CSE 190A Research Proposal: Super-Resolution

Eric Christiansen
echristiansen@cs.ucsd.edu

Abstract

Super-resolution seeks to create a high-quality image of a scene from several low-quality images of the same scene. Though a literature of super-resolution applications exists, implementing super resolution successfully is apparently quite difficult. In fact, a student unsuccessfully attempted to implement super-resolution in a previous offering of this course. I will attempt the same task, but using a different approach that should yield better compartmentalization of bugs and inaccuricies than was enjoyed by the unsuccessful student.

1 Project description

Multiple-image super-resolution seeks to combine the information from several low-quality photographs of a scene to create a single high-quality picture. The field is mature enough that there exists an iPhone application, ClearCam by Occipital, that takes 6 images in quick succession, and then uses them to perform super resolution with about 40 seconds of processing on the iPhone's processor. Another approach to super resolution uses a single image and prior information about the scene to guess the values of missing pixels. This approach is called "hallucination", and will not be addressed in this project. In multiple-image super-resolution, the low quality pictures may be of low resolution, be subject to blurring caused by poor focus, and have noisy pixel values. David Capel's thesis [2] and a follow-up paper [1] describe a framework for performing super-resolution. In the first step, image registration, corresponding points in the low quality images are identified. This algorithm is described concisely in Figure 4 of [1]. These corresponding points are used to compute geometric and photometric transformations between the images. Here, a geometric transformation accounts for a change in the point of view of the camera, and a photometric transformation accounts for changes in illumination and exposure between shots. Once the corresponding points are identified, the main step in super-resolution is performed. The pixel values from the low quality images are represented as the results of a process that creates low quality images from image scenes, and inference is performed to find the ML or MAP value of pixels in the image scene given the low quality pixels.

In the winter 2006 offering of CSE 190A, Jesse Cirimele attempted to implement super-resolution [3]. He first implemented image registration as a prerequisite to performing super-resolution on real images, and appears not to have succeeded in implementing super-resolution by the end of the quarter. Given Jesse's lack of success, and the general difficulty of the problem as described to me by Serge Belongie, I plan to start with a simple subproblem of super-resolution before tackling the full problem. In particular, I plan to start with synthetic images, the corresponding points of which I control, so that I may test my super-resolution implementation independent of a potentially unreliable image registration subroutine. Once this simpler problem is solved, I will extend the implementation to deal with real images.

1.1 Dataset

Initially, I will create synthetic images similar to the synthetic images (the test cards) that Capel used in his thesis. These high-quality images will be warped and degraded according to the generative process assumed by the super-resolution technique. If I get super-resolution working with synthetic images, I will photograph a UCSD landmark with my iPhone and attempt to perform super-resolution on those images. Throughout the project, I will assume any two images are related by a homography. This simplifying assumption is appropriate for 2D objects and 3D objects viewed at a distance.

1.2 Evaluation

In the initial stages of the project, I will not be interested in quantitative evaluations of my super-resolution implementation. When super-resolution works, the effects are dramatic, as illustrated in Capel's thesis, and don't require a numerical analysis to identify. If I require a quantitative measure of success, I may use the performance of an OCR algorithm on super-resolved pictures of text as a proxy for the performancy of my implementation.

2 Milestones

I would like to have met the following milestones at the ends of the following weeks.

- Week 1: Decide on a particular super-resolution model. Capel has continued to publish on super-resolution, and other authors have done work in the area since Capel wrote his thesis. I will investigate some of the newer models, and choose one to use for the project. In particular, I plan to read [4], [6], and [5].
- Week 2: Create a synthetic dataset according to the generative model of the super-resolution technique I choose.
- Week 4: Implement and verify the super-resolution model.
- Week 5: Collect real images and investigate image registration techniques. Capel describes a technique in his thesis, but the intervening years may have yielded better approaches.
- Week 6: Implement an image registration technique.
- Week 7: Combine the image registration and super-resolution implementations, and verify with the real images.
- Week 10: Write up the results.

3 Qualifications

I am a PhD student in computer science, and have taken CSE 250A and CSE 250B, Saul's CSE 191 on high dimensional data analysis, stochastic processes (MATH 285), and convex optimization (ECE 273). Additionally, I attended the 2009 Machine Learning Summer School in Cambridge, England, which exposed me to techniques for performing efficient approximate inference in Bayesian models. As an undergrad, I worked for Gary Cottrell on cognitive models of vision processing.

References

[1] D. Capel and A. Zisserman. Computer vision applied to super resolution. *IEEE Signal Processing Magazine*, 20(3):75–86, 2003.

[2] David Capel. *Image Mosaicing and Super-Resolution (Cphc/Bcs Distinguished Dissertations.)*. SpringerVerlag, 2004.

[3] Jesse Cirimele. Cse 190a: Super resolution. `http://jec190a.blogspot.com/`.

[4] S. Farsiu, M. Elad, and P. Milanfar. A practical approach to super-resolution. In *Proceedings of SPIE*, volume 6077, pages 24–38. Citeseer, 2006.

[5] L.C. Pickup, D.P. Capel, S.J. Roberts, and A. Zisserman. Bayesian methods for image super-resolution. *The Computer Journal*, 52(1):101, 2009.

[6] M. Tipping and C. Bishop. Bayesian image super-resolution. *Advances in Neural Information Processing Systems*, 15:1303–1310, 2003.