Shape descriptors and statistical classification on areal topography data for tile inspection in tessellated surfaces

N. Senin[1,2,], M. Moretti[1] and R.K. Leach[2]*

1 Department of Engineering, University of Perugia, Italy
2 Department of Mechanical, Materials and Manufacturing Engineering, University of Nottingham, UK

**corresponding author: nicola.senin@unipg.it*

Abstract
Verification of conformance to design specifications in production, and identification of defects related to wear or other damage during maintenance, are key metrological aspects that must be addressed for micro-scale tessellated surfaces. A new algorithmic approach is presented that operates on topography datasets as obtained by areal topography instruments. The approach combines segmentation algorithms with a novel implementation of the angular radial transform, originally adopted by the MPEG-7 standard, to implement shape descriptors and associated similarity metrics. Applications to the inspection and verification of laser-manufactured micro-embossing topographies are illustrated. The topographies are first segmented to extract the individual tiles; the tiles are then encoded through shape descriptors. Principal component analysis and cluster analysis are used to investigate the behaviour of the angular radial transform coefficients. Finally, an algorithmic classifier based on supervised learning (k-nearest neighbours) is implemented and shown to be effective at identifying defects and at discriminating between defect types.

Keywords: *Surface metrology*, *Tessellated surfaces*, *Areal surface topography*, *Shape descriptors for encoding topography data*

1 Introduction

1.1 Structured and tessellated surfaces

Structured surfaces are surfaces whose micrometric or sub-micrometric texture is characterised by a deterministic pattern of (often high aspect-ratio) features, designed to achieve a specific functional role [1], [2]. The topography of a structured surface is explicitly defined by design specification, as opposed to conventional, unstructured surfaces, where topography is only partially and indirectly defined through compliance to provided texture descriptors (for example, areal surface texture field parameters as defined in ISO 25178-2 [3], [4]). "Tessellated" is a term commonly used to refer to a particular class of structured surface whose topography can be thought of as comprised of a pattern unit, or tile, replicated multiple times in order to create a regular, periodic pattern [1]. Tessellated surfaces have increasingly found successful uses

in a wide array of industrial applications. Typical examples include optical surfaces (retroreflectors, Fresnel lenses, etc.) [5], mechanical surfaces (for example, low-friction patterns [6]), and biocompatible/biomimetic surfaces (prosthetic implants, microfilters, high-adhesion, hydrophilic, etc.) [7]. Designed topographies are constantly evolving along with the high-precision micro- and nanomanufacturing processes needed to generate them. In this scenario, the importance of providing a dedicated and effective solution for metrological inspection and verification is paramount.

1.2 Test case

The test case used in this work consists of micro-embossing patterns obtained by means of different manufacturing processes. Two example surfaces are shown in Figure 1; albeit fabricated from the same nominal specification, they feature significant topography differences, thus obtaining a quantitative assessment of such differences can be used as a means to obtain information on the performance of the two manufacturing processes.
As typically happens for tessellated surfaces [8], the small size of the topographic features that need to be measured and verified (i.e. checked for compliance to specifications, according to ISO 17450-1 [9]), makes areal topography instruments, such as 3D digital profilometers and 3D digital microscopes, ideal candidates for acquiring quantitative information useful for obtaining 3D reconstructions of measured topography [10], [11]. Therefore, the test topographies were measured with an interferometric probe based on conoscopic holography [8], [12] operating in single-point, raster scanning mode. Aside from the differences in the two topographies, each tile is dale-shaped, with an approximately 280 µm × 280 µm square footprint, and a 50 µm nominal depth.

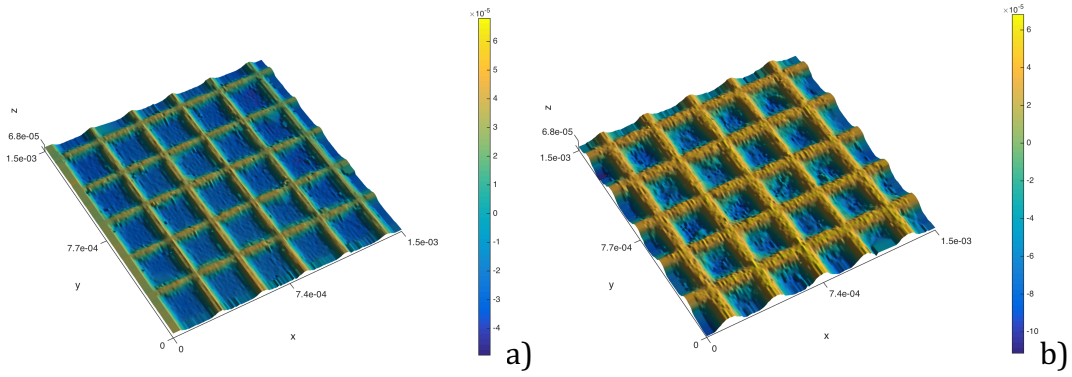

Figure 1. Example reconstructed digital topographies of micro-embossing patterns; a) pattern obtained by micro-milling; b) pattern obtained by pulsed-laser texturing. Both surfaces: field size 1.48 mm × 1.48 mm; 128 pixels × 128 pixels; height-based colouring.

1.3 Metrological characterisation of tessellated surfaces

Typical analysis solutions based on computing surface texture parameters, as provided by current surface metrology literature [13], [14] and standards (in particular ISO 25178-2 [3]), are often not capable of capturing some of the defining properties of tessellated surfaces, such as their degree of regularity/periodicity. Many researchers in surface metrology have thus begun

exploring alternative approaches. For example, some authors have explored the possibility of devising new parameters that better capture the correlation between topography and functional performance (see the work in [15] for an application to anilox printing rollers). Several researchers have been approaching the characterisation problem as a two-step process: firstly the tiles (or the notable topographic features within the tiles) are identified and extracted as separate geometric entities; then dedicated, custom parameters are devised aimed at capturing properties pertaining to the shape of the tiles, and/or their layout over the surface. Notable examples have been shown for hard disk drives [13], abrasives and optical depixelator surfaces [1] and microlens arrays [5], [16]. The same two-step, tile-centred approach has been adopted by the authors of this paper. Two alternative characterisation routes have been investigated. The first route aims at computing dimensional and geometric attributes from measured topography, so that a one-to-one mapping with design specifications is achieved (thus allowing for direct verification akin to common practice in dimensional metrology for standard-sized mechanical parts). Examples of this approach can be found in previous publications [17], [18], and have been specifically applied to tessellated surfaces [19]. The second characterisation route, which is the subject of this paper, is about computing shape descriptors; that is, fast transforms turning topography data into a finite series of numbers useful for shape encoding. Given the high degree of similarity between the mathematical representation of topography data and conventional, digital intensity images (and also equivalently, range images), a large number of techniques, originally developed in computer vision/image processing, can be adapted to generate shape descriptors that efficiently operate on surface topography data [18]. The underlying premise is that the overall size of the tessellated surface is in general large, in comparison to the size of the unit tile; which implies that hundreds, if not thousands of tiles may need to be inspected/verified in an industrial application. Therefore, processing speed becomes a primary issue, and tile characterisation approaches that favour this aspect, sometimes at the expense of a less accurate depiction of tile topography, are given priority.

2 Tile characterisation via the angular radial transform

2.1 The angular radial transform

The angular radial transform (ART) is a moment-based description method adopted by the MPEG-7 standard for shape encoding in video frames [20]–[23]. The ART is defined on a unit disk and based on complex orthogonal sinusoidal basis functions in polar coordinates. The ART coefficients $F_{n,m}$ of order n and m are given by:

$$F_{n,m} = \int_0^{2\pi} \int_0^1 V_{n,m}(\rho,\theta) f(\rho,\theta) \rho d\rho d\theta \quad , \tag{1}$$

where $f(\rho,\theta)$ is the image function in polar coordinates ($\rho \in [0,1]$, $\theta \in [0,2\pi]$) and $V_{n,m}(\rho,\theta)$ is the basis (complex) function separable along the angular and radial directions:

$$V_{n,m}(\rho,\theta) = A_m(\theta)R_n(\rho), \tag{2}$$

where:

$$A_m(\theta) = \frac{1}{2\pi}e^{jm\theta}, \tag{3}$$

$$R_n(\rho) = \begin{cases} 1 & \text{if } n = 0 \\ 2\cos(n\pi\rho) & \text{if } n \neq 0 \end{cases}. \tag{4}$$

2.2 ART as a shape descriptor

The set of coefficients forming the shape descriptor (referred to as the ART descriptor) can be computed from the $F_{n,m}$ coefficients as follows

$$S_{n,m} = \frac{|F_{n,m}|}{|F_{0,0}|}. \tag{5}$$

The $S_{m,n}$ coefficients are real numbers; the modulus of the $F_{n,m}$ complex coefficients is used to obtain rotation invariance, while dividing by $F_{0,0}$ ensures that the final coefficients are approximately independent of the number of pixels used to encode the original shape. In the MPEG-7 standard, three radial functions (n = 0,1,2) and twelve angular functions (m = 0,1,⋯,11) are combined to obtain 36 complex basis functions $V_{n,m}(\rho, \theta)$ [21]; their real parts are illustrated in Figure 2. This choice ensures that a reasonable amount of spatial frequencies are captured by the ART descriptor. Notice that the basis functions in Figure 2 are computed on rectangular (more precisely, square) domains. The purpose of the ART descriptor is to encode the contents of an image (or a rectagular portion of one), therefore, the basis functions are typically computed on discrete regions of matching size and resolution. To compute a basis function onto a rectangular, discrete domain, it is sufficient to solve the integral in equation 1 for the ρ, θ values featured by the pixels comprised within the region (i.e. the portion of a polar space that fits within the rectangular domain of interest).

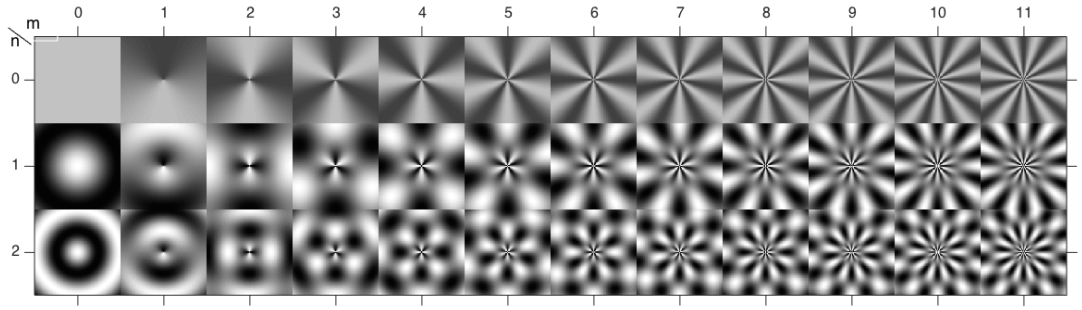

Figure 2. Real parts of 36 basis functions $V_{n,m}(\rho,\theta)$ typically used to compute the ART descriptor (adapted from [23]*).*

2.3 Encoding and comparing tiles through the ART shape descriptor

An example computation of the ART descriptor on an individual tile of one of the test micro-embossing surfaces is shown in Figure 3. The 36-coefficient ART descriptor is conveniently visualised as a 3D bar diagram. Notice that the first descriptor $S_{0,0}$ is always equal to unity due to normalisation and it is, therefore, shape-invariant. It is included only for completeness, but it is not useful for shape differentiation.

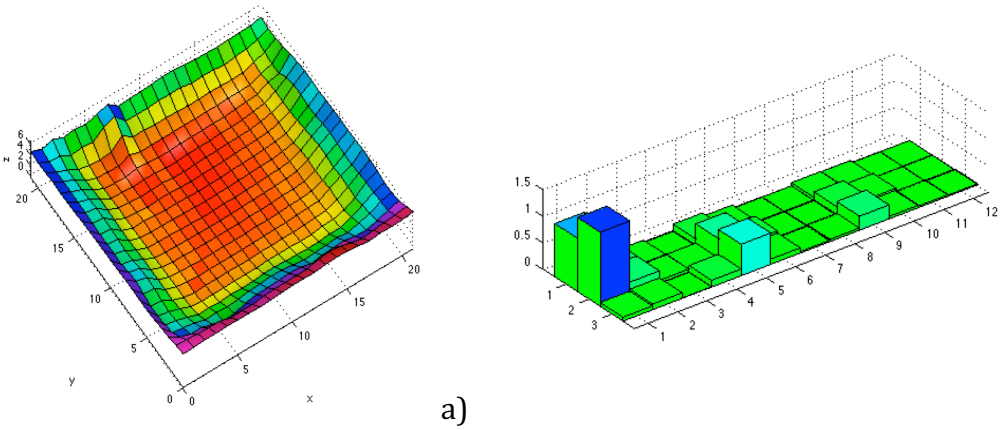

a) b)

Figure 3. Example individual tile and related ART shape descriptor; a) individual tile topography extracted from one of the micro-embossing test surfaces (lateral axes are in pixel units); b) ART shape descriptor with 36 coefficients.

The distance between two images (tile topographies in this case) encoded by the ART descriptor, simply referred to as the ART distance from now on, is typically calculated using the Minkowski formulation:

$$d_{ART}(Q,T) = \left(\sum_{i=0}^{m-1} \sum_{j=0}^{n-1} \left| S_{i,j}(Q) - S_{i,j}(T) \right| \right)^{\frac{1}{p}}, \tag{6}$$

where Q and T are the two images, $S_{i,j}(Q)$ and $S_{i,j}(T)$ are the shape descriptors computed for the two images, and $p = 1$ (Manhattan) and $p = 2$ (Euclidean) are the most popular choices for computing their distance.

2.4 Processing the tessellated surface with the ART shape descriptor

In order to use the ART shape descriptor to encode individual tiles, surface topography needs to be pre-processed so that the individual tiles can be identified and extracted. Several data processing techniques are available for this purpose, many being the subject of active research (for example, see [24]). For the specific test case involving micro-embossing topographies, a simple morphological segmentation into dales, as defined in ISO 25178-2 [3], [25], is perfectly adequate, as the tiles are essentially shaped as dales. The result of segmentation is shown in Figure 4. The identified tiles are individually labelled for further processing through ART encoding.

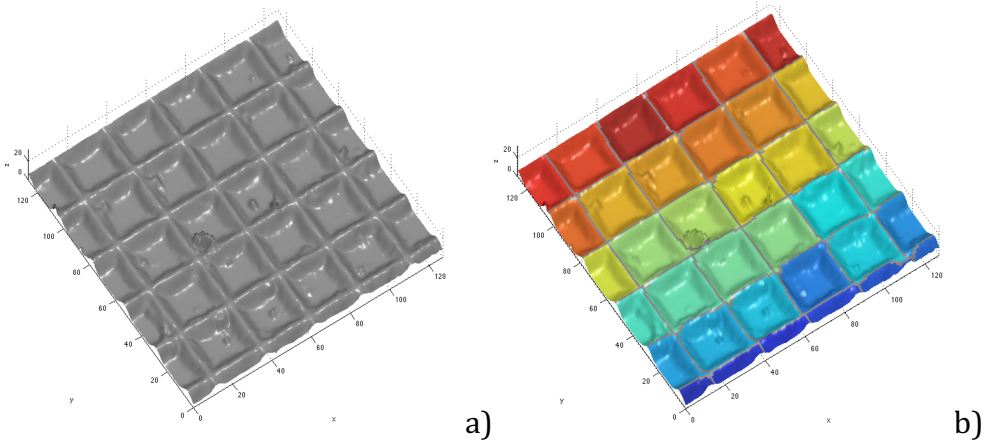

a) b)

Figure 4. Tile identification and extraction for ART encoding in the test case: a) example micro-embossing topography (lateral axes are in pixel units); b) result of morphological segmentation into dales (according to ISO 25178-2 [3]*). Colours indicate different tiles.*

3 Tile inspection and verification with ART distance to nominal

Assuming the tile nominal geometry is available, which is usually the case with a manufactured structured surface, the simplest way to use ART shape descriptors for tile inspection is based on comparing the topography of each manufactured tile with the nominal reference. The ART distance between the two ART encoding results provides a quantitative indication of cumulative shape error between the manufacturing process and the nominal specification. For the example test case, the nominal topography of the tile was available as an STL model. In order to encode it through the ART and to ensure consistency with the real tiles, the STL model must be sampled into a height image with the same number of pixels and resolution as the actual measured tiles (see Figure 5).

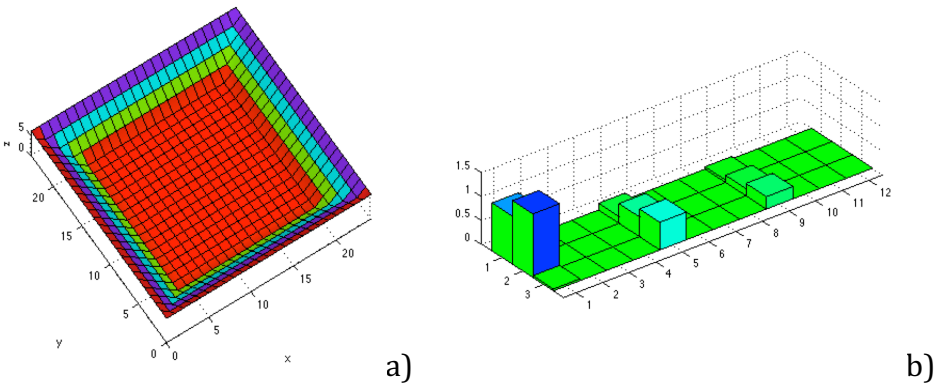

a) b)

Figure 5. Geometric model of the nominal tile and result of ART encoding; a) original geometry sampled into a height image; b) ART shape descriptor.

Processing the tiles extracted from the test surfaces against the nominal template leads to results such as those shown in Figure 6, where the Euclidean distance is used to evaluate the difference between the ART descriptors of each tile and the nominal reference. Visual inspection reveals which tiles feature the largest amount of cumulative error with respect to the nominal template. This method can also be used to study the error distribution over the surface, and to compare the performance of different manufacturing processes.

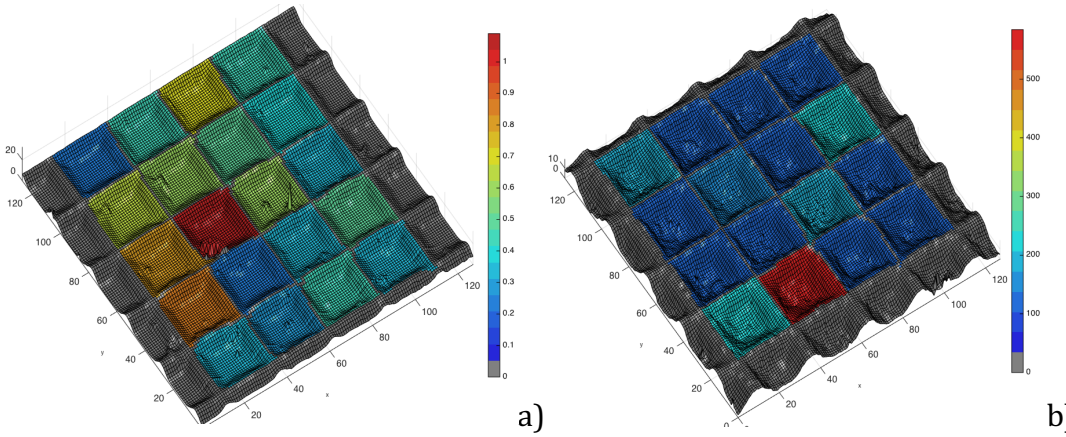

a) b)

Figure 6. Tiles coloured on ART distance from nominal. Euclidean distance in the n × m ART descriptor space was used for this example. Colours shifting towards red indicate a larger amount of cumulative error (i.e. larger departure from nominal); incomplete tiles are discarded (grey) as not containing enough shape information to allow for reliable encoding; a) micro-milled surface; b) pulsed laser textured surface. Lateral axes are in pixel units. Colours are not comparable across images (mapped to different value ranges).

As it is based on simple ART distance to the nominal reference, the inspection of a tessellated surface becomes the process of analysing an individual random variable (the ART distance value). This opens up a large number of techniques already available for univariate statistical process control [26]. An even simpler binary classifier can be set up based on global thresholding. However, an appropriate threshold value on ART distance must be identified, in order to detect any deviation not compliant with the dimensional and/or geometric tolerances imposed by the tile design specifications. An example application of a threshold-based binary classifier is shown in Figure 7.

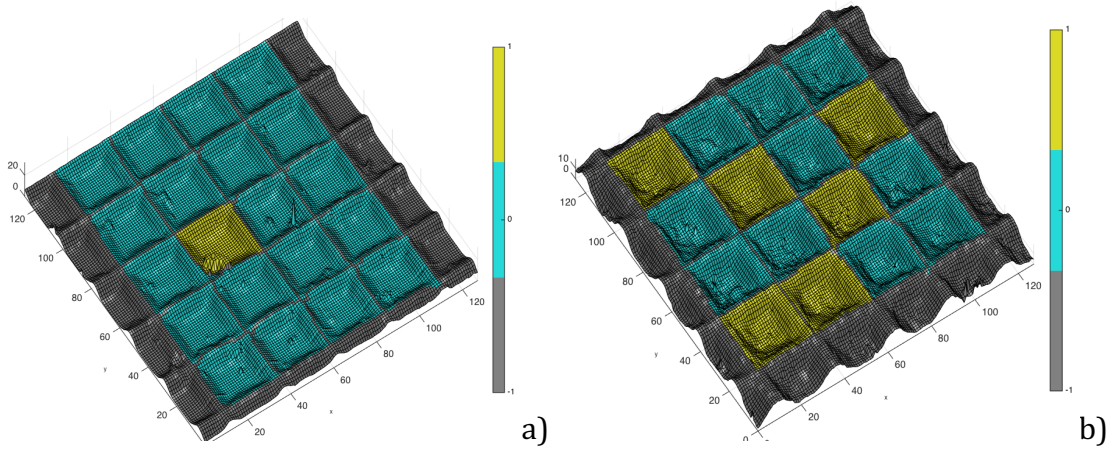

a) b)

Figure 7. Binary classifier based on global thresholding on ART distance from the nominal tile. Tiles whose distance is larger than the threshold (i.e. classified as not compliant to design specifications) are highlighted in yellow; a) micro-milled; b) pulsed laser textured surface. Lateral axes are in pixel units. Thresholds are not comparable across images (different values).

4 Investigation of the discriminating power of the ART shape descriptor

In addition to the difficulty in setting an appropriate threshold value, the main drawback of using statistical process control in univariate space with simple ART distance is that different types of departures from nominal topography (for example, different types of tile defects) may not be properly discriminated if their distances to the template are similar. In order to better understand how the ART shape descriptor really behaves, it is necessary to go back to analysing the descriptor in full $n \times m$ space, and observe how its individual components ($F'_{n,m}$) react to different types of tile topography alterations; that is, the shape-selective sensitivity of the ART coefficients must be determined. This can be accomplished through a multi-step approach: firstly it is necessary to collect a sufficient amount of tile instances, representative of each condition that needs characterising and comparing (for example, the main defect types, as well as the non-defected/acceptable tile conditions), then, the discriminative power of the descriptor can be assessed by preparing different aggregations of tile instances belonging to different classes, and subjecting them to cluster analysis (for example, by k-means clustering) in full ART $n \times m$ space. Further knowledge can be gained by investigating ART descriptor behaviour in a reduced-dimensionality space via principal component analysis (PCA).

4.1 Collection of tile instances

A sufficient coverage of the main types of tile shape alterations must be guaranteed in order to study the discriminating power of the ART descriptor. This implies being able to collect a significant amount of experimental evidence pertaining each class of defects (i.e. multiple instances for each class).
In the specific case of micro-embossing, collecting experimental data is particularly challenging because it implies the availability of multiple printing roller surfaces at different states of damage/wear, and additional others featuring a representative array of manufacturing defects. In absence of appropriate in-line measurement solutions, the need to access these surfaces, potentially removing them from production (or at least temporarily stopping them, in order to acquire physical replicates of local topography), must be factored in as an additional cost. For this research proejct, the number of measured tiles belonging to the available specimens did not guarantee sufficient coverage for all possible types of alterations; therefore, it was decided to resort to simulation. Simulation also guarantees a more ideal coverage of all the possible alteration states the topography may be found in, thus allowing for a more reliable statistical modelling of the problem at stake. In Figure 8, example simulated tile instances are shown, representative of the five principal conditions identified for the micro-embossing test case.

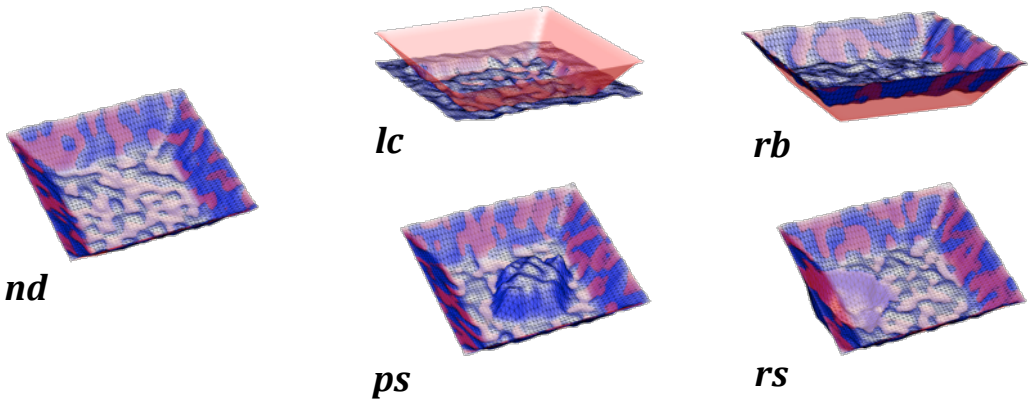

Figure 8. Example simulated tile instances, representative of the main classes identified for the micro-embossing test case. The reference tile nominal geometry is shown in transparent red; nd) non-defective: the amount of topography alterations is within acceptable limits; lc) lowered-crests: the basin ridges are lowered, for example, due to wear; rb) raised-bottom: the bottom of the basin is raised, for example, due to accumulated material; ps) protruded singularity: for example, a leftover particle/debris in the basin; rs) recessed singularity: for example, a pit/cavity which has formed in the basin.

Tile conditions are not classified based on the originating cause (for example, manufacturing error or functional event), instead they are only classified purely in terms of types of shape alteration, which is consistent with using a shape descriptor to discriminate between them.

The simulation model developed to generate the instances relies on a superposition approach, where parametric topography alterations affected by random variations are combined. The starting topography is given by the tile nominal geometry (invariant for all instances), sampled into a discrete function $z = z(x, y)$, defined over a regular *x,y* grid by simulating an ideal range imaging process. Image resolution and pixel width are chosen to match the available measured data acquired from real specimens. Randomly generated instances of topography features, typical of each class of defect, are added, with varying size, shape, localisation and orientation. For each class of topographic features, acceptable random variation is constrained within case-specific, pre-set boundaries. Finally, measurement error is added at the pixel level, defined as $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, i.e. uncorrelated Gaussian noise with zero bias (assuming a calibrated measuring instrument). Noise variance σ_ε^2 is assumed proportional to local slope. Albeit evidently simplified with respect to real-life tile measurement results, the model guarantees the possibility to generate tile instances with varying types and amounts of defects, reproducing both abrupt and gradual transitions between acceptable and defective tile states.

4.2 Investigations with k-means clustering

Clustering experiments on simulated topographies featuring different aggregations of defective and non-defective tiles can be used to investigate the discriminative power of the shape descriptor in full ART $n \times m$ space in the presence of different combinations and amounts of defect types. An example of such a topography is shown in Figure 9a.

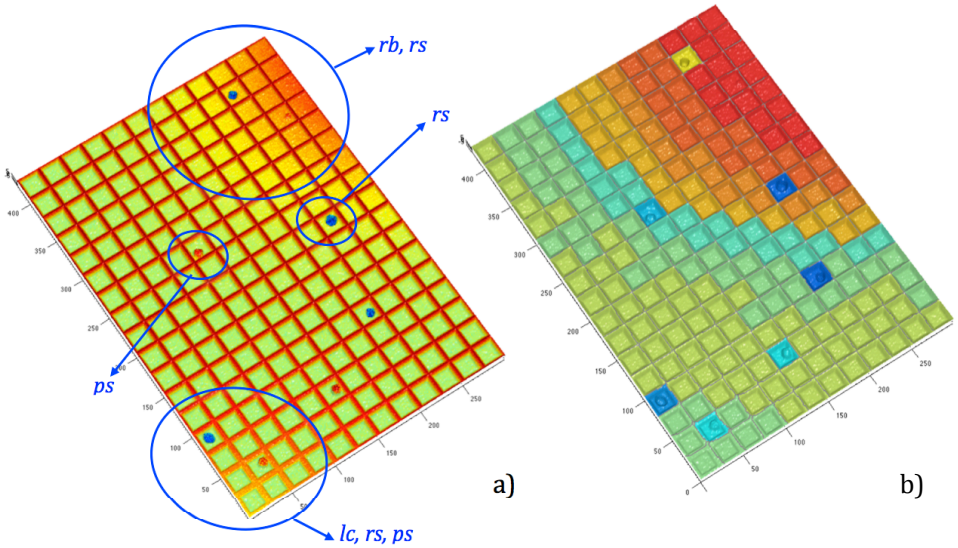

Figure 9. Test for ART descriptor discriminating power; a) simulated topography with defected tiles (some example defects are identified by labels); b) result of k-means clustering on n × m ART descriptor space.

The specific example includes defects that spread over multiple tiles with increasing gravity (for example, lowered crests and raised bottoms – i.e. regions *lc* and *rb* in Figure 9a) and defects localised to individual tiles (for example, protruded and recessed singularities – i.e. tiles *ps* and *rs* in Figure 9a). Some tiles also feature multiple overlapping defects (for example, lowered crests and protruded/recessed singularities). Clustering with the k-means algorithm, Euclidean distance in $n \times m$ ART space and $k = 12$ leads to results such as those in Figure 9b. In Figure 9b, the descriptor is showing a higher discriminative power for some defect types: for example, it is capable of recognising multiple levels for the raised-basin (*rb*) condition (that is, it distinguishes between tiles with different amounts of the defect), but lower discriminative power for others: for example, it discriminates between lowered-crest (*lc*) and non-defective tiles, but does not recognise multiple levels for the *lc* defect, and fails to discriminate between lowered-crest and the less-severe levels of the raised-bottom (*rb*) defect. Protruded and recessed singularities (*ps* and *rs*) are well discriminated from the rest, but again, while the co-presence of a protruded singularity and a raised bottom condition is correctly detected as a new state for the tile, the same does not happen with singularities and lowered-crests, leading to only one defect driving the classification result; the singularity (either *ps* or *rs*) typically prevailing over the lowered-crest. Consistent results can be obtained by running k-means clustering on several different simulated test topographies. Evidently, the ART shape descriptor is more sensitive to some types of shape alterations. This seems to be dependent on the number of pixels involved in the alteration, and also on their localisation in the polar coordinate space, as illustrated in more detail in the discussion section.

4.3 Investigations in principal-component space

An equivalently-important and complementary insight into the behaviour of the shape descriptor can be obtained by visually investigating how the ART

coefficients, computed from a sufficiently large amount of test tile instances, are distributed in principal component space, i.e. in the 2D or 3D Cartesian space identified by the first principal components obtained as the eigenvectors of the covariance matrix, that is, through principal component analysis (PCA). Recent literature findings show that, in fact, PCA and k-means clustering are closely related [27]. In Figure 10, an example PCA result is shown, computed from a test set involving 100 tile instances, with 20 instances per class type (the classes illustrated in Figure 8), including the non-defective class (*nd*). For the specific example, the first three principal components explain approximately 94 % of the total variance.

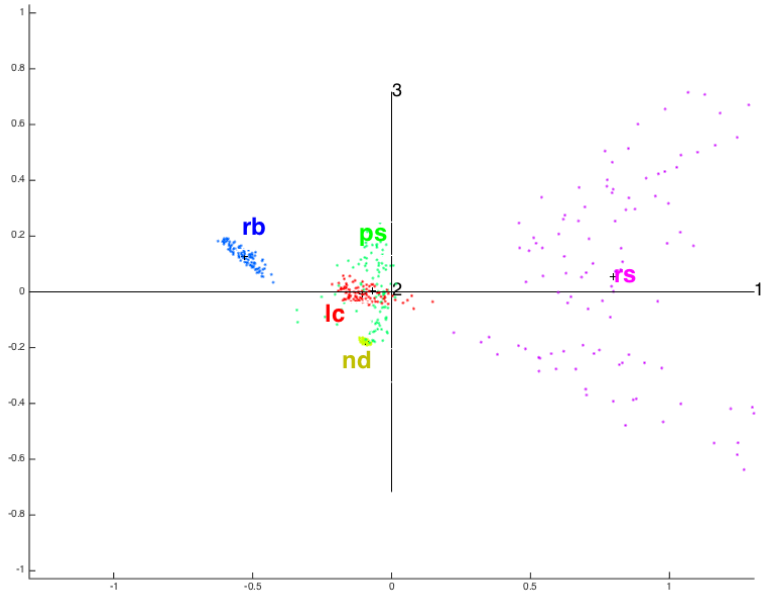

Figure 10. Visualisation of the ART coefficients in principal component space, from a test case of 100 simulated tiles (1 and 3 principal axes shown). In the figure: nd (yellow); lc (red); ps (green), rb (blue) and rs (purple).

Consistent with the results of k-means clustering, the ART descriptor is showing better discriminative power for specific defect types, which in this case is visually conveyed by a smaller within-class scattering of the instances, and larger separation between classes. Of course, the intrinsic nature of the defect types plays a significant role in determining behaviour in principal component space: for a type of defect that exists as a gradual transition from acceptable state to defected, it is natural that its appearance in principal-component space shows a similar degree of continuity. On the other hand, defects that tend to appear abruptly on the surface, i.e. generate a sudden, evident change of tile topography, may correspond to equivalently abrupt jumps in principal-component space, as long as the ART descriptor is sensitive to that defect type. Overall, the results of the investigation in principal-component space confirm that relying on clustering analysis for discriminating between tile conditions may lead to mixed results, and needs to be verified on a case-by-case basis.

5 Development of a classifier based on supervised learning

Regardless of whether the ART shape descriptor is sufficiently capable of discriminating between defect types by means of simple multivariate clustering statistics, an effective performance at identification and discrimination of tile states can still be obtained by devising alternative means of classification, for example, based on supervised learning approaches. Amongst the wide array of supervised learning alternatives available from the literature, for the specific micro-embossing test case, it was found that the k-nearest-neighbour (*knn*) method [28], paired again with Euclidean distance in $n \times m$ ART descriptor space, would provide good results.

5.1 The *knn* classification model

In the learning stage, new observations (i.e. tile instances, either measured or simulated) are fed into the system. A manual classification of each observation is carried out by an expert operator. Notice that multiple classes could be associated to the same tile in case multiple defect types are present. At this stage, the algorithmic classifier simply maintains a growing collection of tile instances, each appropriately flagged as belonging to one of multiple classes, as manually determined by the operator.

At the usage stage, the classifier is fed with an unclassified observation t. Following the *knn* model, the k nearest neighbours to t are initially collected from the database of instances maintained by the classifier and grouped into the temporary set N. Assuming that n_j is the subset of N assigned to class c_j at training, a ranking is computed based on the conditional probabilities $P(c_j|t) = n_j/k$, i.e. the probabilities of t actually belonging to each class c_j depend on how many instances in N were previously classified as c_j. Probabilities are returned in order of decreasing magnitude, making the first class the most-probable match for t. High probability values for multiple classes are an indication that t may feature multiple defects.

5.2 Validation

An assessment of the efficacy of the *knn* classifier can be carried out by means of n-fold cross-validation. For the micro-embossing test case, a dataset of 500 simulated tile instances of known classification, with 100 instances per tile class (i.e. non-defected and four defect types) was generated, and split into five sets of 100 instances, each set featuring internally the same distribution of class types (i.e. 20 instances per class). At each trial run, four such sets were collected into the training set in rotation, and the remaining set was used as the test set, leading to a total of five trial runs. The efficacy of the classifier can be assessed by computing classification accuracy, that is: *number of correct classifications*/*number of test instances* × 100 for each trial run. The results of the experiment are shown in Table 1 as minimum, maximum and mean accuracy over five trial runs.

Table 1. Classification accuracy over five trial runs. Accuracy is shown as computed for each specific class type, and as an aggregated score for all the five class types. Minimum (Min), maximum (Max) and mean values refer to the five trial runs of cross-validation.

	Min	*Max*	*Mean*
nd	100 %	100 %	100 %
lc	100 %	100 %	100 %
rb	100 %	100 %	100 %
ps	85 %	100 %	96 %
rs	100 %	100 %	100 %
total	85 %	100 %	**99.2 %**

Cross-validation shows that the classifier based on supervised learning was highly effective for the test case, with minor discriminating problems for the *ps* (protruded singularity) class; providing an indication that with *knn* classification, good performance can be obtained even with sub-optimal cluster separability in the $n \times m$ ART descriptor space. Moreover, if the system is used as a binary classifier (i.e. lumping all defect types into a generic "defected" class), classification accuracy results at 99 % for a generic defect and 100 % for the non-defected class (percentages refer to true positives in both cases).

6 Discussion

6.1 Computational speed against quality of shape information

In this work, the choice of encoding tile topography through a general-purpose shape descriptor (such as the one based on the ART, presented in this work), against adopting dedicated algorithmic procedures aimed at directly computing specific geometric/dimensional attributes of the tile (for example, depth, footprint area, etc.) is fundamentally driven by computational speed. For example, for detecting a recessed singularity (pit), one could develop a case-specific, dedicated algorithm for identifying pit-like topographic features and for measuring their relevant geometric attributes. While the latter approach may result in more useful feedback for designers and manufacturers [18], [19], because it is likely to provide information in their "language", it involves data processing steps that are typically slower (for example, see [18]) than applying a generalist (i.e. non-dedicated) transform like the ART, specifically designed to be simple and fast to compute. In fact, once the image resolution is fixed, the ART basis functions (equations (2) to (4) in Section 2.1) can be conveniently pre-computed; moreover, the generation of the ART coefficients in discrete space is reduced to a combination of simple sums and multiplications. This speed is appealing for surface inspection applications requiring short execution times, a typical example being inline inspection when the surface is being manufactured. Moreover, a typical tessellated surface is comprised of hundreds, if not thousands of tiles, making inspection speed even more important.

The effectiveness of a classifier based on the ART transform has been illustrated in Section 5.2. However, the investigation on the micro-embossing test cases has revealed that obtaining a shape descriptor through the ART is not as fast as initially thought. While the ART is very fast to compute, it is not invariant to shape translation in the *x,y* plane, because it depends on where the origin is located for the polar coordinate system. Moreover it is also sensitive to translation along *z*, because pixel intensities (i.e. heights) are not normalised. These aspects are typically not an issue when encoding conventional frames of a video stream, because images are generally quantified in the same positive intensity (grey scale) space, and are also typically equalised with respect to each other (that is, *z*-alignment is not necessary). Moreover, for conventional images, alignment in the *x,y* plane is often resolved by placing the origin of the polar coordinate system onto the image centroid, and the centroid is not supposed to change significantly between frames. Things are not so simple for areal topography data. Height (*z*) coordinates may not necessarily be in positive space (in particular, after levelling) and placing the origin of the polar coordinate system on the centroid may not be as reliable, especially considering that localised defects may shift the position of the centroid in significant ways. In this work, it was found that acceptable results from the ART shape descriptor can only be obtained if the topographies to be compared are properly aligned with each other by a rigid rotation and translation transform in 3D space. The most reliable approach is to align them all to the same nominal (template) geometry, as shown in Figure 11.

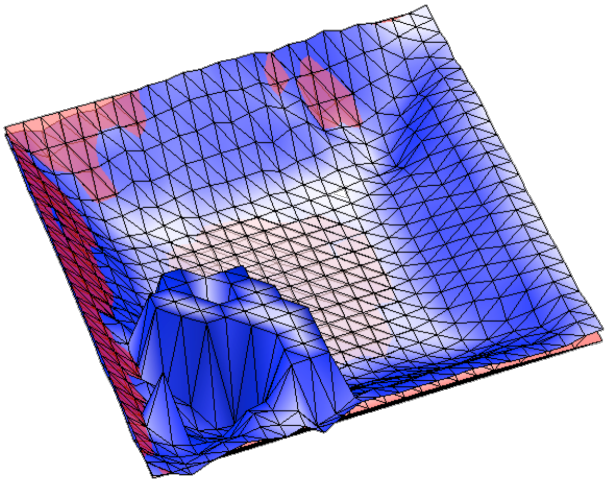

Figure 11. Measured tile registered to the nominal geometry (transparent red). Registration is implemented following previous work [8], [29].

Registration implies a computational effort that in most cases is far from negligible; this increases significantly the overall time needed to compute the ART descriptor onto the individual tiles, thus reducing the theoretical advantage of the descriptor speed. Only the adoption of modified, or entirely new, descriptors, invariant to tile localisation and orientation, may allow for the removal of the registration step, therefore, this is to be considered a research priority to be pursued in this type of approach.

6.2 ART descriptive power and sensitivity

k-means clustering experiments and analyses in principal-component space have clearly shown that not all defect types make the ART shape descriptor react in the same way. Investigations have shown that ART sensitivity is significantly influenced by how many pixels are affected by the shape change in the image (the more pixels, the greater the reaction of the ART coefficients). Also, the localisation of the affected pixels in the image counts, particularly in combination with the frequencies and localisations covered by the selected ART basis functions. In this work, multiple configurations for the basis functions were tested, in addition to the conventional 3 × 12 adopted by the MPEG-7 standard. Albeit no significant differences were observed for the test case, it is believed that higher radial and angular resolutions may help handle shape alterations defined at smaller scales.
Also, the result of segmentation onto a real surface typically leads to tiles with irregular boundaries. When confronted with the ART need to operate on rectangular images (see explanation in section 2.2), either padding (to the smallest image that encloses all the tiles to be compared) or cropping (to the largest image entirely enclosed within all the tiles to be compared) must be adopted. In both cases, this has consequences on the results of ART encoding, thus must be handled with care. For the specific test case, it was found that the effect of padded pixels was significant, thus it was preferred to crop all the tile topographies to the minimum size, which could be safely represented within all the tiles. The disadvantage of this approach is that shape alterations specifically located in proximity of the image boundaries are at risk of being poorly represented.

6.3 Correlation between ART coefficients

Since the ART coefficients are originated by basis functions that capture partially overlapping aspects of the topography, they are not entirely independent of each other and show correlation. This makes it more difficult to investigate ART behaviour in correspondence to shape changes, because typically multiple ART coefficients react to each shape alteration. PCA is only partially able to isolate independent linear combinations of ART coefficients, through principal components, and classification performance is slightly degraded if *knn* operates in principal-component space. If correlation cannot be avoided, alternative distance metrics specifically designed to handle correlated components in multidimensional spaces should be preferentially investigated; the Minkowsky distance (equation (6)) being clearly sub-optimal for this type of situation.

6.4 Classifier issues

Any classifier architecture that relies on manual training runs into the risk of misclassified instances driving down overall classification performance. Aside from prescribing good practice protocols for manual classification, this may become a significant issue in particular in the presence of borderline cases, which are often difficult to classify even for an expert. It is recommended that training focuses on easy classifiable observations, leaving borderline cases to the classifier itself. An additional problem is how to handle the presence of multiple

defects in a tile. The approach adopted in this work is to allow the co-existence of multiple classes that can be associated to the same tile. To handle this, the classifier needs to be able to generate multiple classification outcomes for any observation, which in this work is accomplished by having the classifier return a series of ranked classification results. The challenge in this case is to determine whether additional classification results returned after the first are to be considered as an indication of multiple classes, or should be discarded as not sufficiently representative. An alternative approach is to create a new class for each significant combination of multiple defect types. This approach was found to produce an even greater amount of overlapping in principal-component space for the test case, thus being a potential obstacle to discrimination approaches based on clustering.

6.5 Layout information for tessellated surfaces

Regardless of the specific details involving the shape descriptor, similarity metric and classifier implementation, an intrinsic limitation of the proposed approach is that it is exclusively dedicated to analysis properties pertaining to the individual tiles. In order to completely characterise a tessellated surface, an additional investigation must be carried out on the tile layout, i.e. on how the tiles are spread over the entire surface. This problem has been extensively investigated in the literature [1], [5], [16], [19] and, at the current state of this research, needs to be implemented as an add-on procedure, for example, by reconstructing the pattern lattice through topological connection of the tile centroids.

6.6 Measurement uncertainty

No surface topography characterisation tool is complete without a thorough assessment of its contributions to overall measurement uncertainty. The process of computing the ART descriptors and related similarity metrics is entirely deterministic, thus no additional random error components are added by the encoding and comparison process, unless k-means clustering is used, which typically implies the introduction of a random effect in cluster initialisation. Regardless, the systematic and random error components originally associated to the input image should still be considered, as they propagate through the ART encoding and comparison process, ultimately affecting the final characterisation result. The assessment of the propagation and effects of measurement uncertainty is one of the main subjects of future work, as it is fundamental to further assess the discriminative capabilities of the classification approach, since an increased scatter in the descriptor values, as likely introduced through measurement error, may degrade the performance of the classifier.

6.7 Algorithm configuration and the toolbox approach

The proposed approach is currently implemented as a configurable toolbox, with many editable parameters controlling its behaviour and performance. Some of the most notable parameters include those controlling the encoding resolution for computing the ART coefficients, those controlling the normalisation and registration of topographies; the selection of the similarity metric, the selection of clustering/classification mechanisms and related parameters, and so forth.

While in this work, fundamental research was carried out to understand the effect of the most relevant parameters, it is likely that no unique optimal configuration exists for the toolbox to cope with all real-life application scenarios. In such circumstances, it is envisioned that an algorithm "calibration" stage may be necessary for effective application to each new test case. In parallel, knowledge-based systems are being developed to support the user in scenarios presenting similarities to others encountered in the past.

7 Conclusions and future work

The characterisation of tessellated surfaces presents unique challenges. On one hand, being structured surfaces, tessellated surfaces are not easily specified in terms of conventional surface texture parameters, but instead are suitable to be subjected to dimensional and geometric verification against a nominal topography specification, akin to what happens in verification of standard-sized manufactured components. On the other hand, tessellated surfaces are intrinsically defined as comprised of a typically small-scale unit topography (the tile) replicated hundreds or thousands of times over the surface, ultimately occupying a much larger area; their functional role typically resulting from the concerted interaction of many tiles. Returning to a characterisation approach based on developing synthetic descriptors of the overall topography (i.e. going back to something similar to texture parameters) may indeed be more efficient at describing the surface.

Accuracy and granularity of geometric inspection, targeting individual attributes, clash with the requirements of inspection speed, which could be ideally accomplished by simply computing a texture parameter (for example, the arithmetical mean height *Sa*) over the entire surface. An in-between position may be represented by those methods that adopt partitioning of the topography to identify and separate individual tiles; yet in the end compute synthetic descriptors of topographic properties over the tiles, instead of focusing on individual dimensional and geometric attributes.

The work presented in this paper follows the latter conceptual approach. Tiles are identified and extracted as separate entities and a shape descriptor based on the angular radial transform (ART) is used to encode tile topography. Clustering tests and principal-component analysis are used to investigate the behaviour of the descriptor, and an algorithmic classifier based on supervised learning (*knn* approach) is implemented, showing good capability at recognising defective tiles and at discriminating between defect types. The main drawback of the ART is the need to geometrically register tile topography in order to obtain a meaningful descriptor, which somewhat hampers its main advantage, i.e. being fast to compute. Future work will, therefore, be dedicated to finding ways to improve the descriptor by making it independent of geometric registration needs, while still maintaining the classification and discrimination powers already demonstrated.

Further work will also be devoted to the identification of test cases for which a larger amount of experimental data is available, in order to confirm the current results, primarily obtained through simulation. The availability of a larger amount of experimental data will also be necessary to investigate the

propagation and effects of measurement uncertainty, both through the development of dedicated mathematical models, and through the execution of repeated measurements.

Acknowledgements

N.S. and R.K.L. would like to thank the EC for supporting this work through the grant: FP7-PEOPLE-IEF-METROSURF. R.K.L. would also like to thank the EPSRC (grant: EP/M008983/1). N.S. and M.M. would also like to aknowledge the grant: UM12024L002, POR Umbria FSE 2007-2013, awarded by Regione Umbria.

8 References

[1] X. J. Jiang and D. J. Whitehouse, "Technological shifts in surface metrology," *CIRP Ann. - Manuf. Technol.*, vol. 61, pp. 815–836, 2012.

[2] K. J. Stout and L. Blunt, "A contribution to the debate on surface classifications -- random, systematic, unstructured, structured and engineered," *Int. J. Mach. Tools Manuf.*, vol. 41, no. 13–14, pp. 2039–2044, 2001.

[3] International Organization for Standardization (ISO), "25178 part 2 Geometrical Product Specification (GPS) - surface texture: areal - part 2: terms, definitions and surface texture parameters ." 2012.

[4] F. Blateyron, "The Areal Field Parameters," in *Characterisation of Areal Surface Texture*, R. K. Leach, Ed. Heidelberg: Springer, 2013.

[5] L. B. Kong, C. F. Cheung, X. Q. Jiang, W. B. Lee, S. To, L. Blunt, and P. Scott, "Characterization of surface generation of optical microstructures using a pattern and feature parametric analysis method," *Precis. Eng.*, vol. 34, no. 4, pp. 755–766, 2010.

[6] L. Wang, "Use of structured surfaces for friction and wear control on bearing surfaces," *Surf. Topogr. Metrol. Prop.*, vol. 2, no. 4, p. 43001, 2014.

[7] A. Malshe, K. Rajurkar, A. Samant, H. N. Hansen, S. Bapat, and W. Jiang, "Bio-inspired functional surfaces for advanced applications," *CIRP Ann. - Manuf. Technol.*, vol. 62, no. 2, pp. 607–628, 2013.

[8] N. Senin, S. Pini, and R. Groppetti, "Identification of Microtopographic Surface Features and Form Error Assessment," in *Geometric tolerances: impact on product design, quality inspection and statistical process monitoring*, B. M. Colosimo and N. Senin, Eds. London: Springer, 2010, pp. 159–187.

[9] International Organization for Standardization (ISO), "ISO 17450 - Geometrical product specifications (GPS) -- General concepts -- Part 1: Model for geometrical specification and verification." International Organization for Standardization (ISO), 2011.

[10] R. K. Leach, "Optical measurement of surface topography." Springer, Berlin, 2011.

[11] R. Leach, *Fundamental Principles of Engineering Nanometrology*, 2nd ed. William Andrew, 2014, p. 384.

[12] G. Sirat and F. Paz, "Conoscopic probes are set to transform industrial metrology," *Sens. Rev.*, vol. 18, no. 2, pp. 108–110, 1998.

[13] X. Jiang, P. Scott, D. Whitehouse, and L. Blunt, "Paradigm shifts in surface metrology. Part II. The current shift," *Proc. R. Soc. A*, vol. 463, no. 2085, pp. pp. 2071–2099, 2007.

[14] F. Blateyron, "3D Parameters and New Filtration Techniques," *Proc. Japanese Soc. Prod. Eng.*, pp. 21–27, 2006.

[15] A. Weckenmann and W. Hartmann, "Function-oriented method for the definition and verification of microstructured surfaces," *Precis. Eng.*, vol. 37, no. 3, pp. 684–693, Jul. 2013.

[16] A. Burla, "Fourier descriptors for defect indication in a multiscale and multisensor measurement system," *Opt. Eng.*, vol. 50, no. 4, p. 043603, Apr. 2011.

[17] N. Senin, L. A. Blunt, and M. Tolley, "Dimensional metrology of micro parts by optical three-dimensional profilometry and areal surface topography analysis," *Proc. Inst. Mech. Eng. Part B J. Eng. Manuf.*, vol. 226, no. 11, pp. 1819–1832, 2012.

[18] N. Senin and L. Blunt, "Characterisation of Individual Areal Features," in *Characterisation of Areal Surface Texture*, R. K. Leach, Ed. Heidelberg: Springer, 2013, pp. 179–216.

[19] N. Senin, G. MacAulay, C. Giusca, and R. K. Leach, "On the characterisation of periodic patterns in tessellated surfaces," *Surf. Topogr. Metrol. Prop.*, vol. 2, no. 2, p. 25005, 2014.

[20] W. Y. Kim and Y. S. Kim, "A new region-based shape descriptor," in *Mpeg Meeting, TR 15-01*, 1999.

[21] S. Jeannin and A. Divakaran, "MPEG7 visual motion descriptors," *IEEE Trans. CSVT*, 2001.

[22] M. Bober, "MPEG-7 Visual Shape Descriptors," *IEEE Trans. Cir. Sys. Video Technol.*, vol. 11, no. 6, pp. 716–719, Jun. 2001.

[23] J. Ricard, D. Coerjolly, and A. Baskurt, "Generalizations of angular radial transform for 2D and 3D shape retrieval," *Pattern Recognit. Lett.*, vol. 25, pp. 2174–2186, 2005.

[24] G. D. MacAulay, N. Senin, C. L. Giusca, R. K. Leach, and Atanas Ivanov, "Review of feature boundary identification techniques for the characterisation of tessellated surfaces," *Surf. Topogr. Metrol. Prop.*

[25] F. Blateyron, "The Areal Feature Parameters," in *Characterisation of Areal Surface Texture*, R. K. Leach, Ed. Heidelberg: Springer, 2013, pp. 45–65.

[26] D. C. Montgomery, *Introduction to statistical quality control*, 7th ed. John Wiley & Sons, 2012, p. 768.

[27] C. Ding and X. He, "K -means clustering via principal component analysis," in *Twenty-first international conference on Machine learning - ICML '04*, 2004, p. 29.

[28] G. Shakhnarovic, T. Darrell, and P. Indyk, *Nearest-Neighbor Methods in Learning and Vision*. MIT Press, 2006, p. 280.

[29] N. Senin, B. M. Colosimo, and M. Pacella, "Point set augmentation through fitting for enhanced ICP registration of point clouds in multisensor coordinate metrology," *Robot. Comput. Integr. Manuf.*, vol. 29, pp. 39–52, 2013.