

Minimization and optimization

Introduction

Minimization is a problem of finding the minimum (or the maximum) of a given, generally non-linear, real valued function $f(\mathbf{x})$ of an n -dimensional argument $\mathbf{x} \doteq \{x_1, \dots, x_n\}$.

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

Minimization is related to root-finding as at the minimum all partial derivatives of the objective function vanish,

$$\frac{\partial f}{\partial x_i} = 0. \quad (1)$$

One can alternatively solve this system of non-linear equations.

Local minimization

Newton's methods

Newton's method is based on the quadratic approximation of the objective function f in the vicinity of the suspected minimum,

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

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

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

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

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

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

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

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

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

The original Newton's method is simply the iteration,

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

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 Newton's method is modified to include a smaller step size, $\lambda \Delta\mathbf{x}$, with $0 < \lambda < 1$. The step-size 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*,

$$f(\mathbf{x} + \lambda \Delta\mathbf{x}) < f(\mathbf{x}) + \alpha \nabla f(\mathbf{x})^T \lambda \Delta\mathbf{x}, \quad (8)$$

is satisfied. The parameter α can be chosen as small as 10^{-4} .

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

$$\nabla f(\mathbf{x} + \mathbf{s}) = \nabla f(\mathbf{x}) + (\mathbf{H} + \delta\mathbf{H})\mathbf{s}, \quad (9)$$

called *secant equation*, which is the Taylor expansion of the gradient. This secant condition 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 quasi-Newton methods one often starts with a unity matrix as the zeroth approximation for the Hessian matrix and then applies the updates. In this case one can actually use the inverse Hessian matrix—thus avoiding the need to solve the linear equation at each iteration—and apply the updates directly to the inverse matrix.

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

$$\delta\mathbf{H} = \mathbf{c}\mathbf{s}^T. \quad (10)$$

where the vector \mathbf{c} is found from the condition (9). The resulting update is

$$\mathbf{H} \rightarrow \mathbf{H} + \frac{\mathbf{y} - \mathbf{H}\mathbf{s}}{\mathbf{s}^T\mathbf{s}}\mathbf{s}^T, \quad (11)$$

where $\mathbf{y} \doteq \nabla f(\mathbf{x} + \mathbf{s}) - \nabla f(\mathbf{x})$. The update for the inverse matrix is given as¹

$$\mathbf{H}^{-1} \rightarrow \mathbf{H}^{-1} + \frac{(\mathbf{s} - \mathbf{H}^{-1}\mathbf{y})\mathbf{s}^T\mathbf{H}^{-1}}{\mathbf{y}^T\mathbf{H}^{-1}\mathbf{s}}. \quad (12)$$

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

$$\delta\mathbf{H} = \mathbf{u}\mathbf{u}^T, \quad (13)$$

where the vector \mathbf{u} is found from the condition (9). The resulting update is

$$\mathbf{H} \rightarrow \mathbf{H} + \frac{(\mathbf{y} - \mathbf{H}\mathbf{s})(\mathbf{y} - \mathbf{H}\mathbf{s})^T}{(\mathbf{y} - \mathbf{H}\mathbf{s})^T\mathbf{s}}. \quad (14)$$

And for the inverse matrix

$$\mathbf{H}^{-1} \rightarrow \mathbf{H}^{-1} + \frac{(\mathbf{s} - \mathbf{H}^{-1}\mathbf{y})(\mathbf{s} - \mathbf{H}^{-1}\mathbf{y})^T}{(\mathbf{s} - \mathbf{H}^{-1}\mathbf{y})^T\mathbf{y}}. \quad (15)$$

DavidonFletcherPowell (DFP) update

$$\delta\mathbf{H}^{-1} = \frac{\mathbf{y}\mathbf{y}^T}{\mathbf{y}^T\mathbf{s}} - \frac{\mathbf{H}^{-1}\mathbf{s}\mathbf{s}^T\mathbf{H}^{-1}}{\mathbf{s}^T\mathbf{H}^{-1}\mathbf{s}} \quad (16)$$

¹using the Sherman-Morrison formula,

$$(A + uv^T)^{-1} = A^{-1} - \frac{A^{-1}uv^TA^{-1}}{1 + v^TA^{-1}u}.$$

Table 1: Downhill simplex (Nelder-Mead) algorithm

```

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

```

Downhill simplex method

The *downhill simplex method* (also called Nelder-Mead method or amoeba method) 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 relatively 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: $f(\mathbf{p}_{\text{hi}}) = \max_k f(\mathbf{p}_k)$.
- Lowest point: the vertex, \mathbf{p}_{lo} , with the lowest value of the function: $f(\mathbf{p}_{\text{lo}}) = \min_k f(\mathbf{p}_k)$.
- 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. $\mathbf{p}_{k \neq \text{lo}} \rightarrow \frac{1}{2}(\mathbf{p}_k + \mathbf{p}_{\text{lo}})$.

A C-implementation of the simplex operations is listed in Table 4.

Finally, Table 1 shows one possible algorithm for the downhill simplex method.

Global optimization

Global optimization is the problem of locating (a good approximation to) the global minimum of a given objective function in a search space large enough to prohibit exhaustive enumeration.

When only a small sub-space of the search space can be realistically sampled the stochastic methods usually come to the fore.

Table 2: Simulated annealing algorithm

```

state ← start_state
T ← start_temperature
energy ← E(state) + random_thermal_energy(T)
REPEAT :
  new_state ← neighbour(state)
  new_energy ← E(new_state) + random_thermal_energy(T)
  IF new_energy < energy :
    state ← new_state
    energy ← new_energy
  reduce_temperature_according_to_schedule(T)
UNTIL terminated

```

A good local minimizer converges to the nearest local minimum relatively fast, so one possible global minimizer can be constructed by simply starting the local minimizer from different random starting points.

In the following several popular global minimization algorithms are shortly described.

Simulated annealing

Simulated annealing is a stochastic metaheuristic algorithm for global minimization. The name and inspiration come from annealing — heating up and cooling slowly — in material science. The slow cooling allows a piece of material to reach a state with "lowest energy".

The objective function in the space of states is interpreted as some sort of potential energy and the states — the points in the search space — are interpreted as physical states of some physical system. The system attempts to make transitions from its current state to some randomly sampled nearest states with the goal to eventually reach the state with minimal energy — the global minimum.

The system is attached to a thermal reservoir with certain temperature. Each time the energy of the system is measured the reservoir supplies it with a random amount of thermal energy sampled from the Boltzmann distribution,

$$P(E) = Te^{-E/T}. \quad (17)$$

If the temperature equals zero the system can only make transitions to the neighboring states with lower potential energy. In this case the algorithm turns merely into a local minimizer with random sampling.

If temperature is finite the system is able to climb up the ridges of the potential energy — about as high as the current temperature — and thus escape from local minima and hopefully eventually reach the global minimum.

One typically starts the simulation with some finite temperature on the order of the height of the typical hills on the potential energy surface, letting the system to wander almost unhindered around the landscape with a good chance to locate if not the best then at least a good enough minimum. The temperature is then slowly reduced following some annealing schedule which may be supplied by the user but must end with $T = 0$ towards the end of the allotted time budget.

Table 2 lists one possible variant of the algorithm.

The function `neighbour(state)` should return a randomly chosen neighbour of the given state.

Downhill simplex method can incorporate simulated annealing by adding the stochastic thermal energy to the values of the objective function at the vertices.

Quantum annealing

Quantum annealing is a general global minimization algorithm which — like simulated annealing — also allows the search path to escape from local minima. However instead of the thermal jumps over the potential barriers quantum annealing allows the system to tunnel through the barriers.

In its simple incarnation the quantum annealing algorithm allows the system to attempt transitions not only to the nearest states but also to distant states within certain "tunneling distance" from the current state. The transition is accepted only if it reduces the potential energy of the system.

Table 3: Quantum annealing algorithm

```

state ← start_state
energy ← E(state)
R ← start_radius
REPEAT :
  new_state ← neighbour_within_radius(state, R)
  new_energy ← E(new_state)
  IF new_energy < energy :
    state ← new_state
    energy ← new_energy
  reduce_radius_according_to_schedule(R)
UNTIL terminated

```

At the beginning of the minimization procedure the tunnelling distance is large — on the order of the size of the region where the global minimum is suspected to be located — allowing the system to explore the region. The tunneling distance is then slowly reduced according to a schedule such that by the end of the allotted time the tunnelling distance reduces to zero at which point the system hopefully is in the state with minimal energy.

Evolutionary algorithms

Unlike annealing algorithms, which follow the motion of only one point in the search space, the evolutionary algorithms typically follow a set of points called a population of individuals. A bit like the downhill simplex method which follows the motion of a set of points — the simplex.

The population evolves towards more fit individuals where fitness is understood in the sense of minimizing the objective functions. The parameters of the individuals (for example, the coordinates of the points in multi-dimensional minimization of a continuous objective function) are called genes.

The algorithm proceeds iteratively with the population in each iteration called a generation. In each generation the fitness of each individual — typically, the value of the objective function — is evaluated and the new generation is generated stochastically from the gene pool of the current generation through crossovers and mutations such that the genes of more fit individuals have a better chance of propagating into the next generation.

Each new individual in the next generation is produced from a pair of "parent" individuals of the current generation. The use of two "parents" is biologically inspired, in practice more than two "parents" can be used as well. The parents for a new individual are selected from the individuals of the current generation through a fitness based stochastic process where fitter individuals are more likely to be selected.

The "child" individual shares many characteristics of its "parents". In the simplest case the "child" may get its genes by simply averaging the genes of its parents. Then a certain amount of mutations — random changes in the genes — are added to the "child's" genes.

Generation of "children" continues until the population of the new generation reaches the appropriate size after which the iteration repeats itself.

The algorithm is terminated when the fitness level of the population is deemed sufficient or when the allocated budget is exhausted.

Implementation in C

Table 4: C implementation of simplex operations

```

#include <math.h>
#define float double
void reflection
    (float* highest, float* centroid, int dim, float* reflected){
    for(int i=0;i<dim;i++) reflected[i]=2*centroid[i]-highest[i];
}
void expansion
    (float* highest, float* centroid, int dim, float* expanded) {
    for(int i=0;i<dim;i++) expanded[i]=3*centroid[i]-2*highest[i];
}
void contraction
    (float* highest, float* centroid, int dim, float* contracted){
    for(int i=0;i<dim;i++)
        contracted[i]=0.5*centroid[i]+0.5*highest[i];
}
void reduction( float** simplex, int dim, int lo){
    for(int k=0;k<dim+1;k++) if(k!=lo) for(int i=0;i<dim;i++)
        simplex[k][i]=0.5*(simplex[k][i]+simplex[lo][i]);
}
float distance(float* a, float* b, int dim){
    float s=0; for(int i=0;i<dim;i++) s+=pow(b[i]-a[i],2);
    return sqrt(s);
}
float size(float** simplex, int dim){
    float s=0; for(int k=1;k<dim+1;k++){
        float dist=distance(simplex[0], simplex[k], dim);
        if(dist>s) s=dist; }
    return s;
}

```

```

#define float double
void simplex_update(float** simplex, float* f_values, int d,
    int* hi, int* lo, float* centroid)
{
    *hi=0; float highest=f_values[0];
    *lo=0; float lowest =f_values[0];
    for(int k=1;k<d+1;k++) {
        float next=f_values[k];
        if(next>highest){ highest=next; *hi=k;}
        if(next<lowest) {lowest=next; *lo=k;}
    }
    for(int i=0;i<d;i++) {
        float s=0; for(int k=0;k<d+1;k++) if(k!=*hi) s+=simplex[k][i];
        centroid[i]=s/d; }
}
void simplex_initiate(
    float fun(float*), float** simplex, float* f_values, int d,
    int* hi, int* lo, float* centroid)
{
    for(int k=0;k<d+1;k++) f_values[k]=fun(simplex[k]);
    simplex_update(simplex, f_values, d, hi, lo, centroid);
}

```

Table 5: C implementation of downhill simplex algorithm

```

#include <simplex.h>
#include <stdlib.h>
#include <stdio.h>
int downhill_simplex(
    double F(double*), double**simplex, int d, double simplex_size_goal)
{
    int hi, lo, k=0; double centroid[d], F_value[d+1], p1[d], p2[d];
    simplex_initiate(F, simplex, F_value, d, &hi, &lo, centroid);
    while( size(simplex, d) > simplex_size_goal){
        simplex_update(simplex, F_value, d, &hi, &lo, centroid);
        reflection(simplex[hi], centroid, d, p1); double f_re=F(p1);
        if(f_re < F_value[lo]){ // reflection looks good: try expansion
            expansion(simplex[hi], centroid, d, p2); double f_ex=F(p2);
            if(f_ex < f_re){ // accept expansion
                for(int i=0; i<d; i++) simplex[hi][i]=p2[i]; F_value[hi]=f_ex;
            }
            else{ // reject expansion and accept reflection
                for(int i=0; i<d; i++) simplex[hi][i]=p1[i]; F_value[hi]=f_re;
            }
        }
        else{ // reflection wasn't good
            if(f_re < F_value[hi]){ // ok, accept reflection
                for(int i=0; i<d; i++) simplex[hi][i]=p1[i]; F_value[hi]=f_re;
            }
            else{ // try contraction
                contraction(simplex[hi], centroid, d, p1); double f_co=F(p1);
                if(f_co < F_value[hi]){ // accept contraction
                    for(int i=0; i<d; i++) simplex[hi][i]=p1[i]; F_value[hi]=f_co;
                }
                else{ // do reduction
                    reduction(simplex, d, lo);
                    simplex_initiate(F, simplex, F_value, d, &hi, &lo, centroid);
                }
            }
        }
        k++;
    }
    return k;
}

```