

## Super-resolution

One would like to have an intelligent method for expanding the resolution of an image. It should keep edges which are implicitly described in the low resolution image sharp. It should make intelligent guesses about the details of textures.



We have developed a Bayesian method to estimate the detail components of a scene, given the low resolution components. The result of our method is shown at left. The input image was the 70x70 pixel image at top left. A standard method to zoom the image to 280x280 resolution, cubic spline interpolation, gives the result at the top right. Textures and edges are necessarily made blurry. The 280x280 output of our new method is shown at the bottom left. Note that sharp, linear details in the hair are properly guessed, as is the eyelid. The overall impression of sharpness seems more like that of the actual 280x280 version of the image, shown at the bottom right. This new method for super-resolution may have applications to many sorts of digital image manipulations: reproduction of photographs, television display, image enhancement, etc.

**Background and objectives:** There is much still and moving image content at a low resolution. Much current NTSC programming may be desired to play on future high definition television (HDTV) players. By the time HDTV sets are commonplace, consumers may have come to expect that level of resolution quality. (Just as consumers expect to see color images now, instead of the old black and white ones).

**Technical discussion:** We use a training based approach. We examine many pairs of high resolution, and low resolution versions of the same image data. We divide each image into patches, both high resolution and low-resolution patches. We describe the patches as vectors in a continuous space, and model the probability densities as gaussian mixtures. (We reduced the dimensionality of the scene and image data within each patch by principal components analysis). We had approximately 20,000 patch samples from our training data, and typically used 9 dimensional representations for both the low-resolution patches (7x7 pixels) and the high resolution patches (3x3 pixels).

Each patch of the low and high resolution images is a node in a Markov network. Given some new image, we seek to infer the corresponding high resolution image components. During inference, we evaluate the prior and conditional distributions of the high resolution data, given the low resolution observation. The high resolution components are a sampling of those high resolution components which correspond to the observed low resolution components at that node. We think of it as a "lineup of suspects". Each node has its own set of suspects. Each scene in a node's lineup has in common the fact that it renders to the low-resolution observation at that node. We evaluate the likelihoods by a set of belief propagation equations. The computation converges in just 3 iterations. The iterations themselves take about 5 seconds each. However, the set-up time prior to beginning the computation takes about 1 hour. We hope to reduce that time with future research.

[TR2001-30](#)

*Example-based super-resolution* (a summary of our work in this area)

William T. Freeman, Thouis R. Jones, and Egon C. Pasztor.

[TR2000-05](#)

*Learning low-level vision*

(longer journal version, Intl. Journal of Computer Vision, 40(1), pp. 25-47, 2000) William T. Freeman, Egon C. Pasztor, and Owen T. Carmichael

[TR99-12](#)

*Learning low-level vision*

(conference version, Intl. Conf. on Computer Vision, Corfu, Greece, 1999) William T. Freeman, Egon C. Pasztor

[TR99-08](#)

*Markov networks for low-level vision*

William T. Freeman, Egon C. Pasztor

[TR99-05](#)

*Learning to estimate scenes from images*

William T. Freeman, Egon C. Pasztor,  
Neural Information Processing Systems 11, 1998.