

Texture classification by statistical learning from morphological image processing: application to metallic surfaces

A. CORD*, F. BACH† & D. JEULIN‡

*UniverSud, LIVIC, INRETS–LCPC, Versailles, France

†Laboratoire d'Informatique de l'École Normale Supérieure (CNRS/ENS/INRIA UMR 8548), Paris, France

‡Centre de Morphologie Mathématique (CMM), Mathématiques et Systèmes, Fontainebleau CEDEX, France

Key words. Mathematical morphology, rough surfaces, statistical learning, texture classification.

Summary

A classification method based on textural information for metallic surfaces displaying complex random patterns is proposed. Because these kinds of textures show fluctuations at a small scale and some uniformity at a larger scale, a probabilistic approach is followed, considering textural variations as realizations of random functions. Taking into account information of pixel neighbourhoods, the texture for each pixel is described at different scales. By means of statistical learning, the most relevant textural descriptors are selected for each application. The performance of this approach is established on a real data set of steel surfaces.

Introduction

Texture description in image processing may have different objectives: extraction of characteristics of the scene, defect detection, surface inspection, image classification. For defect inspection on materials displaying complex patterns but appearing visually periodic on a larger scale, automated softwares are very successful by their ability to make fast accurate and repetitive measurements (Bennamoun & Bodnarova, 2003; Murino *et al.*, 2004; Mak *et al.*, 2005). A variety of statistical techniques were investigated, such as histogram-based texture analysis techniques corresponding to the use of co-occurrence matrices (Ivarinen, 2000; Latif-Amet *et al.*, 2000), structural approach (Chen & Jain, 1988), texture modelling (Cohen *et al.*, 1991), filtering approach (Meylani

et al., 1996; Kumar & Pang, 2002) and wavelet transformation of images (Lambert & Bock, 1997; Karras & Mertzios, 2002; Serdaroglu *et al.*, 2006). However, on materials displaying complex patterns that are random in appearance (i.e. not periodic), detecting subtle local defects turns out to be difficult (Xie & Mirmehdi, 2005). This kind of textured surfaces shows fluctuations at a small scale and some uniformity at a larger scale. To handle it, we rely on a probabilistic approach, considering textural variations as realizations of random functions. A theoretical way to characterize them is to evaluate their Choquet capacity, that describes their properties on sets of compact support. Therefore, we describe the texture for each pixel of images accounting for information in its neighbourhood at different scales, and thus follow the texton approach of Malik *et al.* (1999).

For textural description, we applied both linear filtering, relying on curvelet transformations, and nonlinear filtering by morphological transformations of images. To compare their performance, statistical learning is used on documented databases. It is based on linear discriminant analysis and allows us to evaluate the textural descriptors performance for the considered application.

Following a first validation of this approach on synthetical random texture images (Cord *et al.*, 2007), we present here applications on a real data set of scrapping metallic surface images. We propose an approach for detecting small textural anomalies drowned in a textural background.

This paper is organized as follows: in Section 'Materials and methods', the textural descriptors are proposed and the statistical learning tools are introduced, in Section 'Metallic surface application', an application on a real data set is presented, in Section 'Results on metallic surfaces at the image level', the application results are extended to a 800 images database.

Correspondence to: D. Jeulin, Centre de Morphologie Mathématique (CMM), Mathématiques et Systèmes, Mines ParisTech 35 rue Saint Honoré, 77305 Fontainebleau CEDEX, France. e-mail: Dominique.Jeulin@ensmp.fr

Materials and methods

Pixel texture description

The presence of fluctuations at a small scale suggests the use of a probabilistic approach to characterize random textures. Therefore, from a theoretical point of view, a random texture is completely known from its Choquet capacity.

In defect detection studies, image descriptors are generally evaluated on small images extracted from larger ones (Xie & Mirmehdi, 2005). Here, we exploit the theoretical approach to characterize textural properties and calculate them for each pixel, taking into account local properties of its neighbourhood at varying scales.

Theoretical approach of random textures. In this framework, we consider textures as realizations of random functions (RF). Morphological and probabilistic information on RF is obtained by means of lower semi continuous (l.s.c.) test functions g with a compact support K in R^n : it can be shown (Matheron, 1969, 1975; Jeulin, 1992, 2000) that an upper semi-continuous RF $Z(x)$ with real values and defined in R^n is completely characterized by the Choquet capacity functional $T(g)$, defined, for all l.s.c. functions g and for all compact supports K , by

$$T(g) = P\{x \in D_Z(g)\} = 1 - P\{x \in D_Z(g)^c\} \quad (1)$$

where P is the probability of the event $\{\}$ and

$$D_Z(g)^c = \{x, Z(x+y) < g(y), \forall y \in K\}. \quad (2)$$

$D_Z(g)^c$ is the complementary set of $D_Z(g)$ in R^n . Note that in Jeulin (1992) an extension to multivariate RF $Z_i(x)$ (like colour images, or more generally multispectral images) is given.

As particular cases, when the compact set K is a point x and $g(x) = z$ ($g(y) = +\infty$, for $y \neq x$), the cumulative distribution function is obtained from $1 - T(g) = P\{Z(x) < z\}$. When using the two points $\{x, x+h\}$ and the function defined by $g(x) = z_1$ and $g(x+h) = z_2$, we can derive the bivariate distribution $F(z_1, z_2, h) = P\{Z(x) < z_1, Z(x+h) < z_2\}$. More generally, for a function defined on n points, we recover the multivariate spatial distribution of the RF. Using $g(x) = z$ for $x \in K$ and $g(x) = +\infty$ for $x \notin K$, we obtain the distribution function of $Z(x)$ after a change of support according to the supremum over any compact set K . This corresponds to the morphological grey-level dilation of Z by K (Serra, 1982, chapter XII). In addition to dilation and erosion, we use in this study openings and closings applied to grey-level images. The opening, obtained by an erosion followed by a dilation by a given structuring element, preserves bright parts of the image which can contain the structuring element, whereas the closing operation (dilation followed by an erosion) preserves dark parts. Applying a succession of openings (respectively closings) by structuring elements of increasing sizes makes disappear one by one the characteristics of the image.

To address the problem of segmentation of pixels from texture properties and of textural anomaly detection, a classification of pixels must be performed. This requires a local characterization, that can be made in different ways:

1. As described earlier, transform the grey-level image by opening or closing by K_i , and generate a multimodal image from the collection of K_i . It corresponds to filters, such as granulometries, as already proposed by Sivakumar & Goutsias (1997); Aubert *et al.* (2000); Fricout & Jeulin (2004). These morphological operations, which are the basis to the probabilistic characterization of random functions through the Choquet capacity are good candidates as texture descriptors. In this study, different structuring elements are chosen: squares, vertical and horizontal segments. The structuring elements sizes are $[2, 4, 8, 16, 32] \times 2 + 1$ pixels, that are adapted to our application. For a given type of structuring element and the list of sizes, all opening and closing images are evaluated. The descriptor is obtained by calculating the difference between the open images at sizes n and $n+1$ as well as between the closed images at sizes $n+1$ and n . Therefore, each pixel is described by a vector with 30 morphological components, where $30 = \text{number of sizes (5)} \times \text{number of structuring elements (3)} \times \text{number of operations (opening/closing)}$.
2. Consider a neighbourhood $B(x)$ of each pixel, and use a local estimate of $T(g)$ inside $B(x)$. From the estimates, generate a multimodal image from the collection of g_i . It uses a local estimate in B of the Choquet capacity and requires the appropriate choice of B . This approach developed in Cord *et al.* (2007), will not be used here.
3. An alternative approach to the Choquet capacity and to the dilations/erosions, is to use measures μ_i with a compact support K_i and estimate $\mu_i(Z)$, generating a multimodal image from the collection of K_i . This approach is less general than the morphological approach, because it is limited to linear operators. Particular cases are given by various types of linear filters, like multiscale convolution by Gaussian kernels, wavelets, curvelets, ... The application of this approach is presented in Section 'Linear filtering using curvelets'.

Linear filtering using curvelets. The curvelet transform is a higher-dimensional generalization of the wavelet transform, designed to represent images at different scales and different angles (Candes & Donoho, 1999). The use of this tool to characterize the texture in an image is recent (Elad *et al.*, 2005). Curvelets have very interesting properties in the context of object detection, in particular curved singularities can be well approximated with very few coefficients. This makes the curvelet coefficients for pixels belonging to a particular object very specific.

The curvelet filter bank is in essence a set of bandpass filters with range and orientation selective properties. Typically, we

apply a linear filtering of each 100×100 neighbourhood of every pixel by curvelets with different frequencies and orientations. The filter bank is decomposed into four sets of frequencies containing, respectively, 1, 8, 16 and 1 filters of varied orientations from the smallest frequency to the largest one. The features used are the magnitudes of each of the 26 filters results.

Synthesis of pixel texture description. To minimize border effects on descriptor calculation, images were extended by neutral values for the corresponding feature (max for the erosion, min for the dilation and mean for the linear filtering).

Combining linear and nonlinear filtering for textural properties description provides us with a three-dimensional data set having two spatial dimensions and a descriptor dimension. Each pixel is then described by a vector. This allows us to follow the texton approach of Malik *et al.* (1999).

The calculation is evaluated using an image sizing 256×256 on a PC intel CPU 3.4 GHz with 1 GB of RAM running Matlab 7.1. The morphological descriptors having 30 dimensions are calculated in 0.58 s. The 26 curvelets are calculated in 2.4 s. The optimization of the calculation time is not an issue in this paper, and we focus on the following on prediction performance.

Statistical learning

From a practical point of view, using pixel texture descriptions provides us with a large amount of descriptors for each pixel of each image in the database. We deliberately use a large number of descriptors, that are automatically selected for each application by statistical learning. In this way, the approach is generic and the pattern and orientations that are chosen are problem dependent. This approach was successfully applied to the classification of images in the standard case when there is a single texture per field of view (Aubert *et al.*, 2000; Fricout & Jeulin, 2004).

The statistical learning approach described here could be decomposed in three steps.

1. A dimensional reduction of the set of independent variables by principal component analysis.
2. A supervised learning using linear discriminant analysis.
3. A variable selection based on forward selection.

This approach will help us to know which set of descriptors can best determine the pixel classes and what classification rule should be applied to best separate the groups of pixels.

Principal component analysis (PCA) is a vector space transform often used to reduce multidimensional data sets to lower dimensions. PCA is defined as an orthogonal linear transformation that maximize the variance of the projected data. In our applications, PCA is used for dimensionality reduction in a data set by retaining those characteristics of the data set that contribute most to its variance, by keeping lower-order principal components and ignoring higher-order ones.

A fixed percentage of the variance, typically 99% is conserved to minimize the possible correlation existing between textural descriptors.

Linear discriminant analysis (LDA) The purpose of linear discriminant analysis is to classify objects (here, pixels) into one of two or more groups based on a set of features that describe the objects (here, textural properties). It assumes that the groups can be separated by a linear combination of features. It is a supervised learning method, meaning that some object, that has an identified label, should be extracted from the database. More details may be found in Hastie *et al.* (2001). Other methods, as SVM, could be used to produce this classification, but at the expense of a higher computational cost.

The applied method is a combination of a PCA followed by an LDA on selected axis (that represent 99% of the global variance). It produces new variables, as linear combinations of the initial features that maximizes the ratio of between-class variance to the within-class variance.

Cross-validation is the statistical practice of partitioning data (pixels in the present case) into subsets such that the analysis is initially performed on all except one subsets, whereas the last subset is retained for subsequent use in validating the initial analysis. In this work, the original data is partitioned into five subsets. Of the five subsets, a single subset is retained as validation data for testing the model, and the remaining four subsets are used as training data. The cross-validation process is then repeated five times, with each of the five subsets used exactly once as validation data. We calculate both the mean of the five results to produce an estimation of the learning error, and the standard deviation to evaluate the expected precision of this error.

Variable selection is made using the classical approach named forward selection (Guyon & Elisseeff, 2003). It is a method to find an efficient combination of variables by starting with a single variable, and progressively incorporating variables into larger and larger subsets. The method starts by first selecting the variable which results in the lowest classification error. Next, this variable is used to test all combinations with the remaining variables to find the best pair of variables. In all further steps, additional variables are added until all variables are used up. It is then possible to select the best subset of variables corresponding to the step where the error reaches a minimum. Note that the forward selection does not necessarily find the best combination of variables (out of all possible combinations). However, it will result in a combination which comes close to the optimum solution.

Metallic surface application

In automatic surface inspection, the conventional method of detection/classification relies on two steps. First, fast real-time algorithms pull out 'objects' of interest. Then powerful

classification calculations adequately reject pseudo-defects and sort the pertinent defects in appropriate families. However, on materials displaying complex patterns that are random in appearance, detecting subtle local defects or describing textural variation turns out to be a rather difficult task.

The main idea of this application is to detect defects on this kind of materials. For this purpose, we calculate for each pixel of the image if its belonging to a defect or not (Section 'Metallic surface application'). Then we implicitly gather the defect pixels to reinforce the detection (Section 'Results on metallic surfaces at the image level').

Then, we propose to describe pixels textures using the advanced descriptors presented earlier and demonstrate their potential in the evaluation of light defect detection in strongly textured surfaces.

The study was made possible thanks to the 'digital coil recording' feature equipping VAI MT SAS systems, permitting the uncompressed recording of the video stream of an entire coil as a full image, when desired. Large amounts of images were thus processed with this aim.

Training database extraction. For this application, we use a database containing 800 images of scraping steel surface factory, provided by VAI MT SAS. Those images typically have a size of 240×296 pixels, with a 1 mm scale resolution, and all of them contain different defects on a nonuniform background showing fluctuations of texture. Our objective is to distinguish standard steel surface, versus defects.

Among the 800 images, we randomly pick 50 of them for the training of the system. The rest will be used to validate the results at the very end. On these 50 images,

1. all the defect pixels are manually localized, and
2. we measure all the 56 texture descriptors that are described in Sections 'Theoretical approach of random textures' and 'Linear filtering using curvelets' (30 morphological granulometries and 26 curvelets linear filtering).

The number of images for the training is sufficient because the objects that are classified during the learning procedure are pixels. Many of them can be extracted from these 50 images. Indeed, we randomly select 100 000 pixels belonging to the defect class and 100 000 pixels outside, corresponding to the background. All the descriptor vectors of those pixels are placed in a database, that is used in the following learning procedure to evaluate and select variables.

Classification using all the descriptors. The first step of the analysis is meant to evaluate the efficiency of both the descriptors and the learning procedure in the task of classifying pixels between the two classes 'defect' versus 'background'. We run the learning procedure combining PCA and LDA with a cross-validation step as described in Section 'Statistical learning'. The use of the cross-validation is essential to evaluate the performance of the learning procedure. Without it, our performance would be too optimistic because of potential overfitting. The histogram of the pixels descriptors




Fig. 1. Projection of the test data on the first linear discriminant axis for one subset of the cross-validation. The x-axis is the coordinate of the projection on the first linear discriminant axis. The y-axis is the proportion of pixels. The background data are plotted in black. The defect data are plotted in white.

projection on the first discriminant axis for one subset of the cross-validation is presented in Fig. 1. All other subsets are very similar.

In this figure, the separation between the two classes looks good. It is confirmed by the evaluation of the following statistical criteria for the test data:

1. The accuracy (number of pixels correctly classified/total number of pixels) is 94.4% (with a standard deviation of 0.2%).
2. The sensitivity (percentage of defect pixels classified as defect) is 96.3% (0.3%).
3. The specificity (percentage of normal pixels classified as normal) is 92.5% (0.3%).
4. The positive predictive value (percentage of classified as defect pixels who are correctly classified) is 92.8% (0.2%).
5. The negative predictive value (percentage of classified as normal pixels who are correctly classified) is 96.2% (0.2%).

There was no significant difference between the performances on the training and the test data. This indicates that the learning procedure does not overfit the data, and that the number of selected pixel in the database (100 000 per class) is sufficient compared to the number of descriptors (56). It thus confirms that the choice of extracting only 50 images for the 800 available images is relevant.

The accuracy of 94.4% shows that the descriptors we propose are valid for the classification task. The small classification error may be explained by the fact that our pixel database is not perfect. Indeed, pixels located near a defect border may not be easy to classify.

Relying on these results, we consider the possibility of obtaining such a score using only a subset of variables.




Fig. 2. Classification test error versus number of selected variables. The error bars correspond to error standard deviation.

Variable selection. Using forward selection of variables as detailed in Section ‘Statistical learning’, we calculate the mean test error and its standard deviation as a function of the number of remaining variables in Fig. 2.

It shows that the error decreases as the number of variables increases and reaches a minimum for 20 selected variables. The standard deviation keeps a very low level whatever the number of selected variables. It assures us of the stability

of the learning procedure. Keeping the selected subset of 20 descriptors leads to a classifier as efficient as having the whole set of descriptors.

Those 20 variables, sorted as function of their classification power, are presented in Table 1 with the associated classification error and standard deviation.

From this Table, it appears that curvelet linear filters are the best descriptors for the classification task in the present application. Indeed, the curvelets are the three first selected variables and they represent 13 of the 20 selected variables. We found both low- and high-frequency filters. It confirms the relevance of those descriptors for the presented application, in particular the choice of filter frequencies.

However, the curvelets are not sufficient by themselves. Indeed, a quick test shows that keeping only these descriptors leads to an error of 11.4%. It appears that the morphological granulometry cooperates well with the curvelets. The structuring elements corresponding to the selected descriptor presents a good representation of the available ones, both in terms of forms (Horizontal lines, vertical lines and square) and sizes (65, 33, 17, 9 and 5 pixels).

In our approach, the relevant descriptors are selected for the current application by statistical learning. In this way, the method is generic and the pattern and orientations of the filters that are chosen are problem dependent.

Threshold setting. On the selected subset of descriptors, we calculate the projection matrix that corresponds to a linear combination of the 20 selected descriptors. To find a mapping

Table 1. Classification test error with respect to the selected variables (cf. Fig. 2).

	Descriptor	Error (%)	SD ($\times 0.01$)
1	Curvelet low-frequency angle from vertical of 112.5°	24.9	0.316
2	Curvelet high-frequency angle from vertical of 78.75°	20.64	0.151
3	Curvelet low-frequency angle from vertical of 135°	17.59	0.054
4	Opening with horizontal line of size 65	14.6	0.105
5	Closing with square of size 33	12.99	0.192
6	Curvelet high-frequency angle from vertical of 90°	11.65	0.175
7	Curvelet high-frequency angle from vertical of 123.75°	10.42	0.153
8	Closing with square of size 65	9.53	0.211
9	Curvelet low-frequency angle from vertical of 45°	8.69	0.165
10	Curvelet high-frequency angle from vertical of 0°	7.83	0.094
11	Closing with square of size 17	7.47	0.148
12	Opening with square of size 65	7.10	0.109
13	Curvelet low-frequency angle from vertical of 22.5°	6.71	0.104
14	Curvelet high-frequency angle from vertical of 157.5°	6.53	0.097
15	Curvelet high-frequency angle from vertical of 67.5°	6.33	0.245
16	Curvelet high-frequency angle from vertical of 112.5°	6.19	0.117
17	Curvelet high-frequency angle from vertical of 135°	6.04	0.099
18	Opening with vertical line of size 5	5.89	0.098
19	Closing with square of size 9	5.79	0.103
20	Curvelet low-frequency angle from vertical of 90°	5.66	0.095




Fig. 3. ROC curve for the metallic defect detection. The black point corresponds to the selected threshold.

of defects on a global image, we project every pixels on the first discriminant axis. We need to find a threshold over which a pixel is considered as a defect.

The Receiver Operating Characteristics (ROC) curve (Egan, 1975), plotted in Fig. 3, is produced on the 200 000 pixels from the database by varying the threshold corresponding to the limit between background and defect.

In our application, we need to minimize the false detection rate. Indeed, we only need to find some pixels in a global defect to produce an alert. Then we seek on the ROC curve the threshold that corresponds to a false detection rate of 1%. It leads to a high missed detection rate of 35%, in terms of pixels. However, as it is presented in the following, it mainly corresponds to pixels located on defects borders.

Results on metallic surfaces at the image level

In this section, we translate the detection from the pixel scale to the image scale, by implicitly gathering detected defect pixels. Because a defect is composed of numerous pixels, using this information enhances the statistical results. The method was evaluated on available images. The process of any incoming image is the following:

1. Calculation of the 20 selected descriptors.
2. Projection on the first discriminant axis.
3. Application of the threshold, determined using the ROC curve.

We applied this process on the 800 available images. Some results are presented in Fig. 4. The two top rows illustrate the variability of defects that can be detected by our approach. On the third row, the thinnest part of the defect is missed, but there is still a detection on its larger part. On the last row, the defect located at the center of the image is too thin to be detected.

All images contain both defect and background. The classified images were all manually scanned to evaluate the following results. First, with the selected threshold, there is not a single false detection (meaning a detection occurring outside the defects in an image). It corresponds to a specificity and a positive predictive value of 100%. Only 20 images present a partially missed defect, as shown in the third row in Fig. 4. Four images present a totally missed defect, as shown in the last row in Fig. 4. It then corresponds to a sensitivity of 99.5%, a negative predictive value of 97% and an accuracy of 98.5%. It shows that defects are well located by the approach, whatever their shape and structure. A complementary description of this shape and structure could be used to determine the defect category.

As a conclusion on this example, the pixel texture description allows us to characterize pixel neighbourhood properties and to distinguish defects and a standard textured background. We could extend this study by testing some other texture descriptors, or applying some post-processing on detected area to improve the results.

Discussion and conclusion

The methodology presented in this paper is slightly different from typical classification tasks on images. The originality is to describe all pixels of each image by some local descriptors and to use the learning process at the pixel scale. In this way, we have access to an quasi-unlimited number of samples to train and test the statistical model. The gathering of the classified pixel enhance the detection at the final step.

Morphological image transformations and curvelets provide efficient descriptors for pixels classification according to the local texture. In the tested application, they allow to point out the textural variations existing in images, and to discriminate between specific characteristics. We could add some other descriptors, as wedge or wavelet for instance, to generalize the approach and may be to improve the results. However, the complementarity of linear (curvelets) and nonlinear (morphological) filtering presented here leads to a satisfactory performance on the considered example.

We deliberately use a very large number of descriptors, which could be selected for each application by statistical learning. In this way, the approach is generic and the pattern and orientations that are chosen are problem dependent. Indeed, the pertinence of a subset of image transformations for texture classification depends on the application, and can be ascertained from statistical learning techniques. Then, the method could be applied to a large set of problems relying on pattern recognition in random textured images.

Other learning approaches could be used in the same kind of study, as SVM, LARS/Lasso for instance. We deliberately choose to use the combination of simple tools as PCA and LDA because they are easy to handle, fast to compute and produce satisfactory results for our application.




Fig. 4. Examples of defect detection. On the left, the original image. In the middle, the projection on the first LDA axis. On the right, the detection after applying the threshold. From top to bottom, two good detections, a partially missed defect and a totally missed defect.

Acknowledgements

The authors are grateful to Siemens VAI MT SAS (France) for the steel plate images. This study was supported by Siemens VAI MT SAS (France). The authors are very grateful to the reviewers for their valuable reviews and careful reading of earlier versions, which helped improving this paper.

References

- Aubert, A., Jeulin, D. & Hashimoto, R. (2000) Surface texture classification from morphological transformations. In *Proceedings of the ISMM'2000, Mathematical Morphology and its Applications to Image and Signal Processing* (ed. by I.J. Goutsias, L. Vincent & D.S. Bloomberg), pp. 253–252. Kluwer Academic Publishers, Palo Alto, USA.
- Bennamoun, M. & Bodnarova, A. (2003) Digital image processing techniques for automatic textile quality control. *Syst. Anal. Model. Simul.* **43**(11), 1581–1614.
- Candes, E.J. & Donoho, D.L. (1999) Curvelets – a surprisingly effective non-adaptive representation for objects with edges. In *Curve and Surface Fitting*, vol. P879110 (ed. by A. Cohen, C. Rabut & L.L. Schumaker), pp. 1–16. Vanderbilt University Press, Nashville, TN, Saint-Malo Proceedings.
- Chen, J. & Jain, K. (1988) A structural approach to identify defects in textured images. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, vol. 1, pp. 29–32. IEEE, Beijing.
- Cohen, F., Fan, Z. & Attali, S. (1991) Automated inspection of textile fabrics using textural models. *IEEE TPAMI* **13**(8), 803–808.
- Cord, A., Jeulin, D. & Bach, F. (2007) Segmentation of random textures by morphological and linear operators. In *Mathematical Morphology and its Applications to Signal Processing* (ed. by G. Bannon, e.a.e.), pp. 397–398.
- Egan, J.P. (1975) *Signal Detection Theory and ROC Analysis*. Series in Cognition and Perception, Academic Press, New York.
- Elad, M., Starck, J., Donoho, D. & Querre, P. (2005) Simultaneous cartoon and texture image inpainting using morphological component analysis (mca). *ACHA* **19**, 340–358.
- Fricout, G. & Jeulin, D. (2004) Analisis de imagen y morfologia matematica para la caracterizacion en linea del aspecto de las superficies pintadas. *Pinturas y Acabados XLVI*(292), 6–13.
- Guyon, I. & Elisseeff, A. (2003) An introduction to variable and feature selection. *J. Mach. Learn. Res.* **3**, 1157–1182.
- Hastie, T., Tibshirani, R. & Friedman, J. (2001) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, Springer.
- Iivarinen, J. (2000) Surface defect detection with histogram-based texture features. In *Intelligent Robots and Computer Vision XIX: Algorithms, Techniques, and Active Vision*, vol. 4197, pp. 140–145. SPIE, Boston, MA, USA.
- Jeulin, D. (1992) Multivariate random image models. *Acta Stereologica* **11**, 59–66.
- Jeulin, D. (2000) Random texture models for materials structures. *Stat. Comp.* **10**, 121–131.
- Karras, D.A. & Mertzios, B.G. (2002) Improved defect detection using novel wavelet feature extraction involving principal component analysis and neural network techniques. In *Australian Joint Conference on Artificial Intelligence*, pp. 638–647. Springer-Verlag, London, UK.
- Kumar, A. & Pang, G. (2002) Defect detection in textured materials using Gabor filters. *Indus. Appl.* **38**(2), 425–440.
- Lambert, G. & Bock, F. (1997) Wavelet methods for texture defect detection. In *Proceedings of the 1997 International Conference on Image Processing*, vol. 3, pp. 201–204. IEEE Computer Society, Washington, DC, USA.
- Latif-Amet, A., Ertüzün, A. & Erçil, A. (2000) An efficient method for texture defect detection: sub-band domain co-occurrence matrices. *Im. Vis. Comp.* **18**(6), 543–553.
- Mak, K., Peng, P. & Lau, H. (2005) Optimal morphological filter design for fabric defect detection. In *Indus. Technol., ICIT*, pp. 799–804.
- Malik, J., Belongie, S., Shi, J. & Leung, T. (1999) Textons, contours and regions: cue integration in image segmentation. In *Proceedings of the Seventh International Conference on Computer Vision (ICCV'99)*, vol. 2, p. 918.
- Matheron, G. (1969) *Théorie des Ensembles Aléatoires*. Ecole des Mines de Paris.
- Matheron, G. (1975) *Random Sets and Integral Geometry*. Wiley, New York.
- Meylani, R., Ertüzün, A. & Erçil, A. (1996) Texture defect detection using the adaptive two-dimensional latticefilter. In *Proceedings of International Conference on Image Processing*, vol. 3, pp. 165–168.
- Murino, V., Bicego, M. & Rossi, I. (2004) Statistical classification of raw textile defects. In *ICPR*, vol. 4, pp. 311–314.
- Serdaroglu, A., Ertüzün, A. & Erçil, A. (2006) Defect detection in textile fabric images using wavelet transforms and independent component analysis. *J. Pat. Rec. Im. Anal.* **16**(1), 61–64.
- Serra, J. (1982) *Image Analysis and Mathematical Morphology*. Academic Press, London.
- Sivakumar, K. & Goutsias, J. (1997) Discrete morphological size distributions and densities: estimation techniques and applications. *J. Elect. Imaging* **6**(1), 31–53.
- Xie, X. & Mirmehdi, M. (2005) Texture exemplars for defect detection on random textures. In *ICAPR*, vol. 3687, pp. 404–413. Springer-Verlag, Berlin Heidelberg, Germany.