

1 Minimization

1.1 Introduction

Minimization (maximization) is the problem of finding the minimum (maximum) of a given—generally non-linear—real valued function $\phi(\mathbf{x})$ of an n -dimensional argument $\mathbf{x} \doteq \{x_1, \dots, x_n\}$. The function is often called the *objective function* or the *cost function*.

Minimization is a simpler case of a more general problem—*optimization*—which includes finding the best available values of the objective function within a given domain and/or subject to given constraints.

Minimization is not unrelated to root-finding: at the minimum all partial derivatives of the objective function vanish,

$$\frac{\partial \phi}{\partial x_i} = 0 \Big|_{i=1 \dots n}, \quad (1)$$

and one can alternatively solve this system of (non-linear) equations.

1.2 Local minimization

Local minimization refers to a group of algorithms that move from one candidate solution to another candidate solution by applying local changes and moving “downhill” until a solution deemed optimal is found (or the allotted time is elapsed).

1.2.1 Newton’s method

Newton’s method is based on the quadratic approximation of the objective function $\phi(\mathbf{x})$ in the vicinity of the suspected minimum,

$$\phi(\mathbf{x} + \Delta \mathbf{x}) \approx \phi(\mathbf{x}) + \nabla \phi(\mathbf{x})^\top \Delta \mathbf{x} + \frac{1}{2} \Delta \mathbf{x}^\top \mathbf{H}(\mathbf{x}) \Delta \mathbf{x}, \quad (2)$$

where the vector $\nabla \phi(\mathbf{x})$ is the gradient of the objective function at the point \mathbf{x} ,

$$\nabla \phi(\mathbf{x}) \doteq \left\{ \frac{\partial \phi(\mathbf{x})}{\partial x_i} \right\}_{i=1 \dots n}, \quad (3)$$

and $\mathbf{H}(\mathbf{x})$ is the *Hessian matrix*—a square matrix of second-order partial derivatives of the objective function at the point \mathbf{x} ,

$$\mathbf{H}(\mathbf{x}) \doteq \left\{ \frac{\partial^2 \phi(\mathbf{x})}{\partial x_i \partial x_j} \right\}_{i,j \in 1 \dots n}. \quad (4)$$

The minimum of the quadratic form (2), as function of $\Delta \mathbf{x}$, is found at the point where its gradient with respect to $\Delta \mathbf{x}$ vanishes,

$$\nabla \phi(\mathbf{x}) + \mathbf{H}(\mathbf{x}) \Delta \mathbf{x} = 0. \quad (5)$$

This gives an approximate step towards the minimum, called the *Newton's step*,

$$\Delta \mathbf{x} = -\mathbf{H}(\mathbf{x})^{-1} \nabla \phi(\mathbf{x}). \quad (6)$$

The original Newton's method is simply the iteration,

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \mathbf{H}(\mathbf{x}_k)^{-1} \nabla \phi(\mathbf{x}_k), \quad (7)$$

where at each iteration the full Newton's step is taken and the Hessian matrix is recalculated. In practice, instead of calculating \mathbf{H}^{-1} one rather solves the linear equation (5).

Usually the Newton's method is modified to take a smaller step \mathbf{s} ,

$$\mathbf{s} = \lambda \Delta \mathbf{x}, \quad (8)$$

with $0 < \lambda < 1$. The factor λ can be found by a backtracking algorithm similar to that in the Newton's method for root-finding. One starts with $\lambda = 1$ and then backtracks, $\lambda \leftarrow \lambda/2$, until the *Armijo condition*,

$$\phi(\mathbf{x} + \mathbf{s}) < \phi(\mathbf{x}) + \alpha \mathbf{s}^T \nabla \phi(\mathbf{x}), \quad (9)$$

is satisfied (or the minimal λ (say, $1/1024$) is reached, in which case the step is taken unconditionally). The parameter α can be chosen as small as 10^{-4} .

1.2.2 Numerical calculation of gradient and Hessian matrix

The *forward* finite-difference approximation to the derivative $f'(x)$ of a function $f(x)$ is given as

$$f'(x) \approx \frac{f(x + \delta x) - f(x)}{\delta x}, \quad (10)$$

where δx is usually chosen as $(1 + |x|)\sqrt{\epsilon}$ where ϵ is the machine precision (for "double" numbers $\sqrt{\epsilon} = 2^{-26}$). This gives the following approximation for the gradient,

$$\nabla \phi(x)_i = \frac{\partial \phi(x)}{\partial x_i} = \frac{\phi(x_1, \dots, x_i + \delta x_i, \dots, x_n) - \phi(x_1, \dots, x_n)}{\delta x_i}, \quad (11)$$

where $\delta x_i = (1 + |x_i|)\sqrt{\epsilon}$.

Similarly, the Hessian matrix is given as second forward finite difference,

$$\frac{\partial^2 \phi}{\partial x_i \partial x_j} = \frac{\partial}{\partial x_i} \nabla \phi_j = \frac{\nabla \phi(x_1, \dots, x_i + \delta x_i, \dots, x_n)_j - \nabla \phi(x_1, \dots, x_n)_j}{\delta x_i} \quad (12)$$

where the step-size δx_i however must be larger than in the first finite difference, for example, $\delta x_i = (1 + |x_i|)\epsilon^{1/4}$.

The *central* (possibly better) finite difference formula for the first derivative is given as,

$$f'(x) \approx \frac{f(x + \delta x) - f(x - \delta x)}{2\delta x}. \quad (13)$$

This gives the following approximation for the gradient,

$$\nabla\phi(x)_i = \frac{\partial\phi(x)}{\partial x_i} = \frac{\phi(x_1, \dots, x_i + \delta x_i, \dots, x_n) - \phi(x_1, \dots, x_i - \delta x_i, \dots, x_n)}{2\delta x_i}, \quad (14)$$

and the following approximation for the Hessian matrix,

$$\mathbf{H}_{jk} = \frac{\partial^2\phi}{\partial x_j \partial x_k} \approx \frac{\phi(\mathbf{x} + \delta\mathbf{x}_k + \delta\mathbf{x}_j) - \phi(\mathbf{x} + \delta\mathbf{x}_k - \delta\mathbf{x}_j) - \phi(\mathbf{x} - \delta\mathbf{x}_k + \delta\mathbf{x}_j) + \phi(\mathbf{x} - \delta\mathbf{x}_k - \delta\mathbf{x}_j)}{4\delta x_k \delta x_j}. \quad (15)$$

where $\delta\mathbf{x}_k$ is a vector in the direction k with the length $\delta x_k = (1 + |x_k|)\sqrt{\epsilon}$ where ϵ is the *machine epsilon*. In practice the calculation of the gradient should reuse the ϕ -values from the calculation of the Hessian matrix.

1.2.3 Quasi-Newton methods

Quasi-Newton methods are variations of the Newton's method which attempt to avoid recalculation of the Hessian matrix at each iteration, trying instead certain updates based on the analysis of the gradient vectors. The update $\delta\mathbf{H}$ is usually chosen to satisfy the condition

$$\nabla\phi(\mathbf{x} + \mathbf{s}) = \nabla\phi(\mathbf{x}) + (\mathbf{H} + \delta\mathbf{H})\mathbf{s}, \quad (16)$$

called *secant equation*, which is the Taylor expansion of the gradient.

The secant equation is under-determined in more than one dimension as it consists of only n equations for the n^2 unknown elements of the update $\delta\mathbf{H}$. Various quasi-Newton methods use different choices for the form of the solution of the secant equation.

In practice one typically uses the inverse Hessian matrix (often—but not always—denoted as \mathbf{B}) and applies the updates directly to the inverse matrix thus avoiding the need to solve the linear equation (5) at each iteration.

For the inverse Hessian matrix the secant equation (16) reads

$$(\mathbf{B} + \delta\mathbf{B})\mathbf{y} = \mathbf{s}, \quad (17)$$

or, in short,

$$\delta\mathbf{B}\mathbf{y} = \mathbf{u}, \quad (18)$$

where $\mathbf{B} \doteq \mathbf{H}^{-1}$, $\mathbf{y} \doteq \nabla\phi(\mathbf{x} + \mathbf{s}) - \nabla\phi(\mathbf{x})$, and $\mathbf{u} \doteq \mathbf{s} - \mathbf{B}\mathbf{y}$.

One usually starts with the identity matrix as the zeroth approximation for the inverse Hessian matrix and then applies the updates.

If the minimal λ (say, $1/1024$) is reached during the backtracking line-search—which might be a signal of lost precision in the approximate (inverse) Hessian matrix—it is advisable to reset the current inverse Hessian matrix to identity matrix.

Table 1.2.3 lists one possible algorithm of the quasi-newton method with updates.

Table 1: Quasi-newton minimisation algorithm with updates.

```

set the inverse Hessian matrix to unity,  $B = 1$ 
repeat until converged (e.g.  $\|\nabla\phi\| < \text{tolerance}$ ) :
    calculate the Newton's step  $\Delta\mathbf{x} = -B\nabla\phi$ 
    if not  $\nabla\phi \cdot \Delta\mathbf{x} < 0$  :
        reset  $B = 1$  and  $\Delta\mathbf{x} = -\nabla\phi$ 
    do linesearch starting with  $\lambda = 1$  :
        if  $\phi(\mathbf{x} + \lambda\Delta\mathbf{x}) < \phi(\mathbf{x})$  or  $\lambda < \lambda_{\min}$  :
            accept the step  $\mathbf{x} = \mathbf{x} + \lambda\Delta\mathbf{x}$ 
            update  $B = B + \delta B$ 
            break linesearch
    backtrack  $\lambda / = 2$ 

```

Broyden's update The Broyden's update is chosen in the form

$$\delta B = \mathbf{c}\mathbf{s}^T. \quad (19)$$

where the vector \mathbf{c} is found from the condition (18),

$$\mathbf{c} = \frac{\mathbf{u}}{\mathbf{s}^T\mathbf{y}}. \quad (20)$$

Sometimes the dot-product $\mathbf{s}^T\mathbf{y}$ becomes very small or even zero which results in serious numerical difficulties. One can avoid this by only performing update if the condition $|\mathbf{s}^T\mathbf{y}| > \epsilon\|\mathbf{s}\|\|\mathbf{y}\|$ is satisfied where ϵ is a small number, say 10^{-6} .

Symmetric Broyden's (aka BFGS) update The Broyden's update (19) is not symmetric (while the Hessian matrix should be) which is an obvious drawback. Therefore a better approximation might be the symmetric Broyden's update,

$$\delta B = \mathbf{a}\mathbf{s}^T + \mathbf{s}\mathbf{a}^T. \quad (21)$$

The vector \mathbf{a} is again found from the condition (18),

$$\mathbf{a} = \frac{\mathbf{u} - \gamma\mathbf{s}}{\mathbf{s}^T\mathbf{y}}, \quad (22)$$

where $\gamma = (\mathbf{u}^T\mathbf{y})/(\mathbf{s}^T\mathbf{y})$.

Again one only performs the update if $|\mathbf{s}^T\mathbf{y}| > \epsilon\|\mathbf{s}\|\|\mathbf{y}\|$.

This update is generally known as Broyden–Fletcher–Goldfarb–Shanno (BFGS) update.

SR1 update The symmetric-rank-1 update (SR1) is chosen in the form

$$\delta B = \mathbf{v}\mathbf{v}^T, \quad (23)$$

where the vector \mathbf{v} is again found from the condition (16), which gives

$$\delta B = \frac{\mathbf{u}\mathbf{u}^T}{\mathbf{u}^T\mathbf{y}}. \quad (24)$$

Again, one only performs the update if denominator is not too small, that is, $|\mathbf{u}^T\mathbf{y}| > \epsilon\|\mathbf{u}\|\|\mathbf{y}\|$.

Other popular updates The wikipedia article “Quasi-Newton method” list several other popular updates.

1.2.4 Downhill simplex method

The *downhill simplex method* [1] (also called “Nelder-Mead” or “amoeba”) is a commonly used minimization algorithm where the minimum of a function in an n -dimensional space is found by transforming a simplex—a polytope with $n+1$ vertexes—according to the function values at the vertexes, moving it downhill until it converges towards the minimum.

The advantages of the downhill simplex method is its stability and the lack of use of derivatives. However, the convergence is reatively slow as compared to Newton’s methods.

In order to introduce the algorithm we need the following definitions:

- Simplex: a figure (polytope) represented by $n+1$ points, called vertexes, $\{\mathbf{p}_1, \dots, \mathbf{p}_{n+1}\}$ (where each point \mathbf{p}_k is an n -dimensional vector).
- Highest point: the vertex, \mathbf{p}_{hi} , with the highest value of the function: $\phi(\mathbf{p}_{\text{hi}}) = \max_k \phi(\mathbf{p}_k)$.
- Lowest point: the vertex, \mathbf{p}_{lo} , with the lowest value of the function: $\phi(\mathbf{p}_{\text{lo}}) = \min_k \phi(\mathbf{p}_k)$.
- Centroid: the center of gravity of all points, except for the highest: $\mathbf{p}_{\text{ce}} = \frac{1}{n} \sum_{(k \neq \text{hi})} \mathbf{p}_k$

The simplex is moved downhill by a combination of the following elementary operations:

1. Reflection: the highest point is reflected against the centroid, $\mathbf{p}_{\text{hi}} \rightarrow \mathbf{p}_{\text{re}} = \mathbf{p}_{\text{ce}} + (\mathbf{p}_{\text{ce}} - \mathbf{p}_{\text{hi}})$.
2. Expansion: the highest point reflects and then doubles its distance from the centroid, $\mathbf{p}_{\text{hi}} \rightarrow \mathbf{p}_{\text{ex}} = \mathbf{p}_{\text{ce}} + 2(\mathbf{p}_{\text{ce}} - \mathbf{p}_{\text{hi}})$.
3. Contraction: the highest point halves its distance from the centroid, $\mathbf{p}_{\text{hi}} \rightarrow \mathbf{p}_{\text{co}} = \mathbf{p}_{\text{ce}} + \frac{1}{2}(\mathbf{p}_{\text{hi}} - \mathbf{p}_{\text{ce}})$.
4. Reduction: all points, except for the lowest, move towards the lowest points halving the distance. $\mathbf{p}_{k \neq \text{lo}} \rightarrow \frac{1}{2}(\mathbf{p}_k + \mathbf{p}_{\text{lo}})$.

Table 2 shows one possible algorithm for the downhill simplex algorithm.

Table 2: Downhill simplex (Nelder-Mead) algorithm

```

REPEAT :
  find highest, lowest, and centroid points of the simplex
  try reflection
  IF  $\phi(\text{reflected}) < \phi(\text{lowest})$  :
    try expansion
    IF  $\phi(\text{expanded}) < \phi(\text{reflected})$  :
      accept expansion
    ELSE :
      accept reflection
  ELSE :
    IF  $\phi(\text{reflected}) < \phi(\text{highest})$  :
      accept reflection
    ELSE :
      try contraction
      IF  $\phi(\text{contracted}) < \phi(\text{highest})$  :
        accept contraction
      ELSE :
        do reduction
UNTIL converged (e.g. size(simplex) < tolerance)

```

1.2.5 Gauss-Newton algorithm

The Gauss-Newton algorithm is designed to minimize an objective function $\phi(\mathbf{c})$ that is given as a sum of squares of several (non-linear) functions $r_i(\mathbf{c})$,

$$\phi(\mathbf{c}) = \sum_{i=1}^n r_i^2(\mathbf{c}), \quad (25)$$

where $\{c_{k=1\dots m}\}$ is the set of parameters of the objective function. In particular, the algorithm can be used to solve a non-linear least squares curve fitting problem where the function to minimize is given as

$$\chi^2(\mathbf{c}) = \sum_{i=1}^n \left(\frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \right)^2, \quad (26)$$

where $\{x_i, y_i \pm \delta y_i\}$ is the set of data to fit and $f(\mathbf{c}, x)$ is the fitting function that depends on a set of parameters \mathbf{c} .

The algorithm can also be used to find an approximate solution to an overdetermined (if $n > m$) system of non-linear equations

$$\mathbf{r}(\mathbf{c}) = 0. \quad (27)$$

Just like the Newton's method the algorithm relies on the Taylor expansion of the objective function in the vicinity of the suspected minimum,

$$\phi(\mathbf{c} + \Delta\mathbf{c}) \approx \phi(\mathbf{c}) + \mathbf{g}^\top \Delta\mathbf{c} + \frac{1}{2} \Delta\mathbf{c}^\top \mathbf{H} \Delta\mathbf{c}, \quad (28)$$

where the (size- m) gradient \mathbf{g} is given as

$$g_k = 2 \sum_{i=1}^n r_i \frac{\partial r_i}{\partial c_k} \quad (29)$$

and the (square $m \times m$) Hessian matrix \mathbf{H} is given as

$$H_{jk} = 2 \sum_{i=1}^n \left(\frac{\partial r_i}{\partial c_j} \frac{\partial r_i}{\partial c_k} + r_i \frac{\partial^2 r_i}{\partial c_j \partial c_k} \right). \quad (30)$$

Now, in the Gauss-Newton method one ignores the second-derivative term in (30) which results in the following approximation for the Hessian matrix,

$$\mathbf{H}_{jk} \approx 2\mathbf{J}^T \mathbf{J}, \quad (31)$$

where \mathbf{J} is the (tall $n \times m$) Jacobian matrix of the $\{r_i\}$ functions,

$$J_{ik} = \frac{\partial r_i}{\partial c_k}. \quad (32)$$

The approximation (31) may be valid in two cases,

1. The functions r_i are small in the vicinity of the minimum;
2. The functions r_i are only slightly non-linear such that the second derivatives are small in magnitude.

Using the Jacobian matrix the gradient of ϕ can be written as

$$\mathbf{g} = 2\mathbf{J}^T \mathbf{r}. \quad (33)$$

From here the algorithm proceeds as in the usual Newton's method: one finds the Newton's step,

$$\Delta \mathbf{c} = -\mathbf{H}^{-1} \mathbf{g} \approx -(\mathbf{J}^T \mathbf{J})^{-1} (\mathbf{J}^T \mathbf{r}), \quad (34)$$

and then does the backtracking line-search.

Note that $(\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T$ is the left pseudo-inverse of the matrix \mathbf{J} . Therefore the Newton's step of the Gauss-Newton method for the objective function $\sum_{i=1}^n r_i^2$ is equivalent to the step of the root-finding Newton's method for the system of equations $\mathbf{r}(\mathbf{c}) = 0$.

1.3 Uncertainties of nonlinear least squares fit parameters

The *non-linear least squares* fit is the process of fitting a curve (a mathematical function, $f(\mathbf{c}, x)$, where \mathbf{c} is the set of fitting parameters) to a set of data points with uncertainties, $\{x_i, y_i \pm \delta y_i\}$. However, unlike the ordinary least squares fit, where the fitting parameters enter linearly, in the non-linear least squares fit the parameters enter essentially non-linearly.

The fit is achieved by minimizing the sum of squares (hence the name) of the deviations of the curve from the data (called χ^2 in physics),

$$\chi^2(\mathbf{c}) = \sum_{i=1}^n \left(\frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \right)^2 \equiv \sum_{i=1}^n r_i^2(\mathbf{c}), \quad (35)$$

where

$$r_i(\mathbf{c}) \doteq \frac{f(\mathbf{c}, x_i) - y_i}{\delta y_i} \quad (36)$$

are the (weighted) residuals.

The χ^2 can be minimized in the space of the fitting parameters either using any of the general minimization algorithms or using the Gauss-Newton algorithm which is specifically designed for an objective function in the form of the sum of squares of some residuals.

The uncertainties of the fitting parameters can be estimated by i) Taylor expansion of χ^2 around the minimum; ii) linearizing the problem; iii) calculating the uncertainties using the same technique as for the ordinary least squares fit. In other words, we apply the Newton's method to find the solution of the minimization problem and then determine the uncertainties of the fitting parameters from the last Newton's step (the one that brings us to the minimum).

1.3.1 Linearization of nonlinear problem at minimum

The Newton's step $\Delta \mathbf{c}$ toward the minimum is found from the second order Taylor expansion of χ^2 ,

$$\chi^2(\mathbf{c} + \Delta \mathbf{c}) \approx \chi^2(\mathbf{c}) + \mathbf{g}^\top \Delta \mathbf{c} + \frac{1}{2} \Delta \mathbf{c}^\top \mathbf{H} \Delta \mathbf{c}, \quad (37)$$

where the gradient \mathbf{g} is given as

$$g_k = \frac{\partial \chi^2}{\partial c_k} = \sum_{i=1}^n 2r_i \frac{\partial r_i}{\partial c_k}. \quad (38)$$

and the Hessian matrix \mathbf{H} is given as

$$H_{jk} = \frac{\partial^2 \chi^2}{\partial c_j \partial c_k} = \sum_{i=1}^n \left(2 \frac{\partial r_i}{\partial c_j} \frac{\partial r_i}{\partial c_k} + 2r_i \frac{\partial^2 r_i}{\partial c_j \partial c_k} \right). \quad (39)$$

Close to the minimum the term with the second derivative can be (hopefully) neglected (since at the minimum $f(\mathbf{c}, x_i) \approx y_i$).

Introducing the (tall $n \times m$) Jacobian matrix \mathbf{J} of the residuals,

$$J_{ij} = \frac{\partial r_i}{\partial c_j} = \frac{1}{\delta y_i} \frac{\partial f(\mathbf{c}, x_i)}{\partial c_j}, \quad (40)$$

one can rewrite the gradient as

$$\mathbf{g} = 2\mathbf{J}^T \mathbf{r} \quad (41)$$

and the Hessian matrix as

$$\mathbf{H} = 2\mathbf{J}^T \mathbf{J} . \quad (42)$$

The corresponding Newton's step is determined by the equation

$$\mathbf{H}\Delta\mathbf{c} = -\mathbf{g} , \quad (43)$$

or

$$\mathbf{J}^T \mathbf{J} \Delta\mathbf{c} = -\mathbf{J}^T \mathbf{r} , \quad (44)$$

which gives the Newton's step to the minimum as

$$\Delta\mathbf{c} = -\mathbf{J}^{-1} \mathbf{r} \quad (45)$$

where

$$\mathbf{J}^{-1} = (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T \quad (46)$$

is the pseudo-inverse of the matrix \mathbf{J} .

1.3.2 Uncertainties of the fit parameters

Equation (45) defines $\Delta c_{k=1\dots m}$ as function of $y_{i=1\dots n}$. The question is, if y_i are determined with uncertainties δy_i , what are the uncertainties of Δc_k ?

The answer is given by the *propagation of uncertainty* rule which says that the (co)variances $\delta c_k \delta c_j$ are given as

$$\delta c_k \delta c_j = \sum_i \frac{\partial \Delta c_k}{\partial y_i} \frac{\partial \Delta c_j}{\partial y_i} \delta y_i \delta y_i = \sum_i (\mathbf{J}^{-1})_{ki} (\mathbf{J}^{-1})_{ji} . \quad (47)$$

In matrix notation the covariance matrix $\Sigma_{kj} = \delta c_k \delta c_j$ is given as

$$\Sigma = \mathbf{J}^{-1} \mathbf{J}^{-T} = (\mathbf{J}^T \mathbf{J})^{-1} . \quad (48)$$

The uncertainties of the fitting parameters are then given as the square roots of the diagonal elements of the covariance matrix,

$$\delta c_k = \sqrt{\Sigma_{kk}} . \quad (49)$$

Notice that within the approximation (42) (that should work well at the minimum) the covariance matrix is given via the inverse of the Hessian matrix, \mathbf{H}^{-1} , at the minimum,

$$\Sigma = (\mathbf{J}^T \mathbf{J})^{-1} = 2\mathbf{H}^{-1} , \quad (50)$$

which is the canonical textbook result. It can also be obtained from the Taylor expansion of the variation of χ^2 with respect to fit parameters at the minimum, where the gradient is zero,

$$\delta\chi^2 = \frac{1}{2} \delta\mathbf{c}^T \mathbf{H} \delta\mathbf{c} = \text{trace} \left(\frac{1}{2} \Sigma \mathbf{H} \right) . \quad (51)$$

The uncertainties of fit parameters are determined by a unit variation of χ^2 per degree of freedom (that is, per fit parameter). That is, the matrix inside the trace operator must be the unit $m \times m$ matrix. This gives

$$\Sigma = 2H^{-1} . \tag{52}$$

References

- [1] J.A. Nelder and R. Mead, "A simplex method for function minimization", Computer Journal 7 (1965) 308