Non-parametric synthesis of laminar volumetric textures from a 2D sample

Radu Urs[1, 2]
radu.urs@ims-bordeaux.fr
Jean-Pierre Da Costa[1, 2]
jean-pierre.dacosta@ims-bordeaux.fr
Jean-Marc Leyssale[3]
leyssale@lcts.u-bordeaux1.fr
Gérard Vignoles[3]
vinhola@lcts.u-bordeaux1.fr
Christian Germain[1, 2]
christian.germain@ims-bordeaux.fr

[1] University of Bordeaux, IMS, UMR 5218, F-33400 Talence, France
[2] CNRS, IMS, UMR 5218, F-33400 Talence, France
[3] CNRS, LCTS, UMR 5801 CNRS-UB1-CEA-Safran, F33600 Pessac, France

Abstract

The goal of this paper is to evaluate non-parametric algorithms that achieve 3D texture synthesis from a single 2D sample. The algorithms under study are variants of the algorithm proposed by Wei and Levoy [1]. Several authors have proposed different algorithms that intend to better reproduce, in the output texture, the diversity learned in the input sample. Hence, we turn our attention to the improved algorithm proposed by Kopf et *al.* [2] and the particular histogram matching approach of Chen and Wang [3]. In addition, we propose to visit the voxels during synthesis according to the 3D extension of space filling curves. We investigate the algorithms capability to reproduce anisotropic textures. In particular we are interested in laminar textures, i.e. textures made of anisotropic sheets stacked along a given direction. Examples of such textures are snapshots of dense carbons observed by high resolution transmission electronic microscopy (HRTEM). Beyond the traditional subjective evaluation of the synthesized textures, we propose a genuine quantitative benchmark for the analysis of the synthesized textures which consists in comparing input and output gray level statistics and patterns morphology.

1 Introduction

Many tools in image processing, vision and computer graphics involve image analysis and synthesis, and especially textured images. For this reason, the disciplinary field of texture synthesis is particularly dynamic and very productive with notable applications in image compression, inpainting, image extrapolation or texture mapping. This research field has led to the development of many 2D synthesis techniques, but their extension to the 3D environment remains unstable proving itself as a very complex and computational issue. 3D textures are mainly used for texturing volumetric objects trying to increase the realism of the 3D scenario, but they can also be observed in 3D vision when exploring for instance material structure or seismic data.

 http://dx.doi.org/10.5244/C.26.54

The interest of the scientific community for solid texture synthesis gave birth to various families of algorithms. The procedural family of 3D synthesis algorithms models a solid texture by adjusting some mathematical functions. The range of applicability of the procedural approaches [6, 9, 10] remains sparse, leading mostly to algorithms hard to optimize [6, 22-24]. In opposition to the global growth procedure involved in the procedural methods, the local methods [1-4, 11-13] generate the texture one voxel/patch at a time maintaining the coherence of the local texture with its vicinity. Some of the most underlined synthesis methods based on local optimization are the exemplar-based synthesis algorithms. These algorithms require only one sample texture as input image and synthesize the solid texture guaranteeing the best possible similitude with the sample. Allowing a smoother control during the synthesis process, the pixel by pixel synthesis algorithms [1-3, 8, 12] are among the most successful techniques mainly because of their simplicity. To assure a reasonable computational complexity, most of the approaches propose efficient searching algorithms combined with a multi-scale implementation capable of capturing patterns at different scales and not increasing the computational load.

In this article we are interested in the synthesis of anisotropic volumetric textures, in particular textures with laminar structure. Some 2D examples of laminar textures from quite different domains are provided in Fig. 1. The synthesis of such textures from a unique sample is however delicate. It means to infer 3D information from 2D information only. This inference is possible only under some assumptions of anisotropy. A recent 2D/3D extension of Portilla and Simoncelli's parametric algorithm [9] was proposed by Da Costa and Germain [10], relying on the 2D to 3D inference of the 1st and 2nd order statistics and on an original 3D pyramidal decomposition algorithm. The authors showed that the 2D/3D inference was relevant and that their algorithm was able to produce convincing 3D synthetic textures from a band-pass filtered image sample. However, their approach revealed some limitations on raw (unfiltered) images, while showing high frequency artefacts on filtered input samples.

In the following, we will investigate a few non-parametric approaches from the literature for the synthesis of 3D textures from a unique 2D sample. We will carry out an experimental study to evaluate their capability to reproduce both low-level pixel statistics and pattern morphology.

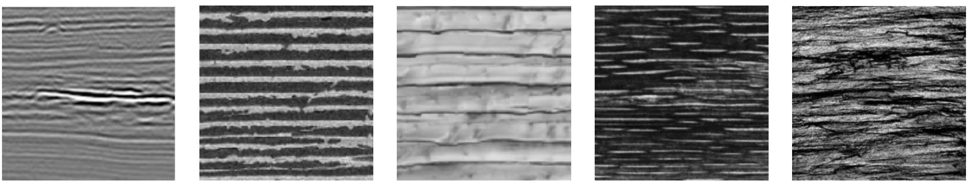

a *b* *c* *d* *e*

Figure 1: Examples of laminar textures: *a* – seismic data, *b* – ceramic material with polymer infiltration (image from [18]), *c* – wood profile, *d* – brodatz (D68 image from [19]) and *e* – sedimentary rock.

2 Non-parametric volumetric synthesis techniques

In this chapter we are briefly exposing the non-parametric approaches on which we focused our attention. The first one, the algorithm of Wei and Levoy [1, 4], will serve as a reference. The second one, namely Kopf et *al.* [2], is employed for the improvements brought by its non-parametric global texture optimization procedure. Finally, the last one is

Chen and Wang's algorithm [3], proposing a high quality texture synthesis by integrating in the texture optimization process two new kinds of matching histograms.

2.1 Overview of Wei and Levoy's algorithm

Based on the Markov field hypothesis, this method relies on texture locality (every pixel is predictable from the few pixels in its neighbourhood) and stationarity (spatial statistics invariance by translation), assumptions stated in [4]. The 2D/3D extension that we put up here is an adaptation of the algorithm of texture synthesis from multiple sources [1] using a single 2D image as source of synthesis. Starting from a random initialization, we synthesize voxel by voxel by examining the 2D neighbourhood of the current voxel from two orthogonal views of the 3D block (front and side view) as shown in Fig. 2. This phase implies a TSVQ (Tree-Structured Vector Quantization) search of the best match for each of these two neighbourhoods in the input image. The output pixel corresponds to the average of the two found voxels. The process is reiterated several times until reaching the same neighbourhoods after 2 consecutive iterations. This 2D/3D extension raises different downsides being constrained by the neighbourhood size and shape or the synthesis scan type. However, by including a multi-resolution scheme it proves to be a well adaptable and largely applicable synthesis method. In spite of its versatility and its good computational properties, this algorithm is also known to produce output textures that are more regular than the examples it tries to mimic.

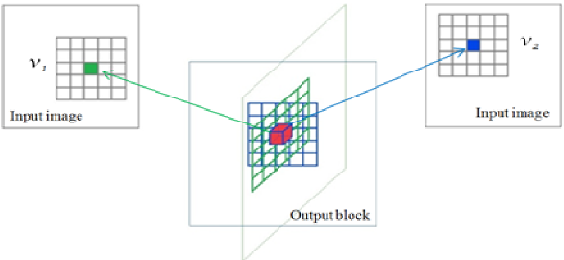

Figure 2: Principle of non-parametric synthesis: extract two neighbourhoods in the output block (front and side view), search the best two neighbourhoods in the input image, combine their central pixel values and modify the output voxel.

2.2 Overview of Kopf's algorithm

A critical issue of Wei and Levoy's algorithm resides in the way of combining information from the front view and the side view. Simple averaging leads to a loss of dynamics between the original and the synthesized texture. One possible answer is to replace the average with a better combination [2, 5] and, in the same time, adding a colour histogram matching mechanism [2, 6] in the texture optimization procedure. Kopf et *al.* [2] propose to give weights to each of the two found voxels and to mix them according to the robust optimization proposed by Kwatra et *al.* [5]. The colour histogram matching works by adjusting the weight of each voxel engaged in the weighted average that could bring differences between the colour histogram of the 3D result and the colour histogram of the 2D exemplar. In this way the global statistics are preserved and the synthesis does not rely only on local decisions. The quality of the results delivered with this technique exceeds that of many other methods. Despite this, computing the results by using the weighted average may produce blurry results if the variance of the exemplar voxels is too large [2].

2.3 Overview of Chen and Wang's algorithm

Unfortunately, most of the non-parametric methods are affected by blurring, missing textural patterns or mismatched input/output histograms. To ameliorate these drawbacks, Chen and Wang [3] propose a new modus operandi for texture optimization approaches. Mainly, it consists in using two new kinds of histograms – position and index histograms [3] and in updating the output voxel using the discrete solver as Han et *al.* [7]. Prior to synthesis, for every pixel in the exemplar, the *k* best matches are found in the exemplar, idea inspired from the natural texture synthesis algorithm proposed by Ashikmin [8]. The set of candidates being constructed, a synthesis process is started, similar to Kopf et *al.* [2]. However, instead of the colour histogram, the position histogram is used to adjust the weights. The position histogram records the number of occurrences in the solid texture of the corresponding pixel in the 2D exemplar. The weights are adjusted to get a histogram as uniform as possible. Once the position histogram matching is achieved, the colour histogram matching is achieved also. Index histogram counts the frequency of the voxels candidates in the volume texture. During the search phase, the choice of pixels is modulated according to their index histogram frequency. Having all the candidates uniformly distributed means having a well preserved texture.

2.4 Improving the pixel scan order

Whatever the method, a problematic issue is the choice of the neighbourhood system.

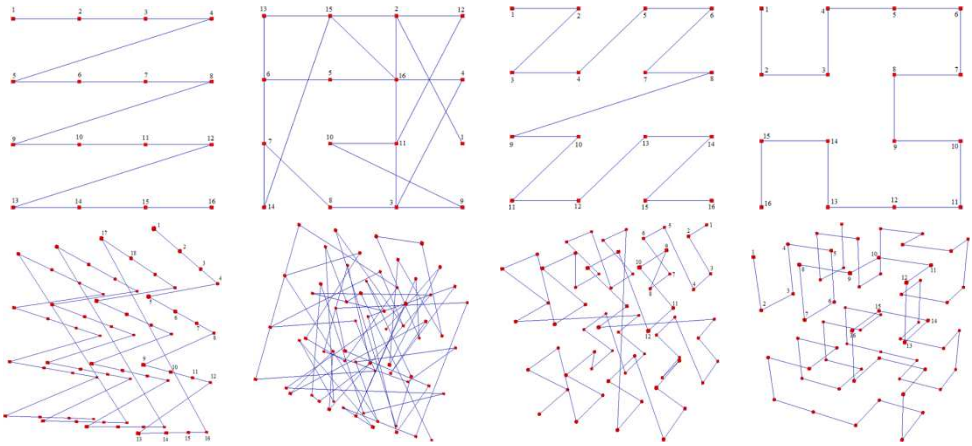

a *b* *c* *d*

Figure 3: Illustration of different scanning types, in 2-dimensions (the top column) and their corresponding extensions to 3-dimensions (the bottom column): from left to right, *a* - lexicographic path, *b* - random walk, *c* - Z-curve and *d* - Hilbert curve.

The choice of a causal neighbourhood makes the synthesis of a pixel totally dependent on previous pixels and leads in the case of lamellar textures to a higher regularity of synthesized results. The choice of a non-causal neighbourhood can partly overcome this determinism, but now the choice of the scan type may be challenged. Replacing the lexicographical scan type with a completely random walk allows the synthesis of a pixel by freeing itself from its past and so multiplying the possible configurations. However, the convergence time can become prohibitive in this case. To reconcile the deterministic part and the randomness character we propose an alternative scan type, namely the use of space filling curves [17] extended to 3D, like Morton code (Z-curve) or fractal curves (e.g. Hilbert curve) illustrated in Fig. 3.

3 Experimentation background

We were able to implement different variants of the non-parametric synthesis algorithm and apply them to the volumetric texture synthesis starting from a single 2D sample. Beyond the traditional subjective evaluation of the synthesized textures, we propose a genuine quantitative benchmark for the analysis of the synthesized textures which consists in comparing input and output image characteristics.

We have provided some solid results in Fig. 4. The results are obtained after 10 iterations, using the optimization approach proposed by Kopf, 5 pyramidal levels, a neighbourhood of size 7×7 pixels and random scanning. Input images have a size of 128×128 pixels and the produced blocks have a size of 128×128×128 pixels. It proves the capacity of the non-parametric algorithm to synthesize different types of textures, showing the potential of this method in different domains. More or less, the synthesized results resemble the source textures conserving the dynamics and the structure.

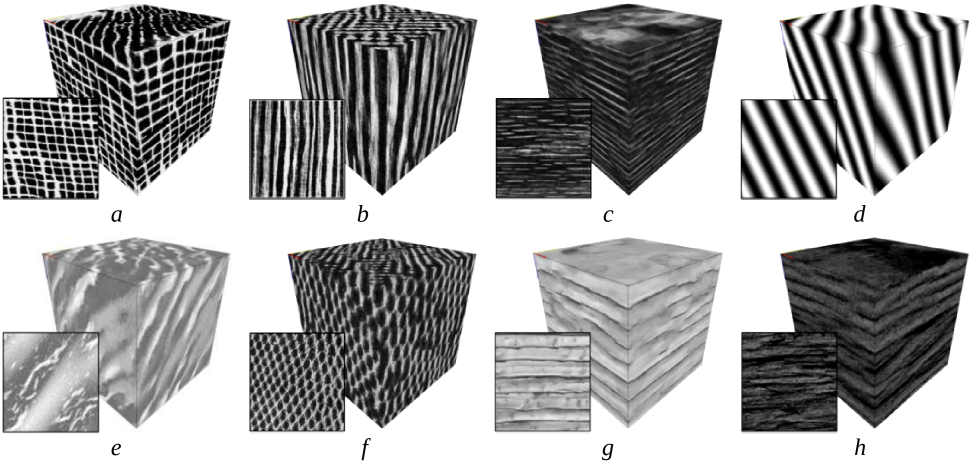

Figure 4: Several solid results, using as sources for synthesis images from Brodatz [19] database (*a, b, c, d, e, f*), wood wall profile (*g*) and sedimentary rock texture (*h*).

But we have focused our attention to the snapshots of dense carbons observed by microscopy [14] on which we have conducted our experimental procedure. Four examples of such laminar textures called lattice fringe images are found in Fig. 5a and 5c. They show locally periodic patterns together with high and low frequency artefacts. Neither high frequency nor low frequency components provide useful information about the atomic structure of the observed material. A good compromise for analyzing them consists in using their filtered versions (Fig. 5b and 5d) by applying a radial and directional band-pass filter in the frequency domain [20]. But in the following, we have carried out the quantitative experimentation on the raw HRTEM images.

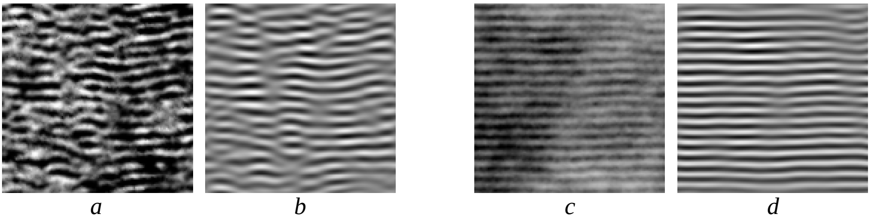

Figure 5: Samples of lattice fringe images of dense pre-graphitic carbons: HRTEM images (*a*, *c*) and their filtered versions (*b*, *d*). Sample sizes are 128×128 pixels.

The near periodic structure observed in such texture images can be investigated in 2D, but cannot be investigated in 3D for many technical reasons related to sample preparation and HRTEM technology. The only way to access the three dimensional structure is to infer it from 2D by image synthesis [10].

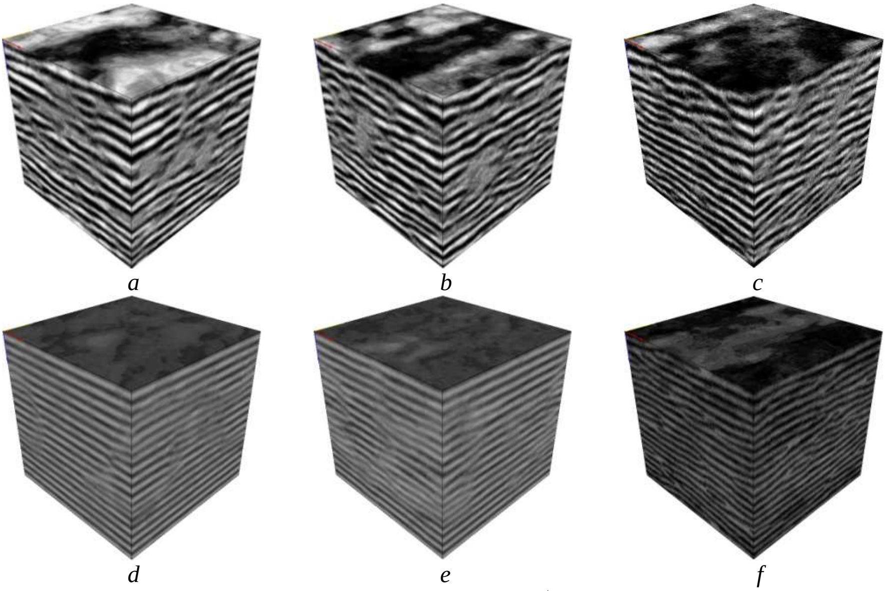

Figure 6: Volumetric results: from top to bottom, the 1st line represents the 3D views of the solid textures obtained after 10 iterations from the raw HRTEM sample in Fig. 5a using, from left to right, *a* - Wei and Levoy, *b* - Kopf, *c* - Chen and Wang; the 2nd line (*d*, *e*, *f*) shows the blocks after applying the same algorithms on the sample in Fig. 5c.

Fig. 6 contains some volumetric results obtained from a single HRTEM texture using non-parametric approaches. Fig. 6a and 6d show the volumetric textures for Wei and Levoy's approach [1, 4] while Fig. 6b and 6e show the results obtained using the backbone of Wei and Levoy but optimized by using the weighted average and the colour histogram matching technique proposed by Kopf et *al.* [2]. In Fig. 6c and 6f are represented the 3D blocks obtained by using the discrete solver plus the index and the position histogram matching technique of Chen and Wang [3]. The experimentation framework consists in using 5 pyramidal levels, non-causal neighbourhoods of size 7×7 pixels, Hilbert curve scan type and, for Chen and Wang, a set of 15 candidates for each pixel.

3.1 Subjective evaluation

The visual quality of the results is relatively convincing, in terms of dynamics and structure. Measuring the results based on our human perception, we can say that the 3D results and the original images seem to have been created by the same underlying process (see the image from Fig. 5a and the results in Fig. 6a-c for analogy). The perceptual comparison of the results in Fig. 6d-f with the one in Fig. 5c reveals that, in some cases, the algorithms are not capable to preserve the gray levels even when the algorithms involve histogram adjustments, while the structure is preserved. In addition, taking into account the size of the structures in determining the number of synthesis levels and the neighbourhood size is a critical point in terms of synthesis quality.

3.2 Quantitative evaluation

The visual interpretation of the synthesis results remains subjective, so we are constrained to compare them objectively. This original quantitative study evaluates the results taking into account the results gray level dynamics (first order statistics) and the morphological properties like elongated pattern (fringe) lengths, tortuosity and orientation. It aims to evaluate the algorithms capacity to reproduce a 3D texture respecting the statistical properties of the exemplar.

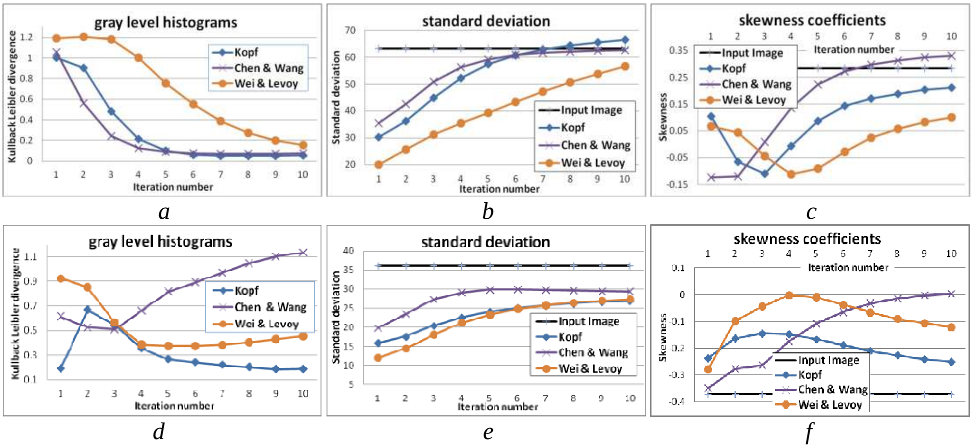

Figure 7: Different indicators of the objective comparison of the 3D textures obtained by synthesizing the HRTEM texture in Fig. 5a (the 1st line) and the sample in Fig. 5c (the 2nd line): gray level histograms (*a*, *d*) and 1st order statistics (*b*, *e* - standard deviation; *c*, *f* - skewness). The procedure consist in comparing input image statistics with the ones obtained after a multi-2D solid blocks analysis (*b*, *c*, *e*, *f*) or by computing Kullback-Leibler divergence between the histograms of the exemplar and the output block (*a*, *d*).

The quantitative indicators from Fig. 7, Fig. 9 and Fig. 10 reveal the 3D textures tendency towards the same structure observed on the HRTEM images. The comparison of the results dynamics from Fig. 7 confirms the visual conclusion attributed to some of the algorithms not being able to conserve the contrast. The plots in Fig.7a-c show that in general the approach of Wei and Levoy takes more time to reach the statistics of the input image because of the averaging operation when combining the two orthogonal views. This objective comparison of the evolution of the 3D blocks statistics towards the ones from the

2D sample strengthens the visual assumptions relative to the significant improvements brought by the histogram matching techniques. However, the indicators in Fig. 7d-f show that using the index and the position histograms, proposed by Chen and Wang, does not always improve the dynamics, by contrary, it deviates the statistics from the target.

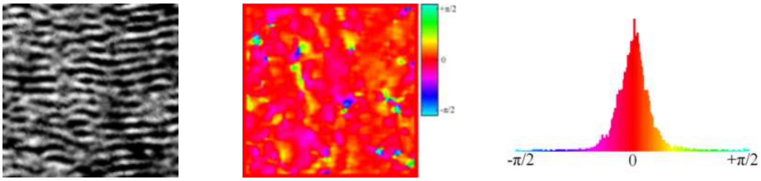

a *b* *c*

Figure 8: Texture orientation: *a* – the original raw HRTEM sample from Fig. 5a; *b* - the correspondent orientation field and the colour bar giving the correspondence between colours and orientations; *c* – the orientation histogram computed from the orientation field.

As for the structural properties of the 3D blocks, we have compared in Fig. 9 the local orientations of the multi-2D block relative to the input sample based on a structure tensor estimated on a 7×7 neighbourhood inspired from Bigun et *al.* [15]. Fig. 8 exemplifies the image orientation field and its associated orientation histogram. Fig. 9 confirms the impression that in some cases the proposed algorithms tend to produce textures more regular than that of the input image and iterating more doesn't resolve the reproduction of patterns. This phenomenon may be related to the deterministic character of the algorithms, based on local decisions, which tends to produce these very regular or repetitive textures.

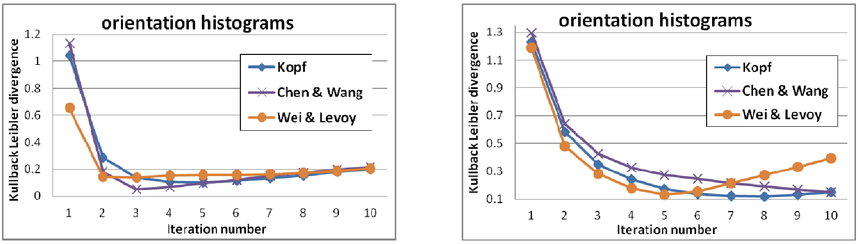

a *b*

Figure 9: Spatial structure variation indicators: the Kullback - Leibler divergence between the orientation histograms of the exemplar and the output blocks obtained by applying the algorithms on the sample in Fig.5a (plot *a*) and on the sample in Fig.5c (plot *b*).

The comparison of the morphological properties consists in tracking the texture level curves and describing the elongated patterns contained within the texture. We address the orientation and pattern length measurement techniques as described by Da Costa et *al.* [16] and Raynal et *al.* [21] to extract textural features (fringes) and compute their lengths and tortuosity. The procedure consists first in filtering [20] the raw HRTEM images and then counting the frequency of appearance of the fringe lengths and classifying them as a function of their lengths. Similarly, we built tortuosity classes and compare the tortuosity values distribution. The results from Fig. 10 show that the produced textures contain longer fringes than the input image (plot 10a) and in the same time they are more regular (plot 10b). Using larger images on input and synthesizing bigger blocks on output should impose a sufficient number of fringes capable of describing better the morphology.

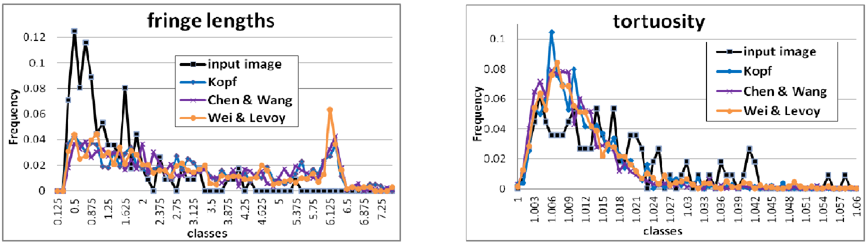

a *b*

Figure 10: Plots of the morphological structure indicators: *a* - the distribution of fringe lengths in the input raw HRTEM image from Fig. 5a and in the multi-2D synthesized volumetric textures from Fig. 6a-6c; *b* – the distribution of the tortuosity values corresponding to the same input image and output blocks used in *a*. The fringes were retrieved by manipulating the front and the side slices in the 3D textures.

Comparing the different types of pixel scanning during synthesis, it seems that the orientation histograms of the produced textures are closer to the one of the input image when using the random scan or the space filling curve (as shown by Fig. 11a), while the lexicographical walk produces textures more regular than the input. Fig. 11b shows that the scan type does not have a significant influence on the gray level dynamics.

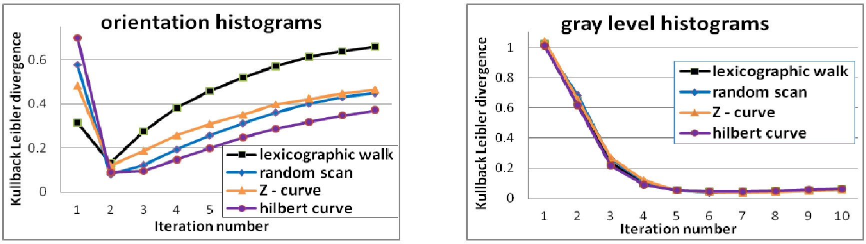

a b

Figure 11: Comparing the synthesis scan types: *a* – the orientation histograms indicator and *b* – the gray level histograms indicator, corresponding to the synthesis of the image from Fig.5a, using Kopf's approach.

4 Conclusions

In this paper, we have brought about several non-parametric algorithms that achieve with success solid texture synthesis starting from a 2D texture exemplar. Results interpretation was accompanied by an original experimental procedure marking a qualitative and objective evaluation. Wei and Levoy's approach is not able to preserve gray levels while the learning phase associated to the set of candidates of Chen and Wang is time consuming and increasing with the size of the set. It proves to be more efficient to use Kopf's approach and a good compromise between the quality of the results and the synthesis time is to use Hilbert space filling curve as scan type. As an application, we have addressed these algorithms to the synthesis of laminar textures, and more precisely the anisotropic HRTEM images of dense carbons. Apart from certain limitations, the obtained results are convincing in term of preserving texture gray levels and texture structure properties.

Acknowledgements

This work has been carried out within the project PyroMaN (Pyrocarbon Matrices at the Nanoscale, http://pyroman.dr15.cnrs.fr) funded by the French National Research Agency (ANR) and Aerospace Valley, World Competitiveness Cluster in Aeronautics, Space and Embedded Systems of Midi-Pyrenees and Aquitaine. The company Snecma Propulsion Solide (Safran group) has to be also acknowledged for providing the image samples.

References

[1] L.-Y. Wei and M. Levoy. Texture synthesis from multiple sources. *Proceedings of ACM SIGGRAPH* 2003 Sketches & Applications, New York.

[2] J. Kopf, C.W. Fu, D. Cohen-Or, O. Deussenn, D. Lischinski and T.T. Wong. *Solid Texture Synthesis from 2D Exemplars*. ACM SIGGRAPH, TOG, vol. 26, issue 3, 2007.

[3] J. Chen and B. Wang. *High quality solid texture synthesis using position and index histogram matching*. The Visual Computer, vol.26, pp. 253-262, 2010.

[4] L.-Y. Wei and M. Levoy. *Fast texture synthesis using tree-structured vector quantization*. Proceedings of ACM SIGGRAPH 2000, pp. 479-488, 2000.

[5] V. Kwatra, A. Schodl, I. Essa, G. Turk and A. Bobick. *Graphcut textures: image and video synthesis using graph cuts*. ACM SIGGRAPH, pp. 277-286, 2003.

[6] D. J. Heeger and J.R. Bergen. *Pyramid-based texture analysis/synthesis*. Proc. 22nd annual Conf. on Computer graphics and interactive techniques, pp. 229–238, SIGGRAPH 1995.

[7] J. Han, K. Zhou, L.-Y. Wei, M. Gong, H. Bao, X. Zhang, B. Guo. *Fast example-based surface texture synthesis via discrete optimization.* Visual Computer, 22(9), pp. 918–925, 2006.

[8] M. Ashikhmin. *Synthesizing natural textures*. ACM Symposium on Interactive 3D Graphics, Research Triangle Park, North Carolina, pp. 217-226, 2001.

[9] J. Portilla and E.P. Simoncelli. *A parametric texture model based on joint statistics of complex wavelet coefficients*. International Journal of Computer Vision, 40(1): 49–70, 2000.

[10] J.-P. Da Costa and C. Germain. *Synthesis of solid textures based on a 2D example: application to the synthesis of 3D carbon structures observed by transmission electronic microscopy*. Proc. of SPIE, vol. 7538, Image Processing: Machine Vision Applications III, 2010.

[11] R. Paget and I.D. Longstaff. *Texture synthesis via a noncausal nonparametric multiscale markov random field*. IEEE Trans. on Image Processing, 7(6): 925–931, 1998.

[12] A.A. Efros and T.K. Leung. *Texture synthesis by non-parametric sampling*. Proc. of International Conference on Computer Vision (ICCV), pp. 1033-1038, 1999.

[13] S. Lefebvre and H. Hoppe. *Appearance-space texture synthesis*. ACM Transactions on Graphics, 25(3):541-548, 2006.

[14] J.-M. Leyssale, J.-P. Da Costa, C. Germain, P. Weisbecker and G. Vignoles. *An image-guided atomistic reconstruction of pyrolytic carbons*. Applied Physics Letters, vol. 95, 2009.

[15] J. Bigun, G. Granlund, J.Wiklund. *Multidimensional orientation estimation with applications to texture analysis and optical flow*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 13(8): 775-789, 1991.

[16] J.-P. Da Costa, C. Germain and P. Baylou. *Level curve tracking algorithm for textural feature extraction*. 15th International Conference on Pattern Recognition. Vol.3, pp. 909-912, 2000.

[17] H. Sagan. *Space-Filling Curves*. Springer-Verlag, New York, 1994. ISBN : 0-387-94265-3.

[18] E. Munch, M.E. Launey, D.H. Alsem, E. Saiz, A.P. Tomsia, R.O. Ritchie. *Tough, Bio-Inspired Hybrid Materials.* Science Magazine, Vol. 322 no. 5907 pp. 1516-1520, December 20008.

[19] P. Brodatz. *A Photographic Album for Artists and Designers*. Dover, New York, 1966

[20] J.-P. Da Costa, C. Germain, P. Baylou and M. Cataldi. *An image analysis approach for the structural characterization of pyrocarbons.* Proceedings of Composite Testing, University Press, Bristol, UK, 2004.

[21] P. I. Raynal, M. Monthioux, J.-P. Da Costa O. Dugne. "*Multi-scale quantitative analysis of carbon structure and texture: III. Lattice fringe imaging analysis*". Carbon 2010 Conference, Clemson, SC, 11-16 July 2010, M. Thies ed. (2010) ref. 627.

[22] D. R. Peachey. *Solid texturing of complex surfaces*. Computer Graphics, 19:279–286, 1985.

[23] K. Perlin. *An image synthesizer*. In Proceedings of SIGGRAPH ’85, pages 287–296, ACM Press, New York, USA, 1985.

[24] A. Witkin and M. Kass. *Reaction-diffusion textures*. SIGGRAPH Computer Graphics, 25(4):299–308, 1991.