

## REFERENCES AND FURTHER READING:

- Acton, Forman S. 1970, *Numerical Methods That Work* (New York: Harper and Row), p. 55, pp. 454-458.
- Brent, Richard P. 1973, *Algorithms for Minimization without Derivatives* (Englewood Cliffs, N.J.: Prentice-Hall), p. 78.

## 10.4 Downhill Simplex Method in Multidimensions

With this section we begin consideration of multidimensional minimization, that is, finding the minimum of a function of more than one independent variable. This section stands apart from those which follow, however: All of the algorithms after this section will make explicit use of the one-dimensional minimization algorithms of §10.1, §10.2, or §10.3 as a part of their computational strategy. This section implements an entirely self-contained strategy, in which one-dimensional minimization does not figure.

The *downhill simplex method* is due to Nelder and Mead (1965). The method requires only function evaluations, not derivatives. It is not very efficient in terms of the number of function evaluations that it requires. Powell's method (§10.5) is almost surely faster in all likely applications. However the downhill simplex method may frequently be the *best* method to use if the figure of merit is "get something working quickly" for a problem whose computational burden is small.

The method has a geometrical naturalness about it which makes it delightful to describe or work through:

A *simplex* is the geometrical figure consisting, in  $N$  dimensions, of  $N + 1$  points (or vertices) and all their interconnecting line segments, polygonal faces, etc. In two dimensions, a simplex is a triangle. In three dimensions it is a tetrahedron, not necessarily the regular tetrahedron. (The *simplex method* of linear programming also makes use of the geometrical concept of a simplex. Otherwise it is completely unrelated to the algorithm that we are describing in this section.) In general we are only interested in simplexes that are nondegenerate, i.e. which enclose a finite inner  $N$ -dimensional volume. If any point of a nondegenerate simplex is taken as the origin, then the  $N$  other points define vector directions that span the  $N$ -dimensional vector space.

In one-dimensional minimization, it was possible to bracket a minimum, so that the success of a subsequent isolation was guaranteed. Alas! There is no analogous procedure in multidimensional space. For multidimensional minimization, the best we can do is give our algorithm a starting guess, that is, an  $N$ -vector of independent variables as the first point to try. The algorithm is then supposed to make its own way downhill through the unimaginable

complexity of an  $N$ -dimensional topography, until it encounters an (at least local) minimum.

The downhill simplex method must be started not just with a single point, but with  $N + 1$  points, defining an initial simplex. If you think of one of these points (it matters not which) as being your initial starting point  $P_0$ , then you can take the other  $N$  points to be

$$P_i = P_0 + \lambda e_i \quad (10.4.1)$$

where the  $e_i$ 's are  $N$  unit vectors, and where  $\lambda$  is a constant which is your guess of the problem's characteristic length scale. (Or, you could have different  $\lambda_i$ 's for each vector direction.)

The downhill simplex method now takes a series of steps, most steps just moving the point of the simplex where the function is largest ("highest point") through the opposite face of the simplex to a lower point. These steps are called reflections, and they are constructed to conserve the volume of the simplex (hence maintain its nondegeneracy). When it can do so, the method expands the simplex in one or another direction to take larger steps. When it reaches a "valley floor," the method contracts itself in the transverse direction and tries to ooze down the valley. If there is a situation where the simplex is trying to "pass through the eye of a needle," it contracts itself in all directions, pulling itself in around its lowest (best) point. The routine name *amoeba* is intended to be descriptive of this kind of behavior; the basic moves are summarized in Figure 10.4.1.

Termination criteria can be delicate in any multidimensional minimization routine. Without bracketing and with more than one independent variable, we no longer have the option of requiring a certain tolerance for a single independent variable. We typically can identify one "cycle" or "step" of our multidimensional algorithm. It is then possible to terminate when the vector distance moved in that step is fractionally smaller in magnitude than some tolerance  $tol$ . Alternatively, we could require that the decrease in the function value in the terminating step be fractionally smaller than some tolerance  $ftol$ . Note that while  $tol$  should not usually be smaller than the square root of the machine precision, it is perfectly appropriate to let  $ftol$  be of order the machine precision (or perhaps slightly larger so as not to be diddled by roundoff).

Note well that either of the above criteria might be fooled by a single anomalous step that, for one reason or another, failed to get anywhere. Therefore, it is frequently a good idea to *restart* a multidimensional minimization routine at a point where it claims to have found a minimum. For this restart, you should reinitialize any ancillary input quantities. In the downhill simplex method, for example, you should reinitialize  $N$  of the  $N + 1$  vertices of the simplex again by equation (10.4.1), with  $P_0$  being one of the vertices of the claimed minimum.

Restarts should never be very expensive; your algorithm did, after all, converge to the restart point once, and now you are starting the algorithm already there.




Figure 10.4.1. Possible outcomes for a step in the downhill simplex method. The simplex at the beginning of the step, here a tetrahedron, is drawn with solid lines. The simplex at the end of the step (drawn dashed) can be either (a) a reflection away from the high point, (b) a reflection and expansion away from the high point, (c) a contraction along one dimension from the high point, or (d) a contraction along all dimensions toward the low point. An appropriate sequence of such steps will always converge to a minimum of the function.

Consider, then, our  $N$ -dimensional amoeba: .

```
#include <math.h>

#define NMAX 5000           The maximum allowed number of function evaluations, and three
#define ALPHA 1.0          parameters defining the expansions and contractions.
#define BETA 0.5
#define GAMMA 2.0

#define GET_PSUM for (j=1;j<=ndim;j++) { for (i=1,sum=0.0;i<=mpts;i++)\
    sum += p[i][j]; psum[j]=sum;}

void amoeba(p,y,ndim,ftol,funk,nfunk)
```

All the minimization methods in this section and in the two sections following fall under this general schema of successive line minimizations. In this section we consider a class of methods whose choice of successive directions does not involve explicit computation of the function's gradient; the next two sections do require such gradient calculations. You will note that we need not specify whether `linmin` uses gradient information or not. That choice is up to you, and its optimization depends on your particular function. You would be crazy, however, to use gradients in `linmin` and *not* use them in the choice of directions, since in this latter role they can drastically reduce the total computational burden.

But what if, in your application, calculation of the gradient is out of the question. You might first think of this simple method: Take the unit vectors  $e_1, e_2, \dots, e_N$  as a *set of directions*. Using `linmin`, move along the first direction to its minimum, then *from there* along the second direction to *its* minimum, and so on, cycling through the whole set of directions as many times as necessary, until the function stops decreasing.

This dumb method is actually not too bad for many functions. Even more interesting is why it is bad, i.e. very inefficient, for some other functions. Consider a function of two dimensions whose contour map (level lines) happens to define a long, narrow valley at some angle to the coordinate basis vectors (see Figure 10.5.1). Then the only way "down the length of the valley" going along the basis vectors at each stage is by a series of many tiny steps. More generally, in  $N$  dimensions, if the function's second derivatives are much larger in magnitude in some directions than in others, then many cycles through all  $N$  basis vectors will be required in order to get anywhere. This condition is not all that unusual; by Murphy's Law, you should count on it.

Obviously what we need is a better set of directions than the  $e_i$ 's. All *direction set methods* consist of prescriptions for updating the set of directions as the method proceeds, attempting to come up with a set which either (i) includes some very good directions that will take us far along narrow valleys, or else (more subtly) (ii) includes some number of "non-interfering" directions with the special property that minimization along one is not "spoiled" by subsequent minimization along another, so that interminable cycling through the set of directions can be avoided.

### Conjugate Directions

This concept of "non-interfering" directions, more conventionally called *conjugate directions*, is worth making mathematically explicit.

First, note that if we minimize a function along some direction  $u$ , then the gradient of the function must be perpendicular to  $u$  at the line minimum; if not, then there would still be a nonzero directional derivative along  $u$ .

Next take some particular point  $P$  as the origin of the coordinate system with coordinates  $x$ . Then any function  $f$  can be approximated by its Taylor




Figure 10.5.1. Successive minimizations along coordinate directions in a long, narrow "valley" (shown as contour lines). Unless the valley is optimally oriented, this method is extremely inefficient, taking many tiny steps to get to the minimum, crossing and re-crossing the principal axis.

series

$$f(\mathbf{x}) = f(\mathbf{P}) + \sum_i \frac{\partial f}{\partial x_i} x_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 f}{\partial x_i \partial x_j} x_i x_j + \dots \quad (10.5.1)$$

$$\approx c - \mathbf{b} \cdot \mathbf{x} + \frac{1}{2} \mathbf{x} \cdot \mathbf{A} \cdot \mathbf{x}$$

where

$$c \equiv f(\mathbf{P}) \quad \mathbf{b} \equiv -\nabla f|_{\mathbf{P}} \quad [\mathbf{A}]_{ij} \equiv \left. \frac{\partial^2 f}{\partial x_i \partial x_j} \right|_{\mathbf{P}} \quad (10.5.2)$$

The matrix  $\mathbf{A}$  whose components are the second partial derivative matrix of the function is called the *Hessian matrix* of the function at  $\mathbf{P}$ .

In the approximation of (10.5.1), the gradient of  $f$  is easily calculated as

$$\nabla f = A \cdot x - b \quad (10.5.3)$$

(This implies that the gradient will vanish — the function will be at an extremum — at a value of  $x$  obtained by solving  $A \cdot x = b$ . This idea we will return to in §10.7!)

How does the gradient  $\nabla f$  change as we move along some direction? Evidently

$$\delta(\nabla f) = A \cdot (\delta x) \quad (10.5.4)$$

Suppose that we have moved along some direction  $u$  to a minimum and now propose to move along some new direction  $v$ . The condition that motion along  $v$  not *spoil* our minimization along  $u$  is just that the gradient stay perpendicular to  $u$ , i.e. that the change in the gradient be perpendicular to  $u$ . By equation (10.5.4) this is just

$$0 = u \cdot \delta(\nabla f) = u \cdot A \cdot v \quad (10.5.5)$$

When (10.5.5) holds for two vectors  $u$  and  $v$ , they are said to be *conjugate*. When the relation holds pairwise for all members of a set of vectors, they are said to be a conjugate set. If you do successive line minimization of a function along a conjugate set of directions, then you don't need to redo any of those directions (unless, of course, you spoil things by minimizing along a direction that they are *not* conjugate to).

A triumph for a direction set method is to come up with a set of  $N$  linearly independent, mutually conjugate directions. Then, one pass of  $N$  line minimizations will put it exactly at the minimum of a quadratic form like (10.5.1). For functions  $f$  which are not exactly quadratic forms, it won't be exactly at the minimum; but repeated cycles of  $N$  line minimizations will in due course converge *quadratically* to the minimum.

### Powell's Quadratically Convergent Method

Powell first discovered a direction set method which does produce  $N$  mutually conjugate directions. Here is how it goes: Initialize the set of directions  $u_i$  to the basis vectors,

$$u_i = e_i \quad i = 1, \dots, N \quad (10.5.6)$$

Now repeat the following sequence of steps ("basic procedure") until your function stops decreasing:

- Save your starting position as  $P_0$ .
- For  $i = 1, \dots, N$ , move  $P_{i-1}$  to the minimum along direction  $u_i$  and call this point  $P_i$ .
- For  $i = 1, \dots, N - 1$ , set  $u_i \leftarrow u_{i+1}$ .
- Set  $u_N \leftarrow P_N - P_0$ .
- Move  $P_N$  to the minimum along direction  $u_N$  and call this point  $P_0$ .

Powell, in 1964, showed that, for a quadratic form like (10.5.1),  $k$  iterations of the above basic procedure produce a set of directions  $u_i$  whose last  $k$  members are mutually conjugate. Therefore,  $N$  iterations of the basic procedure, amounting to  $N(N + 1)$  line minimizations in all, will exactly minimize a quadratic form. Brent (1973) gives proofs of these statements in accessible form.

Unfortunately, there is a problem with Powell's quadratically convergent algorithm. The procedure of throwing away, at each stage,  $u_1$  in favor of  $P_N - P_0$  tends to produce sets of directions that "fold up on each other" and become linearly dependent. Once this happens, then the procedure finds the minimum of the function  $f$  only over a subspace of the full  $N$ -dimensional case; in other words, it gives the wrong answer. Therefore, the algorithm must not be used in the form given above.

There are a number of ways to fix up the problem of linear dependence in Powell's algorithm, among them:

1. You can reinitialize the set of directions  $u_i$  to the basis vectors  $e_i$  after every  $N$  or  $N + 1$  iterations of the basic procedure. This produces a serviceable method, which we commend to you if quadratic convergence is important for your application (i.e. if your functions are close to quadratic forms and if you desire high accuracy).

2. Brent points out that the set of directions can equally well be reset to the columns of any orthogonal matrix. Rather than throw away the information on conjugate directions already built up, he resets the direction set to calculated principal directions of the matrix  $A$  (which he gives a procedure for determining). The calculation is essentially a singular value decomposition algorithm (see §2.9). Brent has a number of other cute tricks up his sleeve, and his modification of Powell's method is probably the best presently known. Consult his book for a detailed description and listing of the program. Unfortunately it is rather too elaborate for us to include here.

3. You can give up the property of quadratic convergence in favor of a more heuristic scheme (due to Powell) which tries to find a few good directions along narrow valleys instead of  $N$  necessarily conjugate directions. This is the method which we now implement. (It is also the version of Powell's method given in Acton, from which parts of the following discussion are drawn.)

### *Powell's Method Discarding the Direction of Largest Decrease*

The fox and the grapes: Now that we are going to give up the property of quadratic convergence, was it so important after all? That depends on the function that you are minimizing. Some applications produce functions

with long, twisty valleys. Quadratic convergence is of no particular advantage to a program which must slalom down the length of a valley floor that twists one way and another (and another, and another, ... - there are  $N$  dimensions!). Along the long direction, a quadratically convergent method is trying to extrapolate to the minimum of a parabola which just isn't (yet) there; while the conjugacy of the  $N - 1$  transverse directions keeps getting spoiled by the twists.

Sooner or later, however, we do arrive at an approximately ellipsoidal minimum (cf. equation 10.5.1 when  $\mathbf{b}$ , the gradient, is zero). Then, depending on how much accuracy we require, a method with quadratic convergence can save us several times  $N^2$  extra line minimizations, since quadratic convergence *doubles* the number of significant figures at each iteration.

The basic idea of our now-modified Powell's method is still to take  $\mathbf{P}_N - \mathbf{P}_0$  as a new direction; it is, after all, the average direction moved after trying all  $N$  possible directions. For a valley whose long direction is twisting slowly, this direction is likely to give us a good run along the new long direction. The change is to discard the old direction along which the function  $f$  made its *largest decrease*. This seems paradoxical, since that direction was the *best* of the previous iteration. However, it is also likely to be a major component of the new direction that we are adding, so dropping it gives us the best chance of avoiding a buildup of linear dependence.

There are a couple of exceptions to this basic idea. Sometimes it is better *not* to add a new direction at all. Define

$$f_0 \equiv f(\mathbf{P}_0) \quad f_N \equiv f(\mathbf{P}_N) \quad f_E \equiv f(2\mathbf{P}_N - \mathbf{P}_0) \quad (10.5.7)$$

Here  $f_E$  is the function value at an "extrapolated" point somewhat further along the proposed new direction. Also define  $\Delta f$  to be the magnitude of the largest decrease along one particular direction of the present basic procedure iteration. ( $\Delta f$  is a positive number.) Then:

1. If  $f_E \geq f_0$ , then keep the old set of directions for the next basic procedure, because the average direction  $\mathbf{P}_N - \mathbf{P}_0$  is all played out.
2. If  $2(f_0 - 2f_N + f_E)[(f_0 - f_N) - \Delta f]^2 \geq (f_0 - f_E)^2 \Delta f$ , then keep the old set of directions for the next basic procedure, because either (i) the decrease along the average direction was not primarily due to any single direction's decrease, or (ii) there is a substantial second derivative along the average direction and we seem to be near to the bottom of its minimum.

The following routine implements Powell's method in the version just described. In the routine,  $\mathbf{x}_i$  is the matrix whose columns are the set of directions  $\mathbf{n}_i$ ; otherwise the correspondence of notation should be self-evident.

```
#include <math.h>

#define ITMAX 200                Maximum allowed iterations.
static float sqrg;
#define SQR(a) (sqrg=(a).sqrg*sqrg)

void powell(p,xi,n,ftol,iter,frret,func)
float p[],**xi,ftol,*frret,(*func)();
```



```
int n,*iter;
```

Minimization of a function func of n variables. Input consists of an initial starting point p[1..n]; an initial matrix xi[1..n][1..n] whose columns contain the initial set of directions (usually the n unit vectors); and ftol, the fractional tolerance in the function value such that failure to decrease by more than this amount on one iteration signals doneness. On output, p is set to the best point found, xi is the then-current direction set, fret is the returned function value at p, and iter is the number of iterations taken. The routine linmin is used.

```
{
  int i,ibig,j;
  float t,fptt,fp,del;
  float *pt,*ptt,*xit,*vector();
  void linmin(),nrerror(),free_vector();

  pt=vector(1,n);
  ptt=vector(1,n);
  xit=vector(1,n);
  *fret=(*func)(p);
  for (j=1;j<=n;j++) pt[j]=p[j];           Save the initial point.
  for (*iter=1;(*iter)<=ITMAX) {
    fp=(*fret);
    ibig=0;
    del=0.0;                                Will be the biggest function decrease.
    for (i=1;i<=n;i++) {                   In each iteration, loop over all directions in the set.
      for (j=1;j<=n;j++) xit[j]=xi[j][i];   Copy the direction,
      fptt=(*func)(p);
      linmin(p,xit,n,fret,func);           minimize along it.
      if (fabs(fptt-(*fret)) > del) {      and record it if it is the largest decrease
        del=fabs(fptt-(*fret));           so far.
        ibig=i;
      }
    }
    if (2.0*fabs(fp-(*fret)) <= ftol*(fabs(fp)+fabs(*fret))) {
      free_vector(xit,1,n);
      free_vector(ptt,1,n);
      free_vector(pt,1,n);                Termination criterion.
      return;
    }
    if (*iter == ITMAX) nrerror("Too many iterations in routine POWELL");
    for (j=1;j<=n;j++) {                 Construct the extrapolated point and the average
      ptt[j]=2.0*p[j]-pt[j];             direction moved. Save the old starting point.
      xit[j]=p[j]-pt[j];
      pt[j]=p[j];
    }
    fptt=(*func)(ptt);                   Function value at extrapolated point.
    if (fptt < fp) {
      t=2.0*(fp-2.0*(fret)+fptt)*SQR(fp-(fret)-del)-del*SQR(fp-fptt);
      if (t < 0.0) {
        linmin(p,xit,n,fret,func);       Move to the minimum of the new direction,
        for (j=1;j<=n;j++) xit[j][ibig]=xit[j];   and save the new direction.
      }
    }
  }
}

Back for another iteration.
```

### Implementation of Line Minimization

Make no mistake, there is a *right* way to implement linmin: It is to use the *methods* of one-dimensional minimization described in §10.1-§10.3, but to rewrite the programs of those sections so that their bookkeeping is done

on vector-valued points  $P$  (all lying along a given direction  $n$ ) rather than scalar-valued abscissas  $z$ . That straightforward task produces long routines densely populated with "for( $k=1;k<=n;k++$ )" loops.

We do not have space to include such routines in this book. Our `linmin`, which works just fine, is instead a kind of bookkeeping swindle. It constructs an "artificial" function of one variable called `fidim`, which is the value of your function, say, `func`, along the line going through the point  $p$  in the direction  $x_1$ . `linmin` calls our familiar one-dimensional routines `mbrak` (§10.1) and `brent` (§10.2) and instructs them to minimize `fidim`. `linmin` communicates with `fidim` "over the head" of `mbrak` and `brent`, through global (external) variables. That is also how it passes to `fidim` a pointer to your user-supplied function.

The only thing inefficient about `linmin` is this: Its use as an interface between a multidimensional minimization strategy and a one-dimensional minimization routine results in some unnecessary copying of vectors from hither to yon and back again. That should not normally be a significant addition to the overall computational burden, but we cannot disguise its inelegance.

```
#define TOL 2.0e-4
```

```
int ncon=0; /* defining declarations */
float *pcon=0,*xicon=0>(*nrfunc());
```

```
void linmin(p,xi,n,fret,func)
float p[],xi[],*fret,(*func)();
int n;
```

Given an  $n$  dimensional point  $p[1..n]$  and an  $n$  dimensional direction  $xi[1..n]$ , moves and resets  $p$  to where the function `func(p)` takes on a minimum along the direction  $x_1$  from  $p$ , and replaces  $xi$  by the actual vector displacement that  $p$  was moved. Also returns as `fret` the value of `func` at the returned location  $p$ . This is actually all accomplished by calling the routines `mbrak` and `brent`.

```
{
    int j;
    float xx,xmin,fx,fb,fa,bx,ax;
    float brent(),fidim(),*vector();
    void mbrak(),free_vector();

    ncon=n; /* Define the global variables. */
    pcon=vector(1,n);
    xicon=vector(1,n);
    nrfunc=func;
    for (j=1;j<=n;j++) {
        pcon[j]=p[j];
        xicon[j]=xi[j];
    }
    ax=0.0; /* Initial guess for brackets. */
    xx=1.0;
    bx=2.0;
    mbrak(&ax,&xx,&bx,&fa,&fx,&fb,fidim);
    *fret=brent(ax,xx,bx,fidim,TOL,&xmin);
    for (j=1;j<=n;j++) { /* Construct the vector results to return. */
        xi[j] += xmin;
        p[j] += xi[j];
    }
    free_vector(xicon,1,n);
    free_vector(pcon,1,n);
}
```