Exploring the use of wavelet thresholding for denoising images

Imola K. Fodor and Chandrika Kamath
Lawrence Livermore National Laboratory

October 17, 2001

Scientific data sets collected by sensors are generally contaminated by noise. Imperfect instruments, problems with the data acquisition process, and interfering natural phenomena, can all degrade the target data of interest. Cleaning, or denoising, is often a necessary first step in mining such noisy data sets.

A traditional way to remove noise from image data is to employ spatial filters as described in the signal processing literature. Recently, several new techniques have been proposed in various scientific communities that claim to improve on spatial filters by denoising more effectively while better preserving the edges in the data. Despite numerous publications, there seems to be a lack of available software and guidance about the relative advantages and disadvantages of the different methods. We attempt to address this drawback in a small way by exploring various wavelet-based denoising techniques and assessing their performance for different types of data sets and noise distributions.

In this presentation, we concentrate on denoising of images based on thresholding of wavelet coefficients. These methods transform the data into a wavelet basis, threshold the wavelet coefficients, then transform back the thresholded coefficients into the original domain to obtain the denoised data. Under certain conditions, the "large" coefficients in the wavelet domain correspond to the signal, while the "small" ones represent mostly noise. One obtains different denoisers depending on the wavelet transform, the number of multiresolution levels, the thresholding function (which specifies how to apply the threshold to the wavelet coefficients), the thresholding rule (which specifies how to calculate the threshold), and in certain cases, the noise estimate (which specifies how to estimate the usually unknown level of the noise) used. To explore these different options, our software includes various wavelets (e.g. daublets, symmlets, bi-orthogonal), thresholding functions (e.g. hard, soft, semisoft, garrote), noise estimates (e.g. median absolute deviation of the diagonal detail coefficients on the first multiresolution level), and thresholding rules (e.g. universal, top, minimizing false discovery rate, SUREShrink, hypothesis testing, BayesShrink). We extend the denoising methods to images by providing global (one thresholding rule and function for all the coefficients), level-dependent (separate thresholding rules and/or functions for the coefficients on the different multiresolution levels), and subband-dependent (separate thresholding rules and/or functions for the coefficients on the different subbands) implementations.

We report our results on test images contaminated with synthetic noise. In our experiments, the BayesShrink and the SUREShrink methods outperform (in terms of MSE and visual quality) the other wavelet-based techniques we considered. The best wavelet-based methods, however, for most test images and noise levels, resulted in slightly higher MSE values than a combination of traditional spatial filters.

UCRL-JC-145671 ABS. This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract W-7405-ENG.