Unconstrained optimization

Optimization with Applications to Image Processing*

12/04/2012

1 Rosenbrock Banana

The goal of this set of exercises is to compare different optimization techniques, in order to understand the fundamental differences between relaxation, fixed stepsize gradient descent and gradient descent with optimal stepsize. For such a comparison, we consider the *Rosenbrock Banana* function,

$$J(x, y) = (x-1)^2 + 10(x^2 - y)^2.$$

- Is this function convex?
- In what point does J reach its minimum?
- Compute $\nabla J(x, y)$ and $\nabla^2 J(x, y)$.
- Visualize J in 3D using Matlab's function *meshgrid*.
- Visualize J in 2D as an image, using using Matlab's function *imagesc*.
- Comment on this. Explain why function J is not easy to minimize.

2 Preliminaries: optimization in one dimension

The one dimensional case is a special case in optimization, basically because $\mathbb{R}$ is an ordered set. Several optimization methods (the steepest descent with optimal stepsize, for instance) are based on one dimensional optimization. Write a function that takes as argument a point $\mathbf{x} = (x, y)$ and a vector $\mathbf{d}$, and computes the minizer of the function $t \mapsto J(\mathbf{x} + t\mathbf{d})$ using Newton's method.

3 Minimization of J

The goal of this exercise is to compare the following methods, and their performance when applied to function J:

- Relaxation method
- Gradient descent with fixed stepsize
- Gradient descent with optimal stepsize
- Newton's method
- Conjugate gradient method (Polack-Ribire version).

1. Implement these methods, and apply them to function J, starting from point $(-1, 1)$. For each method, show the evolution of the minimization procedure: display a figure with level lines and the trajectory described by the points that are successively computed by the method.
2. Compare the obtained trajectories, and the methods performance (number of iterations, time complexity, computational burden, etc.).

*Exercises by J.-F. Aujol

4 Application to image restoration: Tikhonov regularization

4.1 Discretization

A digital image will be represented as a 2D vector of size $N \times N$. We denote by X the Euclidean space $\mathbb{R}^{N\times N}$, and by $Y = X \times X$. We endow X with the scalar product:

$$< u, v >_X = \sum_{1\le i,j\le N} u_{i,j} v_{i,j}$$

and the norm

$$\|u\|_X = \sqrt{< u, u >_X}.$$

We will consider the following discrete version of the gradient operator. Given $u \in X$, its gradient ∇u is an element of Y given by:

$$(\nabla u)_{i,j} = \left((\nabla u)^1_{i,j}, (\nabla u)^2_{i,j}\right),$$

where

$$(\nabla u)^1_{i,j} = \begin{cases} u_{i+1,j} - u_{i,j} & \text{if } i < N \\ 0 & \text{if } i = N, \end{cases}$$

and

$$(\nabla u)^2_{i,j} = \begin{cases} u_{i,j+1} - u_{i,j} & \text{if } j < N \\ 0 & \text{if } j = N. \end{cases}$$

For the divergence, given $p \in Y$, we define

$$(\text{div}(p))_{i,j} = \begin{cases} p^1_{i,j} - p^1_{i-1,j} & \text{if } 1 < i < N \\ p^1_{i,j} & \text{if } i = 1 \\ -p^1_{i-1,j} & \text{if } i = N \end{cases} + \begin{cases} p^2_{i,j} - p^2_{i,j-1} & \text{if } 1 < j < N \\ p^2_{i,j} & \text{if } j = 1 \\ -p^2_{i,j-1} & \text{if } j = N. \end{cases}$$

As for the Laplacian, we will consider the discrete version given by the composition:

$$\Delta u = \text{div} \nabla u.$$

1. Implement all these discrete operators.
2. To test your implementation, choose an image and display its vertical and horizontal gradients, its gradient norm and its laplacian.

4.2 Solving by PDE

Consider the following problem:

$$\inf_u \left\{\|f - u\|^2_X + 2\lambda \|\nabla u\|^2_X\right\}. \tag{1}$$

Compute the corresponding Euler-Lagrange equation, then find the minimizer using a fixed stepsize gradient descent method.

4.3 Solving by Fourier transform

We recall that the DFT of a digital image $f(m, n)$, $(m, n) \in \{0, 1, \ldots, N - 1\}^2$, and its inverse, are is given by:

$$\mathcal{F}(f)(p, q) = F(p, q) = \sum_{m=0}^{N-1}\sum_{n=0}^{N-1} f(m, n) \exp\left(-i(2\pi/N)pm\right) \exp\left(-j(2\pi/N)qn\right),$$

$$f(m, n) = \frac{1}{N^2}\sum_{p=0}^{N-1}\sum_{q=0}^{N-1} F(p, q) \exp\left(i(2\pi/N)pm\right) \exp\left(j(2\pi/N)qn\right).$$

where $(p, q) \in \{0, 1, \ldots, N - 1\}^2$. We have also Parseval and Plancherel relations:

$$\|\mathcal{F}(f)\|^2_X = N^2\|f\|^2_X, \quad < \mathcal{F}(f), \mathcal{F}(g) >_X = N^2 < f, g >_X.$$

1. Using Parseval relation, show that the solution u of (1) satisfies:

$$\mathcal{F}(u)(p, q) = \frac{\mathcal{F}(f)(p, q)}{1 + 8\lambda(\sin^2 \frac{\pi p}{N} + \sin^2 \frac{\pi q}{N})}.$$

2. Code this new version. Consider extending the image by symmetry, to reduce aliasing effects.

5 Application to image restoration: ϕ regularization

5.1 Denoising

We now consider the following problem: Consider the following problem:

$$\inf_u \left\{ \|f - u\|_X^2 + \lambda \int \phi(|\nabla u|) \right\}. \tag{2}$$

1. Derive the corresponding Euler-Lagrange equation.
2. Solve (3) using a gradient descent method. Consider the following ϕ functions: $\phi(t) = t$, $\phi(t) = t^2/(1+t^2)$, $\phi(t) = \log(1+t^2)$, $\phi(t) = 2\sqrt{1+t^2} - 2$.
3. Solve (3) using a half quadratic minimization for different ϕ functions. Comments?

5.2 Deconvolution

Repeat the previous exercise for the deconvolution problem:

$$\inf_u \left\{ \|f - Au\|_X^2 + \lambda \int \phi(|\nabla u|) \right\}. \tag{3}$$

For the numerical implementation, consider for A a Gaussian kernel.