Example-Based Image Synthesis
MIT 2001-07

Progress Report: January 1, 2002—June 30, 2002

William T. Freeman

Project Overview

The collaboration has grown as we have learned more about each others' research interests. The initial topic was example-based image synthesis: synthesizing image detail using a database of example images. This synthesis can be of texture, as in the Image Quilting technique, or of high-resolution image detail, as in Super-resolution.

After face-to-face meetings, we learned of a common research interest in shape estimation from images. Mr. Sato and Dr. Onozawa of NTT have developed laboratory equipment for analyzing object images under varying optical conditions, useful for shape estimation. Prof. Freeman and Dr. Torralba of MIT have developed a technique to improve shape estimates, using the image information and the rendering parameters learned from the initial shape estimate. These two research components are a great fit. We hope the NTT laboratory can use the technique and code we have developed, and we look forward to applying our shape estimation method to their data.

Progress Through June 2002

On example-based image synthesis: MIT graduate student Bryan Russell is seeking to extend example-based super-resolution methods to video. These methods use a database of examples of high- and low-resolution patches of the same image data. A Markov network model can assign the best high-resolution output patch for any given low-resolution input patch, taking both the training examples, and spatial context into account. Application of this to video data would allow low-resolution moving images to be displayed at high resolution.

The problem for video data: Inconsistent high-resolution details generated for different video frames create objectionable flicker artifacts in the synthesized sequence. Our solution: Mr. Russell and post-doc Dr. Torralba developed an algorithm to gate the synthesized high-resolution details by the amount of image change at each location. This causes static image regions to maintain constant high-resolution details, while the details synthesized in moving regions are allowed to change. This substantially reduces flicker artifacts, while adding high-resolution details to the moving images.

On shape recipes: Prof. Freeman and Dr. Torralba have developed a method to improve the accuracy of shape estimates. The idea is to use the initial shape estimate and the observed image to develop a local formula, or "recipe", for going from high-resolution image detail to high-resolution shape detail. This lets us exploit shape details captured in the image, but not in the initial shape reconstruction. We have illustrated this method by improving upon initial stereo depth estimates for a variety of objects. The stereo is accurate enough to capture low-resolution shape details, but not high-resolution ones. We implicitly learn the material and lighting properties of the scene from the image and the low-resolution shape details, learning a "recipe" to go from the low-resolution image to the low-resolution shape. We then apply that recipe to the high-resolution image details to infer high-resolution shape details. The result adds high-resolution shape details visible in the image, but not captured by the stereo.

We have written-up a conference submission about this work, which we have provided to Mr. Sato and Dr. Onozawa, and we have provided Mr. Sato with computer code.

Research Plan for the Next Six Months

On example-based image synthesis: Mr. Russell plans to incorporate optical flow measurements to create consistent high-resolution image detail even for patches of high image motion. This should improve the image quality of the synthesized high-resolution video data.

Mr. Sato and Dr. Onozawa have expressed interest in the Image Quilting texture synthesis method, for incorporation with their synthetic reality work. We will be able to give them that code on a research basis, with approval from my former company. I am interested to help with any implementation problems that arise.

On shape recipes: At MIT, using computer graphics simulations, we will explore properties of the shape recipes method (when it works, when it breaks down, what order approximation is needed for it to work well under general conditions). For explorations of the method, using real data, we hope to use data from the laboratory set-up of Dr. Onozawa and Mr. Sato of NTT. We hope that Mr. Sato can use our technique and code to further improve the shape estimates he acquires using his optical imaging system.

We expect to plan out these collaboration issues during Mr. Sato's visit to MIT on August 19.