Noise Reduction Through Detection of Signal Redundancy

Jeremy S. De Bonet

The Problem:

The transmission of information usually results in the corruption of the desired signal by noise in the transmitter, channel and receiver. In the general case it is impossible to recover the original signal from only the noisy signal which has received. For many of the types of information about which we care, such as sound and images, we have prior information about the types of signals we expect to see. How to characterize these priors and how to best use them to recover the transmitted signal is a difficult and important problem.

Motivation:

During the transmission of images, noise is often introduced. For analog television transmission noise is introduced in the form of “snow.” For digital imagery, error-checking and correction schemes eliminate most transmission errors, however to achieve this, other more systematic compression artifacts are introduced. In both cases we would like to “clean up” the images and recover a signal which is closer to the truth. Further we want to recover a signal which is simultaneously quantitatively closer to the truth, in terms of its actual pixel values, and which is qualitatively close to the perceptual effect of the original.

Previous Work:

The previous approach to image denoising is to take advantage of the prior belief that the noisy signal actually represents a natural image [4, 5]. In general for natural images, it turns out that the distribution of multiresolution wavelet coefficients is highly kurtotic, which indicates that coefficients are most often zero but in rare cases take on very large values. For white-noise however, this distribution is Gaussian. The difference in these distributions can be used as the basis of noise reduction algorithms, by reducing the wavelet coefficients which are more likely to be noise than signal. Such denoising relies on a generic image model which is weak and will therefore not remove all possible noise.

Approach:

In contrast to earlier methods, which rely on a generic image model, we construct a probability density model from the noisy image itself. To do this we take advantage of the prior belief that within natural images, there are regions which have the same visual appearance. For example image regions within a single solid patch, or regions along border between two patches are examples of regions which are similar using a very simple notion of visual appearance. By averaging these regions, we can improve the strength of the underlying signal.

Using more sophisticated notions of visual appearance, developed in [2, 1] we can extract redundancy from regions within image patches with the same texture, or along borders between such patches. Further, applying this notion recursively any of the visual structure within an image which occurs more than once can be used to mutually reinforce its strength everywhere.

The success of this approach is directly related to the redundancy in the image. If the redundancy in the image is very low, then no reinforcement will occur and the noise will persist. But if there is some redundancy in the image — that might arise from complex reoccuring structures or a simple smoothly varying patch — the signal will be strengthened within these regions.

Difficulty:

In Figure 1 we show results of this denoising approach: (a) the original image; (b) the image corrupted with white gaussian noise (SNR 8.9 dB); (c) the results of denoising using wavelet shrinkage or coring [4, 5] (SNR 9.8 dB); (d) the residual error between the denoised result in (c) and the original (a) — notice that the error contains a great deal of interpretable structure; (e) our denoising approach (SNR 13.2 dB); and (f) the residual error between (e) and (a) — these errors are far smaller than (d) and are much less structured.

In Figure 1 (g) and (h) we show results of using this denoising approach to remove the artifacts caused by extreme JPEG compression. On the left, the Lena image was JPEG compressed to 0.118 bits per pixel, on

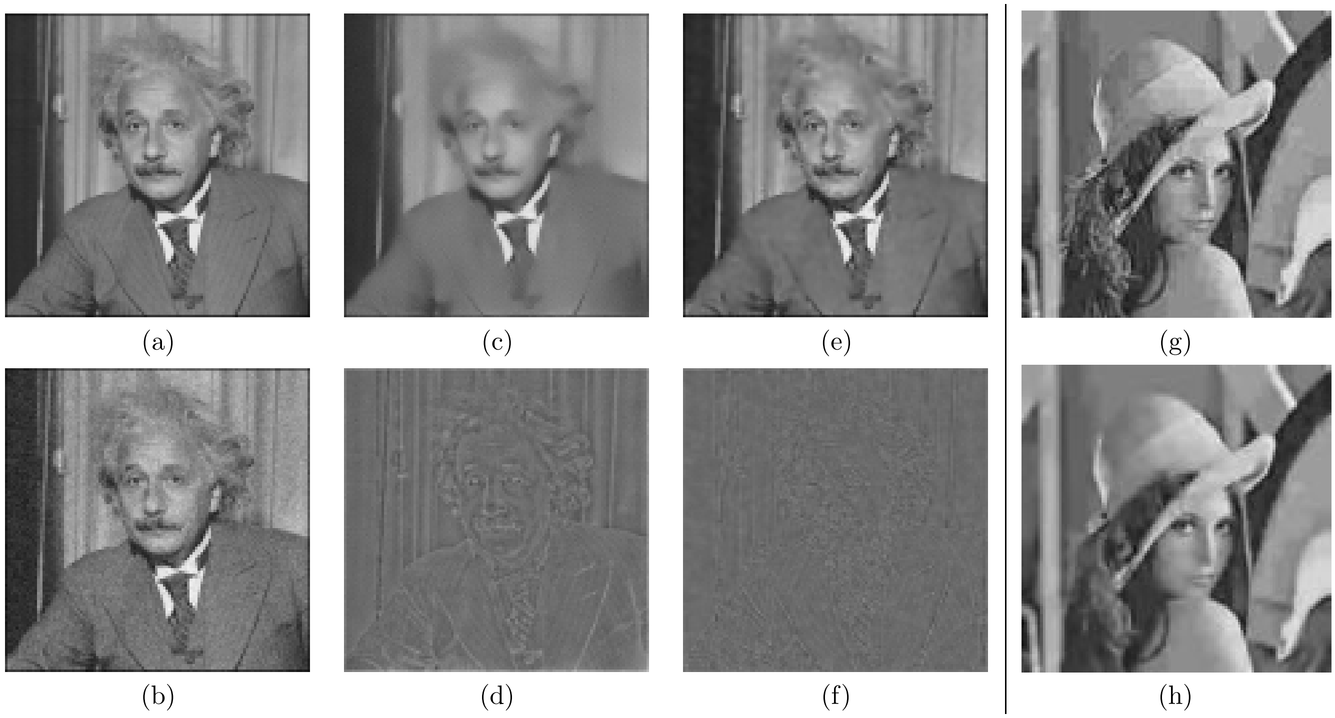

Figure 1: For description see text.

the right, we show the denoised version. Since only the image on the left was used as input, the image on the right can also be encoded with only 0.118 bits per pixel.

Impact:

The explicit use of signal redundancy to remove noise is a new and powerful technique. It is capable of better noise reduction than the other state-of-the-art techniques. Furthermore, it is independent of these other techniques, and can therefore be used *in conjunction* with them, producing results which outperform each technique separately.

Future work:

When considering video streams, there is additional redundancies in the visual structures across different frames. By further developing these techniques to take advantage of this redundancy, we hope to explore video denoising, including ultra-low bit-rate video compression schemes.

Research Support:

Research supported in part by DARPA contract No. N00014-94-1-0128.

References:

[1] J. S. De Bonet. Flexible histograms: Multiresolution texture discrimination model. Unpublished. http://www.ai.mit.edu/people/jsd/Research/Publications/FlexHist.ps, April 1997.

[2] J. S. De Bonet. Multiresolution sampling procedure for analysis and synthesis of texture images. In *Computer Graphics*. ACM SIGGRAPH, 1997.

[4] Donoho, D. L. and Johnstone, I. M.1993. Ideal spatial adaptation by wavelet shrinkage. Technical Report 400, Stanford Statistics Department.

[5] Simoncelli, E. P. and Adelson, E. H.1996. Noise removal via bayesian wavelet coring. In *IEEE Third Int'l Conf on Image Processing*, Laussanne Switzerland. IEEE.