Measurement 277 (2026) 121590

Contents lists available at ScienceDirect

Measurement

journal homepage: www.elsevier.com/locate/measurement

Efficient road roughness characterization using Gaussian process regression and 2D imaging

Raffaele Stefanelli [a],*, Georgios Mavros [b], Miguel Martínez-García [b], James Knowles [b], Francesco Timpone [a], Flavio Farroni [a]

[a] University of Naples Federico II, Department of Industrial Engineering, Via Claudio, 21, Naples, 80125, Italy
[b] Department of Aeronautical and Automotive Engineering, Loughborough University, Epinal Way, Loughborough, LE11 3TU, Leicestershire, United Kingdom

ARTICLE INFO

Keywords:
Surface roughness
Image processing
Virtual surface reconstruction
Machine learning

ABSTRACT

Accurate knowledge of surface roughness is essential for understanding key phenomena such as tire–road friction and wear rate. However, high-resolution surface characterization typically requires scanning devices that are time-consuming to use and expensive. In time-critical contexts like motorsport, where setup windows are narrow, having a faster and more accessible alternative could provide a significant advantage. This paper presents a fast, cost-effective and physically interpretable methodology to reconstruct the three-dimensional roughness of road surfaces directly from grayscale images. The results demonstrate a strong correlation between 2D photometric information and 3D surface features, preserving both the height distribution and spectral content. The proposed framework combines light-intensity features extracted from the image with a Gaussian Process Regression (GPR) model, offering high physical controllability and reduced data requirements compared to deep learning approaches such as convolutional neural networks. The model was validated under various conditions and proved effective even on unconventional textures, including sandpaper and marble samples. It also showed robustness to different image formats. Overall, the method provides a practical and innovative alternative for surface roughness assessment, with strong potential for applications in vehicle dynamics, tire modeling and tribology.

1. Introduction

Road surface characterization is a crucial topic in different engineering fields [1,2]. Indeed, understanding and monitoring road surface conditions is fundamental in the context of transportation, civil engineering and tire performance evaluation as well for contact mechanics analysis [3,4]. Accurate surface characterization is particularly critical as it directly affects rolling resistance, safety, noise, ride comfort and tire friction analysis [5–7].

In the context of vehicle dynamics, road surface roughness has a direct influence on tire and vehicle performance [8,9]. In motorsport, where tires operate at their frictional limits, even small variations in surface texture can significantly alter the tire behavior from both frictional and wear points of view [10,11]. From a contact mechanics perspective, roughness determines both the real area of contact and the pressure distribution at the interface, influencing not only friction but also subsurface stress propagation, energy dissipation and local viscoelastic deformation [12–14]. Indeed, many contact mechanics theories have emphasized the critical role of road surface roughness in accurately modeling tire–road interactions [15–18]. In particular, roughness at multiple scales, ranging from macrotexture to microtexture, affects how tires deform, grip and slide across the pavement, making it a fundamental input for predictive models of tire friction, wear and rolling resistance [19,20].

The measurement of surface texture is a topic relevant to many fields of engineering [21,22]. Traditionally, road texture is measured using laser scanner or optical devices, which, despite their accuracy, are often expensive, require specialized handling and are constrained by low acquisition speed and limited spatial coverage [23]. In light of these limitations, there is a growing demand for accessible and rapid alternatives that can provide surface texture acquisition without compromising accuracy: in civil engineering, fast surface inspection supports asset management strategies and early deterioration detection [24,25]; in motorsport and performance vehicle testing, being able to scan entire track surfaces in minutes rather than hours opens up unprecedented opportunities in predictive modeling, tire compound selection

* Corresponding author.
E-mail addresses: raffaele.stefanelli@unina.it (R. Stefanelli), G.Mavros@lboro.ac.uk (G. Mavros), m.martinez-garcia@lboro.ac.uk (M. Martínez-García), J.A.C.Knowles@lboro.ac.uk (J. Knowles), francesco.timpone@unina.it (F. Timpone), flavio.farroni@unina.it (F. Farroni).

https://doi.org/10.1016/j.measurement.2026.121590
Received 5 November 2025; Received in revised form 30 March 2026; Accepted 21 April 2026
Available online 24 April 2026

and simulation accuracy. Such capabilities are particularly valuable in contexts where real-time decisions are critical and operational budgets or timeframes are constrained.

Recent advances in computer vision and sensor technology have prompted the exploration of alternative, cost-effective solutions for surface reconstruction [26–28]. In particular, the development of novel fast technologies has been a promising field for visual inspection and qualitative surface evaluation. However, the translation from 2D visual information to quantitative 3D surface profiles remains a fundamental challenge, due to the inherent loss of depth information in standard photographs.

The surface texture of roads is governed by the aggregate distribution, binder type and laying process, resulting in characteristic spatial and statistical features [29–31]. While the human eye can qualitatively recognize the differences among a rough or a smooth surface from a photograph, the exact physical relation between surface reflectance and topography is complex. However, light interacts with textured surfaces through mechanisms such as scattering, shadowing and specular reflection, indicating that 2D images may contain relevant information which is correlated with the underlying 3D roughness profile [32,33]. In the current state of the art, an interesting methodology for generating surface topographies with a desired height distribution and power spectral density has been proposed in [34]. Unlike previous approaches, this method allows independent control over the height probability distribution and the spectral content, enabling the synthesis of surfaces that accurately reflect specific statistical and frequency domain characteristics.

Building on this approach, the present research introduces an innovative methodology to reconstruct a 3D surface from a single 2D image. Indeed, this paper builds upon the hypothesis that although a single 2D image cannot fully determine a 3D surface structure, it contains statistical and spectrum features that correlate with surface height information. In this spirit, a model to obtain physical roughness starting from images is presented, enabling a predictive approach. This combines physics-inspired analysis with a machine learning approach. For this reason, the approach relies on traditional signal processing procedures and, to ensure predictive reliability, Gaussian Process Regression (GPR) was adopted as technique. This choice allows for effective learning with limited data, while also providing interpretable uncertainty estimates. Indeed, GPR provides a powerful probabilistic framework for regression that is well-suited to situations where the amount of training data is limited and physical interpretability is essential [35,36].

Convolutional Neural Networks (CNNs) were deliberately avoided due to their significant limitations, including the need for large datasets and the lack of physical interpretability [37,38]. Indeed, the lack of control makes it difficult to ensure that the extracted features related to physically meaningful variables, especially in engineering tasks requiring functional control and uncertainty quantification. In contrast, the primary goal of this research is to develop a fast, physically interpretable and controllable method that captures the correlation between 3D surface roughness and its 2D signature.

In summary, the approach presented in this paper is built on a signal-processing-based analysis of image features derived from light intensity, followed by the reconstruction of surface texture to replicate both statistical properties and frequency content. The main innovation goal is to provide an alternative that eliminates the need for costly surface texture scanners, significantly reducing both equipment costs and acquisition time, as the process relies on simple photographic capture rather than time-consuming roughness scanning.

This paper is organized as follows: Section 2 presents the experimental testing campaign used to collect the data for model development and validation, along with the main pre-processing steps. Section 3 provides a detailed description of the proposed methodology, including the procedure for extracting surface heights from images with the desired height distribution and frequency content and the implementation of the GPR model. The results and a relevant discussion are presented in Section 4, followed by conclusions in Section 5.

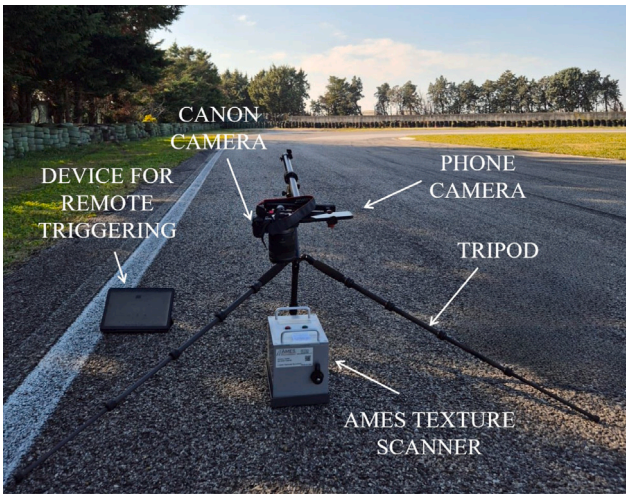

Fig. 1. Experimental setup for data acquisition.

2. Experimental testing campaign

The experimental testing campaign was essential to acquire the roughness and 2D image data required for developing and validating the roughness characterization method. Specifically, both high-resolution photographs and laser scans were collected from real road surfaces, using a dedicated setup and an AMES laser scanner, respectively. The reference scans were useful to calibrate the developed methodology capable of reconstructing the 3D roughness profile of a surface starting from 2D images. In particular, the experimental setup was developed to capture images and perform surface texture scanning on precisely the same surface area. A 26 MP full frame Canon RP camera equipped with an 85 mm macro lens was mounted at a fixed height and orientation relative to the sample surface to ensure repeatability across all tests. To avoid structural vibrations during image capture, particularly those induced by manually pressing the shutter, remote triggering was used to control the camera focus and acquisition, thus preserving image stability and consistency. To ensure consistency in the optical signal, particular attention was paid to the lighting conditions during image acquisition. Images were captured under diffuse lighting and the exposure settings were adjusted to avoid intensity saturation, ensuring that the maximum grayscale value remained within the available dynamic range of the sensor. Furthermore, light intensity values were normalized by the maximum intensity of each image in order to reduce the influence of illumination variability. The area scanned by the AMES texture scanner for each surface was set to a length of 104 mm and a width of 72 mm. Multiple scan lines were acquired for each sample, maintaining a resolution of 6.35E−06m along the scanning direction. Careful alignment between the image frame and the scanned region was maintained throughout the campaign, making it possible to compare, pixel-by-pixel (measuring the physical dimensions as well with a ruler), the optical signal (light intensity-based) and the reference roughness profile (height-based), as required for model calibration and validation. The experimental setup was also designed to perform the acquisition of images using an additional device: alongside the main camera, a smartphone was mounted in a fixed position, allowing the acquisition of photographs from both the devices. A representation of the experimental setup is shown in Fig. 1, illustrating the configuration described above.

Approximately 20 different surfaces were considered for the construction of the dataset used in this study, covering a wide range of surface types such as urban roads, racing tracks and some road samples typically used for indoor friction testing. To further expand the variability in surface roughness characteristics and consider unconventional

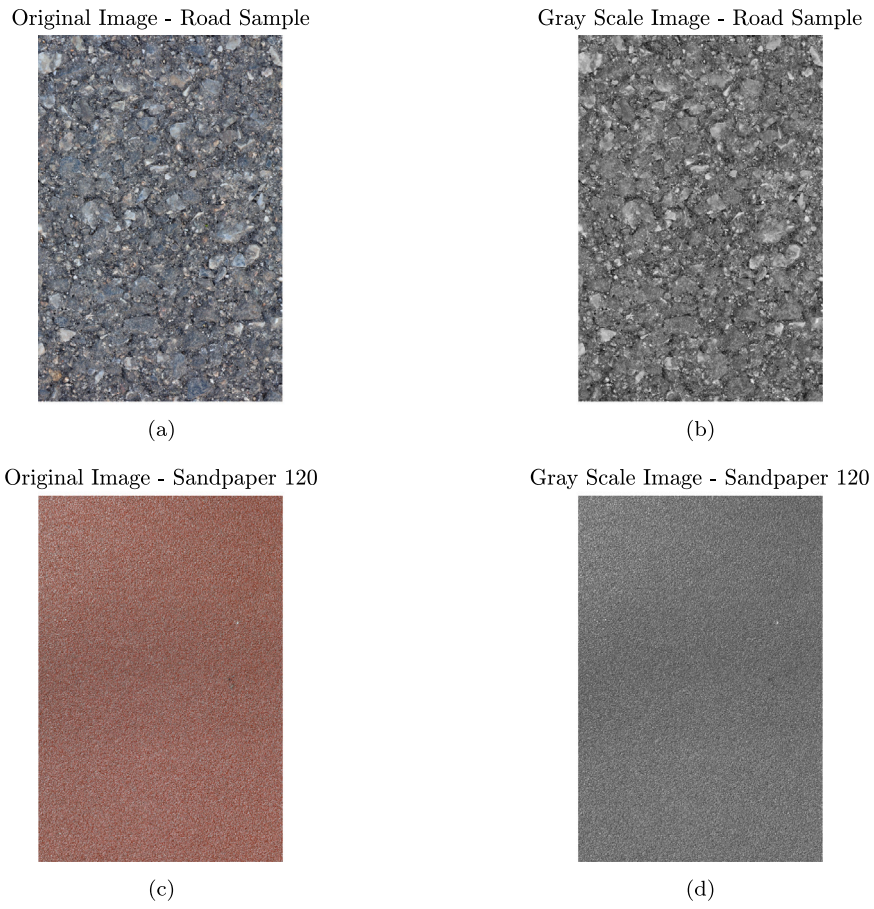

Fig. 2. RGB to grayscale example for both a road sample and a sandpaper: (a) Road sample original image, (b) Road sample grayscale image, (c) Sandpaper grit 120 original image, (d) Sandpaper grit 120 grayscale image.

scenarios, additional specimens such as sandpapers with different grit sizes and a marble sample surface were also included.

2.1. Data pre-processing

To enable a meaningful comparison between image-based information and laser scanner data, a dedicated pre-processing pipeline was developed. The goal of this step was to align the raw image with the AMES roughness data both spatially and in terms of resolution.

First, the RGB image was converted into grayscale, as the light intensity distribution in the gray domain was assumed to correlate with surface morphology due to variations in diffuse and specular reflections. An example of conversion from RBG image into grayscale is reported in Fig. 2 which shows an example of road sample and a sandpaper grit 120.

Subsequently, the physical size of the scanned area was matched through geometric transformations since the dimensions of the images were measured with a ruler in both directions. Indeed, the image was either centrally cropped or rescaled to fit the AMES scan dimensions, based on user-defined physical parameters. This ensured that the image and scan refer to the same surface portion.

Once geometrically aligned, a set of equally-spaced vertical columns was extracted from the image, excluding a few mm on each side to remove marginal effects. In this way, the number of extracted columns corresponded to the number of scan lines selected during acquisition. Each column was then resampled in the vertical direction to match the resolution of the AMES scan using 1D interpolation. The result is a light intensity matrix of the same size as the AMES height matrix, thus allowing matrix-dimensions correspondence.

To prevent aliasing introduced by resampling and interpolation, an anti-aliasing Butterworth low-pass filter is applied along the vertical direction of each interpolated column. The cutoff frequency of the filter is set to match the Nyquist frequency of the original image along the vertical axis, thus preserving the physical bandwidth of the original optical signal. The filtered matrix was then rescaled to match the grayscale range of the original image, ensuring consistency in intensity values linked to surface roughness transitions. All resulting light-intensity data were stored as structured matrices, ready for subsequent height estimation, while the reference heights (AMES scans) were used as ground truth for model calibration and validation.

It should be noted that strict spatial alignment between the image and the laser scan is not required. The proposed methodology does not rely on point-wise correspondence, but on matching statistical and spectral properties. Therefore, it is sufficient that the scanned region is representative of the same surface patch and contained within the photographed area, while small misalignments have negligible impact on the roughness generation.

3. Methodology

A workflow is developed to reconstruct road surface roughness from grayscale images. To ensure consistency with measured ground truth data, matching is attempted both in terms of statistical and spectral properties of the surface. In particular, the methodology is divided into three main stages. First, an analytical transformation is applied to map image light intensity to estimated surface heights, ensuring consistency in the statistical height distribution. Second, the frequency content of the reconstructed surface is adjusted to match the power spectral density (PSD) of the ground-truth laser-scanned surfaces without losing

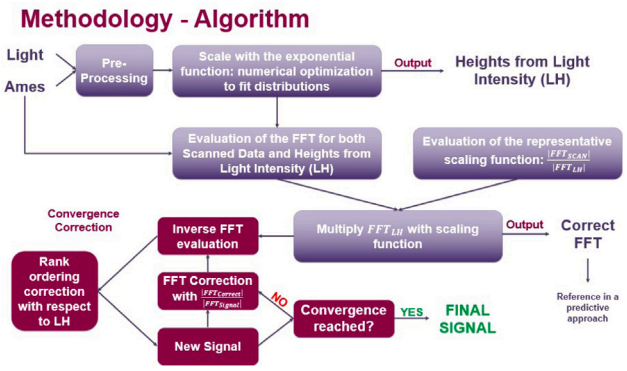

Fig. 3. Algorithm block diagram for virtual surface generation from grayscale image data.

the consistency of the heights distribution. Finally, a model based on Gaussian Process Regression (GPR) is trained to enable a fully predictive approach. Building upon the previous steps, this model allows for the complete generation of surface profiles directly from visual features, thus enabling generalization to unseen surfaces without the need for physical scanning. The ultimate goal of the entire methodology transformation is to generate, starting from the grayscale image, a reconstructed surface profile that not only exhibits a height distribution statistically consistent with that obtained from high-resolution laser scanning (considered the experimental ground truth), but also preserves its characteristic frequency content (i.e., power spectral density).

A block diagram illustrating step by step the procedure that will be described in Sections 3.1 and 3.2 is presented in Fig. 3, providing a clear overview of the entire workflow.

3.1. Height distribution analysis

The first key step of the methodology involves extracting surface heights from grayscale image data, based on an physically-inspired empirical model that relates light intensity to geometric surface features. The fundamental hypothesis underpinning this phase is that the roughness profile of a surface affects its optical appearance through local variations in light scattering and shadowing. These effects modulate the observed grayscale intensity, which is then used as a proxy for depth information. The analytical expression employed to convert grayscale light intensity data — extracted from photographic images of the surface — into numerical height values was designed based on a fundamental physical assumption: the deeper a point lies within a surface depression, the less light it reflects due to occlusion and limited direct illumination. This conceptual framework implies a monotonic relationship between light intensity and elevation, where brighter regions correspond to higher surface features and darker regions are associated with valleys or recessed areas. This assumption would approximately hold for surfaces which are homogeneous at the microscopic level, and of the same color throughout. With asphalt, such an assumption is frequently violated as it is possible to observe dark colored spots at the top of asperities (typically on the surface of stones embedded in the asphalt) as well as brightly colored small stones embedded in the bitumen, at the troughs or road asperities. While these anomalies exist, we will show that the assumed correlation between height and light intensity leads to representative height distributions when the parameters of the empirical function are predicted using GPR, using qualitative information about the type of road.

Starting from the height distribution analysis, the following parametric model is introduced to express the estimated heights as a function of the light intensity (LI):

$$heights_{LH} = h_0 \cdot \left(\frac{1 - e^{-k \cdot \frac{LI}{\max(LI)}}}{1 - e^{-k}} \right) \tag{1}$$

In this expression:

- $heights_{LH}$ denotes the estimated surface heights derived from the grayscale light intensity values;
- LI represents the grayscale light intensity field extracted from the input image;
- h_0 and k are two model parameters to be identified via numerical optimization to fit the experimental height distributions. Specifically, h_0 serves as a scaling factor corresponding to the maximum peak-to-valley distance of the reconstructed surface, while k governs the nonlinear response of the height transformation.

This exponential form is adopted as a simple, empirically motivated mapping that captures the expected monotonic decrease of reflected light with depth, without implying a direct derivation from optical reflection theory. It is also mathematical attractive as it comes with a limited number (two) of parameters with direct physical interpretation.

Road surfaces are composed of a heterogeneous mixture of aggregates and stones that vary significantly in color and texture. This diversity has a direct impact on the grayscale intensity distribution extracted from surface images. For instance, surfaces predominantly composed of dark-colored aggregates typically yield unimodal intensity histograms, with a concentration of pixel values skewed toward the darker end of the grayscale spectrum. On the other hand, when the surface comprises a mix of both light and dark stones, the resulting grayscale distribution often becomes distinctly bimodal, with significant peaks near both the bright and dark ends of the intensity range. A visual illustration of this phenomenon is provided in Fig. 4, where two different surface samples are presented alongside their respective grayscale intensity distributions. The first surface, composed mainly of dark stones, leads to a unimodal intensity histogram. In contrast, the second surface, containing a mixture of bright and dark stones, produces a bimodal distribution with two pronounced peaks — one associated with light tones and the other with dark ones. These differences stem solely from the optical reflectivity properties of the materials and are not necessarily indicative of the actual surface topography.

This inherent variability in the input data presents a significant modeling challenge. While the experimental height distributions obtained from roughness scans are typically unimodal — reflecting the natural randomness of rough surface profiles with a single dominant scale of elevation variation — the corresponding grayscale intensity distributions may not share this property. In cases where the input intensity distribution is also unimodal, the optimization process is relatively straightforward: the algorithm can effectively adjust h_0 and k to match the shape of the experimental height distribution, achieving a high-quality fit.

However, when the input grayscale distribution is bimodal, the situation becomes markedly more complex. The presence of two dominant modes in the intensity signal — resulting from heterogeneous surface coloration rather than topographical variation — introduces distortions that cannot be easily reconciled with the unimodal nature of the true height distribution. As a result, the optimization algorithm struggles to find a pair of parameters that simultaneously account for both intensity modes while still yielding a height histogram consistent with the ground-truth roughness data. This can lead to suboptimal convergence, where the reconstructed height distribution fails to accurately replicate the experimental one, particularly in terms of peak position, width, or overall shape. Indeed, the optimization procedure is effective under conditions of monotonic or unimodal light intensity inputs, while it struggles when faced with bimodal intensity distributions caused by visual heterogeneity of the surface materials.

Due to the potential for real road surfaces to contain bimodal distributions, there is a need to utilize mathematical strategies before starting the optimization procedure for finding the proper h_0 and k to match the shape of the experimental height distribution, that can

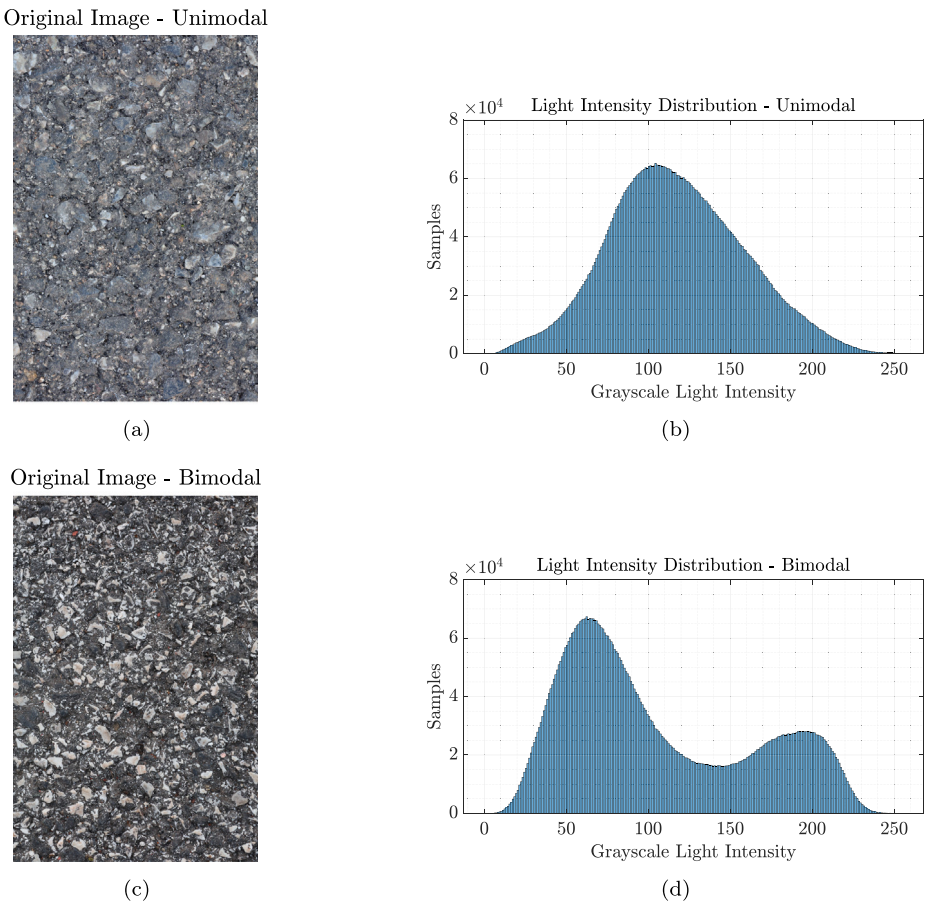

Fig. 4. Grayscale intensity analysis of road surfaces with varying aggregate composition: (a) Image of a surface primarily composed of dark stones, (b) Resulting unimodal intensity histogram concentrated on darker values, (c) Image of a surface with both bright and dark stones, (d) Bimodal intensity histogram reflecting the presence of contrasting materials.

account for reflectivity-driven artifacts in the image data. Hence, a randomized signal processing strategy was introduced prior to the numerical optimization of the model parameters. The rationale behind this approach is to mitigate the effect of local color variations and visual texture, which may distort the statistical correspondence between grayscale values and surface elevation. By doing so, the optimization process is guided by a light intensity signal that preserves the global structure of the data while reducing its dependence on material-specific reflectivity artifacts.

The randomization procedure is based on a specific analysis of the light intensity signal in the frequency domain. First, the grayscale light intensity values extracted from the input image are normalized and a Fourier transform is then applied to this signal, decomposing it into its amplitude and phase components. While the amplitude spectrum — representing the energy content at each spatial frequency — is preserved, the phase information is randomized. Specifically, a set of uniformly distributed random phase shifts is generated and applied to the positive-frequency components of the spectrum (excluding the zero-frequency and Nyquist terms), while ensuring symmetry for the negative-frequency components. This guarantees that the inverse Fourier transform of the modified spectrum remains a real-valued signal.

By randomizing the phase content while maintaining the original spectral amplitude, the signal retains the same global frequency distribution as the original light intensity data but exhibits a different spatial arrangement of grayscale values. This effectively de-correlates the light intensity from localized reflectivity features while preserving the characteristic roughness texture at a statistical level. The resulting randomized signal is then rescaled to ensure non-negativity and used as the input to the parametric model for height estimation (analytical formulation).

This mathematical strategy ensures that the numerical optimization of the model parameters h_0 and k is not influenced by reflectivity-induced bimodality. Instead, it relies on a statistically neutral version of the light intensity signal that better reflects the underlying geometric variations of the surface. The optimization is then carried out by minimizing the statistical discrepancy quantified as the difference between normalized histograms — between the heights estimated from the randomized light intensity and the experimentally measured height distribution. As a result, the obtained parameters are more representative of the physical topography and less sensitive to the optical composition of the surface.

A graphical example illustrating the analytical formulation applied to a representative surface, after numerical optimization of both h_0 and k, is presented in Fig. 5. The resulting height values are plotted as a function of the normalized light intensity, highlighting the characteristic nonlinear relationship imposed by the model.

The functional form was carefully chosen to ensure consistent physical behavior across the light intensity range. In particular, when the light intensity approaches zero (i.e., black pixels), the expression yields a height value close to zero, representing deep valleys or shadowed regions. Conversely, as the intensity approaches its maximum value (i.e., white pixels), the function asymptotically approaches h_0, corresponding to the surface peaks. The normalization by $\max(LI)$ ensures that the transformation remains scale-invariant with respect to lighting conditions. The exponential form ensures a smooth and saturating response, reflecting the diminishing marginal reflectivity changes at higher elevations. The quantities h_0 and k are optimized by minimizing

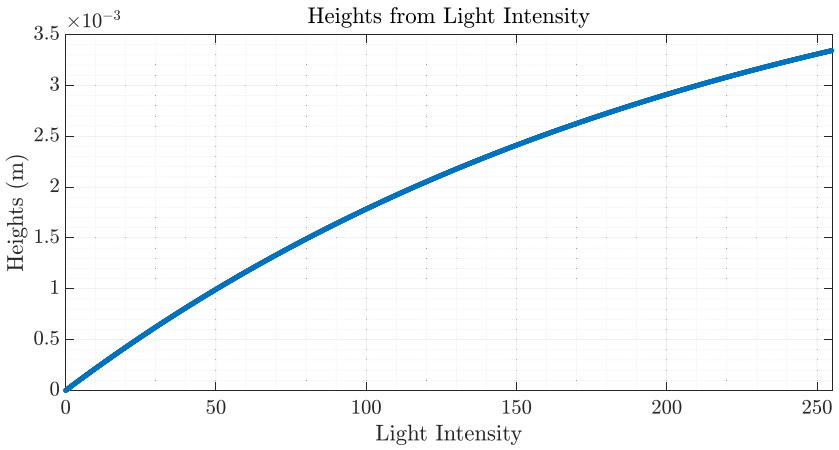

Fig. 5. Exponential function to obtain heights starting from light intensity information.

the statistical discrepancy — measured in terms of histogram distance — between the reconstructed and experimentally measured height distributions, thus ensuring that the reconstructed surface is not only visually plausible but also statistically faithful to the true topography.

Multiple optimization strategies were employed and compared to evaluate their effectiveness in minimizing this discrepancy. Although several optimization algorithms were explored and compared, including `lsqnonlin` combined with a `multistart`, `fmincon` and `surrogateopt`, the final implementation relied on `fminsearch` due to its simplicity, speed and robust performance across all tested surfaces. Its derivative-free nature and low computational cost made it particularly suitable for the low-dimensional parameter space of the proposed model.

Once the optimal values of h_0 and k have been identified through the numerical optimization procedure, the analytical model can be applied to convert grayscale image intensities into quantitative surface height information. By construction, this transformation is designed to reproduce the statistical distribution of the heights measured through the scanner.

Fig. 6 presents a selection of representative distributions (before de-trending) from the optimization process applied to different road samples, including both a racing track (Fig. 6(a)) and an urban asphalt (Fig. 6(b)). It also includes very particular surfaces such as sandpaper (Fig. 6(c)) and marble sample (Fig. 6(d)) to present the response of the model in unconventional texture scenarios. In particular, each subplot of the mentioned figure compares the experimental height distribution obtained from high-resolution laser scanning with the reconstructed distribution generated by the proposed methodology. Despite the diversity of surface textures and light intensity profiles, the optimized model successfully captures the essential statistical shape of the measured height distributions, confirming the robustness and modularity of the approach across varying road textures.

3.2. Frequency domain analysis for spectral consistency

While distribution-based analyses of height profiles capture the statistical amplitude characteristics of a surface, they are not sufficient to fully describe its roughness properties [39]. Indeed, two signals may exhibit identical amplitude distributions but differ significantly in their spectral content — that is, in the spatial frequency composition that governs texture sharpness, wavelength and in general contact and frictional behavior [40,41]. For this reason, a physically accurate reconstruction of surface profiles must not only match the statistical distribution of heights but also preserve their spectral signature. To this end, starting from the general criteria presented in [34], a correction procedure is implemented that combines frequency-domain scaling with iterative rank ordering, enabling simultaneous matching of the power spectral density and amplitude distribution of experimental profiles [42].

The method begins with a frequency analysis of the reference ground-truth surfaces (obtained through laser scanning) and the initial reconstructed heights derived from the light-based estimation. In particular, the evaluation of the Fast Fourier Transform (FFT) is performed, computing both the magnitude spectrum (modulus) and the phase information.

To correct the spectrum of the heights derived from the light-based profiles, a representative scaling function (SF) was computed as the ratio between the average FFT magnitudes of the scanned data and those of the reconstructed signals derived from light intensity (after the optimization process described previously). The formulation of the scaling function is provided in Eq. (2).

$$SF = \frac{|FFT_{SCAN}|}{|FFT_{LH}|} \tag{2}$$

A graphical representation of the scaling functions computed for different surfaces is shown in Fig. 7.

The computation of the scaling function is a critical step in the spectral correction process, as it enables the construction of the corrected FFT that serves as the spectral reference. Specifically, this corrected FFT is obtained by multiplying the scaling function — encoding the amplitude (magnitude) information — with the FFT of the light-based profile after exponential transformation to extract the heights, which retains the original phase information. This formulation ensures that the corrected signal inherits the spectral content of the scanned reference while preserving the phase structure derived from the image-based reconstruction.

As a result, a frequency-domain corrected version of the synthetic signal is obtained. The inverse FFT of this corrected spectrum yields a first estimate of a spatial signal whose spectral content now resembles that of the scanned surface. However, the amplitude distribution of this reconstructed signal no longer matches the target distribution obtained after the exponential model fitting. To enforce distributional consistency, a rank-ordering operation is performed: the samples of the reconstructed signal are sorted in descending order and replaced with the corresponding sorted values from the target signal, i.e., the one obtained after applying the exponential formulation (Eq. (1)). This substitution ensures that the statistical distribution of amplitudes remains consistent with the desired target.

This procedure results in a signal that once again loses the desired frequency content. For this reason, the procedure is embedded within an iterative loop in-line with the approach proposed by Pérez-Ràfols and Almqvist in [34]. At each iteration, the current signal is transformed into the frequency domain, rescaled using the reference spectrum (using the corrected FFT) and then returned to the spatial

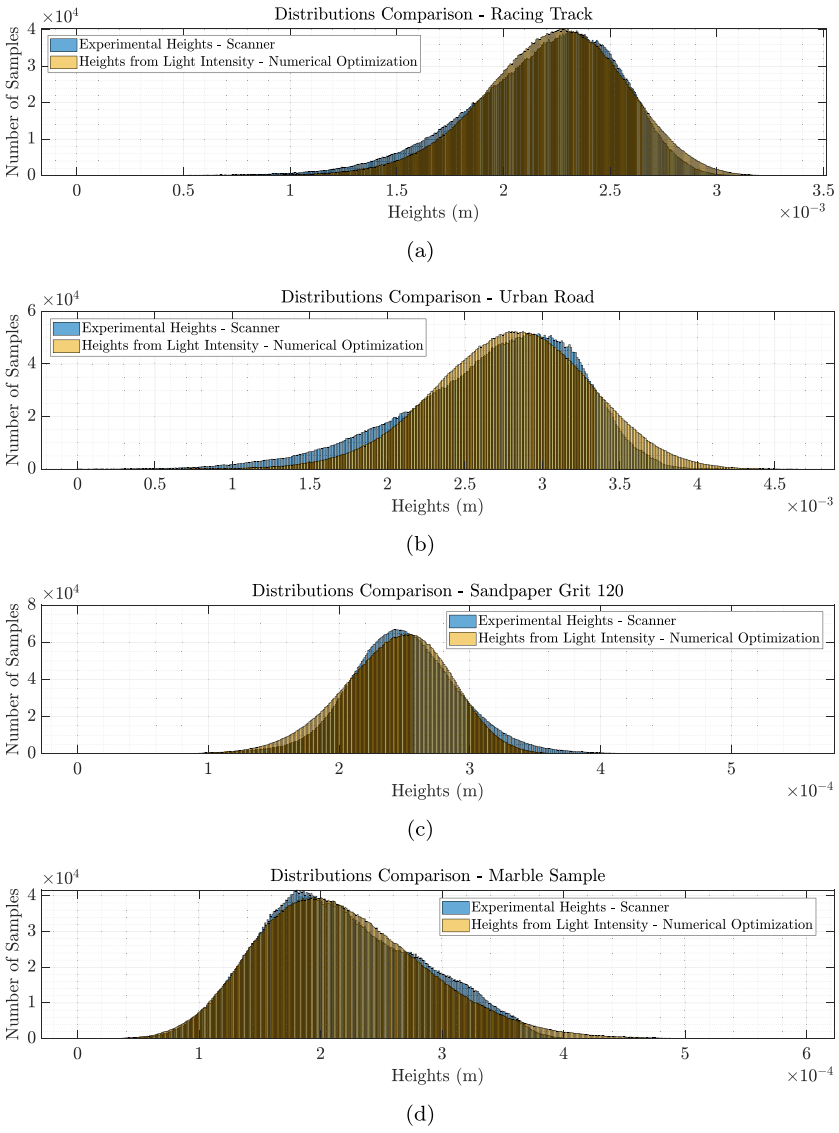

Fig. 6. Comparison between experimental height distributions obtained via laser scanning and the reconstructed distributions derived from image-based light intensity for different surface types: (a) Racing track asphalt, (b) Urban road pavement, (c) Sandpaper with grit size 120, (d) Marble surface.

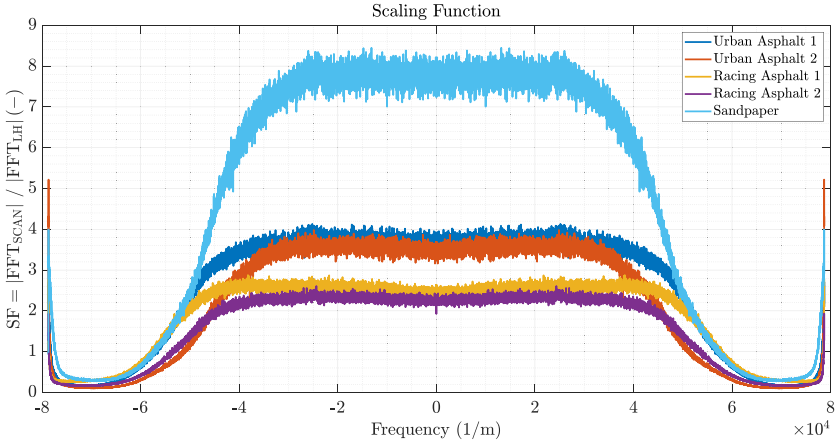

Fig. 7. Scaling function to correct the frequency spectrum for different surfaces.

domain via inverse FFT. After each rescaling, rank-ordering is reapplied to maintain the desired distribution. The iterations continue until convergence, defined by a sufficiently small tolerance that ensures both the desired spectrum and statistical distribution are matched. This

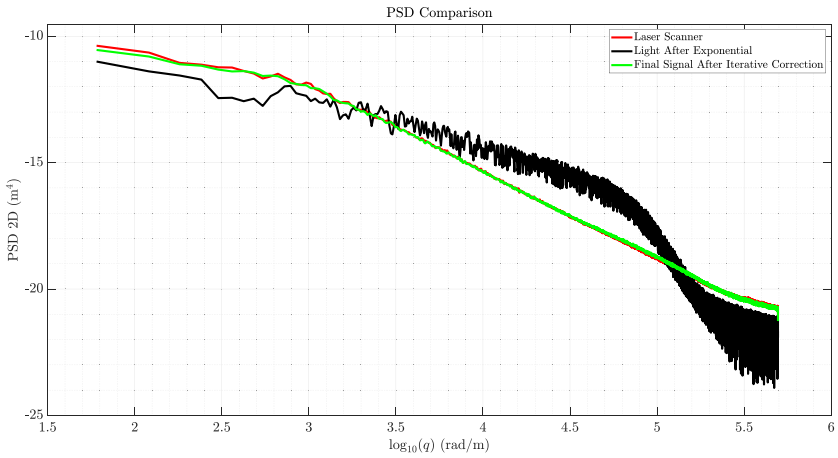

Fig. 8. Power Spectral Density comparison before and after spectral correction.

iterative correction process ensures simultaneous convergence toward the desired power spectral density (PSD) and heights distribution. At the end of the loop, the final reconstructed signal closely replicates both the spectral properties and the statistical distribution of the target experimental profile. This guarantees that the generated surface not only looks statistically similar but also has comparable roughness behavior in both low- and high-frequency components, which are crucial for contact mechanics and tire–road interaction modeling.

To demonstrate that the proposed procedure not only ensures the correct height distribution but also restores the spectral characteristics of the surface, the Power Spectral Density (PSD) of the final reconstructed signal is compared with that of the original scanned surface and with the intermediate signal obtained immediately after the exponential transformation (which required correction). In particular, the 2D PSD was evaluated following the approach described in [21], and the comparison is presented in Fig. 8.

As shown in Fig. 8, the PSD of the final signal obtained after the iterative rank-ordering procedure exhibits excellent agreement with the spectral characteristics of the original scan, without altering the height distribution, which remains consistent with the desired target thanks to the rank-ordering process. This demonstrates that the iterative correction process, by keeping the desired statistical distribution of the surface heights, successfully restores the spatial frequency content of the real surface. In contrast, the intermediate signal (after the exponential transformation to fix distributions but before spectral correction) shows a clear deviation from the desired PSD, confirming the need for the full correction process. The procedure thus ensures that the reconstructed virtual surface preserves both the statistical height distribution and the frequency-dependent roughness characteristics of the physical reference — both of which are key requirements for applications in tribology and contact mechanics.

A graphical representation illustrating some different surfaces as example how the procedure is able to capture differences among the analyzed surfaces — showing, for each surface, both the corresponding associated 2D PSDs — is presented in Fig. 9.

A graphical representation demonstrating the ability of the procedure to replicate the differences among the various analyzed surfaces is presented in Fig. 10. It can be noted that even surfaces with highly distinctive characteristics, such as sandpaper or marble, are realistically reproduced by the procedure, clearly exhibiting substantial differences in roughness properties compared to typical road surfaces. This highlights the robustness of the procedure, which is capable of accurately replicating the physical reality even in the case of highly distinctive and unconventional surfaces — significantly different from typical roads — demonstrating its ability to generalize beyond standard scenarios.

3.3. Predictive modeling approach with Gaussian process regression

To enable predictive surface generation in the absence of a physical scanner, it becomes evident — based on the methodology described thus far — that both the scaling function and the quantities h_0 and k must be accurately estimated. For this purpose, Gaussian Process Regression (GPR) is employed to directly infer these quantities from image-based features, enabling the reconstruction of realistic virtual surfaces that match both the amplitude and the spectral characteristics of the physical references.

GPR is a non-parametric, kernel-based Bayesian regression method that models the target function as a distribution over functions. One of its key advantages is its probabilistic prediction, which yields not only a mean estimate but also a measure of uncertainty, making it especially valuable in contexts where data is limited, noisy, or highly heterogeneous — as is often the case with surface measurements [43–45]. These attributes represent a major innovation compared to alternative solutions such as Convolutional Neural Networks (CNNs), which — despite their popularity — suffer from several limitations in this context. CNNs typically require large labeled datasets, which are often unavailable or impractical to acquire for surface analysis. Moreover, they behave as black boxes, offering limited physical interpretability and they lack intrinsic uncertainty quantification. Their predictions are harder to control and training can be computationally intensive. In contrast, GPR provides a data-efficient, transparent and physically coherent modeling framework, making it far more suitable for engineering applications such as virtual surface reconstruction.

The input space for the Gaussian Process Regression (GPR) model is built from a set of normalized features derived from grayscale image analysis and frequency-domain characterization. To properly train the model, a dedicated dataset was assembled, comprising approximately 20 representative surfaces — including non-road surfaces such as sandpapers — to challenge the model under unconventional conditions. Each surface was processed to extract the following features:

- Spatial frequency values, obtained from FFT analysis;
- Image contrast, to reflect variations in brightness and texture;
- Skewness of the grayscale intensity distribution, providing information on asymmetry in surface reflectance;
- Standard deviation of the randomized intensity signal, capturing local surface irregularities;
- Surface type, encoded using binary flags to provide surface classification details;
- Target values for each surface: the scaling function SF (evaluated at selected frequencies), as well as the optimized quantities h_0 and k.

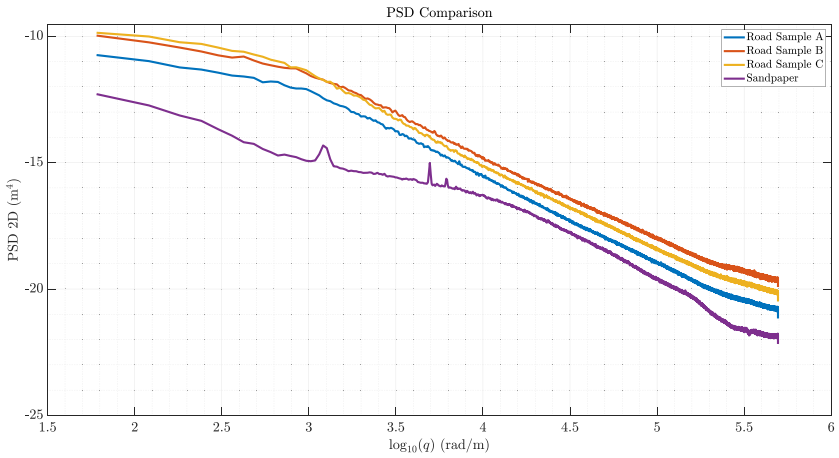

Fig. 9. Power Spectral Density comparison among different surfaces.

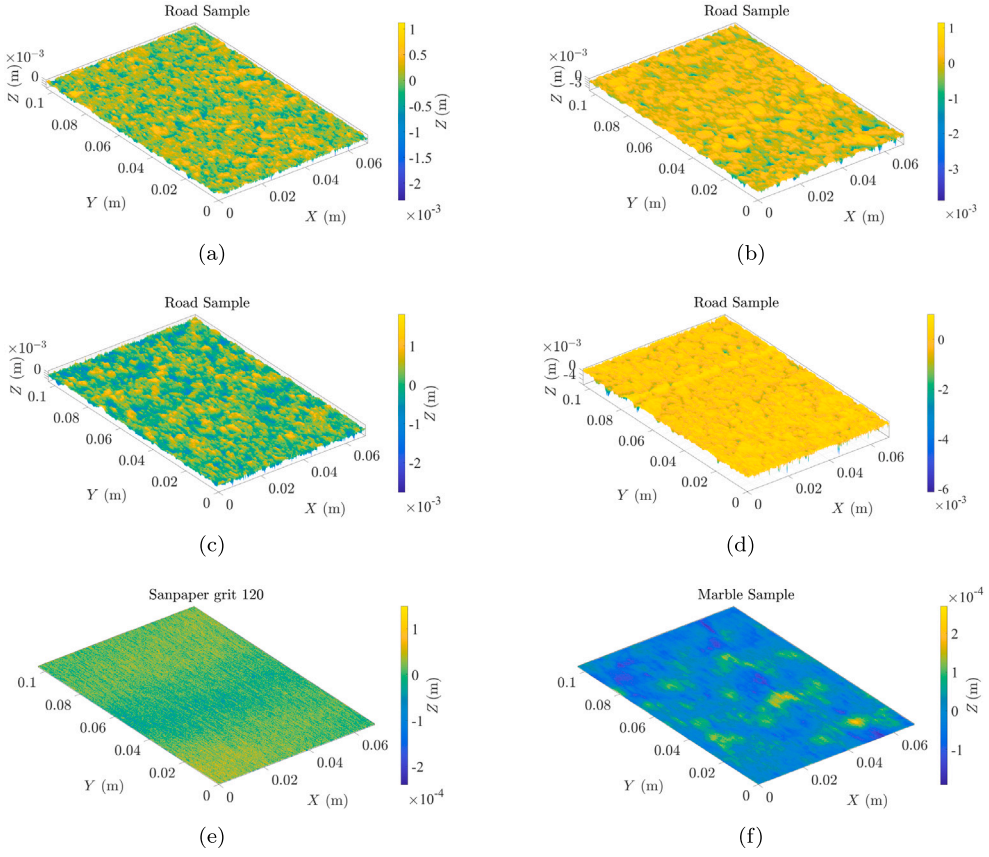

Fig. 10. Virtual surfaces generation from images with target FFT and height distribution: (a) to (d) Different road samples, (e) Sandpaper grit 120, (f) Marble sample.

As mentioned, these features are assembled into a comprehensive dataset that captures both the spatial and spectral characteristics of the surfaces. To construct the spectral component of the input, the Nyquist criterion was considered for selecting the frequency range at which the scaling function is learned. For each frequency point, a local noise estimate is also computed as the standard deviation of the SF within a moving window, which is then used to calibrate the likelihood function of the GPR.

The output of the model contains the surface-specific quantities to be predicted: the value of the scaling function at each frequency point and the quantities h_0 and k, optimized through the exponential fitting process discussed in the previous sections. To provide a clearer understanding of the dataset structure used for GPR training, Table 1 summarizes the variables included in the input–output pairs.

To clarify the notation adopted in Table 1, which serves as a representative example of the GPR training dataset, the following elements are defined:

- f_i denotes the ith frequency point at which the scaling function is evaluated;
- c_j represents the contrast value associated with the jth surface, which remains constant across all frequencies for that surface;

Table 1
Structure of the GPR training dataset, detailing the input features and corresponding output targets.

Surface	Freq.	Cont.	Sk	Std	Asph.	Conc.	Sand.	Other	SF	h_0	k
Surf. 1	f_1	c_1	sk_1	std_1	1	0	0	0	$SF_1(f_1)$	h_1	k_1
Surf. 1	f_2	c_1	sk_1	std_1	1	0	0	0	$SF_1(f_2)$	h_1	k_1
Surf. 1	f_3	c_1	sk_1	std_1	1	0	0	0	$SF_1(f_3)$	h_1	k_1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Surf. 1	f_m	c_1	sk_1	std_1	1	0	0	0	$SF_1(f_m)$	h_1	k_1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Surf. n	f_1	c_n	sk_n	std_n	0	0	1	0	$SF_n(f_1)$	h_n	k_n
Surf. n	f_2	c_n	sk_n	std_n	0	0	1	0	$SF_n(f_2)$	h_n	k_n
Surf. n	f_3	c_n	sk_n	std_n	0	0	1	0	$SF_n(f_3)$	h_n	k_n
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Surf. n	f_m	c_n	sk_n	std_n	0	0	1	0	$SF_n(f_m)$	h_n	k_n

- sk_j and std_j denote the skewness and the standard deviation of the randomized intensity signal, all associated with the jth surface and they remain constant across all frequencies;
- The columns containing binary values (0 or 1) indicate categorical flags used to classify the type of surface (e.g., asphalt, sandpaper, etc.);
- SF, h_0 and k are the output quantities: the scaling function SF varies with both frequency and surface, while h_0 and k depend only on the surface.

An additional remark should be made regarding the use of binary variables, which proved essential to support the training of the GPR model. These variables were introduced primarily to provide a simple yet effective classification of the surfaces represented in the dataset (e.g., asphalt, concrete, sandpaper). At the early stage of development, more conventional unsupervised approaches such as k-means clustering were also explored to automatically group surfaces. However, this approach revealed substantial limitations: the algorithm tended to oversimplify the classification, frequently grouping all road surfaces on one side and non-road materials like sandpapers on the other. Any attempt to refine the number of clusters led to unstable behavior and limited interpretability, making the process less controllable and less meaningful from a physical standpoint. These limitations were overcome with the use of binary flags to enable direct supervision of the classification. This is straightforward to use during prediction, as the surface class is typically known and can be directly provided as an input to the model, avoiding additional uncertainty or error propagation that might result from a separate clustering step. Binary flags were therefore retained to allow direct supervision of surface classification, ensuring consistent and reliable inputs for the GPR model during both training and prediction.

The Gaussian Process Regression (GPR) model adopted in this work is implemented using the GPML Toolbox, a widely used MATLAB library for probabilistic modeling [46], which offers a flexible and modular framework for probabilistic regression. In this work, the GPR is configured to perform multi-output regression in order to jointly predict the previously mentioned quantities (e.g. the scaling function, h_0 and k).

To obtain the desired outputs, the model is configured using a constant mean function, a rational quadratic kernel (RQiso) as covariance function which generalizes the squared exponential (SE) kernel and a Gaussian likelihood function. The rational quadratic kernel is particularly well suited for this application, as it can model multi-scale behaviors by blending different length scales — an essential feature when dealing with surface roughness across a broad frequency spectrum.

The model is trained by minimizing the negative log marginal likelihood, which automatically balances data fitting and model complexity. Hyperparameters — including the characteristic length scale, signal variance and observation noise — are optimized using gradient-based methods provided by the toolbox (conjugate gradient descent). The training process is performed jointly across all outputs, assuming a shared kernel structure and a fixed noise model derived from local FFT statistics. This approach ensures consistency and stability when learning from limited and heterogeneous datasets.

Once trained, the GPR model is used to make probabilistic predictions for a new surface, using image-derived features and frequency points as input. In particular, the model takes as input a set of normalized features extracted from a new surface image such as the frequency values, image contrast, standard deviation of randomized light intensity and surface type flags. The GPR then generates predictive distributions: a mean estimate and associated uncertainty for each output. In this way, the desired quantities can be predicted through a dedicated workflow, enabling the generation of realistic virtual surfaces that retain both the spectral characteristics and statistical properties of the original physical scans, without requiring a physical scanner.

To demonstrate the effectiveness of the training process, the model was then tested on some of the surfaces included in the training dataset. The results for one of the surfaces are illustrated in Fig. 11, where it is evident that the virtually generated surface preserves both the Power Spectral Density (PSD) and the statistical distribution of heights observed in the experimental data obtained from the laser scanner. The PSD after numerical optimization is also shown, representing the PSD obtained in the first step of the process — numerical optimization — to verify that the underlying signal processing method is indeed capable of delivering reliable results. This nomenclature applies to all figures presented hereafter.

The consistency observed in both the statistical distribution and the Power Spectral Density (PSD) of the reconstructed surface can be directly attributed to the accurate prediction of the scaling function and the quantities h_0 and k by the Gaussian Process Regression (GPR) model. While this result is expected — given that in this occasion the tested surface was part of the training dataset — it remains a valuable demonstration of the model's internal consistency (effectiveness of the training process). For the sake of completeness, Fig. 12 illustrates the GPR-based estimation of the quantities required to run the surface reconstruction in fully predictive mode for this particular surface. The plots show an excellent match between the predicted and reference values, confirming the model's ability to capture the underlying physical quantities with high accuracy.

4. Results and discussion

Once the Gaussian Process Regression (GPR) model is trained, it can be used in a fully predictive mode to reconstruct road surface profiles directly from grayscale images. As detailed in the previous sections, once the scaling function and the quantities h_0 and k are predicted by the GPR, the full roughness surface can be generated using the reconstruction algorithm presented previously.

To validate the predictive capabilities of the model, a training set composed of approximately 20 surfaces of varying nature was used, including both conventional road materials and more unconventional textures (e.g., sandpaper samples of different grit sizes).

A simulation was carried out to validate the robustness of the proposed methodology. In this case, the trained GPR model was used

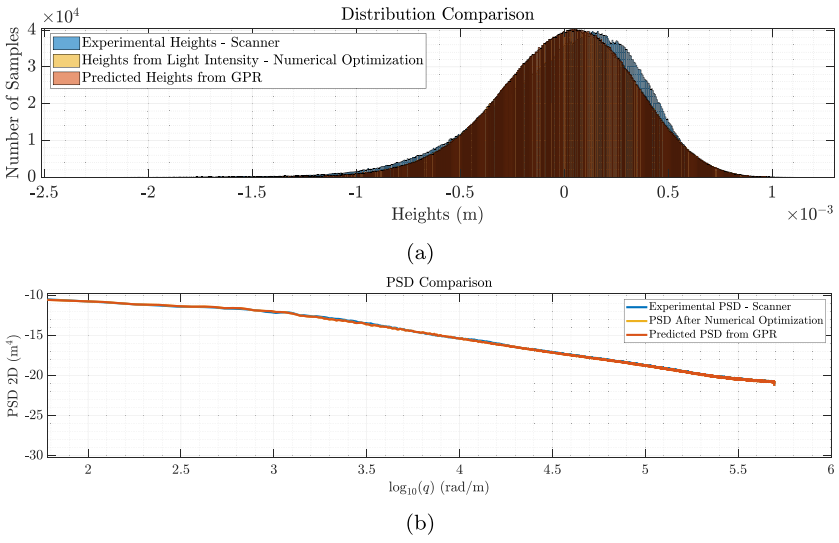

Fig. 11. Testing of the GPR model on a surface included in the training set: (a) Comparison of height distributions, (b) Comparison of Power Spectral Densities (PSD).

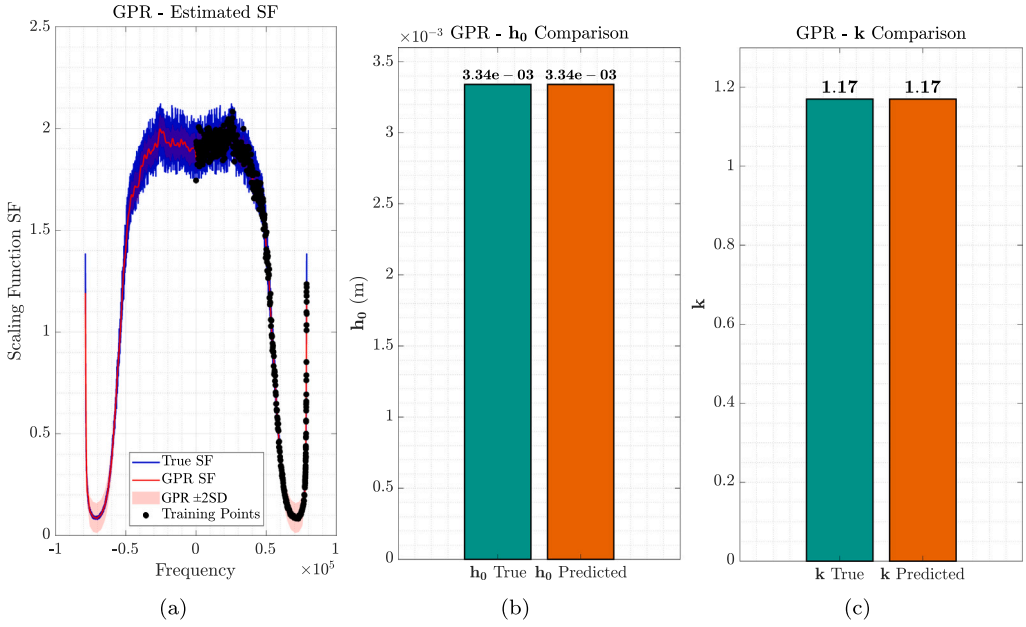

Fig. 12. Evaluation of the GPR model on a surface included in the training set: (a) Estimated scaling function SF with confidence bounds, (b) Predicted vs reference value of h_0, (c) Predicted vs reference value of k.

in a predictive run on a surface already included in the training dataset — but using a completely different image of the same physical surface, one that was never used during training. This scenario tests the model's ability to generalize not just across different surfaces, but also across different image instances of the same surface. The outcome of this experiment is shown in Fig. 13. As can be observed, the reconstructed surface retains a PSD that closely matches the one obtained from laser scanning, along with a highly consistent statistical distribution of surface heights. This result confirms the model's capacity to correctly reconstruct the underlying roughness profile even when presented with new visual data of a known surface.

Such performance highlights the validity and reliability of the overall process, demonstrating that the model does not merely memorize training samples but effectively learns the underlying relationships between image-derived features and physical surface properties. This reinforces the predictive strength and practical applicability of the approach, particularly in scenarios where multiple images of the same surface might be available or captured under different conditions.

In this scenario, the successful reconstruction of the virtual surface — exhibiting both the correct height distribution and Power Spectral Density (PSD) — can be attributed to the accurate prediction of the three key quantities by the Gaussian Process Regression (GPR) model. Specifically, the consistency of the output stems from the precise estimation of the scaling function, h_0 and k, despite the fact that the input image was not used during the training phase. This highlights the generalization capability of the model. The corresponding results are illustrated in Fig. 14, where the predicted scaling function SF, h_0 and k for this unseen image are shown and compared to their respective reference values.

To further challenge the generalization capabilities of the model, the GPR was tested on a surface that was not included in the training dataset at all, just by taking the picture of it. Specifically, one surface was deliberately removed from the dataset prior to training. After

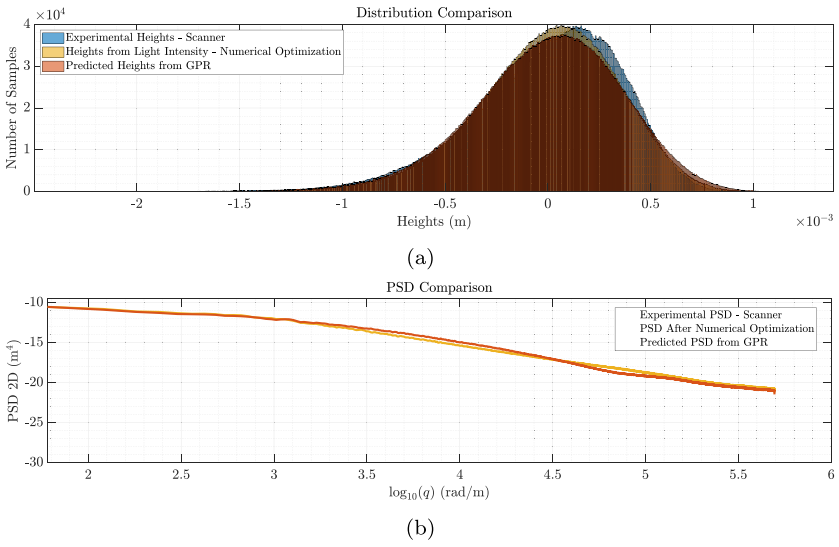

Fig. 13. Predictions of the GPR model on a surface included in the training set but from a different picture used to build the dataset: (a) Comparison of height distributions, (b) Comparison of Power Spectral Densities (PSD).

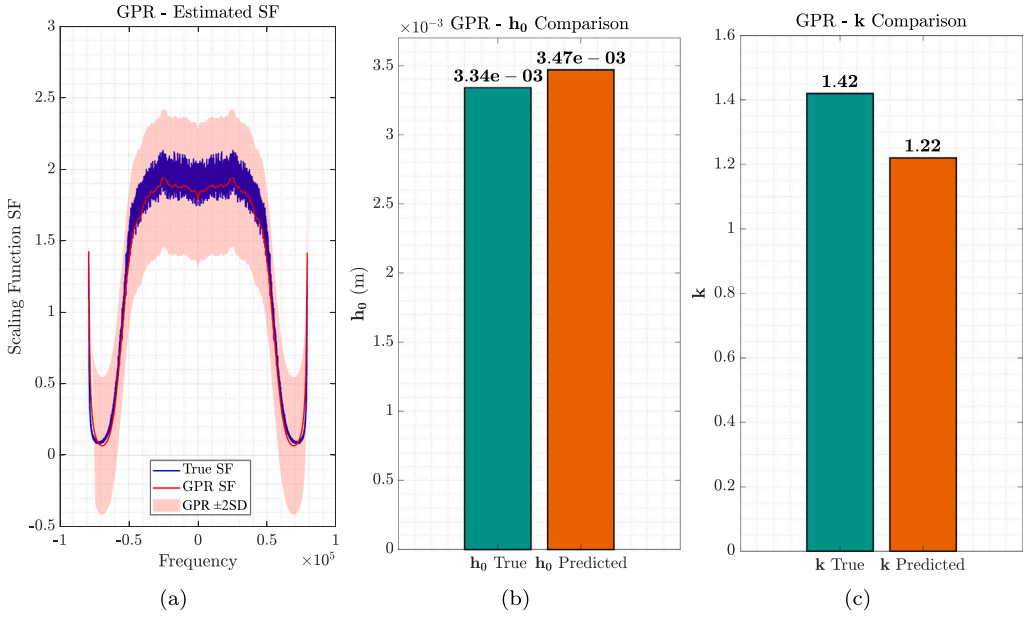

Fig. 14. GPR prediction for an alternative unseen image of a surface included in the training set with a different image: (a) Predicted scaling function SF, (b) Estimated vs true h_0, (c) Estimated vs true k.

training the GPR model on the remaining surfaces, a prediction was carried out on the excluded surface using only a single grayscale image of it. All necessary input features were extracted directly from this image, allowing the model to operate in a fully predictive mode without having ever encountered the surface during training. This represents a realistic deployment condition, in which the trained model is able to generate the surface texture from image-derived features through the proposed algorithm. The results, shown in Fig. 15, demonstrate that the predicted surface closely replicates both the spectral and statistical characteristics of the ground-truth scan, confirming the model's extrapolation potential.

The results shown in Fig. 15 stem from the GPR-based estimation of the scaling function SF and the parameters h_0 and k, which — when combined — enable a close match between the reconstructed and experimental surface in terms of both height distribution and frequency content. The estimates produced by the GPR for these key quantities, in this case of a surface entirely excluded from the training dataset, are presented in Fig. 16.

As illustrated in Fig. 16(a), even for surfaces that were not included in the training dataset, the mean value of the frequency scaling function shows very good correlation with ground truth. This is significant as the surface PSD is the single most important surface metric for the calculation of friction [15,16]. However, in the same figure it is shown that the mean value estimate comes with significant uncertainty in the prediction at certain frequencies. This is a result of the very limited dataset of surfaces. It is expected that as additional surfaces are added in the future, the uncertainty band is likely to shrink accordingly. As any uncertainty in the scaling function is reflected directly in the estimated friction, it is recommended that friction calculations are performed using the mean value but also the limits of the frequency scaling function.

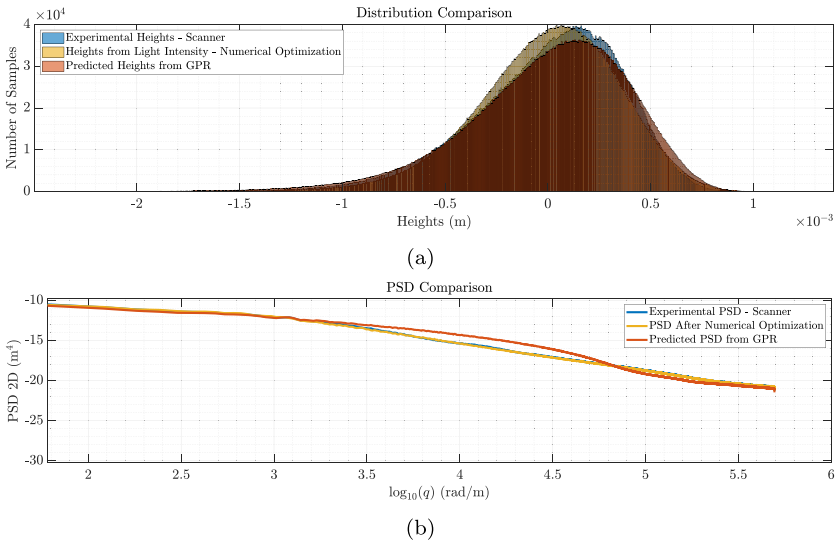

Fig. 15. Predictions of the GPR model on a surface not included in the training dataset: (a) Comparison of height distributions, (b) Comparison of Power Spectral Densities (PSD).

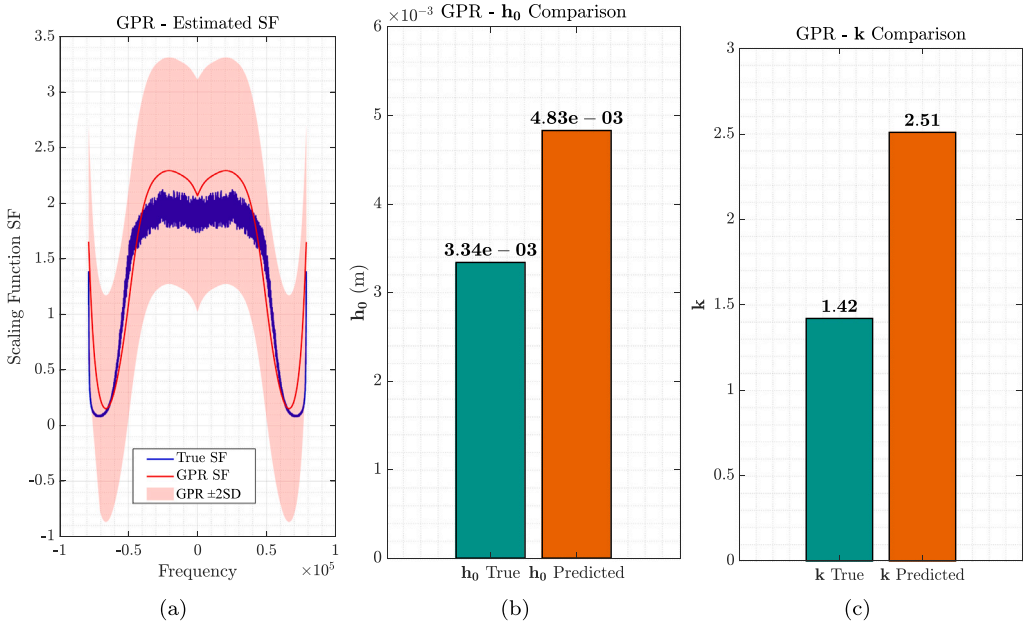

Fig. 16. GPR prediction for a surface not included in the training dataset: (a) Predicted scaling function SF, (b) Estimated vs true h_0, (c) Estimated vs true k.

The computational performance of the proposed methodology was also evaluated. The GPR training required approximately 100 function evaluations and about 1.8 h on a standard 10 years old laptop (Intel Core i7, 9th generation), while the prediction of the quantities via GPR (h_0, k, and SF) together with the complete reconstruction pipeline, including feature extraction, height distribution and spectral corrections, required approximately 100 s. This demonstrates that, once a simple and cost-effective image is acquired, the proposed method enables the reconstruction of surface roughness through an offline computation, providing a practical alternative to traditional scanning techniques, with further reductions in processing time achievable on more advanced hardware.

Additionally, an important investigation was carried out concerning the image format used as input for the predictive model. Under the same conditions previously described — namely, with the target surface removed from the training dataset — the GPR model was fed both a JPG and a RAW version of the same photograph. The goal was to assess whether the internal processing and compression typically applied during format conversion could impact the model's ability to accurately reconstruct the height distribution and frequency spectrum of the surface. The results, shown in Fig. 17, demonstrate that there are no significant differences between the two formats: in both cases, the model successfully replicates the experimental data and the difference between the two formats is minimal. This outcome confirms that the minimal preprocessing performed by the device considered when converting from RAW to JPG does not compromise the model's predictive accuracy.

The proposed methodology can be extended to additional surfaces following the same procedure. To further validate the robustness of the approach, an additional test was carried out by removing a surface characterized by sandpaper grit 120 from the training dataset. Alongside the previously analyzed road surface, this second surface was used as input for the GPR model in fully predictive mode. The final virtually reconstructed surfaces for both cases are shown in Fig. 18, highlighting

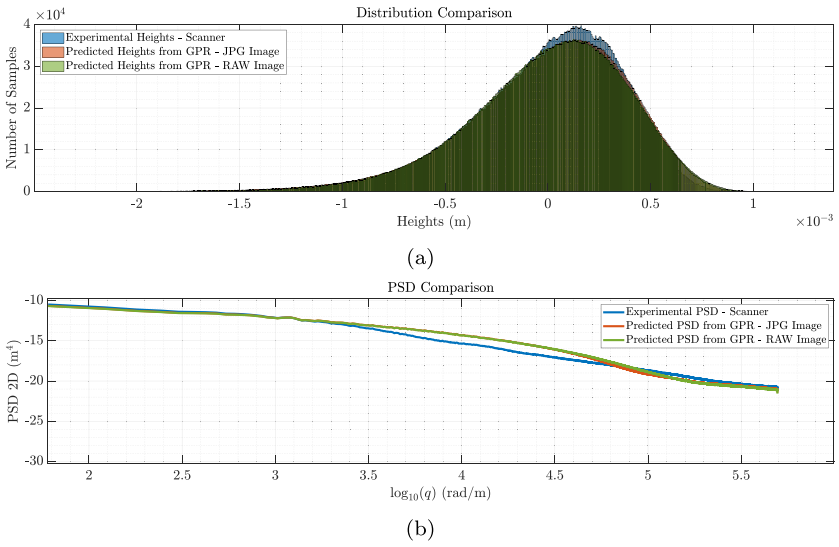

Fig. 17. Effect of input image format on GPR predictions for a surface excluded from the training dataset: (a) Height distribution comparison for RAW and JPG inputs with respect to the experimental heights, (b) Power Spectral Density (PSD) comparison showing consistent reconstruction performance across formats with respect to the experimental PSD.

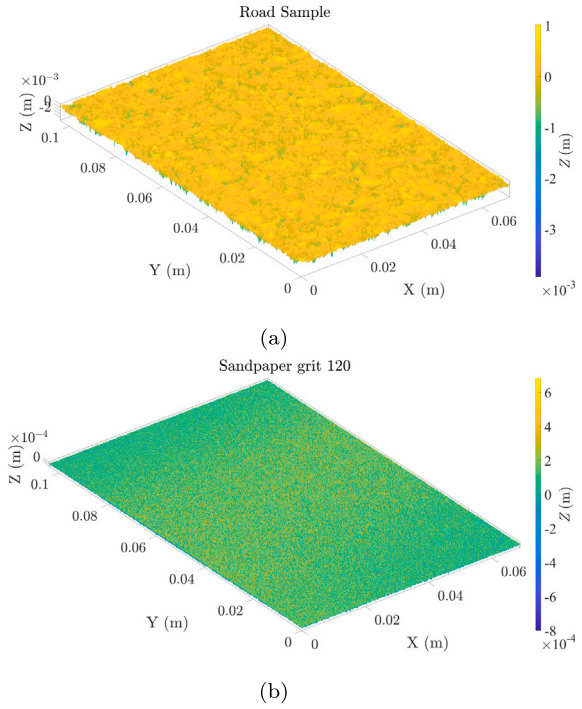

Fig. 18. Reconstructed virtual surfaces using the GPR model in fully predictive mode: (a) Road surface, (b) Sandpaper grit 120.

the model's ability to generalize and accurately reproduce surfaces of different nature.

5. Conclusions

This work introduces a novel and physically interpretable framework for reconstructing three-dimensional road surface roughness directly from grayscale images. Unlike conventional scanning techniques, which are often time-consuming, costly and impractical in real-world scenarios — especially in motorsport applications where speed and flexibility are crucial — the proposed method provides a fast, low-cost and easily deployable alternative. The ability to obtain surface characteristics using only photographic inputs represents a significant breakthrough, enabling surface reconstruction in settings where profilometric tools are not feasible.

The core innovation lies in the integration of a physics-driven model of light-surface interaction with a Gaussian Process Regression (GPR) architecture. This combination not only ensures mathematical consistency and interpretability, but also offers a level of physical controllability that is typically absent in purely data-driven black-box models such as convolutional neural networks. Furthermore, the structure of the dataset allowed the GPR model to inherently incorporate surface classification using binary indicators, further enhancing the model's transparency and explainability.

The robustness and versatility of the approach were validated across both asphalt and artificial textures like sandpaper. The methodology was tested under several conditions: predicting the 3D surface from a photo of a surface included in the training set; predicting the surface using a different photo of a known surface; reconstructing entirely unseen surfaces excluded from training. In all cases, the reconstructed surfaces consistently preserved both the statistical height distribution and the spectral content (PSD), demonstrating the predictive accuracy and physical soundness of the framework. The uncertainty in the predictions is reflective of the relatively small training dataset and is expected to improve with the addition of scanned surface information in the future.

This research activity framework opens several promising directions for future research and practical deployment. One natural extension involves the direct coupling of the reconstructed surface roughness from images with physical models of friction, enabling predictive estimation of the grip based solely on photographic data and known compound properties. By integrating this procedure with viscoelastic or thermodynamic models of the tire, it would be possible to estimate the available friction under various operating conditions, including different temperatures, loads and surface types. Indeed in motorsport, where decision-making must be both rapid and accurate, the ability to quantify grip potential from on-site photographs of the track surface could represent a paradigm shift.

Finally, this methodology holds potential for applications in civil engineering, particularly in the domains of road condition monitoring and predictive maintenance. Its ability to rapidly capture surface roughness makes it well-suited for assessing asphalt degradation and texture evolution, with direct implications for ride comfort, vehicle safety and maintenance planning. Although the proposed approach proved robust

under controlled lighting conditions, extreme illumination scenarios or highly reflective surfaces may introduce local saturation effects in the captured images. Future developments of the methodology may include controlled diffuse illumination setups (e.g., portable shielding systems) to further improve robustness in challenging lighting environments. The model could be integrated into lightweight platforms — such as smartphone-based applications — for scalable deployment across urban and interurban road networks.

CRediT authorship contribution statement

Raffaele Stefanelli: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Georgios Mavros:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Miguel Martínez-García:** Methodology, Resources, Writing – review & editing. **James Knowles:** Writing – review & editing, Resources, Investigation. **Francesco Timpone:** Writing – review & editing, Supervision, Funding acquisition. **Flavio Farroni:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to express their sincere gratitude to the research groups of the Department of Aeronautical and Automotive Engineering at Loughborough University and the Vehicle Dynamics Department of the University of Naples Federico II for providing the experimental facilities and technical resources necessary for conducting this work.

Appendix A. Practical workflow for surface roughness reconstruction from images

This appendix provides a practical step-by-step overview of the methodology presented in this work, illustrating how the proposed framework can be applied in practice, from the creation of the training dataset to the prediction of surface roughness for a new road or racetrack surface using only photographic input.

A.1. Dataset construction and reference measurements

The first step in the proposed methodology consists in building a dataset that links image-based information to physically measured surface roughness. For this purpose, a set of representative surfaces is selected, covering a sufficiently wide range of textures and materials (e.g., urban asphalt, racetrack surfaces, laboratory samples, or artificial textures such as sandpapers).

For each surface included in the dataset, two types of data are acquired:

- High-resolution photographs, captured using the experimental setup described in Section 2. Particular care must be taken to ensure consistent lighting conditions, avoiding saturated reflections and maintaining diffuse illumination when possible.
- Reference roughness measurements, obtained using a surface texture scanner (e.g., AMES scanner). These measurements provide the ground-truth height profiles used to calibrate and validate the reconstruction procedure.

The scanned area should correspond as closely as possible to the physical region captured in the photograph, ensuring spatial consistency between the optical data and the roughness measurements. In practice, an exact overlap between the scanned and photographed regions is difficult to achieve; therefore, careful alignment is performed to ensure that both acquisitions refer to approximately the same surface area.

A.2. Image pre-processing and feature extraction

Once the images and the corresponding roughness scans are acquired, the preprocessing procedure described in Section 2 is applied.

The main steps include:

- Conversion of RGB images to grayscale;
- Geometric alignment between the image and the scanned region;
- Extraction of equally spaced vertical columns corresponding to the scanner profiles;
- Resampling of the intensity signals to match the resolution of the scanning device;
- Application of an anti-aliasing low-pass filter to preserve the physical bandwidth of the signal.

From the processed grayscale image, a set of descriptive features is then extracted for use in the predictive model. These include:

- spatial frequency values obtained from FFT analysis,
- image contrast,
- skewness of the grayscale intensity distribution,
- standard deviation of the randomized intensity signal,
- surface type classification flags.

These features form the input variables for the Gaussian Process Regression (GPR) model.

A.3. Calibration of the analytical height transformation

For each surface included in the dataset, the analytical model introduced in Section 3.1 is calibrated. The grayscale light intensity values are transformed into estimated surface heights through the exponential formulation:

$$heights_{LH} = h_0 \cdot \left(\frac{1 - e^{-k \cdot \frac{LI}{\max(LI)}}}{1 - e^{-k}} \right) \tag{A.1}$$

The quantities h_0 and k are identified through numerical optimization by minimizing the discrepancy between the reconstructed height distribution and the experimental distribution obtained from the scanner.

In cases where the grayscale intensity distribution is affected by reflectivity-driven bimodality, the phase-randomization procedure described in Section 3.1 is applied prior to the optimization process to remove spatial correlations linked to material color rather than topography.

A.4. Determination of the spectral scaling function

Once the exponential transformation parameters are obtained, the frequency-domain correction described in Section 3.2 is applied.

The scaling function SF is computed as the ratio between the magnitude spectra of the scanned surface and the reconstructed light-based surface:

$$SF = \frac{|FFT_{SCAN}|}{|FFT_{LH}|} \tag{A.2}$$

This scaling function represents the spectral correction required to match the power spectral density (PSD) of the reconstructed surface to that of the experimentally measured one.

The final surface reconstruction is obtained through the iterative procedure combining spectral correction and rank-ordering, ensuring simultaneous consistency with both the target height distribution and the reference PSD.

(a) Original photograph of the surface.

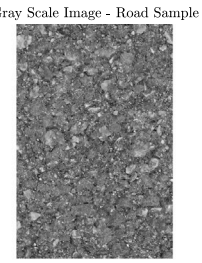

(b) Grayscale intensity field to extract the GPR input from the image.

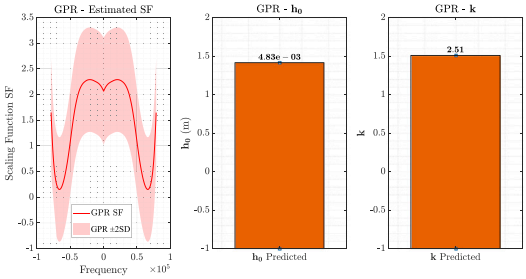

(c) GPR prediction of parameters h_0, k, and scaling function SF from the 2D image input quantities.

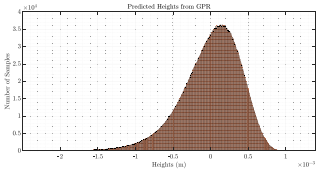

(d) Surface heights estimated using the predicted h_0 and k from GPR.

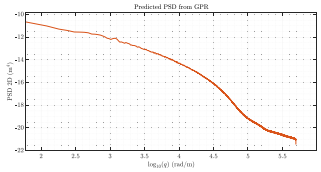

(e) PSD estimated using the predicted quantities from GPR for spectrum correction.

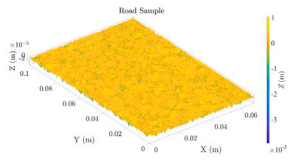

(f) Final reconstructed virtual surface obtained using the proposed methodology.

Fig. A.19. Step-by-step case study showing the complete workflow from image acquisition to final surface reconstruction.

A.5. Training of the Gaussian process regression model

Once the quantities h_0, k, and the scaling function SF are obtained for all surfaces in the dataset, these quantities are used as target variables to train the Gaussian Process Regression model described in Section 3.3.

The GPR model learns the relationship between the image-derived features and the corresponding surface reconstruction parameters. During training, the model estimates the optimal hyperparameters by minimizing the negative log marginal likelihood.

After training, the GPR model is capable of predicting the quantities h_0, k, and the scaling function SF directly from image-based features, without requiring additional roughness measurements.

A.6. Prediction of surface roughness for a new surface

Once the model has been trained, the workflow can be applied to a new surface for which no scanner measurements are available.

The procedure consists of the following steps:

1. Capture a photograph of the surface using the acquisition guidelines described previously.
2. Apply the preprocessing pipeline to obtain the grayscale intensity matrix.
3. Extract the same set of image features used during the training phase.
4. Provide these features as input to the trained GPR model.
5. Obtain the predicted quantities h_0, k, and the scaling function SF.
6. Apply the exponential transformation to convert light intensity into estimated surface heights.
7. Apply the spectral correction and iterative rank-ordering procedure to enforce PSD consistency.

The final result is a reconstructed virtual surface whose statistical and spectral properties closely replicate those of a physically scanned rough surface.

A.7. Practical implications

This workflow enables the generation of physically consistent virtual surfaces using only photographic data. In practical applications, such as motorsport track characterization or road monitoring, this approach can significantly reduce the need for dedicated roughness measurement equipment while still providing surface representations suitable for tribological analysis, contact mechanics simulations, and tire–road interaction modeling.

A summary of the complete step-by-step workflow, illustrated with representative figures for each stage, is presented in Fig. A.19.

Appendix B. Leave-One-Out Cross-Validation (LOOCV)

To further assess the generalization capability of the proposed model and avoid bias from isolated examples, a Leave-One-Out Cross-Validation (LOOCV) procedure was performed across the entire dataset.

For each iteration, one surface was excluded from the training set and used as an unseen test case, while the Gaussian Process Regression (GPR) model was trained on the remaining surfaces. The predicted quantities were then used to reconstruct both the height distribution and the corresponding corrected Power Spectral Density (PSD), which were compared against the experimental extracted quantities.

Fig. B.20 presents a compact summary of the LOOCV results for all available surfaces, showing the comparison between reconstructed and measured distributions. Each subfigure corresponds to a different surface excluded from the training set during the LOOCV procedure. In particular, subfigures (a)–(m) refer to different road surfaces, (n)–(r) correspond to sandpaper samples (grit 40, 80, 120, 500, and 2000), and (s) represents a marble surface.

For completeness, Fig. B.21 also reports the corresponding reconstructed 3D rough surfaces obtained in prediction using the same procedure. These surfaces are generated from the predicted height distributions and PSDs, illustrated in Fig. B.20. The visual comparison highlights the clear distinction between conventional road textures and non-standard surfaces such as sandpapers and marble, confirming the model's ability to capture both statistical and morphological differences.

Overall, the method demonstrates consistent predictive capability across the entire dataset. The LOOCV results provide evidence of the

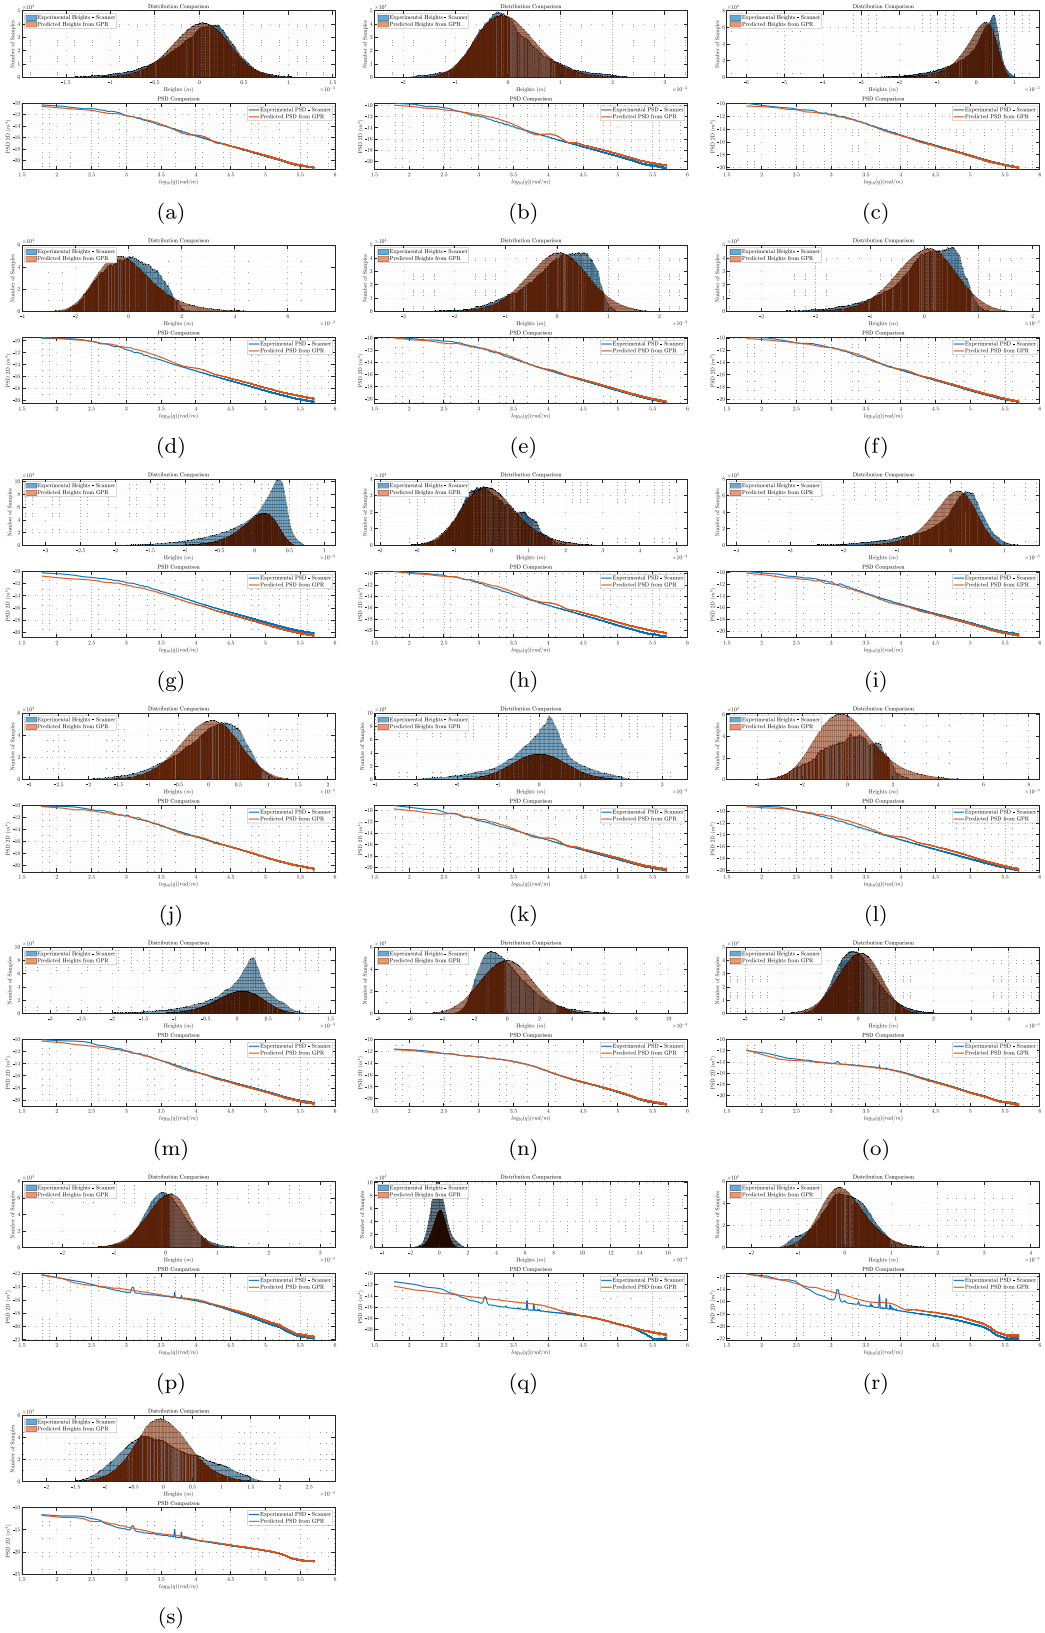

Fig. B.20. Leave-One-Out Cross-Validation results for all surfaces — Height Distributions and PSDs.

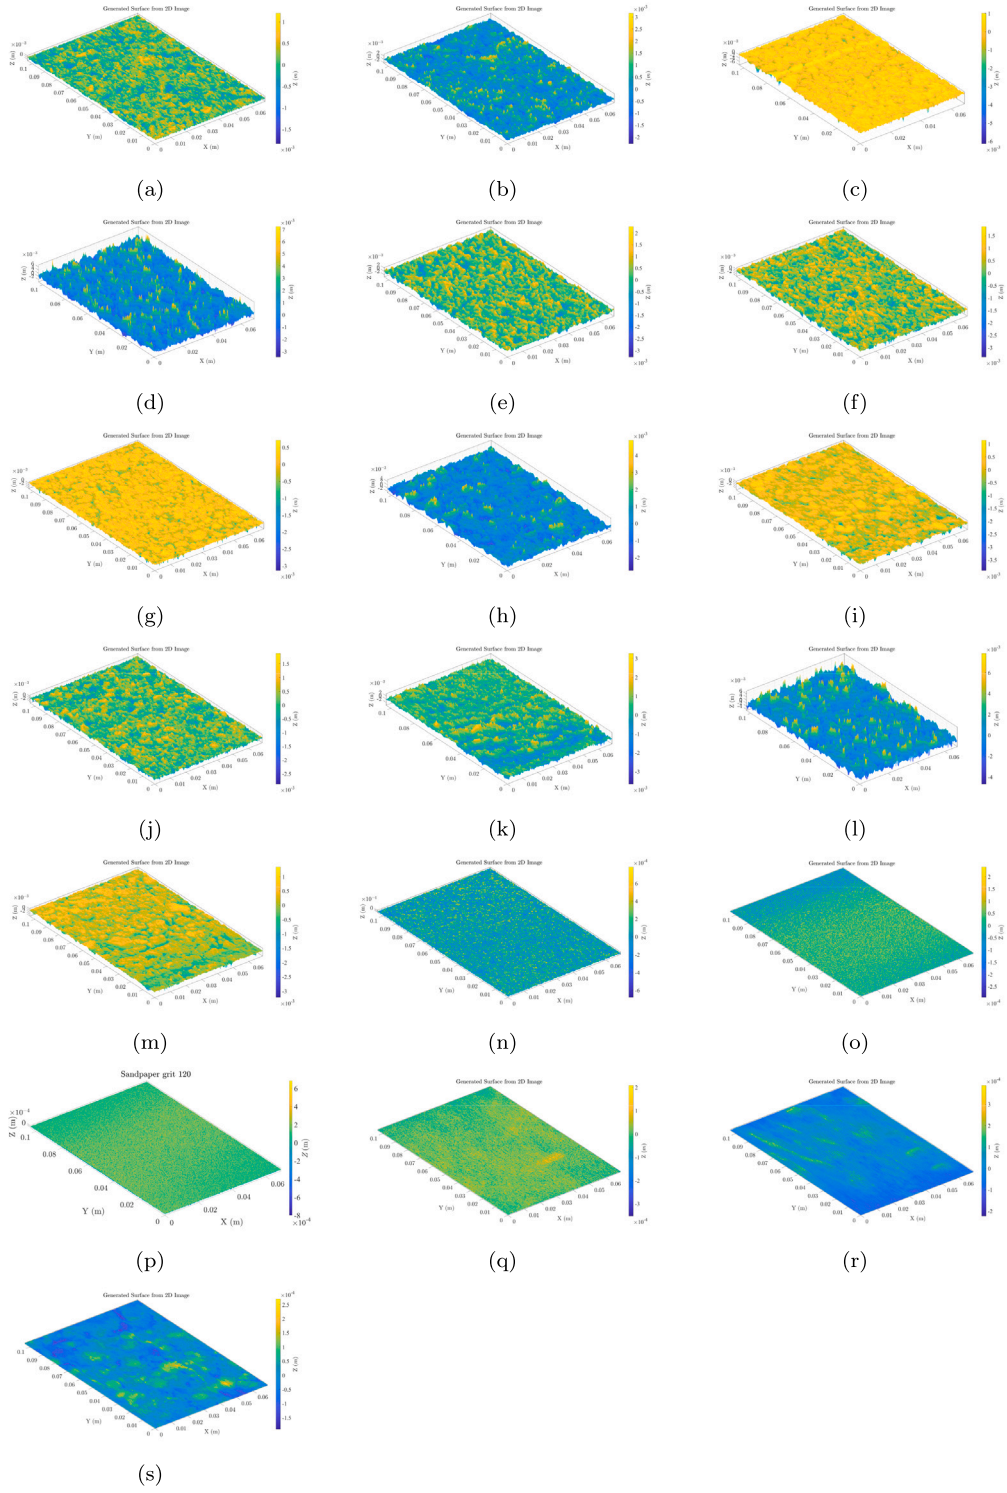

Fig. B.21. Leave-One-Out Cross-Validation results for all surfaces — 3D Surfaces.

robustness and generalization of the proposed framework, even when applied to heterogeneous surface types.

Data availability

The data that has been used is confidential.

References

[1] A. Majumdar, B. Bhushan, Role of fractal geometry in roughness characterization and contact mechanics of surfaces, 1990.
[2] E.S. Gadelmawla, M.M. Koura, T.M. Maksoud, I.M. Elewa, H.H. Soliman, Roughness parameters, J. Mater. Process. Technol. 123 (1) (2002) 133–145.
[3] A. Lang, M. Klüppel, Influences of temperature and load on the dry friction behaviour of tire tread compounds in contact with rough granite, Wear 380 (2017) 15–25.

[4] B. Bhushan, Contact mechanics of rough surfaces in tribology: multiple asperity contact, Tribol. Lett. 4 (1998) 1–35.
[5] K. Grosch, The relation between the friction and visco-elastic properties of rubber, Proc. R. Soc. Lond. Ser. A. Math. Phys. Sci. 274 (1356) (1963) 21–39.
[6] B.N. Persson, O. Albohr, U. Tartaglino, A.I. Volokitin, E. Tosatti, On the nature of surface roughness with application to contact mechanics, sealing, rubberfriction and adhesion, J. Phys.: Condens. Matter. 17 (1) (2004) R1.
[7] H.L. Costa, T. Cousseau, R.M. Souza, Current knowledge on friction, lubrication, and wear of ethanol-fuelled engines—A review, Lubricants 11 (7) (2023) 292.
[8] C. Becker, S. Els, Effect of surface roughness on tyre characteristics, J. Terramech 102 (2022) 27–48.
[9] R. Yong, F. Eiyo, Road surface roughness and tyre performance, J. Terramech 27 (3) (1990) 219–239.
[10] H. Pacejka, Tire and Vehicle Dynamics, Elsevier, 2005.
[11] W.F. Milliken, D.L. Milliken, L.D. Metz, Race car vehicle dynamics, vol. 400, SAE international Warrendale, 1995.
[12] A. Le Gal, M. Klüppel, Investigation and modelling of rubber stationary friction on rough surfaces, J. Phys.: Condens. Matter. 20 (1) (2007) 015007.
[13] V.L. Popov, et al., Contact Mechanics and Friction, Springer, 2010.
[14] H. Ghaednia, X. Wang, S. Saha, Y. Xu, A. Sharma, R.L. Jackson, A review of elastic–plastic contact mechanics, Appl. Mech. Rev. 69 (6) (2017) 060804.
[15] B.N. Persson, Theory of rubber friction and contact mechanics, J. Chem. Phys. 115 (8) (2001) 3840–3861.
[16] M. Kluppel, G. Heinrich, Rubber friction on self-affine road tracks, Rubber Chem. Technol. 73 (4) (2000) 578–606.
[17] K.L. Johnson, Contact Mechanics, Cambridge University Press, 1987.
[18] J.R. Barber, Contact mechanics, vol. 20, Springer, 2018.
[19] A. Jourani, Effect of roughness geometries in contact mechanics, Int. J. Mater. Prod. Technol. 51 (2) (2015) 127–138.
[20] B.N. Persson, Rubber friction: role of the flash temperature, J. Phys.: Condens. Matter. 18 (32) (2006) 7789.
[21] T.D. Jacobs, T. Junge, L. Pastewka, Quantitative characterization of surface topography using spectral analysis, Surf. Topogr.: Metrol. Prop. 5 (1) (2017) 013001.
[22] M.-q. Shao, D. Xu, S.-y. Li, X.-g. Zuo, C.-k. Chen, G.-z. Peng, J.-m. Zhang, X.-c. Wang, Q. Yang, A review of surface roughness measurements based on laser speckle method, J. Iron Steel Res. Int. 30 (10) (2023) 1897–1915.
[23] M. Daneshmand, A. Helmi, E. Avots, F. Noroozi, F. Alisinanoglu, H.S. Arslan, J. Gorbova, R.E. Haamer, C. Ozcinar, G. Anbarjafari, 3D scanning: A comprehensive survey, 2018, arXiv preprint arXiv:1801.08863.
[24] G. Ravindran, Evaluation of new technologies to support asset management of metro systems (Ph.D. thesis), UCL (University College London), 2020.
[25] C. Koch, K. Georgieva, V. Kasireddy, B. Akinci, P. Fieguth, A review on computer vision based defect detection and condition assessment of concrete and asphalt civil infrastructure, Adv. Eng. Inform. 29 (2) (2015) 196–210.
[26] L. Zhou, G. Wu, Y. Zuo, X. Chen, H. Hu, A comprehensive review of vision-based 3D reconstruction methods, Sensors 24 (7) (2024) 2314.
[27] A. Eltner, A. Kaiser, C. Castillo, G. Rock, F. Neugirg, A. Abellán, Image-based surface reconstruction in geomorphometry–merits, limits and developments, Earth Surf. Dyn. 4 (2) (2016) 359–389.
[28] A. Farshian, M. Götz, G. Cavallaro, C. Debus, M. Nießner, J.A. Benediktsson, A. Streit, Deep-learning-based 3-d surface reconstruction—a survey, Proc. IEEE 111 (11) (2023) 1464–1501.
[29] A. Hartley, A review of the geological factors influencing the mechanical properties of road surface aggregates, Q. J. Eng. Geol. 7 (1) (1974) 69–100.
[30] A. Thorpe, R.M. Harrison, Sources and properties of non-exhaust particulate matter from road traffic: a review, Sci. Total Environ. 400 (1–3) (2008) 270–282.
[31] G. Vinogradov, A. Isayev, V. Zolotarev, E. Verebskaya, Rheological properties of road bitumens, Rheol. Acta 16 (1977) 266–281.
[32] A. Bianco, Characterization and Visualization of Reflective Properties of Surfaces (Ph.D. thesis), Philosophisch-Naturwissenschaftliche Fakultät der Universität Basel, 2017.
[33] A.W. Hashmi, H.S. Mali, A. Meena, M.F. Hashmi, N.D. Bokde, Surface characteristics measurement using computer vision: A review., CMES Comput. Model. Eng. Sci. 135 (2) (2023).
[34] F. Pérez-Ràfols, A. Almqvist, Generating randomly rough surfaces with given height probability distribution and power spectrum, Tribol. Int. 131 (2019) 591–604.
[35] E. Schulz, M. Speekenbrink, A. Krause, A tutorial on Gaussian process regression: Modelling, exploring, and exploiting functions, J. Math. Psych. 85 (2018) 1–16.
[36] D. Nguyen-Tuong, M. Seeger, J. Peters, Model learning with local gaussian process regression, Adv. Robot. 23 (15) (2009) 2015–2034.
[37] K. O'shea, R. Nash, An introduction to convolutional neural networks, 2015, arXiv preprint arXiv:1511.08458.
[38] A. Younesi, M. Ansari, M. Fazli, A. Ejlali, M. Shafique, J. Henkel, A comprehensive survey of convolutions in deep learning: Applications, challenges, and future trends, IEEE Access 12 (2024) 41180–41218.
[39] T.S. Sabino, A.C. Carneiro, R.P. Carvalho, F.A. Pires, The impact of non-Gaussian height distributions on the statistics of isotropic random rough surfaces, Tribol. Int. 173 (2022) 107578.
[40] Y. Gong, S.T. Misture, P. Gao, N.P. Mellott, Surface roughness measurements using power spectrum density analysis with enhanced spatial correlation length, J. Phys. Chem. C 120 (39) (2016) 22358–22364.
[41] M.M. Kanafi, A.J. Tuononen, Top topography surface roughness power spectrum for pavement friction evaluation, Tribol. Int. 107 (2017) 240–249.
[42] T. Schreiber, A. Schmitz, Improved surrogate data for nonlinearity tests, Phys. Rev. Lett. 77 (4) (1996) 635.
[43] C. Williams, C. Rasmussen, Gaussian processes for regression, Adv. Neural Inf. Process. Syst. 8 (1995).
[44] C.E. Rasmussen, Gaussian processes in machine learning, in: Summer School on Machine Learning, Springer, 2003, pp. 63–71.
[45] M. Kuss, Gaussian process models for robust regression, classification, and reinforcement learning (Ph.D. thesis), echnische Universität Darmstadt Darmstadt, Germany, 2006.
[46] C.E. Rasmussen, H. Nickisch, Gaussian processes for machine learning (GPML) toolbox, J. Mach. Learn. Res. 11 (2010) 3011–3015.