

9. MINIMIZATION PROBLEMS

We consider two classes of problems: the first is the unconstrained minimization problem (UCMP) and the second is nonlinear least squares (NLLS).

UCMP: Given $f : \mathbb{R}^N \rightarrow \mathbb{R}$, find $\mathbf{x} \in \mathbb{R}^N$ which minimizes $f(\mathbf{x})$.

NLLS: Given $\mathbf{F} = (f_1, \dots, f_m)^T : \mathbb{R}^N \rightarrow \mathbb{R}^m$, with $m \geq N$, find $\mathbf{x} \in \mathbb{R}^N$ which minimizes $\phi(\mathbf{x}) = (1/2) \sum_{k=1}^m [f_k(\mathbf{x})]^2$.

Note that the second problem is a special case of the first, but with more structure.

9.1. Newton's method and steepest descent. To solve UCMP, we can look for \mathbf{x}^* at which

$$\nabla f = (\partial f / \partial x_1, \dots, \partial f / \partial x_N) = \mathbf{0}$$

and the Hessian matrix $H_f = (\partial^2 f / \partial x_i \partial x_j)$ is symmetric and positive definite. Thus the problem becomes one of solving a nonlinear systems of equations $\mathbf{F}(\mathbf{x}) = 0$, where $\mathbf{F} = \nabla f$, i.e., $\mathbf{F}_i = \partial f / \partial x_i$.

If we apply Newton's method, we get the iteration

$$\mathbf{x}^{n+1} = \mathbf{x}^n - J_{\nabla f}(\mathbf{x}^n)^{-1} \nabla f(\mathbf{x}^n) = \mathbf{x}^n - H_f(\mathbf{x}^n)^{-1} \nabla f(\mathbf{x}^n),$$

since $(J_{\nabla f})_{ij} = \partial F_i / \partial x_j$ and hence

$$(J_{\nabla f})_{ij} = \frac{\partial}{\partial x_j} \frac{\partial f}{\partial x_i} = \frac{\partial^2 f}{\partial x_i \partial x_j} = H_f.$$

Note that H_f is symmetric, so for a minimum at \mathbf{x}^* , it will be sufficient to have $H_f(\mathbf{x}^*)$ to be positive definite.

In the special case of UCMP where $f = (1/2)\mathbf{x}^T A \mathbf{x} - \mathbf{x}^T b$, we considered methods of the form $\mathbf{x}^{n+1} = \mathbf{x}^n + \alpha_n \mathbf{p}^n$, where \mathbf{p}^n is a search direction and α_n is a scalar. In the method of steepest descent, we choose $\mathbf{p}^n = -\nabla f(\mathbf{x}^n)$. With this choice and the special form of f , we can solve explicitly for the best choice of α_n . In the more general case, we would consider methods of the form

$$\mathbf{x}^{n+1} = \mathbf{x}^n - \alpha_n \nabla f(\mathbf{x}^n),$$

where α_n is chosen to guarantee that $f(\mathbf{x}^{n+1}) < f(\mathbf{x}^n)$. The advantage of this approach is that we do not need to compute the Hessian. However, the method converges slowly. To compromise, we could consider methods of the form

$$\mathbf{x}^{n+1} = \mathbf{x}^n - \alpha_n B_n^{-1} \nabla f(\mathbf{x}^n),$$

where B_n is symmetric and positive definite and α_n is chosen to insure that $f(\mathbf{x}^{n+1}) < f(\mathbf{x}^n)$. For example, let $B_n = (H_f(\mathbf{x}^n) + \mu_n I)$, with $\mu_n > 0$ chosen so that B_n is positive definite. For large values of μ_n this method behaves like steepest descent and for small values of μ_n , it behaves like Newton's method. The difficult part is deciding how to choose the parameter μ_n .

9.2. Quasi-Newton methods. Analogous to the case of quasi-Newton methods for nonlinear equations, we now wish to avoid computation of the Hessian at each iteration. In this case, we want to generate a sequence of symmetric, positive definite matrices B_n such that B_n approximates $H_f(\mathbf{x}^n)$, but can be computed easily from B_{n-1} . Since we are now solving the system $\nabla f(\mathbf{x}) = \mathbf{0}$, the appropriate Taylor series expansion is:

$$\nabla f(\mathbf{x}^n) = \nabla f(\mathbf{x}^{n+1}) + H_f(\mathbf{x}^{n+1})(\mathbf{x}^n - \mathbf{x}^{n+1}) + O(\mathbf{x}^n - \mathbf{x}^{n+1})^2.$$

Thus, we want the approximation B_{n+1} to $H_f(\mathbf{x}^{n+1})$ to satisfy the quasi-Newton equation

$$\nabla f(\mathbf{x}^n) = \nabla f(\mathbf{x}^{n+1}) + B_{n+1}(\mathbf{x}^n - \mathbf{x}^{n+1}).$$

To simplify notation, let

$$\mathbf{y}^n = \nabla f(\mathbf{x}^{n+1}) - \nabla f(\mathbf{x}^n), \quad \mathbf{s}^n = \mathbf{x}^{n+1} - \mathbf{x}^n,$$

so the quasi-Newton equation is $B_{n+1}\mathbf{s}^n = \mathbf{y}^n$. If we look for a symmetric, single rank (the maximum number of linearly independent rows is one) update satisfying the quasi-Newton equation, then provided $(\mathbf{y}^n - B_n\mathbf{s}^n)^T\mathbf{s}^n \neq 0$, the only one is given by

$$B_{n+1} = B_n + \frac{(\mathbf{y}^n - B_n\mathbf{s}^n)(\mathbf{y}^n - B_n\mathbf{s}^n)^T}{(\mathbf{y}^n - B_n\mathbf{s}^n)^T\mathbf{s}^n}.$$

It turns out that this method does not work well, so we look for a double rank update. It can be shown that the general symmetric rank 2 update is given by:

$$B_{n+1} = B_n + \frac{(\mathbf{y}^n - B_n\mathbf{s}^n)(\mathbf{c}^n)^T + \mathbf{c}^n(\mathbf{y}^n - B_n\mathbf{s}^n)^T}{(\mathbf{c}^n)^T\mathbf{s}^n} - \frac{(\mathbf{y}^n - B_n\mathbf{s}^n)^T\mathbf{s}^n\mathbf{c}^n(\mathbf{c}^n)^T}{[(\mathbf{c}^n)^T\mathbf{s}^n]^2},$$

where \mathbf{c} is an arbitrary vector such that $(\mathbf{c}^n)^T\mathbf{s}^n \neq 0$.

If we choose $\mathbf{c}^n = \mathbf{s}^n$, we get the Powell symmetric Broyden update. Since we would also like to have the property that B_n positive definite implies that B_{n+1} is positive definite, a better choice is $\mathbf{c}^n = \mathbf{y}^n$ (called the Davidon-Fletcher-Powell method).

Finally, as in the case of nonlinear equations, instead of updating B_n and then having to solve a linear system of equations at each step, we can update the inverse directly. If we let $H_n = B_n^{-1}$, we then get the iteration $\mathbf{x}^{n+1} = \mathbf{x}^n - \alpha_n H_n \nabla f(\mathbf{x}^n)$, where

$$H_{n+1} = H_n + \frac{(\mathbf{s}^n - H_n\mathbf{y}^n)(\mathbf{s}^n)^T + \mathbf{s}^n(\mathbf{s}^n - H_n\mathbf{y}^n)^T}{(\mathbf{s}^n)^T\mathbf{y}^n} - \frac{(\mathbf{s}^n - H_n\mathbf{y}^n)^T\mathbf{y}^n\mathbf{s}^n(\mathbf{s}^n)^T}{[(\mathbf{s}^n)^T\mathbf{y}^n]^2}.$$

This method in which the inverse is updated directly is known as the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method.

9.3. Nonlinear least squares. Finally, we consider the nonlinear least squares problem of minimizing $\phi(\mathbf{x}) = (1/2) \sum_{k=1}^m [f_k(\mathbf{x})]^2$. Let $\mathbf{F} = [f_1, \dots, f_m]^T$. Then J_F , the Jacobian matrix of partial derivatives of \mathbf{F} is given by $[J_f(\mathbf{x})]_{kj} = (\partial f_k / \partial x_j)$.

Now

$$\frac{\partial \phi}{\partial x_j} = \sum_{k=1}^m f_k(\mathbf{x}) \frac{\partial f_k}{\partial x_j} = \{[J_f(\mathbf{x})]^T \mathbf{F}(\mathbf{x})\}_j, \quad (H_\phi)_{ij} = \frac{\partial^2 \phi}{\partial x_i \partial x_j} = \sum_{k=1}^m \left[f_k(\mathbf{x}) \frac{\partial^2 f_k}{\partial x_i \partial x_j} + \frac{\partial f_k}{\partial x_i} \frac{\partial f_k}{\partial x_j} \right].$$

Hence, $\nabla\phi(\mathbf{x}) = J_F(\mathbf{x})^T \mathbf{F}(\mathbf{x})$ and

$$H_\phi(\mathbf{x}) = J_F(\mathbf{x})^T J_F + \sum_{k=1}^m f_k(\mathbf{x}) H_{f_k}(\mathbf{x}).$$

Newton's method for the system $\nabla\phi(\mathbf{x}) = J_F(\mathbf{x})^T \mathbf{F}(\mathbf{x}) = 0$, is then given by

$$\begin{aligned} \mathbf{x}^{n+1} &= \mathbf{x}^n - [H_\phi(\mathbf{x}^n)]^{-1} \nabla\phi(\mathbf{x}^n). \\ &= \mathbf{x}^n - [J_F(\mathbf{x}^n)^T J_F + \sum_{k=1}^m f_k(\mathbf{x}^n) H_{f_k}(\mathbf{x}^n)]^{-1} [J_F(\mathbf{x}^n)]^T \mathbf{F}(\mathbf{x}^n). \end{aligned}$$

Since this is fairly complicated to compute, we seek a simpler method. One approach is to linearize $f_k(\mathbf{x})$ about \mathbf{x}^n and minimize the resulting quadratic functional instead. Using Taylor series, we approximate

$$f_k(\mathbf{x}) \approx f_k(\mathbf{x}^n) + [\nabla f_k(\mathbf{x}^n)]^T (\mathbf{x} - \mathbf{x}^n),$$

since we expect the remainder to be small if $\mathbf{x} - \mathbf{x}^n$ is small. Inserting this approximation,

$$\phi(\mathbf{x}) \approx (1/2) \sum_{k=1}^m [f_k(\mathbf{x}^n)]^2 + \sum_{k=1}^m f_k(\mathbf{x}^n) [\nabla f_k(\mathbf{x}^n)]^T (\mathbf{x} - \mathbf{x}^n) + (1/2) \sum_{k=1}^m \{[\nabla f_k(\mathbf{x}^n)]^T (\mathbf{x} - \mathbf{x}^n)\}^2.$$

Then

$$\nabla\phi(\mathbf{x}) \approx \sum_{k=1}^m f_k(\mathbf{x}^n) \nabla f_k(\mathbf{x}^n) + \sum_{k=1}^m \nabla f_k(\mathbf{x}^n) [\nabla f_k(\mathbf{x}^n)]^T (\mathbf{x} - \mathbf{x}^n).$$

Since we want to find \mathbf{x} such that $\nabla\phi(\mathbf{x}) = \mathbf{0}$, we choose \mathbf{x}^{n+1} to satisfy

$$\sum_{k=1}^m \nabla f_k(\mathbf{x}^n) [\nabla f_k(\mathbf{x}^n)]^T (\mathbf{x}^{n+1} - \mathbf{x}^n) = - \sum_{k=1}^m f_k(\mathbf{x}^n) \nabla f_k(\mathbf{x}^n).$$

This may be written in the form

$$J_F^T(\mathbf{x}^n) J_F(\mathbf{x}^n) (\mathbf{x}^{n+1} - \mathbf{x}^n) = -J_F^T(\mathbf{x}^n) \mathbf{F}(\mathbf{x}^n).$$

Hence, this approximation amounts to dropping the term $\sum_{k=1}^m f_k(\mathbf{x}) H_{f_k}(\mathbf{x})$ in Newton's method. We would expect this to be small if the minimum of ϕ is near zero. This method is known as the Gauss-Newton method.

In practice, a modified version of Gauss-Newton, known as the Levenberg-Marquardt method, is used. This method is given by the iteration

$$[\alpha_n I + J_F^T(\mathbf{x}^n) J_F(\mathbf{x}^n)] (\mathbf{x}^{n+1} - \mathbf{x}^n) = -J_F^T(\mathbf{x}^n) \mathbf{F}(\mathbf{x}^n),$$

where $\alpha_n \geq 0$ is an appropriately chosen scalar. The idea is that for α_n large, the method behaves like steepest descent (to ensure that the function $\phi(\mathbf{x})$ is being decreased), while for α_n small, it behaves more like Newton's method, which will converge faster. In general, we would increase α_n if we take a step that increases $\phi(\mathbf{x})$, and decrease α_n if $\phi(\mathbf{x})$ is decreasing.