Improving texture optimization with application to visualizing meat products

Line Clemmensen & Lasse Farnung Laursen
Department of Informatics and Mathematical Modelling
Technical University of Denmark
{lhc, lfla}@imm.dtu.dk

May 17, 2011

Abstract

When inspecting food quality, CT Scanning is among the primary tools used to gain insight. It provides valuable volumetric data using a process, which leaves the product unspoiled and untouched. However, volumetric data is merely a measure of density and therefore contains no appearance information (such as color, translucency, reflective properties). One way of reintroducing this lost information back to the volume data is to synthesize an appropriate texture and apply this to the volume data.

A recent method within the field of texture synthesis is called Texture Optimization presented by Kopf et al. in 2007. This method accepts a number of 2D input exemplars, from which it generates a solid texture volume. The volume is iteratively improved via an expectation maximization algorithm. The bottleneck of Texture Optimization occurs during a nearest neighbor search, between texture patches from the 2D input exemplars and the generated texture volume. We examine the current procedures for minimizing the bottleneck and present a novel approach which increases the speed of the synthesis algorithm while minimizing loss of quality.

The nearest neighbor search is performed in a high dimensional space. Applying a principal component analysis on the texture patches originating from the synthesized solid accelerates the process. These patches are then reduced in dimensionality until ”only” 95% of their original variance remains. This usually results in a dimension reduction from 192 to about 60-80. The reduction in dimensionality speeds up the convergence of the Texture Optimization method considerably.

We examine the impacts of reducing the dimensionality further by tweaking the parameters as well as introducing an alternative method to reducing the dimensionality. Additionally, we study the possibility of selecting only a subsample of the neighborhoods available from the input exemplar without significantly impacting the overall synthesis quality.

Keywords

Volumetric Rendering, Texture Synthesis, Dimension Reduction, Principle Component Analysis, Non-Negative Matrix Factorization, Subset Selection

1 Introduction

Adding appearance information to volumetric data is commonplace. In most cases this is achieved via transfer functions [HKRs$^+$06], but a viable alternative is to use solid textures [LEB11]. Solid textures and volume data are almost a perfect match, since they occupy the same dimensions, making the application straightforward. One drawback of solid textures is that they are not easily acquired. Therefore, there exists a large body of research concerned with synthesizing solid textures from existing 2D textures, also called input exemplars. A recent method is called Texture Optimization[KFCO$^+$07]. In brief terms, this method starts with synthesizing a solid volume based on random samples from the input exemplars. The volume is then iteratively improved on a rough scale using an Expectation Maximization method. Once it has converged, the solid texture is scaled up and iteratively improved until convergence is once again reached. Eventually the synthesis process is complete at the highest resolution of 128x128x128 voxels. Higher resolution synthesis is not performed due to the computational complexity involved.

Local correlation between the synthesized solid and the input exemplars is achieved by comparing small 8 by 8 texture patches, from here on also referred to as neighborhoods. These neighborhoods are essentially vectors in a 192 dimensional space, due to the three-color channels. For each neighborhood in the synthesized solid, the best match is found among the neighborhoods from input exemplars. Figure 1 shows how a number of neighborhoods from the synthesized solid are matched to neighborhoods originating from the input exemplar. This is basically a nearest neighbor search in a high dimensional space. This search is the bottleneck for the entire Texture Optimization method. To lessen the impact of this expensive search, a principle component analysis is performed on the input exemplars. The neighborhoods are then projected into a space where 95% of the variance is retained. This usually results in a dimension reduction from 192 to around 60-80 dimensions.

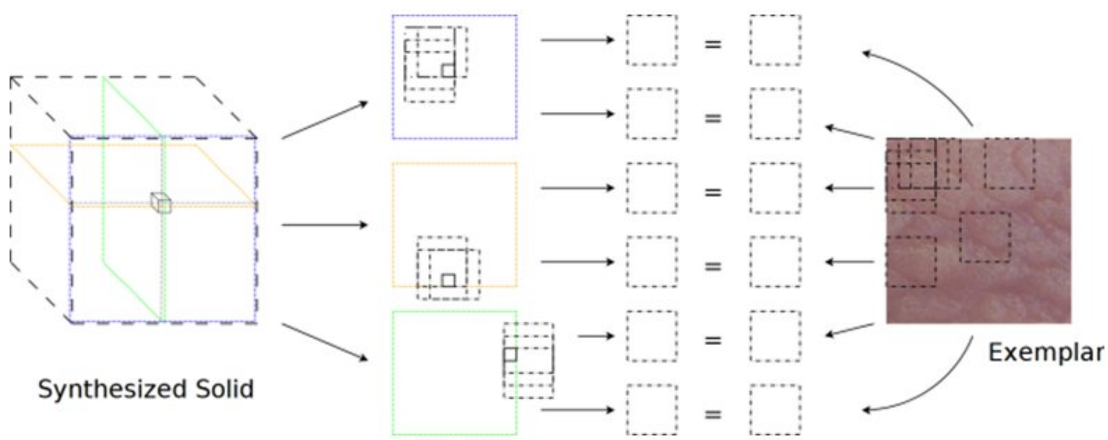

Figure 1: Neighborhoods on axis-aligned planes related to a single voxel are found. These neighborhoods are then compared to all existing neighborhoods in an input exemplar. For each, the optimal match is found.

We intend to show how optimizing the process in two different ways can further reduce this bottleneck. First, we examine whether we can further reduce the dimensionality by applying an alternative dimension reduction method as well as using fewer dimensions in the PCA. Second, we attempt to reduce the number of contributing neighborhoods based on a similarity measure. Eliminating similar neighborhoods from the pool provided by the input exemplar should speed up the synthesis algorithm while only slightly diminishing the texture quality.

2 Preliminary results

Our preliminary research has been performed on the input exemplars shown in figure 2, which originate from pig muscle.

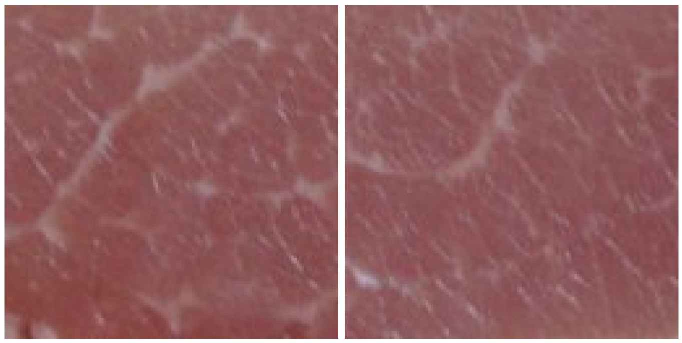

Figure 2: Input Exemplars used to obtain preliminary results.

Figure 3 shows results from the original texture synthesis algorithm, which uses a PCA reducing the number of dimensions to 80 while 95% of the variance is retained.

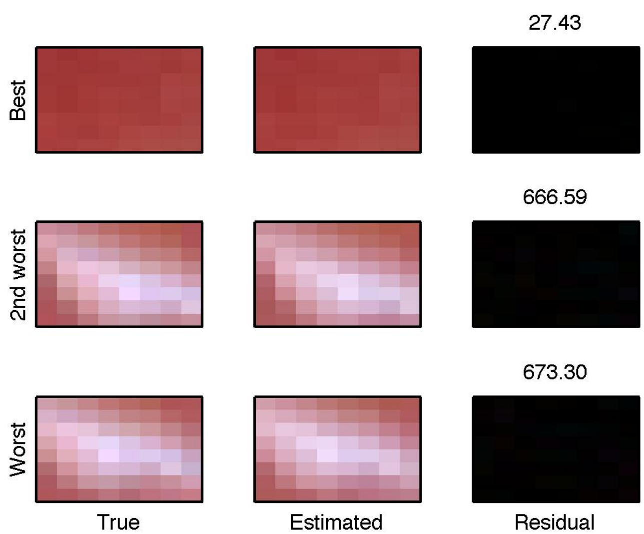

Figure 3: The true and estimated patches using a PCA dimension reduction of size 80. The last column shows the residual between the true and estimated patch and the associated number gives the SSE of the reconstruction. The rows, from top to bottom, represent the best, the 2nd worst and the worst reconstructions.

We apply Non-Negative Matrix Factorization (NMF), which is an alternative decomposition method [LS99]. NMF is known by learning by parts and instead of finding components, which describe maximal variance in data like PCA; it finds components, which describes localized features in data. Figure 4 shows the results from applying the NMF to our input exemplars.

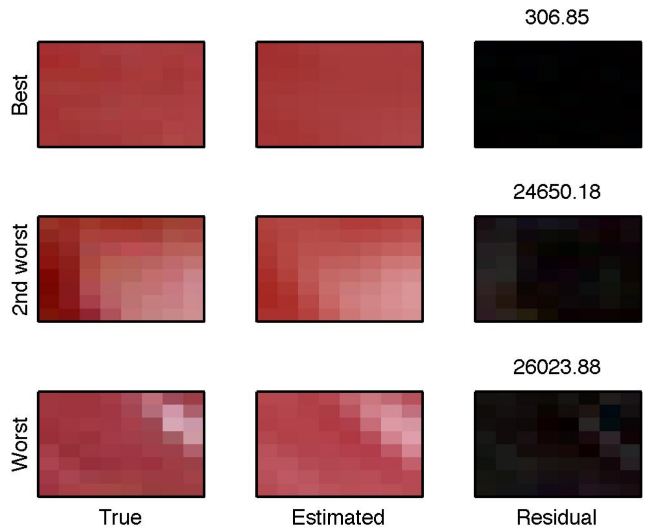

Figure 4: The true and estimated patches using NMF decompostion with a dimensionality of 15. The last column shows the residual between the true and estimated patch and the associated number gives the SSE of the reconstruction. The rows, from top to bottom, represent the best, the 2nd worst and the worst reconstructions.

The NMF reduces the dimensionality to 15, while still retaining similarity, even among the worst estimated neighborhoods.

Figure 5 shows the results of reducing the dimensionality to 15 via PCA, which contains 83% of the variance. Although the sum squared error is better when compared to NMF, the visual result is not noticeably improved. It is difficult at this point to ascertain the global impact of these changes to the final result of the texture synthesis process. The lack of perceptible difference warrants further study.

Figure 6 shows a comparison of a subset of neighborhoods extracted from the input exemplar, after a PCA dimension reduction to 80 dimensions. We measure Euclidian distance between the neighborhoods. At a Euclidian distance of 100, shown in the bottom right of figure 6, the difference between two random patches is visible, yet small. Neighborhoods with an even shorter distance will look increasingly similar. This supports our hypothesis that a number of neighborhoods are redundant and can be discarded with minimal impact to the final synthesized result.

3 Future work

Achieving the optimal balance between speed and quality will require further study. We have briefly covered the main areas within which we believe we can improve the Texture Optimization algorithm. By comparing original results from the method with results produced by our own modified methods, we should be able to achieve considerable increases in speed with hopefully minimal impact to the image quality.

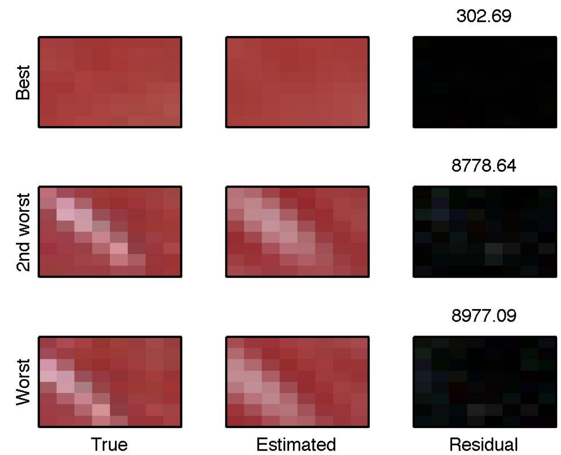

Figure 5: The true and estimated patches using PCA decompostion with a dimensionality of 15. The last column shows the residual between the true and estimated patch and the associated number gives the SSE of the reconstruction. The rows, from top to bottom, represent the best, the 2nd worst and the worst reconstructions.

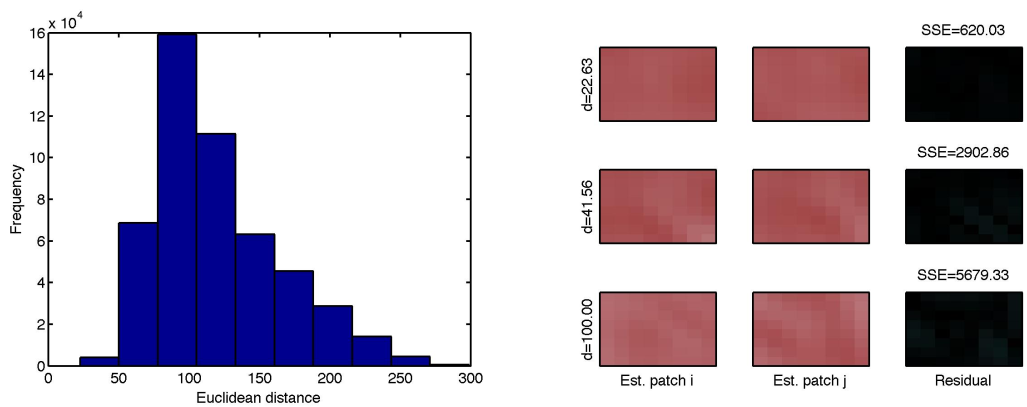

Figure 6: Left: Histogram of the Euclidian distances between the dimension reduced patches using a PCA reduction of 80 dimension. Only a subset of the patches were included in this analysis. Right: Examples of Euclidian distances (d), and an illustration of the corresponding two estimated patches, plus the residual between the two estimated patches and the SSE between the two estimated patches.

Acknowledgements

This research was supported in part by the Danish Meat Research Institute. The input exemplars were acquired with the help of the Danish Meat Research Institute.

References

[HKRs+06] Markus Hadwiger, Joe M. Kniss, Christof Rezk-salama, Daniel Weiskopf, and Klaus Engel, *Real-time volume graphics*, pp. 81–102, A. K. Peters, Ltd., Natick, MA, USA, 2006.

[KFCO+07] Johannes Kopf, Chi-Wing Fu, Daniel Cohen-Or, Oliver Deussen, Dani Lischinski, and Tien-Tsin Wong, *Solid texture synthesis from 2d exemplars*, ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007) **26** (2007), no. 3, 2:1–2:9.

[LEB11] Lasse Farnung Laursen, Bjarne Kjær Ersbøll, and Jakob Andreas Bærentzen, *Anisotropic 3d texture synthesis with application to volume rendering*, 2011.

[LS99] Daniel D. Lee and H. Sebastian Seung, *Learning the parts of objects by non-negative matrix factorization*, Nature **401** (1999), 788–791.