24-505/705: Variational Image Processing
Assignment 3: Image Denoising – Computation
Due 27 October 2004

In the second assignment, you studied the problem of removing noise from an image without blurring sharp edges. The problem setup is repeated here for convenience. Recall that this can be formulated as an infinite-dimensional minimization problem. Given a possibly noisy initial image $u_0(x, y)$, we would like to find the improved image $u(x, y)$ that is closest in the L_2 sense, i.e. we want to minimize

$$\mathcal{F}_{LS} := \frac{1}{2} \int_\Omega (u - u_0)^2 \, d\boldsymbol{x},$$

while also removing noise, which is assumed to comprise very "rough" components of the image. This latter goal can be incorporated as an additional "penalty" term in the objective, either as the Tikhonov regularization

$$\mathcal{R}_{TN} := \beta \int_\Omega \nabla u \cdot \nabla u \, d\boldsymbol{x},$$

or the total variation (TV) regularization

$$\mathcal{R}_{TV}^{\varepsilon} := \frac{\beta}{2} \int_\Omega (\nabla u \cdot \nabla u + \varepsilon)^{\frac{1}{2}} \, d\boldsymbol{x},$$

where β is a constant that controls how strongly we impose the penalty, and ε is a small positive constant added to insure differentiability of $\mathcal{R}_{TV}^{\varepsilon}$ wherever $\nabla u = \boldsymbol{0}$.

We wish to study the performance of the two denoising functionals $\mathcal{F}_{TN}$ and $\mathcal{F}_{TV}^{\varepsilon}$, where

$$\mathcal{F}_{TN} := \mathcal{F}_{LS} + \mathcal{R}_{TN},$$

$$\mathcal{F}_{TV}^{\varepsilon} := \mathcal{F}_{LS} + \mathcal{R}_{TV}^{\varepsilon},$$

with the homogeneous Neumann boundary condition $\nabla u \cdot \boldsymbol{n} = 0$,

1. Using the minimal surface code that I demoed in class as a guide (which you can download from the class website, `http://www-2.cs.cmu.edu/~42-705`; look for `ms.cpp`, `ms-h.cpp`, `ms-hh.cpp`, and `ms-hh-picard.cpp` under *News*), and the expressions you derived on the last assignment for the Newton step, implement a Newton solver for this problem. You can use Alex's utilities for reading/writing an image in `.pnm` format. See the class website for instructions on how to obtain the utilities. Your output should include the enhanced image (which you should print out in your assignment) as well as the L_2 norm of the "error" between the original (not the noisy) image and your enhanced image.

2. Test your code on the *Four Circles* image using a 64×64 mesh. You can download the images from the class website. Compare the use of both Tikhonov and TV regularization. In both cases try to recover the sharpest image you can by experimenting with β until you find the "best" value. You can fix $\varepsilon = 0.001$ for this purpose. What is the effect of varying β and ε on the convergence of Newton's method and the quality of the enhanced image?

3. Test you code on the *Lena* image. There are seven different noisy images available; everyone should attempt the 20% noise level image (`lena20`) and at least one other image. Again try to find the sharpest image you can by varying β, and fix $\varepsilon = 0.001$.

4. Implement a so-called *Picard* method to solve this problem, in which the Hessian of the TV regularization functional is approximated by dropping the tensor part of diffusion coefficient $\mathcal{A}$ (see previous assignment). This generally leads to improved numerical behavior, since the approximate Hessian is much better conditioned than the exact one. Note that you cannot rely on Sundance's linearization to create the Newton equation for you, since it will generate the exact Hessian; you must do that manually. You can use my Picard minimal surface code as a starting point.