

Advancing product surface defect detection and quality inspection with Gaussian Splatting: insights and preliminary experimental evidence

Original

Advancing product surface defect detection and quality inspection with Gaussian Splatting: insights and preliminary experimental evidence / Trombini, M., Maisano, D.A.F., Franceschini, F.. - In: PRODUCTION & MANUFACTURING RESEARCH. - ISSN 2169-3277. - STAMPA. - 14:1(2026), pp. 01-22. [10.1080/21693277.2026.2710381]

Availability:

This version is available at: 11583/3013822 since: 2026-08-01T16:44:28Z

Publisher:

Taylor & Francis

Published

DOI:10.1080/21693277.2026.2710381

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Advancing product surface defect detection and quality inspection with Gaussian Splatting: insights and preliminary experimental evidence

Mattia Trombini , Domenico A. Maisano & Fiorenzo Franceschini

To cite this article: Mattia Trombini , Domenico A. Maisano & Fiorenzo Franceschini (2026) Advancing product surface defect detection and quality inspection with Gaussian Splatting: insights and preliminary experimental evidence, Production & Manufacturing Research, 14:1, 2710381, DOI: [10.1080/21693277.2026.2710381](https://doi.org/10.1080/21693277.2026.2710381)

To link to this article: <https://doi.org/10.1080/21693277.2026.2710381>



© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 01 Aug 2026.



Submit your article to this journal [↗](#)




View related articles [↗](#)



View Crossmark data [↗](#)

RESEARCH ARTICLE



Advancing product surface defect detection and quality inspection with Gaussian Splatting: insights and preliminary experimental evidence

Mattia Trombini, Domenico A. Maisano and Fiorenzo Franceschini

Department of Management and Production Engineering (DIGEP), Politecnico di Torino, Torino, Italy

ABSTRACT

Scheduled inspection is a key element of industrial maintenance, enabling wear and deterioration to be detected before leading to severe failures in machinery or manufactured products. However, inspection records are typically constrained to viewpoints captured during acquisition, limiting the possibility of retrospectively investigating defects that were not adequately documented. This paper proposes the use of a 'scene', a high-fidelity digital representation of an object at a specific point in time, reconstructed through Gaussian Splatting (GS). Generated from a set of images, the scene acts as a digital archive of the object's physical condition and can be interactively explored from arbitrary viewpoints, including perspectives not originally acquired. This capability enhances inspection flexibility and enables more effective retrospective assessment of defect occurrence, contributing to Industry 5.0 goals, including waste reduction and improved maintenance. Successive inspection stages are preserved as navigable scenes and systematically compared to identify surface anomalies and monitor their progression.

ARTICLE HISTORY

Received 11 June 2026
Accepted 23 July 2026

KEYWORDS

Quality inspection; Gaussian splatting; surface defect detection; 3D scene; Industry 5.0

1. Introduction

The quality inspection of machinery and manufactured objects subject to long service times and progressive wear (e.g. aircraft, railway vehicles, trucks and industrial equipment) has long been driven by the need to increase accuracy, improve efficiency, and ensure safety (Ouzineb et al., 2018; Pollock et al., 2021; Sause & Jasiūnienė, 2022). For example, for wind turbine blades, technical conditions must be consistently maintained in order to avoid performance drops or a reduction in power generation (Sareen et al., 2014; Stałowska et al., 2025). Traditionally, photogrammetry is particularly widespread among maintenance companies, as it does not require physical contact with the component (Trombini et al., 2025), thereby reducing the risk of damage and limiting hazards to operators (Sareen et al., 2014), especially in manufacturing and with large-volume structures (Franceschini et al., 2009; 2011; Franceschini et al., 2016).

Laser scanning and large-volume metrology systems have also recently proven effective in the ship-building industry (Maisano et al., 2022, 2023). Maisano et al. (2022) presented a structured comparison of three centralised measuring instruments from both qualitative and quantitative perspectives. Johnson et al. (2004) applied photogrammetry to the plate-burning process; Guo et al. (2017) employed a constellation of R-LAT, while Liu and Wang (2010) proposed a methodology to address the measurement range limitations of laser trackers.

In the automotive industry (Kholkhujayev et al., 2022), inspection of key dimensional and geometrical tolerances is usually carried out with a coordinate measuring machine (CMM) (Martínez et al., 2010; Turley et al., 2014), although laser tracking systems have more recently been adopted due to their measurement speed and non-contact full-field capability (Turley et al., 2014). Vision-based techniques have also been effectively applied to the detection of defects in car bodies (Molina et al., 2017).

In recent years, deep learning has become an increasingly established paradigm for industrial surface defect detection (Maculotti et al., 2025), progressively replacing traditional machine-vision pipelines based

CONTACT Fiorenzo Franceschini  fiorenzo.franceschini@polito.it  Department of Management and Production Engineering (DIGEP), Politecnico di Torino, 10129 Torino, Italy

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

on hand-crafted features and classical classifiers. Early convolutional approaches already showed promising results for classifying defects on industrial surfaces, such as steel and rail defects (Faghih-Roohi et al., 2016; Masci et al., 2012). Later works extended deep architectures towards defect localisation and segmentation, including fully convolutional and segmentation-based networks for surface anomalies and cracks (Rački et al., 2018; Tabernik et al., 2020). More recently, industrial anomaly detection has been strongly influenced by benchmark datasets such as MVTec AD and by methods such as PatchCore and DR $\bar{E}$ M, which aim to detect and localise anomalous regions from image data (Bergmann et al., 2021; Roth et al., 2022; Zavrtnik et al., 2021).

Although inspection techniques are widely used across industrial domains, current approaches still produce representations that are static and difficult to compare across different temporal inspection times (Trombini et al., 2025). Point-cloud or 2D outputs provide useful geometric data, but they lack navigability and do not allow inspectors to change viewpoint or interactively examine the object. More importantly, they do not ensure *back-in-time inspection*, defined as the capability to revisit any past state of the object from any desired perspective and viewing direction, thereby limiting the building of reliable digital archives, tracking defect progression and supporting rapid decision-making in the event of sudden damage (Trombini et al., 2025). There is therefore a clear need for an inspection framework that delivers high-fidelity, full-field 3D measurements that are both temporally consistent and fully navigable, allowing inspectors to revisit past states, analyse defects from different perspectives and gain deeper insights into the condition of industrial assets.

By introducing the concept of a ‘3D scene’ (Sitzmann et al., 2019; Tewari et al., 2022), the limitations of traditional inspection methods can be mitigated. A 3D scene is a digital high-fidelity representation of an object or environment at a specific point in time. Scene generation requires the collection of geometric data enriched with information on colour, texture, shading, lighting and camera perspectives to achieve the highest possible level of realism (Alatan et al., 2007). The 3D scene is fully navigable, enabling users (e.g. inspectors, engineers) to examine the object from alternative digital viewpoints. In addition, new digital images, i.e. renders, can be generated from previously unexplored perspectives, thereby broadening the range of visual analysis and image processing. Furthermore, the scene introduces the innovative capability of remote inspection, allowing inspectors to perform defect detection without being physically present on site.

Among the most relevant methodologies for generating 3D scenes, radiance fields, which describe the distribution of light in 3D space and the way it interacts with surfaces and materials, are attracting increasing attention. The current state of the art is the Neural Radiance Field (NeRF) (Mildenhall et al., 2021), capable of producing highly accurate 3D scenes by processing images and video recordings but requiring extensive computational resources. More recently, Kerbl et al. (2023) introduced 3D Gaussian Splatting (GS), a method for real-time rendering that delivers performance comparable to NeRF while significantly reducing computational time (Kerbl et al., 2023; Mallick et al., 2024). GS models the scene using 3D anisotropic ellipsoids, which are projected onto the 2D image plane to generate renders. In Trombini et al. (2025), a GS-based scene has been exploited to reconstruct the past state of an aeronautical object, generating renders to be integrated into the photogrammetric workflow in order to enhance acquisition coverage and better characterise a misreconstructed geometric defect.

Since the concept of scene is not currently employed for inspection tasks, this study proposes the innovative integration of one of the most recent computer vision techniques, i.e. GS, into manufacturing. This paper discusses a semi-automatic procedure to analyse 3D scenes of the same object at different points in time in order to assess its effectiveness for defect detection and tracking over time. At the current stage, the proposed methodology cannot be considered as a fully automatic inspection technique since the supervision of the operator is still mandatory to establish whether a difference between two scenes corresponds to a physical anomaly (i.e. a defect) rather than a false positive. The methodology is verified through a case study on an automotive brake disc, where both a pristine (defect-free) and a defective state are reconstructed as 3D GS scenes. By systematically comparing the two reconstructions, the presence and morphology of surface defects can be identified. Furthermore, a set of quantitative indicators is defined, and statistical analyses are applied to demonstrate the significance of the observed differences, highlighting the potential of scene-based inspection for reliable defect detection. For the sake of clarity, the present study only focuses on chromatic defects, leaving the geometric aspects to be treated as future developments.

of this methodology. In addition, the main step forward compared with previous research (Trombini et al., 2025) lies in the capability of using GS as an independent and standalone tool for quality inspection, without relying on photogrammetry.

The paper is organised as follows: [Section 2](#) introduces the concept of scene for inspection activities, [Section 3](#) examines the characteristics of scene generation and GS, [Section 4](#) explains the *back-in-time inspection*, [Section 5](#) outlines the methodology, [Section 6](#) describes the case study, and [Section 7](#) presents the results. Finally, [Section 8](#) includes the discussion and [Section 9](#) provides the conclusions and future developments.

2. Scene-based inspection

The adoption of scenes in manufacturing is largely unexplored. Originating from computer vision and computer graphics, a scene is defined as a high-fidelity and realistic representation of an object or environment. This leads to a fundamental research question: *how can a scene become an essential tool in quality inspection?* Addressing this question requires first identifying the existing gap that justifies the use of scenes. For example, in the aeronautical sector, scheduled and comprehensive inspections are indispensable to ensure safety and airworthiness. Consider a surface inspection intended to detect corrosion, paint loss, scratches or more general chromatic and geometric defects on a section of the fuselage (Trombini et al., 2025). The following tools are currently employed to perform defect detection (Sause & Jasiūnienė, 2022):

- Visual inspection;
- Liquid penetrant testing;
- Magnetic particle testing;
- Eddy current testing;
- Thermography.

However, despite the benefits of each technique, a general overview of its main drawbacks can be identified (see [Table 1](#)).

Among the described techniques, visual inspection is the most versatile, as it can be performed using images collected with affordable high-resolution RGB cameras currently widely available and reliable. Their integration as drone payloads further enhances applicability by overcoming accessibility constraints and areas with occlusions. Furthermore, this simple example of surface inspection represents a typical case of in-service quality control and therefore limits the applicability of statistical approaches, since the object under investigation is a unique and isolated case, making the process non-repeatable. A further limitation concerns the temporal dimension of the available techniques. When a new defect is detected, the inspector may be interested in determining whether it was already present during earlier inspections but remained undetected or may require a more detailed retrospective characterisation of the anomalous area. However, each technique is static-in-time: while relevant data are provided, they are restricted to the specific moment of acquisition. Accordingly, the limitations lead to the following research question: *what if an inspector could perform a digital yet realistic walk around the object at a specific moment in the past?*

The scene addresses this gap by providing a fully navigable environment in which the user is immersed. Within the scene, the observer can navigate freely and examine the object from novel viewpoints, simulating the movement around it in physical space. This capability transforms the inspection record from a static collection of images into an environment that can be explored interactively. Unlike a photograph, which conveys a fixed and immutable view captured at a single position in space, the

Table 1. The main limitations of techniques currently employed for defect detection.

Technique	Limitation
Visual inspection	Highly operator-dependent and limited by lighting and accessibility
Liquid penetrant testing	Requires very clean and properly prepared surfaces
Magnetic particle testing	Applicable only to ferromagnetic materials and requiring close contact with the surface
Eddy current testing	Sensitive to geometry, conductivity variations and surface roughness
Thermography	Strongly affected by environmental conditions

scene constitutes an immersive and navigable representation of the object, reflecting the same temporal state of the physical acquisitions. When interacting with the scene, the inspector conducts activities including:

- Close-up inspection;
- Comparative analysis;
- Simulation of inspection conditions;
- Documentation and reporting;
- Tracking defects over time.

In this way, the scene transforms the inspection process from a static assessment into an interactive activity in which defects can be examined from multiple perspectives and with greater detail. To conclude, in a sustainable manufacturing context, the scene is not only a visualisation tool but also a reusable digital asset for manufacturing inspections. Once generated, the same scene can be re-analysed and compared against future or past states without requiring the entire physical acquisition to be repeated. This creates a link between defect detection and lifecycle decision-making, since the digital tool can support the selection of repair actions and the preparation of documentation for digital and reusable archives.

3. Fundamentals of GS

Among the available approaches for constructing 3D scenes, GS has recently emerged as a significant breakthrough (Kerbl et al., 2023), representing a notable shift from state-of-the-art radiance field techniques (Mildenhall et al., 2021). GS generates a 3D scene by processing a set of input images (i.e. ground-truth images) and representing it with learnable 3D anisotropic Gaussian ellipsoids. The ellipsoids are defined by their position, scale, opacity, and colour, enabling a detailed and adaptable reconstruction of the scene. Structure-from-Motion (SfM) (Snavely et al., 2006) is employed to perform camera calibration and generate a sparse point cloud, which provides the baseline for the initial placement of the Gaussians. Typically, SfM is carried out in COLMAP (Schönberger & Frahm, 2016), a general-purpose SfM and Multi-View Stereo (MVS) pipeline with graphical and command-line interfaces.

During training, the parameters of the Gaussian ellipsoids are iteratively refined so that the rendered projections match the original images. In this process, new ellipsoids can also be introduced in regions where a more refined scene description is required. As a result, the set of Gaussians progressively evolves into a coherent and photorealistic model of the scene. In a somewhat unconventional interpretation, GS resembles an ‘impressionist’ technique, where colour blending creates a perceptually accurate representation of the scene. Compared to alternative radiance field techniques (e.g. NeRF), GS provides several advantages (Kerbl et al., 2023):

- *Computational efficiency*: GS enables real-time rendering and interactive navigation, overcoming the high training and inference costs typical of radiance field methods.
- *Adaptive representation*: the use of anisotropic Gaussian ellipsoids allows detailed local refinement, enhancing the reconstruction of fine details.
- *Photorealistic fidelity*: GS achieves high spatial coherence and extremely accurate visual realism.

Novel viewpoints can then be generated by re-projecting the optimised ellipsoids onto the 2D plane of the image, ensuring spatial coherence and visual realism. The viewpoints generated within the scene are not limited to virtual observations but constitute digital images (i.e. renders) that can be extracted and treated as photographs.

The output of GS is a 3D scene that can be condensed into a dense point cloud enriched with multiple parameters. Unlike conventional point clouds obtained through laser scanning or photogrammetry, the points in GS not only carry 3D spatial coordinates but also correspond to the ellipsoid centres and encode additional attributes such as colour, orientation, size, shape, extension and opacity (Kerbl et al., 2023). Colour is defined by means of the so-called spherical harmonics (SH) (Green, 2003; Kerbl et al., 2023).

Kerbl et al. (2023) use harmonics up to degree 3, and consequently, the total number of scalar coefficients is equal to 48 (16 for each channel).

During an inspection phase, the operator can navigate different scenes of the object of interest, each captured at distinct time instances. These scenes can then be compared to detect variations that may indicate anomalies, and therefore, to support this process, an alignment procedure is required, at least in a semi-automated form. Scene alignment allows renders to be generated from consistent perspectives, thereby enabling image-based analysis for defect detection. COLMAP generates a self-consistent sparse point cloud, but the scale is arbitrary, and the spatial location depends on the seed pair, i.e. the pair of images with the highest shared information (Schönberger & Frahm, 2016).

4. Back-in-time inspection

The introduction of the 3D scene concept into manufacturing provides the novel capability of freezing temporal information from an inspection task and revisiting it whenever required. As anticipated in Section 2, current machine vision technologies, such as photogrammetry, allow spatial reproducibility, since predetermined camera positions can be repeated by fixing the observation directions. However, traditional methods lack the temporal dimension because they are constrained by the initial image acquisition. The main limitation is that if an unexpected defect is detected or if a specific portion of the object requires further observation from different viewpoints, conventional methods do not allow the addition of new observation points. In addition, new acquisitions may not be feasible for some reasons, including (i) inaccessibility of the object, For instance when it is located remotely, (ii) rapid changes such as vibration or thermal variation that may already have altered the object condition, and (iii) any new physical acquisition (i.e. new photographs) would reflect the condition of the object in a different temporal state.

A 3D scene, by contrast, is navigable and enables digital exploration of the object also from novel viewpoints. Technically, it can be regarded as a virtual walk around the object in a specific temporal state. This leads to the concept of *back-in-time inspection*, defined as the ability to revisit any past state of the object recorded during inspection and to observe it from new and relevant perspectives.

In the present study, the capability of superimposing scenes is exploited to detect defects emerging at different temporal states. This *back-in-time* functionality enables the recovery of past object states and supports verification of whether suspicious anomalies were already present during previous inspections. To this end, *back-in-time inspection* transforms industrial quality maintenance inspection by embedding retrospective inspection capabilities into routine processes. It ensures full traceability of asset conditions over time, strengthens compliance with industrial quality standards and enables informed decision-making for maintenance scheduling and process optimisation.

Furthermore, back-in-time inspection also introduces sustainability dimensions. By preserving previous states as reusable digital assets, the same inspection record can be re-analysed when new anomalies, stricter acceptance criteria or unexpected operating events emerge. This capability may reduce repeated physical acquisitions, support condition-based decisions and strengthen the continuity of maintenance knowledge when assets are geographically distributed or difficult to access.

5. Methodology

Before delving into the details of the methodology for surface analysis using GS, it is useful to outline the operational procedure to be adopted during the inspection phase:

- Time t_0 : the object is inspected, for example, due to scheduled maintenance activity. The inspection activity also involves the collection of images (i.e. ground-truth images), which are used to create the 3D GS-based scene. The collection of images is an *on-site* activity.
- Time t_1 : after being put back into service, the object is subjected to a new stage of scheduled maintenance activity that involves the generation of the GS scene (i.e. the scene at time t_1). During the inspection, various unexpected situations may arise, including (i) the detection of an unexpected defect, (ii) an increase in inspection severity, leading to the reclassification of certain anomalies as more

critical and (iii) previously unexamined areas exhibiting suspicious discolouration or geometric changes. In such cases, the inspector may need to revisit a specific region of the component to determine whether any anomalies were already present but remained undetected or were absent. With the scene, any past state of the object can be revisited by navigating the scene from a specific moment in time. The activity can involve superimposing different scenes or digitally observing the object from novel perspectives to perform a visual check.

Regarding the acquisition setup, the object of interest is assumed to remain stationary while a camera (e.g. a smartphone) moves around it during acquisition, ensuring that all relevant viewpoints are covered. The acquisition process, using photographs or videos, can be performed by a non-specialist operator, since the inspector (e.g. the engineer) can analyse the scene even remotely. If the acquisition is made via video, a set of frames is extracted by sampling the video, and they are used to reconstruct the 3D scene through GS. During an inspection phase, the operator can navigate different scenes of the object of interest, each captured at distinct time instances. These scenes can then be compared to detect variations that may indicate anomalies, and therefore, to support this process, an alignment procedure is required, at least in a semi-automated form. Scene alignment allows renders to be generated from consistent perspectives, thereby enabling image-based analysis for defect detection. COLMAP generates a self-consistent sparse point cloud, but the scale is arbitrary, and the spatial location depends on the seed pair, i.e. the pair of images with the highest shared information (Schönberger & Frahm, 2016). The registration task is not straightforward:

- xyz coordinates of Gaussian centre-points can be rotated, translated and scaled.
- The covariance matrix is rotated and scaled.
- The SH coefficients, which encode the view-dependent appearance of each Gaussian, must also be rotated. This step is particularly delicate because, unlike positions or vectors, SH coefficients transform through specific rotation matrices (Wigner-D matrices) that ensure the underlying spherical function (i.e. the directional appearance) rotates consistently with the Gaussian orientation (Green, 2003).

Alignment enables identical camera poses to be used to observe corresponding regions in both scenes, allowing direct analysis of potential differences in texture or surface properties. Since the scenes are aligned, renders can be compared at the pixel level using metrics designed to highlight colour differences. In the scenario of chromatic defects, the metric adopted for render comparison is the gradient (Gonzalez & Woods, 2006) (see Appendix A for a proper definition).

5.1. Indicators

The central aim is the identification of a robust methodology to compare two scenes. The presence of a defect induces a variation in the distribution and geometry of ellipsoids, and therefore, two indicators can be used to assess the differences. The first indicator is the *density*, defined as the number of ellipsoid centres within a reference surface, defined according to Equation (1). This concept directly extends the notion of density commonly used in point cloud processing, where it supports the evaluation of reconstruction quality, guides resampling, and describes local surface properties (Rusu et al., 2009). In the present context, a higher density indicates that the surface requires a larger number of ellipsoids to achieve a more accurate representation, thereby reflecting increased geometric complexity or irregularity. A lower density corresponds to smoother and more homogeneous surface regions. The second indicator is the *average volume*, computed in Equation (2) as the mean volume of ellipsoids within a reference area. This quantity relates to the concept of multiscale representation in point-based rendering (Rusinkiewicz & Levoy, 2001) and has been identified in recent work on 3D GS (Kerbl et al., 2023) as an indicator that balances detail preservation and efficient surface coverage. Large volume indicates that the surface is represented by fewer and coarser ellipsoids, while small volume corresponds to finer detail and higher local variability.

$$Density = \frac{N_{Gaussians}}{Area} \quad (1)$$

$$Average \ volume = \frac{1}{N_{Gaussians}} \sum_{i=1}^N V_{Gaussian}^i = \frac{1}{N_{Gaussians}} \sum_{i=1}^N \frac{4}{3} \pi a_i b_i c_i \quad (2)$$

where a_i , b_i and c_i are the three main semiaxis of the i -th ellipsoid, while $N_{Gaussians}$ is the number of ellipsoids whose centre point is inside the reference area.

However, the scene does not contain only geometric data but also embeds colours. To assess chromatic differences between two scenes, the following procedure is adopted:

- For a selected portion of the scene (i.e. a reference area), all SH scalar coefficients associated with the Gaussians are extracted from both scenes, resulting in two matrices of size $N \times 48$ and $M \times 48$, where N and M denote the number of ellipsoids within the reference area. For the sake of clarity, as written in [Section 3](#), 48 is the total number of SH coefficients.
- The Jensen–Shannon divergence (JSD) is then used to compare, for each coefficient, the distributions within the two scenes. The square root of the Jensen–Shannon divergence is the Jensen–Shannon distance (JSd) and is a proper metric that satisfies the triangle inequality and can be adopted as a valid measure for the statistical analysis deployed in the present work (Lin, 1991).
- Expanding the evaluation of the Jensen–Shannon distance along the columns of the matrices, we obtain a vector of 48 JSd values for each portion of the scene. Differences in the distribution of JSd across the portion can detect significant chromatic deviations.

Considering two discrete probability distributions P and Q over a common domain X , JSD is defined as follows:

$$JSD(P||Q) = \frac{1}{2} D_{KL}(P||M) + \frac{1}{2} D_{KL}(Q||M) \quad (3)$$

with

$$M = \frac{1}{2} (P + Q) \quad (4)$$

and

$$D_{KL}(P||M) = \sum_X P \cdot \log \frac{P}{M} \quad (5)$$

M is a mixture distribution of P and Q , and D_{KL} is the Kullback–Leibler (KL) divergence (Kullback & Leibler, 1951). The Jensen–Shannon distance is therefore defined as:

$$Js d(P, Q) = \sqrt{JSD(P||Q)} \quad (6)$$

Other than JSD, several indicators are commonly employed for comparison tasks between distributions. The Kullback–Leibler divergence (Kullback & Leibler, 1951) is widely used in information theory to quantify differences between distributions, while the Wasserstein distance (Givens & Shortt, 1984) originates from optimal transport theory and measures the cost of transforming one distribution into another. [Table 2](#) summarises the main properties in terms of symmetry and boundedness that help justify the use of JSD.

Table 2. Summary of symmetry and range properties of Kullback–Leibler divergence, Jensen–Shannon divergence and Wasserstein.

Indicator	Symmetry	Range
Kullback–Leibler (KL) divergence (Kullback & Leibler, 1951)	Asymmetric	Unbounded
Jensen–Shannon divergence (JSD) (Lin, 1991)	Symmetric	$\in [0,1]$
Wasserstein (Givens & Shortt, 1984)	Symmetric	Unbounded

Table 2 motivates the adoption of JSD due to its symmetry and its finite value for any pair of distributions. The JSD is used since it provides a symmetric and always finite measure of dissimilarity between distributions, while JSd is a true metric and therefore better suited for quantitative comparison.

5.2. Computational overhead

GS is entirely code-based in Python using the PyTorch framework and custom CUDA[®] for rasterisation. The 3D scene can be visualised by employing an interactive viewer using the open-source System for Image-Based Rendering (SIBR) software, complemented with the procedure of extracting the renders from novel viewpoints. In this work, the SfM and scene-generation phases were performed using an NVIDIA GeForce RTX 5080 GPU. Table 3 reports the corresponding completion time and the final scene size in terms of the number of ellipsoids. In this paper, each render required approximately 3 s to be extracted from SIBR. According to Kerbl et al. (2023), the minimum recommended specification for processing GS scenes is 8 GB of VRAM, while 24 GB is suggested for more complete reconstructions. In the present case, 16 GB of VRAM were sufficient to generate the scene at full resolution in approximately 10 min.

6. Case study

The adopted case study concerns the brake disc of a light commercial vehicle, identified as a small van. During braking action, the calliper clamps the pads onto the rotating disc, generating the friction force responsible for vehicle deceleration. The combination of cyclic thermal and mechanical stresses, together with the disc's direct exposure to ambient air, makes it prone to oxidation and corrosion. These factors contribute to fatigue, surface damage and structural deterioration over time, making the brake disc a relevant component for defect analysis and failure prediction. Figure 1 shows two views of the disc, identifying surfaces A and B (see Figure 1a for A and Figure 1b for B), including the portions of surface under investigation (Figure 1c and d).

Image collection was carried out using an iPhone 14 Plus, with the aim of demonstrating the effectiveness of the proposed methodology even when employing a widely available consumer device. The strategy adopted for generating the GS-based scene can be summarised as follows:

- The object is mounted in a convenient position, i.e. a position useful for collecting images and video from all the required views. **On-site**
- An operator records a video or takes images, ensuring that a sufficient number of views of the object are captured. In some cases, coded markers can be placed on the surface to improve tracking and reduce ambiguities in areas with few reference points. No markers have been used in the present study. **On-site**
- The SfM and GS processes are then performed sequentially. **Post-processing**

Let t_0 be the time of the first inspection; no previous condition is available, and no surface defects are detected. At time t_1 , corresponding to the second inspection, oxidation spots appear as small black areas visible on the working surface of the brake disc, as shown in Figure 2. For clarity, nine oxidation spots are identified and used as the reference benchmark.

On surface B at time t_1 , the disc shows an extended scratch over a large portion of the surface, consistent with the damage caused by a worn brake pad, as illustrated in Figure 3.

7. Results

For clarity, the scenes at time t_0 and t_1 are classified as follows:

Table 3. Specifications of GPU employed for the GS-scene generation.

GPU	VRAM [GB]	Time for scene generation (approx.) [s]	Number of ellipsoids in the GS scene
NVIDIA GeForce RTX 5080	16	675	1,071,885




Figure 1. Identification of surfaces A (a) and B (b) of the brake disc considered as a case study. The two images correspond to opposite views of the same component, obtained by flipping the disc. Images (c) and (d) highlight the specific regions of interest that are selected for the surface analysis presented in this work.




Figure 2. Surface A with damage (i.e. oxidation spots). Red circles highlight the physical oxidation spots.



Figure 3. Surface B with damage (i.e. scratch). The red box highlights a part of a physical scratch over the surface.

- Pristine scene: it corresponds to the condition of the object at time t_0 , when no evident defects are present, and the component can be considered defect-free.
- Defective scene: it corresponds to the condition of the object at time t_1 , where both oxidation spots and an extended scratch are detected on the surface.

7.1. Surface A

As shown in Figure 2, surface A contains oxidation spots that are absent in the pristine condition (see Figure 1a). Once the two scenes are generated, they undergo the alignment process described in Section 3 and are then observed from the same perspective, enabling the generation of consistent renders. The chosen perspective (or viewing angle) does not necessarily correspond to a physical acquisition viewpoint but is selected at the operator's discretion. Suitable perspectives can be selected to facilitate the identification of potential anomalies, for example, a perpendicular view centred on a suspicious area. The renders are compared in terms of gradient differences, following the procedure previously described. In this case, the gradient difference map is shown in Figure 4a, while the application of a minimum threshold set by the inspector enables the segmentation of defective regions (see Figure 4b).

Each defect is enclosed within a box, as shown in Figure 5a. Once the boxes are identified as shown in Figure 5b, intrinsic and extrinsic camera parameters (Trombini et al., 2025) enable the superimposition of the scene onto the render image plane, thereby allowing the selection of the ellipsoids responsible for the description of the defect. In the present case, to better visualise the defects, the scene is modified by assigning a fixed red colour to the corresponding ellipsoids, as illustrated in Figure 6a–c.

In the defective scenes, the clusters of ellipsoids that represent the anomalies are used to define spatial masks on the surface. Proper alignment of the scenes allows the application of spatial masks to the pristine scene, enabling the isolation of the same surface portion highlighted in the defective scene. Their application to the pristine case does not imply the presence of defects but only ensures geometric consistency in the areas being analysed. Local comparison can subsequently be performed by evaluating the number of Gaussian centre-points from both scenes that fall within each portion. The results of this analysis are summarised in Table 4, where the 'Defect ID' column refers to the spatial mask built around each defect, following the ID definition in Figures 5 and 6. The 'Defective' column refers to the Gaussians that chromatically describe the anomalies (highlighted as red ellipsoids in Figure 6), whereas the 'Pristine' column refers to the Gaussian centres from the pristine scene that fall within the same geometrical masks.

By observing Table 4, the pristine scene consistently shows a null number of Gaussian centres within each geometric boundary (except for defect 7). This outcome is explained by the chromatic and geometric homogeneity of the pristine surface. In the absence of defects, the representation relies on fewer but larger ellipsoids, and therefore, there is no need for higher-density local refinement. Consequently, the specific




Figure 4. In (a), a gradient-based map of differences between pristine and defective surface A. In (b) the filtered portions of larger differences that identify the oxidation points.




Figure 5. Detection of defective regions on surface A: (a) bounding boxes enclosing each oxidation defect and (b) magnified view of each box.




Figure 6. Visualisation of oxidation defects on surface A by colouring the corresponding defective Gaussians in red.

Table 4. Number of ellipsoid centres located within each defective area, comparing the defective and the pristine 3D scenes. The presence of additional ellipsoids in the defective scene highlights the local refinement introduced by surface anomalies.

Defect id	Number of ellipsoid centres	
	Defective	Pristine
1	7	0
2	44	0
3	7	0
4	5	0
5	1	0
6	2	0
7	4	1
8	49	0
9	17	0

area is described by one or more ellipsoids whose relative centre-point is localised somewhere else in the neighbourhood of the specific spatial mask.

Moreover, the study examines how the scene locally changes depending on whether a defect is present or absent. Within the areas identified by the previously discussed boxes, the local distribution of ellipsoids is analysed for both pristine and defective scenes. Starting from the centre of each box a sphere with increasing radius is drawn and the number of centre-points falling within the volume is counted as a measure of density. To allow meaningful comparison between pristine and defective conditions, both the radius and the density profiles are normalised, and the resulting chart is shown in Figure 7. The radius is scaled relative to its maximum value for each defect, while density values are min-max normalised to the common [0-1] range. Min-max normalisation is adopted since the densities vary across the areas, and without rescaling, the results would not be directly comparable. By mapping the distributions onto the same numerical interval, the analysis enables a consistent comparison between pristine and defective scenes.

By observing the results, some comments can be drawn:

- Gradient-based difference method is effective in detecting variations in surface texture between two GS-based scenes. Specifically, the defective surface can no longer be regarded as homogeneous since localised oxidation creates areas with different local characteristics. In these areas, the GS procedure locates small clusters of Gaussians that are responsible for the chromatic description of the defective areas and therefore the local density increases.
- In the pristine scene, the surface remains intact, and no additional local refinement is required, resulting in a lower Gaussian density. This outcome does not imply that the pristine surface is not properly described, but rather that the chromatic and geometric consistency of the surface does not require denser local refinement. In contrast, when defects are present, the discontinuities in colour and shape produce a higher density of ellipsoids, with several new centre-points placed in the affected areas to refine the surface chromatic and geometric description.




Figure 7. Normalised local density profiles of Gaussian ellipsoids for each detected oxidation defect on surface A. Density values are plotted against increasing radius to compare pristine and defective local ellipsoid density. For defects ID refer to Figures 5 and 6.

7.2. Surface B

As previously shown, when analysing surface defects on an object, the anomaly may appear either extended or clustered. Surface B of the brake disc (see Figure 3) presents an extensive scratch that cannot be regarded as localised. When the inspector encounters this type of defect, the procedure follows the steps described before, although no localised or clustered assessments can be made. According to the proposed methodology, the first task for the inspector is to distinguish intact surfaces from those that are likely to be defective using a gradient-based approach. The renders from both pristine and defective scenes are generated from an overhead view perpendicular to the surface of interest, and Figure 8 illustrates the gradient-based difference map over the surface.

By analysing Figure 8, some brief considerations can be made:

- The gradient-based map represents the primary indicator for identifying potentially damaged areas. In this case, the defect is clearly present, and its extent is greater than that observed in the previously analysed clustered case (i.e. surface A).
- The gradient map provides guidance on the extent of the defect. In this situation, it clearly shows that a cluster-based evaluation is not feasible.

In this case, the inspector may adopt a more global approach, subdividing the entire portion of the object to detect and confirm the presence of the defect. The main questions to be addressed are therefore the following:

- *Is the suspicious region an actual defective area?*
- *What is the extent of the defect?*

To address the first question, the surface of the object is divided into five distinct regions, as shown in Figure 9, with the alignment procedure ensuring that the pristine and defective scenes share the identical subdivision. Within this framework, *defective areas* are defined as those regions where the scratch is observed in the defective scene, while the corresponding regions in the pristine scene remain defect-free. Conversely, *non-defective areas* denote the remaining regions, which are intact in both scenes. The analysis focuses on two aspects:

- In defective areas, the pristine scene is free of defects, whereas the defective scene exhibits a large scratch. *Does this difference in surface condition lead to significant variations in the SH coefficients?*
- In non-defective areas, no defects are present in either scene. *Does the statistical analysis of the SH coefficients confirm the absence of significant differences?*

By analysing the distribution of gradient differences in Figure 8, it is possible to identify defective and intact zones. To characterise the differences between pristine and defective scenes, the indicators in




Figure 8. Gradient-based difference map on surface B of the brake disc.




Figure 9. Subdivision of surface B into five regions of interest (three intact areas and two defective areas) for local statistical analysis.




Figure 10. Normalised density (a) and normalised volume (b) computed across the 5 areas of interest of surface B.

Table 5. Summary of p -values obtained from ANOVA analysis on JSd values.

Levels ANOVA	p -value
Defective area 1 + Defective area 2 + Non-defective area 1 + Non-defective area 2 + Non-defective area 3	<0.001
Non-defective area 1 + Non-defective area 2 + Non-defective area 3	0.887
<i>Global non-defective = Non-defective area 1 + Non-defective area 2 + Non-defective area 3</i>	
Defective area 1 + Global non-defective	<0.001
Defective area 2 + Global non-defective	<0.001

Section 3 are employed. Density corresponds to the number of Gaussians with centre-points falling within a region, normalised by its surface area, while the average volume represents the mean size of the ellipsoids. Both metrics are normalised to the maximum values observed across the ten regions under analysis (five from the pristine scene and five from the defective one). Figure 10(a) and 10(b) respectively present the bar charts of density and average volume, highlighting the local differences.

Following the procedure defined in the Methodology section, the operator can decide to exploit the chromatic properties of ellipsoids to detect defective regions. As already mentioned, the scenes are properly registered. Consequently, colour information can be compared regardless of the observation direction, focusing on differences detectable in the scalar SH coefficients. In line with the proposed approach, SH coefficients over a specific surface portion are treated as discrete probability distributions, and dissimilarities between the scenes are therefore assessed in terms of the Jensen-Shannon distance. After collecting the JSd (see Appendix B for the summary of JSd for each area) from all five areas, ANOVA is applied to determine whether defective areas present JSd values that differ significantly from those of non-defective regions. Table 5 reports the p -values obtained for all the considered cases:

- Comparison of the five areas yields a p -value lower than 0.001, which leads to rejection of the null hypothesis and indicates a significant difference in JSd values.
- The three non-defective areas show no significant differences (p -value 0.887) and can therefore be aggregated into a single variable, referred to as global non-defective.
- The aggregation of the non-defective areas enables pairwise comparisons with each defective region.

Figure 11 presents the boxplot of JSd values for the five investigated areas. The three non-defective regions display a similar distribution, with medians and interquartile ranges that overlap, thereby supporting their aggregation into a single global non-defective category. In contrast, the two defective regions show clearly higher median values and larger variability, confirming the difference with respect to the non-defective areas. In addition, Figure 12 shows the interval plot of the mean JSd values with 95% confidence intervals for the two defective and the three non-defective regions.

The choice of three non-defective areas is arbitrary, although it is relevant to evaluate whether the defect-free portions maintain coherence when a different number of subdivisions is applied, since the operator can also test smaller or larger subdivisions to verify the robustness of the method and statistical coherence. This verification is necessary to ensure that no false defects are detected within the non-defective regions and to confirm that the outcome does not depend on an arbitrary partitioning of the surface. The number of subdivisions in the non-defective portion is varied from 2 to 30 to assess whether




Figure 11. Boxplot of JSd values for the five investigated areas identified in Figure 9.




Figure 12. Interval plot of JSd at 95% CI for the mean for the five investigated areas identified in Figure 9.

Table 6. *p*-Values obtained from ANOVA on non-defective areas when increasing the number of surface subdivisions.

Subdivisions in the non-defective area	<i>p</i> -value (ANOVA)
2	0.837
3	0.887
4	0.713
5	0.888
6	0.785
7	0.942
8	0.838
9	0.910
10	0.773
20	0.355
30	<0.001

Table 7. Overall surface area of surface B, defective areas and non-defective portions.

Portion	Surface [mm ²]
Surface B brake disc	441.9
Defective area 1	149.6
Defective area 2	67.0
Global non-defective	225.3

significant differences emerge when modifying the subsampling strategy, and the *p*-values from ANOVA are summarised in Table 6.

When the scene is divided into a limited number of portions, the ANOVA test confirms that there are no significant differences between the portions, confirming the homogeneity of the intact surface. For finer subdivisions, significant differences emerge since local variations and noise become predominant rather than global effects. As regards the extent of the scratch, total surface areas are computed and reported in Table 7.

Observing the results, some comments can be drawn:

- Gradient-based detection remains effective for an extended scratch, but the morphology of the defect prevents the use of the cluster-based approaches suitable for localised defects. A subdivision of the surface into distinct regions is therefore required to distinguish between defective and intact portions.
- Local metrics, namely Gaussian density and average volume, reveal systematic differences between scratched and pristine regions. Defective areas exhibit higher ellipsoid density, reflecting the need for a finer description of the surface, while average volumes decrease, consistent with a more heterogeneous structure.
- Statistical testing on the JSd indicator highlights significant differences between defective and non-defective areas, supporting the role of colour distribution as a discriminant for extended surface defects.

8. Discussion

A central aspect of this study is the role of colour in detecting and characterising defects. For the chromatic defects considered in this work, surface anomalies mainly manifest as variations in texture, colour and reflectance. The strength of GS lies in its ability to encode colour explicitly through SH coefficients, enabling these variations to be investigated quantitatively. This capability goes beyond visual inspection by providing interpretable descriptors that can be subjected to statistical analysis. From this perspective, colour is not only a natural choice but also a powerful means of linking the visual manifestation of a defect with quantifiable evidence of its presence.

In this work, two classes of defects were analysed, each requiring a different approach. For localised oxidation, where the defect appears as clusters of small spots, gradient-based detection combined with local density analysis is effective. The presence of oxidation spots leads to the local generation of additional

Gaussians, increasing the local ellipsoids' density in the defective scene compared with the pristine one. This behaviour provides a straightforward and interpretable indicator for identifying localised anomalies.

Extended defects, such as scratches, by contrast, are not suitable for cluster-based approaches. Their morphology requires a more global strategy, with colour-based statistical analysis providing a natural means of highlighting differences between defective and intact regions. The statistical analysis confirmed that defective regions exhibit significantly different distributions of JSd values compared with non-defective areas, supporting that the proposed indicator is reliable for defect detection. Such findings emphasise the flexibility of GS-based descriptors and suggest that defect-specific analyses can reveal which parts of the colour spectrum are most informative.

The procedure outlined in this work constitutes part of a semi-automatic inspection workflow, which can be activated by the operator when suspected regions require additional examination. In case of a fully automatic sequence, the procedure still requires manual verification to discard false positives, i.e. detections not corresponding to real defects.

Compared with recent deep learning-based surface defect detection approaches, the proposed GS-based framework may address a complementary inspection problem. Deep learning methods are highly effective for automatically detecting, localising or segmenting defects in available 2D images. However, the analysis remains constrained by the viewpoints acquired during physical inspection (i.e. the captured images). If a surface region was not properly captured, was partially occluded or needs to be examined later from a different perspective, a conventional image-based pipeline cannot recover that missing view without a new acquisition. In contrast, the GS-based approach reconstructs a navigable 3D scene of the inspected object. This allows the inspector to generate renders from selected, repeatable or previously unavailable viewpoints and to compare the same surface region across different inspection times. The proposed GS-based approach is not intended to exclude the use of deep learning. On the contrary, GS and deep learning can be integrated in a combined inspection pipeline, where GS can provide controlled and temporally consistent renders from the reconstructed scene, while deep learning-based detectors can be applied downstream to detect defects.

The scalability of the proposed approach should also be considered. Although the present case study focuses on a brake disc, the proposed methodology could be extended to different industrial assets. As object size and geometric complexity increase, the generation of a reliable scene generally requires a larger number of images, increasing computational time, memory demand and data management complexity. In addition, the acquisition of large assets may reduce the practicality of simple consumer devices (e.g. smartphones), since the operator may not be able to easily access, frame and move around all relevant surfaces. Alternative systems, e.g. drones or robotic platforms, may therefore be required, although their use introduces further constraints, particularly in indoor industrial environments where limited space, safety requirements, lighting variability and flight restrictions can affect image collection. From a practical perspective, scalability may therefore require a multi-scale, region-based strategy, in which the object can be divided into relevant inspection areas, and each area is reconstructed or analysed with an appropriate level of detail. This strategy could reduce the computational burden and support a more effective inspection process.

9. Conclusions

This work has explored the potential of GS for defect detection through the comparison of pristine and defective 3D GS-based scenes. Two distinct defect morphologies were investigated, namely localised oxidation and extended scratches. The results suggest that the choice of descriptors should be tailored to the defect type: gradient- and density-based metrics are effective for clustered anomalies such as oxidation pits, whereas extended defects require colour-based statistical analyses to capture their morphology and extent.

Beyond these specific findings, the study demonstrates that GS provides an early-stage but valid framework for industrial inspection. Its representation of geometry and colour not only enables high-quality visual reconstructions but also produces interpretable metrics that can be subjected to rigorous statistical testing. In the present case study, acquisition was performed using a standard device such as a smartphone, highlighting the accessibility of the proposed approach, although more complex industrial settings may require more structured acquisition systems. In this way, GS creates a bridge between

computer graphics and quality inspection in manufacturing, offering the basis for a systematic approach to defect characterisation in 3D scenes.

Furthermore, the contribution of this work lies in the development of a digital inspection mechanism that can support sustainable manufacturing in the Industry 5.0 context. Specifically, the GS-based scene preserves the inspected condition of an object as a temporal and reusable digital record, allowing defects to be revisited, quantified and compared across multiple inspection stages. This supports the logic of re-use by connecting digital defect detection with later operational decisions on repair, replacement and process correction.

Regarding the limitations, this work focuses on a single case study and addresses two types of chromatic defects. Geometric defects were not included, and their analysis will require methods beyond colour-based descriptors. In addition, the current implementation relies on high-end computational resources, as an A100 GPU in a cloud environment was employed to ensure fast and accurate scene generation. This highlights the need for future optimisation to make the method feasible on standard hardware. Finally, parts of the analytical workflow are still semi-automated and require operator supervision. Additional limitations concern the robustness of the comparison under different acquisition conditions. Since the proposed analysis mainly exploits chromatic and texture-related differences, variations in illumination, shadows, exposure or surface reflections between different acquisitions may generate apparent differences that are not directly associated with actual defects (e.g. false positives). This aspect is particularly relevant for metallic or reflective surfaces and will be addressed in future work through more controlled acquisition protocols, illumination normalisation and suitable radiometric correction strategies. The comparison between different inspection times also depends on the quality of the alignment between the corresponding scenes. Residual alignment errors may influence the localisation of anomalies, thereby affecting the reliability of multi-image and multi-temporal comparisons. Therefore, future developments will investigate the sensitivity of the proposed indicators to alignment uncertainty. Moreover, the current semi-automatic nature of the workflow should be progressively reduced by introducing automatic viewpoint selection and by exploring the synergistic integration of GS with deep learning-based detectors applied to the generated renders.

Several perspectives for future work emerge from this study. A first direction will be the development of quantitative indicators capable of guiding the choice between local and global analysis, depending on the morphology of the defect. A second line of research will concern the optimisation of viewpoint selection: in this study, a top-view render was used as the basis for gradient-based comparisons, but further investigation is needed to determine whether multi-view analyses or alternative perspectives may provide more reliable or sensitive detection of surface anomalies. In addition, the complementarity between GS and deep learning represents a promising direction for future research. The reconstructed scene could be exploited to generate controlled and temporally consistent renders, which may subsequently be analysed by deep learning-based detectors for defect detection, potentially improving the effectiveness of visual inspection in manufacturing.

In the longer term, the aim is to establish a fully integrated framework for surface defect inspection, in which the operator simply records a short video and the inspection procedure generates as output the likely defective regions along with an indication of the defect type. Embedding such a tool within industrial visual inspection routines would represent a key step towards more reliable, interpretable and repeatable surface quality assessment. Within the proposed framework, the procedure is described as semi-automatic, since manual intervention is still required to guide specific analyses or to supervise the outcomes of the performed tests. The long-term research directions will be oriented towards the ambitious goal of achieving full automation of the entire procedure.

Acknowledgements

This study was carried out within the Space It Up project funded by the Italian Space Agency, ASI, and the Ministry of University and Research, MUR, under contract *n.* 2024-5-E.0 - CUP *n.* I53D24000060005.

Author contributions

CRediT: **Mattia Trombini**: Conceptualization, Data curation, Funding acquisition, Investigation, Methodology, Writing – original draft; **Domenico A. Maisano**: Conceptualization, Methodology, Supervision, Writing – review & editing; **Fiorenzo Franceschini**: Conceptualization, Methodology, Supervision.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research was carried out within the Space It Up project funded by the Italian Space Agency, ASI, and the Ministry of University and Research, MUR, under contract *n.* 2024-5-E.0 - CUP *n.* I53D24000060005.

Data availability statement

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Ethical approval

The authors respect the Ethical Guidelines of the Journal.

Acronyms

ANOVA	Analysis of Variance
CMM	Coordinate Measuring Machine
DRÆM	Discriminatively trained reconstruction anomaly embedding model
GPU	Graphics Processing Unit
GS	Gaussian Splatting
JSD	Jensen–Shannon Divergence
JSd	Jensen–Shannon distance
KL	Kullback–Leibler
MVS	Multi-View Stereo
MVTecAD	MVTec Anomaly Detection Dataset
NeRF	Neural Radiance Field
RGB	Red-Green-Blue
SfM	Structure-from-Motion
SH	Spherical harmonics
SIBR	System for Image-Based Rendering
VRAM	Video Random-Access Memory

References

- Alatan, A. A., Yemez, Y., Gudukbay, U., Zabulis, X., Muller, K., Erdem, C. E., Weigel, C., & Smolic, A. (2007). Scene representation technologies for 3DTV-A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 17(11), 1587–1605. <https://doi.org/10.1109/TCSVT.2007.909974>
- Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., & Steger, C. (2021). The MVTec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129, 1038–1059. <https://doi.org/10.1007/s11263-020-01400-4>
- Faghih-Roohi, S., Hajizadeh, S., Núñez, A., Babuska, R., & Schutter, B. D. (2016). Deep convolutional neural networks for detection of rail surface defects deep convolutional neural networks for detection of rail surface defects. In. *2016 International Joint Conference on Neural Networks (IJCNN)* Vancouver, BC, Canada. 2016, 2584–2589. <https://doi.org/10.1109/IJCNN.2016.7727522>
- Franceschini, F., Mastrogiacomo, L., & Pralio, B. (2009). An unmanned aerial vehicle-based system for large scale metrology applications. *International Journal of Production Research*, 48(13), 3867–3888. <https://doi.org/10.1080/00207540902896220>
- Franceschini, F., Galetto, M., Maisano, D., & Mastrogiacomo, L. (2016). Combining multiple large volume metrology systems: Competitive versus cooperative data fusion. *Precision Engineering*, 43, 514–524. <https://doi.org/10.1016/j.precisioneng.2015.09.014>
- Franceschini, F., Galetto, M., Maisano, D., Mastrogiacomo, L., & Pralio, B. (2011). Distributed large-scale dimensional metrology: New insights. Springer Science & Business Media.
- Givens, C. R., & Shortt, M. (1984). A class of Wasserstein metrics for probability distributions. *Michigan Mathematical Journal*, 31(2), 231–240. <https://doi.org/10.1307/mmj/1029003026>
- Gonzalez, R. C., & Woods, R. E. (2006). *Digital image processing* (3rd ed.) ed.). Pearson.

- Green, R. (2003). Spherical harmonic lighting: The gritty details. In: *Proceedings of the Game Developers Conference*, Vol. 56.
- Guo, S., Lin, J., Yang, L., Ren, Y., & Guo, Y. (2017). A method of reconstructing the spatial measurement network by mobile measurement transmitter for shipbuilding. *Journal of Physics: Conference Series*, 870(1), 012009
- Johnson, G. W., Laskey, S. E., Robson, S., & Shortis, M. R. (2004). Dimensional and accuracy control automation in shipbuilding fabrication: An integration of advanced image interpretation, analysis and visualization techniques. In: *ISPRS Congress XX*, 2004.
- Kerbl, B., Kopanas, G., Leimkuehler, T., & Drettakis, G. (2023). 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 1–14. <https://doi.org/10.1145/3592433>.
- Kholkhujayev, J., Maculotti, G., Genta, G., Galetto, M., & Inoyatkhodjaev, J. (2022). Non-contact articulated robot-integrated gap and flushness measurement system for automobile assembly. *IEEE Access*, 10, 86528–86541. <https://doi.org/10.1109/ACCESS.2022.3199066>
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86. <https://doi.org/10.1214/aoms/1177729694>
- Lin, J. (1991). Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, 37(1), 145–151. <https://doi.org/10.1109/18.61115>
- Liu, W. L., & Wang, Z. K. (2010). New method for large-scale dimensional metrology using laser tracker system. *Advanced Materials Research*, 97, 4247–4250. <https://doi.org/10.4028/www.scientific.net/AMR.97-101.4247>
- Maculotti, G., Giorio, L., Genta, G., & Galetto, M. (2025). Traceability and uncertainty of defects automated measurements by CNN-powered machine vision systems. *CIRP Annals*, 74(1), 661–665. <https://doi.org/10.1016/j.cirp.2025.03.023>
- Maisano, D. A., Mastrogiacomo, L., Franceschini, F., Capizzi, S., Pischedda, G., Laurenza, D., Gomiero, G., & Manca, G. (2022). Dimensional measurements in the shipbuilding industry: On-site comparison of a state-of-the-art laser tracker, total station and laser scanner. *Production Engineering*, 17, 625–642. <https://doi.org/10.1007/s11740-022-01170-7>
- Maisano, D. A., Mastrogiacomo, L., Franceschini, F., Capizzi, S., Pischedda, G., Laurenza, D., Gomiero, G., & Manca, G. (2023). Large-volume metrology in shipbuilding: Structured comparison of innovative measuring instruments. *Procedia CIRP*, 118, 1090–1095. <https://doi.org/10.1016/j.procir.2023.06.187>
- Mallick, S. S., Goel, R., Kerbl, B., Steinberger, M., Vicente Carrasco, F., & De La Torre, F. (2024). Taming 3DGS: High-quality radiance fields with limited resources, *SIGGRAPH Asia 2024 Conference Papers* (2nd ed., pp. 1–11). ACM. Article.
- Martínez, S., Cuesta, E., Barreiro, J., & Álvarez, B. (2010). Analysis of laser scanning and strategies for dimensional and geometrical control. *International Journal of Advanced Manufacturing Technology*, 46, 621–629.
- Masci, J., Meier, U., Ciresan, D., Schmidhuber, J., & Fricout, G. (2012). Steel defect classification with max-pooling convolutional neural networks. *The 2012 International Joint Conference on Neural Networks (IJCNN)*, Brisbane, QLD, Australia, 2012, 1–6.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., & Ng, R. (2021). NeRF: representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 1, 99–106.
- Molina, J., Solanes, J. E., Arnal, L., & Tornero, J. (2017). On the detection of defects on specular car body surfaces. *Robotics and Computer-Integrated Manufacturing*, 48, 263–278. <https://doi.org/10.1016/j.rcim.2017.04.009>
- Ouzineb, M., Mhada, F. Z., Pellerin, R., & El Hallaoui, I. (2018). Optimal planning of buffer sizes and inspection station positions. *Production & Manufacturing Research*, 6(1), 90–112. <https://doi.org/10.1080/21693277.2017.1422812>
- Pollock, L., Abdelwahab, A. K., Murray, J., & Wild, G. (2021). The need for aerospace structural health monitoring: A review of aircraft fatigue accidents. *International Journal of Prognostics and Health Management*, 12(3), 1–16.
- Rački, D., Tomažević, D., & Skočaj, D. (2018). A Compact Convolutional Neural Network for Textured Surface Anomaly Detection. *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Lake Tahoe, NV, USA, 2018, pp. 1331–1339.
- Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., & Gehler, P. (2022). Towards total recall in industrial anomaly detection. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, 2022, pp. 14298–14308.
- Rusinkiewicz, S., & Levoy, M. (2001). Efficient variants of the ICP algorithm. *Proceedings Third International Conference on 3-D Digital Imaging and Modeling. Quebec City, QC, Canada*. pp. 145–152.
- Rusu, R. B., Blodow, N., & Beetz, M. (2009). Fast point feature histograms (FPFH) for 3D registration. *2009 IEEE International Conference on Robotics and Automation, Kobe, Japan*. pp. 3212–3217.
- Sareen, A., Sapre, C. A., & Selig, M. S. (2014). Effects of leading edge erosion on wind turbine blade performance. *Wind Energy*, 17, 1531–1542. <https://doi.org/10.1002/we.1649>
- Sause, M. G. R., & Jasiūnienė, E. (2022). Structural health monitoring damage detection systems for aerospace, Springer Aerospace Technology (1st ed.).
- Schönberger, J. L., & Frahm, J.-M. (2016). Structure-from-motion revisited. In: *Proceedings of the. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4104–4113. <https://doi.org/10.1109/CVPR.2016.445>
- Sitzmann, V., Zollhöfer, M., & Wetzstein, G. (2019). Scene representation networks: Continuous 3D-structure-aware neural scene representations. *Advances in Neural Information Processing Systems (NeurIPS 2019)*, 33, 1121–1132.

- Snavely, N., Seitz, S. M., & Szeliski, R. (2006). Photo tourism: exploring photo collections in 3D. *ACM Transactions on Graphics*, 25(3), 835–846. <https://doi.org/10.1145/1141911.1141964>
- Stałowska, P., Suchocki, C., & Zagubień, A. (2025). Terrestrial laser scanning for wind turbine blade defect detection. *Measurement*, 246, 116706. <https://doi.org/10.1016/j.measurement.2025.116706>
- Tabernik, D., Šela, S., Skvarč, J., & Skočaj, D. (2020). Segmentation-based deep-learning approach for surface-defect detection. *Journal of Intelligent Manufacturing*, 31, 759–776.
- Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., Simon, T., Theobalt, C., Nießner, M., Barron, J. T., Wetzstein, G., Zollhöfer, M., & Golyanik, V. (2022). Advances in neural rendering. *Computer Graphics Forum*, 41, 703–735. <https://doi.org/10.1111/cgf.14507>
- Trombini, M., Maisano, D. A., & Franceschini, F. (2025). A synergistic integration of photogrammetry and Gaussian Splatting for back-in-time quality inspection. *Materials Research Proceedings*, 57, 20–28. *Italian Manufacturing Association Conference*. <https://doi.org/10.21741/9781644903735-3>
- Trombini, M., Maculotti, G., Maisano, D. A., Pagani, A., & Franceschini, F. (2025). Scaling photogrammetry: A comparative evaluation and metrological assessment across small- and meso-scale domains. *Precision Engineering*, 92, 124–140. <https://doi.org/10.1016/j.precisioneng.2024.12.002>
- Turley, G. A., Kiraci, E., Olifent, A., Attridge, A., Tiwari, M. K., & Williams, M. A. (2014). Evaluation of a multi-sensor horizontal dual arm coordinate measuring machine for automotive dimensional inspection. *International Journal of Advanced Manufacturing Technology*, 72, 1665–1675. <https://doi.org/10.1007/s00170-014-5737-3>
- Zavrtanik, V., Kristan, M., & Skočaj, D. (2021). DRÆM - A discriminatively trained reconstruction embedding for surface anomaly detection. *2021 IEEE/CVF International Conference on Computer Vision (ICCV), Montreal, QC, Canada, 2021*. pp. 8310–8319.

Appendices

Appendix A

The gradient is derived from edge detection and describes the directional change in pixel intensity along the x and y axes of an image f , as defined in Equation (A1):

$$\nabla f(x, y) = \begin{bmatrix} g_x \\ g_y \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix} = \begin{bmatrix} f(x+1, y) - f(x-1, y) \\ f(x, y+1) - f(x, y-1) \end{bmatrix} \quad (\text{A1})$$

The gradient quantifies local variations of intensity or colour between adjacent pixels. Typically, gradient-based methods are employed for edge-detection tasks and, therefore, they demonstrate to be effective for the detection of clustered defects, since such imperfections typically manifest as small colour discontinuities on the surface. In addition, even extensive defects can be effectively localised since they appear as colour variations. Preliminarily, the gradient of the two renders is computed, and the pixel-to-pixel distance is then evaluated as follows:

$$\Delta_{f_1-f_0}^{ij} = |\nabla f_1(x_i, y_j) - \nabla f_0(x_i, y_j)| \quad (\text{A2})$$

where f_0 is the render of the scene at time t_0 , and f_1 is the render collected from the same position of the scene at t_1 .

Appendix B

Table B1 contains JSd values for each SH scalar coefficient.

Table B1. Summary of JSd values for each scalar coefficient.

	Defective 1	Defective 2	Area 1	Area 2	Area 3
JSd (c_0^0) Red	0.164	0.361	0.181	0.162	0.198
JSd (c_1^{-1}) Red	0.881	0.752	0.157	0.330	0.356
JSd (c_1^0) Red	0.349	0.483	0.205	0.365	0.361
JSd (c_1^1) Red	0.821	0.719	0.719	0.501	0.196
JSd (c_2^{-2}) Red	0.596	0.520	0.409	0.374	0.134
JSd (c_2^{-1}) Red	0.665	0.384	0.070	0.159	0.209
JSd (c_2^0) Red	0.251	0.131	0.357	0.390	0.270
JSd (c_2^1) Red	0.354	0.489	0.381	0.482	0.445
JSd (c_2^2) Red	0.484	0.701	0.174	0.220	0.218
JSd (c_3^{-3}) Red	0.106	0.553	0.202	0.158	0.055
JSd (c_3^{-2}) Red	0.196	0.352	0.155	0.114	0.219
JSd (c_3^{-1}) Red	0.446	0.398	0.119	0.299	0.139
JSd (c_3^0) Red	0.445	0.356	0.183	0.092	0.176
JSd (c_3^1) Red	0.470	0.192	0.248	0.297	0.407
JSd (c_3^2) Red	0.257	0.181	0.158	0.208	0.333
JSd (c_3^3) Red	0.506	0.491	0.213	0.137	0.143
JSd (c_0^0) Green	0.159	0.348	0.181	0.157	0.202
JSd (c_1^{-1}) Green	0.875	0.766	0.149	0.338	0.373
JSd (c_1^0) Green	0.373	0.465	0.222	0.393	0.385
JSd (c_1^1) Green	0.823	0.712	0.726	0.485	0.214
JSd (c_2^{-2}) Green	0.607	0.510	0.425	0.358	0.115
JSd (c_2^{-1}) Green	0.655	0.459	0.061	0.163	0.222
JSd (c_2^0) Green	0.256	0.185	0.350	0.377	0.305
JSd (c_2^1) Green	0.310	0.482	0.363	0.476	0.450
JSd (c_2^2) Green	0.554	0.740	0.219	0.252	0.216
JSd (c_3^{-3}) Green	0.127	0.596	0.209	0.147	0.056
JSd (c_3^{-2}) Green	0.194	0.377	0.144	0.100	0.217
JSd (c_3^{-1}) Green	0.485	0.324	0.145	0.225	0.155
JSd (c_3^0) Green	0.427	0.347	0.193	0.085	0.180
JSd (c_3^1) Green	0.491	0.188	0.236	0.278	0.402
JSd (c_3^2) Green	0.266	0.148	0.168	0.194	0.387
JSd (c_3^3) Green	0.506	0.462	0.242	0.117	0.132
JSd (c_0^0) Blue	0.148	0.373	0.198	0.146	0.207
JSd (c_1^{-1}) Blue	0.886	0.731	0.190	0.371	0.392
JSd (c_1^0) Blue	0.354	0.380	0.239	0.440	0.421
JSd (c_1^1) Blue	0.811	0.748	0.734	0.475	0.232
JSd (c_2^{-2}) Blue	0.581	0.454	0.410	0.333	0.142
JSd (c_2^{-1}) Blue	0.642	0.490	0.118	0.163	0.194
JSd (c_2^0) Blue	0.239	0.223	0.314	0.342	0.335
JSd (c_2^1) Blue	0.262	0.493	0.357	0.463	0.437
JSd (c_2^2) Blue	0.676	0.848	0.293	0.310	0.202
JSd (c_3^{-3}) Blue	0.162	0.628	0.181	0.135	0.083
JSd (c_3^{-2}) Blue	0.204	0.388	0.219	0.101	0.229
JSd (c_3^{-1}) Blue	0.494	0.274	0.109	0.236	0.124
JSd (c_3^0) Blue	0.365	0.342	0.177	0.097	0.176
JSd (c_3^1) Blue	0.502	0.192	0.256	0.261	0.406
JSd (c_3^2) Blue	0.311	0.286	0.163	0.170	0.416
JSd (c_3^3) Blue	0.510	0.430	0.279	0.120	0.097