1 year and 8 months for Tadcaster). Unfortunately, the measurements of Cantiano are not regular or homogeneous during these years, sometimes also with the absence of data for approximately 11 months (e.g., between summer 2020 and summer 2021). Therefore, in order to obtain daily time series for displacements, we perform the same data augmentation procedure we adopted for the Tadcaster bridge (see Section IV-A); however, due to the different nature of Cantiano, the longer monitoring time window, and the sparse measurements, the augmented displacements can be less representative of the true structural behavior of the bridge. Nonetheless, this operation is necessary to analyze this bridge, and it is also useful to understand if the trained model is robust even in such situations.

In order to adapt the trained model to Cantiano, we need to apply one pre-processing operation to the displacements and one post-processing operation to predictions:

- *Displacement pre-processing*: the models have been trained on displacement values generated by perturbations of Tadcaster data; then, they expect to receive displacements characterized by similar statistical properties. For this reason, we transform the Cantiano's displacements to have average and standard deviation equal to the ones of Tadcaster's displacements; i.e.:

$$d_t^{(m)} = \frac{d_t^{(m)} - \mu_{\text{can}}}{\sigma_{\text{can}}} \cdot \sigma_{\text{tad}} + \mu_{\text{tad}} \,, \tag{18}$$

  for each $t = 0, \ldots, T$, and each $m = 1, \ldots, M$, where μ_{can} and σ_{can} are the average and standard deviation of all the displacements for Cantiano, respectively, and μ_{tad} and σ_{tad} are the average and standard deviation of all the displacements for Tadcaster, respectively.
- *Predictions post-processing*: let T_{can} and T_{tad} be the last instants of the time series of Cantiano and Tadcaster, respectively. Since the model has been trained with respect to synthetic bridges that collapse in T_{tad} steps, and $T_{\text{can}} \gg T_{\text{tad}}$, it is necessary to rescale the risk increment predictions to adapt them to the new time window length. Specifically, the risk index predicted for the Cantiano bridge is given by

$$\widehat{r}'(t) := \widehat{r}'_{\text{model}}(t) \frac{T_{\text{tad}}}{T_{\text{can}}} \,, \tag{19}$$

  where $\widehat{r}'_{\text{model}}(t)$ is the "raw" risk prediction returned by the model with respect to the pre-processed displacements of Cantiano.

Given these premises, we apply the trained model to the displacements of the Cantiano bridge, obtaining the results illustrated in Figure 9. It is interesting to notice that the model returns risk predictions that have an almost-linear shape, like the general trend of the ground truth risk $r'(t)$. Probably, the cause of the almost-linear risks (both predicted and ground truth) is in the particular behavior of Cantiano's displacements. Indeed, for Cantiano we observe an irregular, oscillatory behavior of the displacement measurements; nonetheless, the oscillations of the bridge-points are all in a range between $[-10\text{mm}, 10\text{mm}]$, without points characterized by sudden, steeper decrements (like, instead, in the case of the Tadcaster bridge). Summarizing, such oscillations are probably the cause of the linear shape of risk predictions.

However, we observe that the model's prediction performance for Cantiano is analogous to the one obtained for Tadcaster: a good behavior prediction for risk r', with a generic tendency to overestimate the risk (see Figure 9 and Table 4). As we did for the Tadcaster bridge, also for this validation experiment we measure the days t on which the bridge actually reached specific risk levels, as well as the days t on which these risk levels were predicted by the model. Looking at the "early warning" values in Table 4, we notice that for almost all the risk values we have an overestimate of the risk that is of approximately 1.6 years in the average ($\sim$ 14.5% of the monitoring interval of $T = 4124$ days).

These results are extremely important for proving the potentialities of the method, the modelization, and the NN model proposed in this work. Indeed, even just by improving

the amount of available data for training, we can aim to reach even better generalization performances in future models.

TABLE 4. Quantification of the risk overestimate of the model for the Tadcaster bridge.

risk	day t		early warning	
value	pred.	true	days	% (over T)
0.50	1775	2220	-445	-10.79%
0.55	1953	2476	-523	-12.68%
0.60	2130	2724	-594	-14.40%
0.65	2308	2959	-651	-15.79%
0.70	2485	3147	-662	-16.05%
0.75	2663	3377	-714	-17.31%
1.00	3550	4124	-574	-13.92%

V. DISCUSSIONS

The results presented in this study provide a comprehensive assessment of bridge monitoring using InSAR-derived displacement time series. When compared with previous studies, our findings show both consistencies and unique contributions. For instance, Selvakumaran et al. [?]db@bib:24 and Talledo et al. [26] reported that persistent scatterer InSAR (PSI) techniques can reliably detect millimeter-scale deformations on long-span bridges, particularly under favorable radar line-of-sight conditions. Similarly, Quqa et al. [21] highlighted the importance of temporal resolution and data continuity in capturing seasonal and operational variations in bridge behavior. In line with these studies, our results confirm that high-resolution InSAR observations can detect subtle deformation patterns, but they also emphasize the role of synthetic data in testing monitoring frameworks under controlled conditions.

Since bridge inspections are typically carried out on a scheduled or random basis, often without prior insight into the actual structural condition of the bridge, the proposed early-warning system is intended to address this limitation by directing inspection priorities toward bridges that show a higher estimated risk of collapse, rather than attempting to predict the exact moment of failure. The application of the model to Tadcaster and Cantiano bridge showed overestimation of the risk (Table 3 and Table 4). Nevertheless, the observed overestimation corresponds to a conservative bias, which is preferable in safety-critical contexts such as bridge monitoring, where underestimation of risk would be far more catastrophic. The proposed tool can thus serve as a decision-support mechanism, providing valuable intuition for optimizing inspection and maintenance strategies. By enabling more targeted interventions, it has the potential to reduce maintenance costs and, ultimately, to mitigate the frequency of bridge collapses through improved resource allocation.

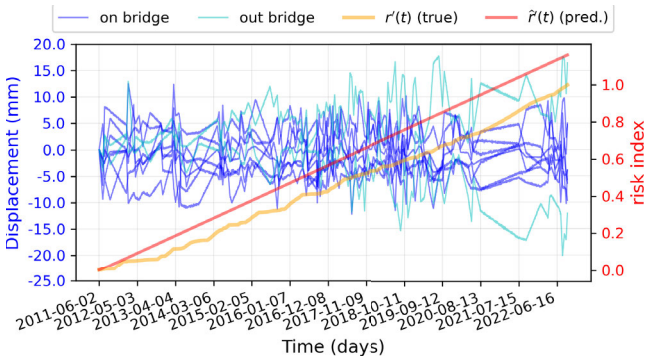

FIGURE 9. Displacements (blue and cyan curves) and $r'(t)$ risk prediction (red curve) and ground truth (orange curve) for the Cantiano bridge. Left-hand side y-axis for the displacements, right-hand side y-axis for the risk indices.

Nevertheless, some limitations must be acknowledged. The synthetic data generation process, while useful for benchmarking, may not fully capture the complex environmental and structural interactions present in real bridges. Differences in bridge typology, including span length, structural system, and material properties, can significantly influence the displacement response and the performance of InSAR-based monitoring. Therefore, while the presented methodology demonstrates promising capabilities, its generalization to diverse bridge types requires careful validation.

Future work should address these limitations by exploring multimodal data integration. Combining InSAR observations with environmental and operational data, such as temperature, traffic load, and strain measurements, could enhance the detection and interpretation of bridge deformations. Such an integrated approach may also facilitate the development of predictive maintenance models, enabling more effective and timely intervention. Furthermore, improvements in synthetic data realism, informed by availability of larger and more informative datasets, would strengthen the testing and calibration of InSAR monitoring frameworks.

Overall, the discussion highlights that while InSAR offers substantial potential for bridge health monitoring, attention to data limitations, structural variability, and complementary information sources is crucial for practical implementation.

VI. CONCLUSION

This work introduced a novel data-driven methodology for predicting bridge failure risk using InSAR displacement data processed through a Deep Learning model based on Multi-Head Attention. A mathematically-defined, quantitative risk index was proposed to formalize the concept of bridge-collapse probability directly from displacement time series, allowing the problem to be framed as a supervised learning task.

Due to the limited availability of real InSAR data for collapsed bridges, a dedicated synthetic dataset was generated through geometrical transformations and perturbations of the Tadcaster bridge displacement data. This synthetic data generation strategy enabled the training of a Neural Network

capable of estimating the risk increment over short time windows, by jointly exploiting the displacements of bridge and surrounding geospatial points.

The trained model was validated on the real Tadcaster bridge data, showing good predictive performance and a consistent tendency to slightly overestimate risk, a desirable behavior for safety-critical applications. A further stress test on the Cantiano bridge confirmed the robustness of the model, as it reproduced the overall temporal behavior of the risk despite the different structural typology, data statistics and duration of monitoring.

Overall, the results demonstrate that InSAR-derived displacements, when combined with a Multi-Head Attention-based Neural Network architecture, can meaningfully contribute to the development of early warning systems for bridge monitoring. The proposed framework thus represents a promising step toward data-driven, scalable approaches for assessing bridge stability using remote sensing information.

Future developments will focus on expanding the available dataset with additional real and synthetic InSAR measurements to improve the generalizability of the model and the reliability of the synthetic data generated. In addition, more advanced architectures will be investigated to handle a larger number of geospatial points, more complex spatial configurations, and environmental data as furhter input features.

APPENDIX A
DETAILS ABOUT TADCASTER'S DATA

The Tadcaster bridge dataset has been recovered from [25]. The data in [25] have been computed by using 48 TerraSAR-X scenes [19], acquired from ascending passes over the bridge. The acquisitions were collected in StripMap mode, with a ground resolution of approximately 3m $\times$ 3m, HH polarization, and a mean incidence angle of about 21.4°. The temporal coverage spans from 9 March 2014 to 26 November 2015. Unfortunately, no acquisitions are available between late November and the collapse of 29 December 2015, meaning that the very last phase of pre-failure evolution was not captured.

The InSAR images were processed following the Small BAseline Subset (SBAS) approach for multi-temporal interferometry (MT-InSAR) [3]. To remove the topographic phase from the interferograms, the Shuttle Radar Topography Mission Digital Elevation Model at 3 arc-second resolution (SRTM 3-arcsec DEM) as the reference was employed as the reference DEM, as explicitly reported in [?]db@bib:24. This processing workflow ultimately produced the dataset analyzed in [25]. For full technical details, the reader is referred to [?]db@bib:24, [?]db@bib:26 and [25].

APPENDIX B
DETAILS ABOUT CANTIANO'S DATA

The raw InSAR scenes of the Cantiano bridge were acquired by the COSMO-SkyMed constellation (Italian Space Agency) in ascending and descending orbital geometries. Then, these raw data have been processed by the NHAZCA S.r.l. company to obtain displacement time series in the vicinity of collapsed bridges. Specifically, two multi-temporal stacks were compiled, covering the period from February 2011 to October 2022 and including a total of 294 Single Look Complex (SLC) images (138 ascending, 156 descending). Processing was performed using NHAZCA's proprietary Advanced Differential SAR Interferometry (A-DInSAR) algorithm, to derive displacement time series and average annual velocities for persistent scatterers. Persistent scatterer candidates (i.e., geospatial points candidates) have been selected based on amplitude stability and spatial coherence thresholds to ensure temporal signal stability. Atmospheric phase screen (APS) contributions have been estimated and removed to improve the accuracy of displacement rates. For topographic phase removal and geocoding, the digital elevation model TINITALY/01 with 10m resolution [29] has been used. Overall, the whole methodology follows similar and recent works in satellite bridge monitoring (e.g., see [5]).

Actually, the interferometric analysis covered four zones between the provinces of Ancona and Pesaro Urbino (Marche Region, Italy). Nevertheless, due to space reasons, the study reported here focuses exclusively on the Cantiano area. This site includes one of the bridges that collapsed during the flood of 15 September 2022 and lies within a built-up area, enabling the identification of reliable persistent scatterers (i.e., geospatial points with reliable displacement timeseries). The selected area allows targeted assessment of pre-collapse ground displacement time series, using only the subset of interferometric results pertaining to Cantiano.

ACKNOWLEDGMENT

Marco Civera is supported by the Centro Nazionale per la Mobilitá Sostenibile (MOST–Sustainable Mobility Centre), Spoke 7 (Cooperative Connected and Automated Mobility and Smart Infrastructures), Work Package 4 (Resilience of Networks, Structural Health Monitoring and Asset Management). This manuscript reflects only the authors' views and opinions; neither the European Union nor the European Commission can be considered responsible for them.

REFERENCES

[1] E. Aktan, I. Bartoli, B. Glišić, and C. Rainieri, "Lessons from bridge structural health monitoring (SHM) and their implications for the development of cyber-physical systems," *Infrastructures*, vol. 9, no. 2, p. 30, Feb. 2024.

[2] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, *arXiv:1607.06450*.

[3] P. Berardino, G. Fornaro, R. Lanari, and E. Sansosti, "A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 11, pp. 2375–2383, Nov. 2002.

[4] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 10, pp. 281–305, 2012.
[5] V. Costantini, B. Chiaia, M. Civera, A. Ciavattone, D. Ambrosio, C. Ranalletta, E. D. Monte, R. Marini, and P. Mazzanti, "ISABHEL (integrated satellite and ground-based monitoring for bridge health lifetime assessment)," in *Proc. 13th Int. Conf. Struct. Health Monit. Intell. Infrastruct. (SHMII-13)*, Graz, Austria, Mar. 2025, pp. 1130–1139.
[6] M. Crosetto, O. Monserrat, M. Cuevas-González, N. Devanthéry, and B. Crippa, "Persistent scatterer interferometry: A review," *ISPRS J. Photogramm. Remote Sens.*, vol. 115, pp. 78–89, Apr. 2015.
[7] M. D'Angelo, M. Civera, P. F. Giordano, P. Borlenghi, F. Ballio, M. P. Limongelli, and B. Chiaia, "Bridge collapses in Italy across the 21st century: Survey and statistical analysis," *Struct. Infrastruct. Eng.*, pp. 1–23, Apr. 2025.
[8] A. Ferretti, C. Prati, and F. Rocca, "Permanent scatterers in SAR interferometry," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 1, pp. 8–20, 2001.
[9] E. Ghaderpour, P. Mazzanti, F. Bozzano, and G. S. Mugnozza, "Ground deformation monitoring in the sacco river valley (Central Italy) using multi-temporal InSAR analysis," *Remote Sens. Appl., Soc. Environ.*, vol. 33, May 2024, Art. no. 101067.
[10] P. F. Giordano, A. M. Zocca, L. Cascini, M. Manunta, R. Lanari, and A. Meda, "Uncertainty propagation in satellite InSAR data analysis for structural health monitoring of bridges," *Automation Construction*, vol. 177, Sep. 2025, Art. no. 106371.
[11] R. F. Hanssen, *Radar Interferometry: Data Interpretation and Error Analysis* (Remote Sensing and Digital Image Processing), vol. 2. Cham, Switzerland: Springer, 2001.
[12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
[13] G. Herrera, R. Tomás, J. M. Lopez-Sanchez, J. Delgado, J. J. Mallorqui, S. Duque, and J. Mulas, "Advanced DInSAR analysis on mining areas: La union case study (Murcia, SE Spain)," *Eng. Geol.*, vol. 90, nos. 3–4, pp. 148–159, Mar. 2007.
[14] A. Hooper, P. Segall, and H. Zebker, "Persistent scatterer interferometric synthetic aperture radar for crustal deformation analysis, with application to Volcán Alcedo, Galápagos," *J. Geophys. Res., Solid Earth*, vol. 112, no. B7, Jul. 2007, Art. no. B07407.
[15] E. Intrieri, F. Raspini, R. Fumagalli, F. Luzi, and N. Casagli, "The contribution of satellite SAR data to landslide forecasting: The poggio baldi landslide, northern Apennines, Italy," *Landslides*, vol. 15, no. 3, pp. 407–420, 2018.
[16] B. M. Kampes, *Radar Interferometry: Persistent Scatterer Technique*. Cham, Switzerland: Springer, 2006.
[17] X. Kang, B. Zhu, Y. Cai, Y. Xiao, N. Liu, Z. Guo, Q.-A. Wang, and Y. Luo, "A concise review of state-of-the-art sensing technologies for bridge structural health monitoring," *Sensors*, vol. 25, no. 17, p. 5460, Sep. 2025.
[18] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. 3rd Int. Conf. Learn. Represent.*, 2014, pp. 1–15.
[19] W. Pitz and D. Miller, "The TerraSAR-X satellite," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 2, pp. 615–622, Feb. 2010.
[20] L. Pulvirenti, G. Squicciarino, E. Fiori, L. Candela, and S. Puca, "Analysis and processing of the COSMO-SkyMed second generation images of the 2022 Marche (central Italy) flood," *Water*, vol. 15, no. 7, p. 1353, Apr. 2023.
[21] S. Quqa, A. Palermo, F. Ubertini, and A. Marzani, "Regional-scale bridge condition monitoring using InSAR displacements and environmental data," *J. Infrastruct. Preservation Resilience*, vol. 24, no. 4, pp. 2271–2291, Jul. 2025.
[22] *Linee Guida Per L'Utilizzo Dei Dati Interferometrici Satellitari AI Fini Dell'interpretazione Del Comportamento Strutturale Delle Costruzioni, 2023*, ReLUIS—Rete dei Laboratori Universitari di Ingegneria Sismica e Strutturale and Dipartimento della Protezione Civile and CNR-IREA, Naples, Italy, 2023.
[23] S. Selvakumaran, S. Plank, C. Geiß, and C. Rossi, "Using InSAR stacking techniques to predict bridge collapse due to scour," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2018, pp. 866–869.
[24] S. Selvakumaran, C. Rossi, A. Marinoni, G. Webb, J. Bennetts, E. Barton, S. Plank, and C. Middleton, "Combined InSAR and terrestrial structural monitoring of bridges," *IEEE Trans. Geosci. Remote Sens.*, vol. 58, no. 10, pp. 7141–7153, Oct. 2020.
[25] S. Selvakumaran, S. Plank, C. Geiß, C. Rossi, and C. Middleton, "Remote monitoring to predict bridge scour failure using interferometric synthetic aperture radar (InSAR) stacking techniques," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 73, pp. 463–470, Dec. 2018.
[26] D. A. Talledo, A. Miano, M. Bonano, F. Di Carlo, R. Lanari, M. Manunta, A. Meda, A. Mele, A. Prota, A. Saetta, and A. Stella, "Satellite radar interferometry: Potential and limitations for structural assessment and monitoring," *J. Building Eng.*, vol. 46, Apr. 2022, Art. no. 103756.
[27] D. A. Talledo and A. Saetta, "A multi-level semi-automatic procedure for the monitoring of bridges in road infrastructure using MT-DInSAR data," *Remote Sens.*, vol. 17, no. 14, p. 2377, Jul. 2025.
[28] R. Tao, A. Lau, M. E. Mossefin, G. Kong, S. Nordal, and Y. Pan, "PS-InSAR for railway anomaly detection: A case study on long-term infrastructure monitoring," *Measurement*, vol. 235, Mar. 2025, Art. no. 115599.
[29] S. Tarquini, I. Isola, M. Favalli, and A. Battistini, "Tinitaly, a digital elevation model of Italy with a 10 meters cell size (version 1.0)," Istituto Nazionale di Geofisica e Vulcanologia (INGV), 2007, vol. 10.
[30] *Early Stopping—Tensorflow Callbacks*, TensorFlow, Mountain View, CA, USA, Oct. 2023.
[31] *Reduce Learning Rate on Plateau—Tensorflow Callbacks*, TensorFlow, Mountain View, CA, USA, Oct. 2023.
[32] D. Tonelli, V. F. Caspani, A. Valentini, A. Rocca, R. Torboli, A. Vitti, D. Perissin, and D. Zonta, "Interpretation of bridge health monitoring data from satellite InSAR technology," *Remote Sens.*, vol. 15, no. 21, p. 5242, Nov. 2023.
[33] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2025, pp. 5998–6008.

FRANCESCO DELLA SANTA received the master's degree in mathematics from the University of Florence and the joint Ph.D. degree in pure and applied mathematics from the University of Turin and Politecnico di Torino. He is currently a Research Associate with Politecnico di Torino. His main scientific research interests include deep learning, surrogate models, uncertainty quantification, and numerical optimization.

FEDERICO SISCI received the bachelor's degree in mathematics for engineering from Politecnico di Torino, in 2021, where he is currently pursuing the master's degree in mathematical engineering and data science, which he is about to complete with a thesis on the application of deep learning to bridge monitoring.

FARBOD KHOSRO ANJOM received the Ph.D. degree in civil and environmental engineering from Politecnico di Torino, Italy, in 2021. Since 2023, he has been an Assistant Professor with the Department of Environment, Land, and Infrastructure Engineering, Politecnico di Torino. His research is focused on machine learning and data-driven solutions for near-surface characterization of the subsurface from seismic waves. He has been involved in many international projects, collaborating with companies, institutions, and universities from Italy, Sweden, Greece, U.K., South Africa, and France. He has published more than 30 papers and conference proceedings on seismic surface wave, seismic exploration, microseismic monitoring of natural hazards, and bridge monitoring. He is an Associate Editor of *Geophysical Prospecting* and contributes to peer review and editorial development within the applied geophysics community.

FLAVIO PINO is currently an Assistant Professor with Politecnico di Torino. His research is mostly focused on the field of digital economics, with a particular emphasis on how the collection and use of consumer data shapes competition. He is working on the competition effects of using data to price discriminate consumers in various settings, such as hybrid marketplaces, and with different data sources, such as data brokers.

MARCO CIVERA received the M.Sc. degree in civil engineering and the Ph.D. degree in aerospace engineering from Politecnico di Torino, Italy, in 2016 and 2021, respectively. Since 2023, he has been an Assistant Professor with the Department of Structural, Geotechnical, and Building Engineering (DISEG), Politecnico di Torino; on a tenure-track, since November 2025. His research focuses on vibration-based assessment of structural changes for aerospace, civil, and mechanical structures. He is a Faculty Member and a Lecturer for the courses of Structural Mechanics, Dynamics of Structures, and Collapse of Structures and Structural Health Monitoring. He is a member of the Associazione Italiana di Meccanica Teorica e Applicata (AIMETA), Società Italiana di Scienza delle Costruzioni (SISCO), and several international associations in bridge maintenance and monitoring, including IABSE, IABMAS, and FIB Working Group TG3.3.