

MICROSTRUCTURE PRESERVING SYNTHESIS OF BIOMEDICAL IMAGES

Shantanu Singh, Kishore Mosaliganti, Raghu Machiraju

Department of Computer Science and Engineering
The Ohio State University, Columbus, Ohio, U.S.A.
{singhsh,mosaligk,raghu}@cse.ohio-state.edu

ABSTRACT

We present an approach for synthesizing biological tissue textures from existing tissue samples. At microscopic resolution, tissues are characterized by a spatial arrangement of nuclei, cytoplasm, red-blood cells (RBCs) and adipose components etc. We employ 2-point correlation functions (*2-pcfs*) to encode the geometrical aspects of component arrangements in the synthesis process. The *2-pcfs* belong to a class of neighborhood density estimators and were recently introduced in microscopic image analysis. We provide examples of their application toward synthesis of histology-based tissue textures. We show that our methods retain properties such as component volume fractions, sizes and density in comparison to standard approaches. Our methods are also shown to improve the performance of segmentation algorithms by automatically generating labeled texture classes.

Index Terms— N -point correlation functions, texture analysis

1. INTRODUCTION

The problem of texture synthesis has been studied extensively in the field of material science [1, 2], computer graphics [3] and image analysis [4]. These techniques, in general, enhance the perceptual similarity between the original and synthesized textures [5, 6] and find applications in creating novel graphics, animation and models etc. Texture synthesis, in the context of biomedical imaging, constitutes an important step towards creating realistic models of biological tissue and imaging systems for the timely diagnosis of disease. Biological texture synthesis was first studied by Brettle *et al.* [7] for conducting psychophysical experiments. Bochud *et al.* [8] improvised by presenting a biological synthesis technique that maintained first and second-order statistics of the original image. They used synthesized images to provide the user with a large number of independent samples, in a controlled environment with analytically tractable properties.

Our work applies the same idea of synthesizing new samples in order to enhance the performance of supervised learning-based techniques. We seek to retain features that characterize a texture class, while providing greater entropy

in the training set. For example, consider the problem of segmenting large histology images consisting of several tissue classes [9]. This approach classifies (using k -nearest neighbor) individual blocks of the image into one of many texture classes. The training phase requires many labeled examples of the tissue class to be available. Labeled samples are expensive to generate because it requires an expert to intervene. Furthermore, such algorithms would need to be frequently re-trained for a new novel dataset for optimal performance. Hence, manual intervention becomes quite significant and cumbersome. Our technique focuses on synthesizing new samples that maintains some microstructural features from the original class. These microstructural features robustly characterize tissue textures.

Earlier, we had characterized a tissue sample as a spatial arrangement of nuclei, cytoplasm, red-blood cells (RBCs) and adipose components etc. The organization of components is related to the definition of a material microstructure. Microstructure may be measured as a collection (ensemble) of points, lines, internal surfaces, and volumes [10]. Each microstructural feature is associated with size, shape, volume, surface area, length and curvature attributes etc. Statistical distributions of such attributes collectively specify the geometric state of a microstructure. These properties of microstructure are formalized by the statistical N -point correlation functions (N -*pcfs*) [11]. In this work, we show the utility of *2-pcfs* for comparing the biological organization in images in terms of component arrangements. We demonstrate synthesis results on histology images of adipose tissue from the mammary gland of a mouse and provide classification results showing the improvement gained by supplementing the training data with synthesized samples of tissue classes.

2. TEXTURE SYNTHESIS

The texture synthesis approach that we employ consists of 3 components: (i) measuring microstructural features that characterize tissues using *2-pcfs*, (ii) a similarity metric (kernel) to compare in the *2-pcf* space and (iii) a search technique using kernel principal components analysis for finding similar texture regions. We now describe each component and finally provide an overview of the overall synthesis approach.

2.1. 2-Point Correlation Functions

The 2-*pcf* describes the spatial distribution of a material component in a multi-phase medium. A detailed description of 2-*pcfs* and other microstructural functions is given in [11]. In brief, we denote by $\mathcal{I}^{(i)}(\mathbf{x})$ the indicator function for the i^{th} phase at point \mathbf{x} in the medium such that

$$\mathcal{I}^{(i)}(\mathbf{x}) = \begin{cases} 1 & \text{if point } \mathbf{x} \text{ is in component } i, \\ 0 & \text{otherwise.} \end{cases}$$

The variable $\mathcal{I}^{(i)}(\mathbf{x})$ is thus a random variable whose distribution characterizes the material. The expected value of the variable $\langle \mathcal{I}^{(i)}(\mathbf{x}) \rangle$ is called the 1-*pcf* and is denoted by $S_1^{(i)}(\mathbf{x})$. Similarly, the 2-*pcf* is defined as

$$S_2^{(i)}(\mathbf{x}_1, \mathbf{x}_2) = \langle \mathcal{I}^{(i)}(\mathbf{x}_1) \mathcal{I}^{(i)}(\mathbf{x}_2) \rangle \quad (1)$$

If the material is assumed to be statistically isotropic, that is, the distribution is rotationally and translationally invariant, then the distance between \mathbf{x}_1 and \mathbf{x}_2 is sufficient to parameterize the two-point probability function. Thus the 2-*pcf* for a statistically isotropic medium is given by $S_2^{(i)}(r)$ where r is the distance separating \mathbf{x}_1 and \mathbf{x}_2 . Similarly, the 1-*pcf* for an isotropic medium is constant everywhere. Thus $S_1^{(i)} = \phi_i$ gives the volume fraction of phase i in the material.

In order to define a similarity measure for the 2-*pcf* feature we look into some of its properties. Consider the matrix Γ where

$$\Gamma_{ij} = S_2(s_i - s_j) \quad (2)$$

where i and j are two arbitrary scales. If we evaluate S_2 at n scales from s_1 to s_n then the resulting $n \times n$ matrix Γ can be shown to be symmetric and positive semi-definite. The eigenvectors of Γ form an orthogonal basis, with the corresponding eigenvalues providing the scale. We use this property to define the similarity between two matrices Γ_1 and Γ_2 . Given the spectral decomposition $\Gamma_1 = E_1 \Lambda_1 E_1^T$ and $\Gamma_2 = E_2 \Lambda_2 E_2^T$, we define a similarity measure

$$k(\Gamma_1, \Gamma_2) = \text{tr}(E_1 \Lambda_1^{-\frac{1}{2}})^T (E_2 \Lambda_2^{-\frac{1}{2}}) \quad (3)$$

Essentially, the similarity is the sum of the dot products of corresponding (scaled) eigenvectors of each matrix.

2.2. Kernel PCA

Kernel Principal Components Analysis (Kernel PCA) is a non-linear extension of the standard PCA algorithm which can be used to extract a non-linear structure from a possibly high dimensional dataset [12]. In brief, kernel PCA finds the eigenvectors of the Gram matrix (dot product matrix) of a dataset that is defined with respect to a kernel. Regular PCA can be shown to be a special case of kernel PCA when the dot product is the canonical Euclidean dot product. With an appropriately defined kernel, using kernel PCA for feature

extraction provides the lowest mean-squared approximation error in representing the data points in the implicit feature space defined by the kernel k . For a given feature point \mathbf{x} , the component of its projection in the i^{th} largest principal component of the feature space is given by:

$$\tilde{x}_i = \sum_{j=1}^m \alpha_j^i k(\mathbf{x}, \mathbf{x}_j) \quad (4)$$

where α_j^i is the i^{th} component of the j^{th} eigenvector of the Gram matrix and \mathbf{x}_j is the j^{th} data point. We use the kernel defined in Equation 3 for performing kernel PCA on the data points. We note that the kernel we define is not positive definite so we use the standard technique of subtracting the smallest eigenvalue from the diagonal of the Gram matrix in order to guarantee positive-definiteness.

2.3. Synthesis Algorithm

The texture synthesis technique iteratively constructs an image that has similar textural properties as that of a specified target image. The similarity measure can be chosen based on the feature that we want to preserve. The technique is based on a general paradigm prevalent in texture synthesis in computer graphics literature - global similarity between two textures is achieved by imposing local similarity between pairs of neighborhoods in the two textures. In other words, if every neighborhood of the synthesized image is similar to some neighborhood of the target image, then the two textures are similar. In order to reduce computational complexity, we consider only a subset neighborhoods but ensure that they have sufficient overlap. This idea is exemplified in [5]. Our technique is modeled along a similar paradigm and seeks to retain the specified features from the target image.

Given a target image T we extract the set N_T of image patches corresponding to local overlapping neighborhoods of T by sliding a window across the image. Each patch captures the local structure of the image. Similarly, we extract N_S , the set of image patches for a randomly initialized synthetic image S . We use kernel PCA with the 2-*pcf* kernel on the set $N_T \cup N_S$ to get a lower dimensional representation that is optimal with respect to the kernel, i.e., has the lowest mean-squared approximation error in the feature space defined by the kernel. With this representation, we can now use the standard (Euclidean) nearest neighbor algorithm on each element of N_S to find the most similar image patch in N_T . We replace each element in N_S with its nearest neighbor from N_T .

The updated image S' is constructed by replacing the original image patches with the corresponding patches from the updated set N_S . Since we used a sliding window to extract the patches, the updated image will contain regions of overlap. The overlapping regions are assigned the median of the values across all overlapping patches in the region. That is, if n tiles overlap at a pixel location \mathbf{x} , then the value

assigned to the pixel is the median of the n values obtained from each of the tiles at the location x . This completes one iteration of the algorithm. The iterations proceed until the number of replaced patches drops below a threshold.

To reiterate, the goal of the algorithm is to construct a new image such that every neighborhood in the image is similar to some neighborhood in the target image. By allowing overlaps between neighborhoods, the local similarity of texture induces a global similarity between the two images. By using the *2-pcf* feature to measure similarity, we coerce the resulting image to have similar *2-pcf* characteristics.

3. EXPERIMENTS AND RESULTS

We present results of the synthesis technique on histology images of adipose tissue from the mammary gland of a mouse. The algorithm proceeds in a coarse to fine fashion - the texture at the coarsest level is synthesized first, and the finer levels of texture are progressively synthesized using the output of the previous level for initialization. We used a target texture at resolution 512×512 pixel and a window width of 32 pixels with a 16 pixel overlap. The lower levels of resolution have proportionate window and overlap dimensions. The coarsest level of the synthesized image was initialized randomly. An approximate nearest neighbor algorithm was used to find the closest match between windows. The target texture along with the synthesized texture at different stages of synthesis are shown in Fig. 1.

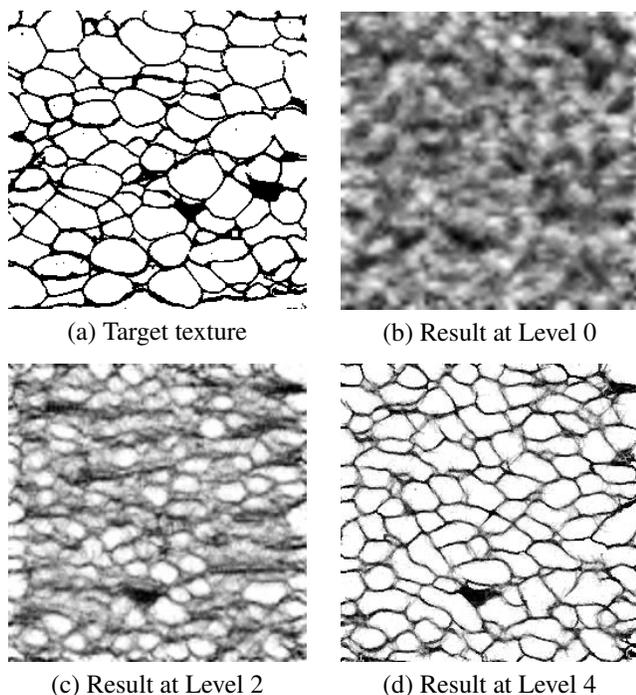


Fig. 1.

We compared the performance of the algorithm using two different similarity measures - the *2-pcf* kernel and sum of squares error (SSE). SSE is used for comparison by virtue of it being a standard measure of error in Euclidean space. Fig 2 shows the resulting textures, demonstrating that the *2-pcf* measure is able to better preserve the characteristics of the target texture. We also compared some seminal functions that are used to characterize microstructures, summarized in Table 1. The area fraction is the ratio of the area covered by a component to the total area of the medium. The average component size the average size of blobs of a component in the medium. *2-pcf*(normalized) is the similarity between the *2-pcf*'s normalized between 0 and 1. The results were obtained using 10 images from each type. The standard deviation of the results is given in parenthesis. The results illustrate that the *2-pcf* similarity measure is able to preserve the statistics that characterize the microstructures in the target image.

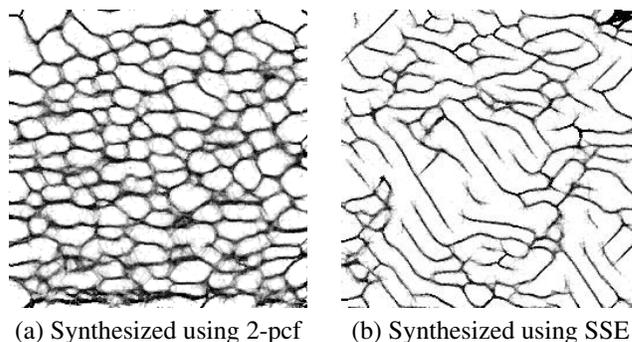


Fig. 2. Synthesis results using *2-pcf* and SSE

	Original	Syn. <i>2-pcf</i>	Syn. SSE
Area Fraction	0.79	0.85 (.01)	0.91(.01)
Avg. Comp. size	451	514 (60)	2550 (550)
<i>2-pcf</i> (Norm.)	1	0.86 (.03)	0.52 (.02)

Table 1. Comparing statistics from synthesis results

We show that the synthesis technique can also be used to boost the performance of segmentation algorithms by providing synthesized ground truth data. As discussed before, most segmentation techniques require many labeled examples of a texture class to be available. The quality of segmentation depends on the number and quality of the labeled data. By synthesizing new examples of a texture class that reliably preserve the characteristics of the original sample, we can improve of the segmentation quality. To illustrate this point, we conduct the following experiment. We train a classifier using a *single* example each of two texture classes. The classification methodology used was similar to that employed in [9]. We add one new synthesized example to each texture

class in the training set, and retrain the classifier. We progressively add new synthesized examples, retrain, and record the performance of the classifier on a testing set. Fig 3 shows the performance of a nearest-neighbor classifier on two texture classes. The experiment was repeated by choosing 10 different sets of training examples. The classification results for a single example are poor as expected. As the number of synthesized examples are increased, the classification results progressively improves.

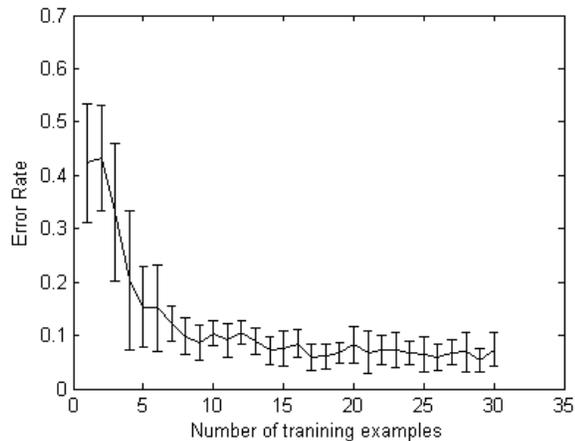


Fig. 3. Performance of texture classification by increasing training set using synthesized examples. (Error bars mark one standard deviation)

These results indicate that by providing just a few labeled examples of texture classes, we will be able to achieve significantly better segmentation results by supplementing our training set with synthesized examples. While a conclusive result that supports this claim would require performing segmentation on an image with multiple textures, these results support the hypothesis that texture synthesis can directly aid texture segmentation.

4. DISCUSSION

The synthesis technique works on a uniform scale at each level. This results in undesired blurring especially when the scale of the texture is smaller than the window. An adaptive version of the algorithm, which identifies the local scale at every level can help alleviate this problem. Further, an interesting extension of the work would be to synthesize 3D textures given 2D slices of the texture, borrowing ideas from stereology.

In conclusion, we have presented a synthesis technique designed to preserve important microstructural characteristics that are used in the analysis of biological textures. The goal of synthesis, in the context we have defined, is to aid analysis. The motivation to do so lies in the fact that generating labeled

data is expensive. We have presented indicative results to support the claim that synthesized data can be used for segmentation of biological tissue textures using only a few labeled examples.

5. REFERENCES

- [1] C. Yeong and S. Torquato, "Reconstructing random media," *Phys. Rev. E*, vol. 57, no. 1, pp. 495–506, Jan 1998.
- [2] L. Graham-Brady and X. Xu, "Stochastic morphological modeling of random multiphase materials," *ASME Journal of Applied Mechanics*, 2006.
- [3] A. Efros and T. Leung, "Texture synthesis by non-parametric sampling," *Intl. Conf. on Comp. Vision*, vol. 02, pp. 1033, 1999.
- [4] A. Efros and W. Freeman, "Image quilting for texture synthesis and transfer," in *SIGGRAPH 2001, Computer Graphics Proceedings*, Eugene Fiume, Ed., 2001, pp. 341–346.
- [5] V. Kwatra, I. Essa, A. Bobick, and N. Kwatra, "Texture optimization for example-based synthesis," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 795–802, 2005.
- [6] A. Zalesny, V. Ferrari, G. Caenen, and L. Van Gool, "Composite texture synthesis," *Int. J. Comput. Vision*, vol. 62, no. 1-2, pp. 161–176, 2005.
- [7] D. Brettell, E. Berry, and M. Smith, "Synthesis of texture from clinical images," *Image Vision Computing*, vol. 21, no. 5, pp. 433–445, 2003.
- [8] F. Bochud, C. Abbey, and M. Eckstein, "Statistical texture synthesis of mammographic images with superblob lumpy backgrounds," *Opt. Express*, vol. 4, no. 1, pp. 33–42, 1999.
- [9] F. Janoos, K. Mosaliganti, O. Irfanoglu, R. Machiraju, and K. Huang, "Multi-resolution image segmentation using the 2-point correlation functions," *International Symposium on Biomedical Imaging (ISBI)*, pp. 300–303, 2007.
- [10] A. Gokhale, "Experimental measurements and interpretation of microstructural n-point correlation functions," *Microscopy and Microanalysis*, vol. 10, pp. 736–737, 2004.
- [11] S. Torquato, *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*, Springer, 2001.
- [12] B. Scholkopf and A. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, Cambridge, MA, USA, 2001.