Image Quilting for Texture Synthesis and Transfer

Alexei A. Efros & William T. Freeman

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: A texture synthesis algorithm should be able to take a sample of texture and generate an unlimited amount of image data which, while not exactly like the original, will be perceived by humans to be *the same texture*. Furthermore, it would be useful to be able transfer the texture from one object to anther (e.g. the ability to cut and paste material properties on arbitrary objects). In this work we present an extremely simple algorithm for texture synthesis. We synthesize new texture by taking patches of existing texture and stitching them together in a consistent way.

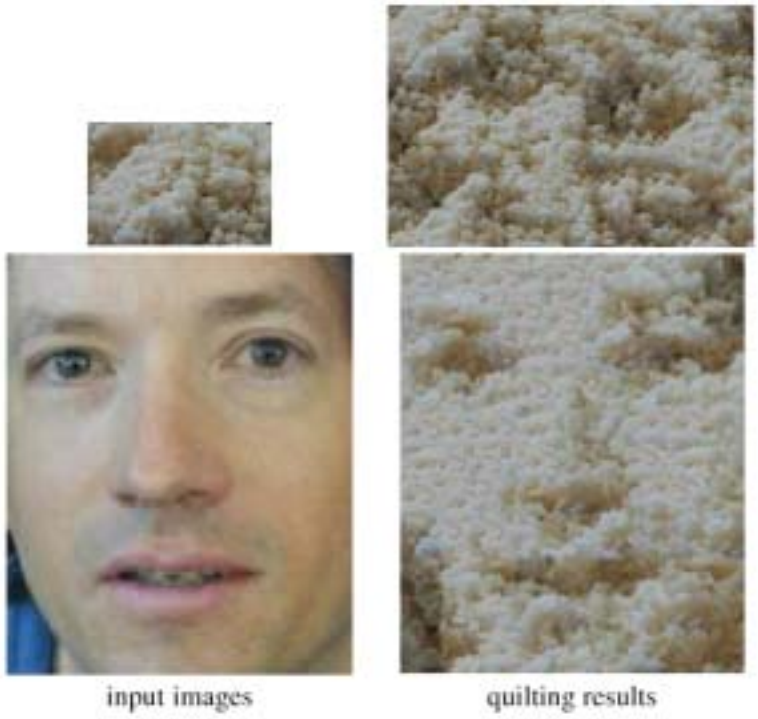

Figure 1: Demonstration of quilting for texture synthesis and texture transfer. Using the rice texture (upper left), we can synthesize more such texture (upper right). We can also transfer the rice texture onto another image (lower left) for a strikingly different result.

Previous Work: With Heeger and Bergen [4] the quality of texture synthesis reached a level acceptable for use in computer graphics. Their main insight was to model texture in terms of histograms of filter responses at multiple scales and orientations. But these marginal statistics fail to capture important relationships across scales and orientations for structured textures. Several attempts have been made to extend this idea to capture a wider range of textures, including De Bonet [1] who samples from conditional distribution over multiple scales, and later Portilla et.al. [5] who match both first and second order properties of wavelet coefficients. While important from a theoretical point of view, neither method is successful at capturing local detail of many structured textures.

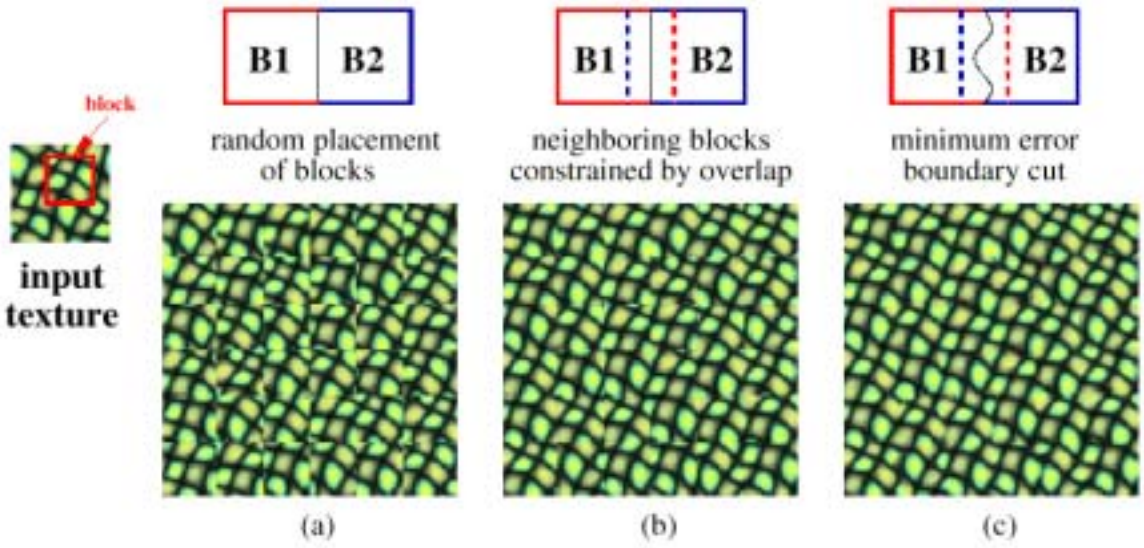

Figure 2: Quilting texture. Square blocks from the input texture are patched together to synthesize a new texture sample: (a) blocks are chosen randomly (similar to [7, 6]), (b) the blocks overlap and each new block is chosen so as to "agree" with its neighbors in the region of overlap, (c) to reduce blockiness the boundary between blocks is computed as a minimum cost path through the error surface at the overlap.

Approach–Image Quilting: Here, we outline our patch-based texture synthesis procedure, image quilting. To synthesize a new texture image, as a first step let us simply tile it with blocks taken randomly from S_B. The result shown on Figure 2(a) already looks somewhat reasonable, but for most structured textures it will be quite obvious that the blocks do not match.

As the next step, let us introduce some overlap in the placement of blocks onto the new image. Now, instead of picking a random block, we search for a block that by some measure agrees with its neighbors along the region of overlap. Figure 2(b) shows an improvement in result, however the edges between the blocks are still quite noticeable.

Finally, we will let the blocks have ragged edges which will allow them to better approximate the features in the texture. Now, before placing a chosen block into the texture we will look at the error in the overlap region between it and the other blocks. We find a minimum cost path through that error surface and declare that to be the boundary of the new block. Figure 2(c) shows the results of this simple modification.

Impact and Future Work: Despite its simplicity, this method [2] works amazingly well when applied to texture synthesis, producing results that are equal or better than the state-of-the-art [3] but at a fraction of the computational cost. We have also extended our method to the novel task of "texture transfer" with promising results. Fig. 1 shows the result of transferring the rice texture onto the man's face. We hope to extend this to higher-level visual objects beyond texture.

Research Support: This research was conducted at Mitsubishi Electric Research Laboratories, where AAE was a summer intern and WTF was a Senior Research Scientist.

References:

[1] J. S. De Bonet. Multiresolution sampling procedure for analysis and synthesis of texture images. In *SIGGRAPH 97*, pages 361–368, 1997.

[2] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. In *ACM SIGGRAPH*, 2001. In *Computer Graphics* Proceedings, Annual Conference Series.

[3] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *International Conference on Computer Vision*, pages 1033–1038, Corfu, Greece, September 1999.

[4] David J. Heeger and James R. Bergen. Pyramid-based texture analysis/synthesis. In *SIGGRAPH 95*, pages 229–238, 1995.

[5] 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):49–71, December 2000.

[6] Emil Praun, Adam Finkelstein, and Hugues Hoppe. Lapped textures. In *SIGGRAPH 00*, pages 465–470, 2000.

[7] Y. Xu, B. Guo, and H.-Y. Shum. Chaos mosaic: Fast and memory efficient texture synthesis. Technical Report MSR-TR-2000-32, Microsoft Research, April 2000.