Image Smoothness and Denoising

Mentor: Yeon Hyang Kim

Introduction to Image Processing: Prepared by Narfi Stefansson

When doing image compression, we take a wavelet description of an image

$$f = \sum c_{\psi_{j,k}} \psi_{j,k}$$

and we alter the coefficients $c_{\psi_{j,k}}$ to get new coefficients $\tilde{c}_{\psi_{j,k}}$ such that:

1. Most of the coefficients $\tilde{c}_{\psi_{j,k}}$ are zero.
2. Those coefficients $\tilde{c}_{\psi_{j,k}}$ that are non-zero are stored with limited finite precision.

After that, one needs to somehow encode the location of the nonzero coefficients $\tilde{c}_{\psi_{j,k}}$ as well as their values and write that information in a binary string in an efficient way. We will not touch this second step in this project. (The 2nd step is treated in the image1 project)

As for the question about how to choose the coefficients $\tilde{c}_{\psi_{j,k}}$, one can investigate the following scheme:

Let N and let $\tilde{c}_{\psi_{j,k}}$ be such that N of the wavelet coefficients are nonzero. Let us choose the non-zero coefficients in such a manner that the L_2 error:

$$\|f - \sum \tilde{c}_{\psi_{j,k}} \psi_{j,k}\|_{L_2}$$

is minimal.

This is easy to do: simply select the N largest coefficients $c_{\psi_{j,k}}$.

Alternatively, one can let $1 < p < \infty$ and choose $\tilde{c}_{\psi_{j,k}}$ such that

$$\|f - \sum \tilde{c}_{\psi_{j,k}} \psi_{j,k}\|_{L_p}$$

is minimal. In this case we would call this "compression in L_p". And it is not much harder than compression in L_2, we rescale the coefficients $c_{\psi_{j,k}}$ in a simple way and select the N largest rescaled coefficients.

Now, we can do this selection of $\tilde{c}_{\psi_{j,k}}$ for various values of N. And then we can look at the behavior of the error

$$E(p, N) := \|f - \sum \tilde{c}_{\psi_{j,k}} \psi_{j,k}\|_{L_p}$$

as N increases. From the numbers $E(p, N)$ we can calculate the so called Besov space smoothness of the image f, and its Besov space norm. These two are believed to characterize the image quite well: The number of edges and corners and the amount of texture in the image as well as the amount of information in the image.

A number of questions arise immediately. One is: How do we know in which L_p we should compress our images? That is something that depends on our perceived visual quality of the compressed image and cannot be quantified in a simple manner.

Another question is: How does this relate to other uses of wavelets in image processing? In particular, how does this relate to image denoising? This is the subject of our next section.

Introduction to image denoising

In practice, image denoising is used when you somehow obtain a noisy image, and you want to remove as many of the speckles from the image as possible, without removing or distorting any features in it.

However, in order to quantify the performance of the various denoising algorithms, you start with an high quality image, f, and you add noise to it. This would be normally distributed noise with standard deviation σ. We'll call the noisy image $\bar{f}$. A denoising algorithm receives the noisy image $\bar{f}$ as input, and maybe σ as well, and outputs an image $\tilde{f}$, which is hopefully close to the original image, f.

Now we have a concrete way of evaluating the performance of denoising algorithms. We can e.g. measure

$$\|f - \tilde{f}\|_{L_2}.$$

Wavelet denoising consists of writing

$$\bar{f} = \sum \bar{c}_{\psi_{j,k}} \psi_{j,k}$$

and then processing the coefficients $\bar{c}_{\psi_{j,k}}$ in some way, usually through soft thresholding:

$$\eta_S(c,t) := \begin{cases} c+t & \text{if } c < -t \\ 0 & \text{if } |c| < t \\ c-t & \text{if } c > t \end{cases}$$

We then threshold all the wavelet coefficients, $\bar{c}_{\psi_{j,k}}$ and let $\tilde{c}_{\psi_{j,k}} := \eta_S(\bar{c}_{\psi_{j,k}}, t)$ and then

$$\tilde{f} := \sum \tilde{c}_{\psi_{j,k}} \psi_{j,k}$$

Now the question is: How should we choose the threshold t? Should we perhaps let the threshold depend on j or k? Perhaps soft thresholding is not optimal?

Hopefully this introduction is sufficient for you to choose among the projects. You will of course get more information once you have made your choice.

Reference:

Antonin Chambolle, Ronald A. DeVore, Namyong Lee, and Bradley J. Lucier, "Nonlinear Wavelet Image Processing: Variational Problems, Compression, and Noise Removal through Wavelet Shrinkage", IEEE Transactions on Image Processing, 7 (1998), 319-335

Level I. Write a MATLAB function that uses the wavelet toolbox and estimates the smoothness of images in Besov spaces and computes their norm in $B^\alpha_\tau(L_\tau)$.

Process a large class of images of various sizes and with various systems and compare the results.

How well does the calculated smoothness compare to smoothness as your eyes perceive it?
Turn in both an electronic(yeon@cs.wisc.edu) and a hard copy of the report.

Level III. In this project, you add noise to images and try to use thresholding to remove it. Chambolle, DeVore, Lee and Lucier suggested a method for choosing the threshold. It is rather recent (published in 1998) and it is not as simple as some of the other choices of threshold. In particular, it requires knowledge of the $B^\alpha_\tau(L_\tau)$ norm of the original image.

Use various systems (the biorthogonal 6/10 and 7/9 are the first candidates, then Haar and Daubechies 4 and 8). You should record a number of visual qualities in the denoised images (say, blurring of edges and quality of texture as well as a few others that you think of). Also record the norm of the difference of the original image and the denoised image in L_2 and a few L_p for $1 < p < 2$. Turn in both an electronic copy (to yeon@cs.wisc.edu) and a hard copy of the report.