

10.3 One-Dimensional Search with First Derivatives

Here we want to accomplish precisely the same goal as in the previous section, namely to isolate a functional minimum that is bracketed by the triplet of abscissas (a, b, c) , but utilizing an additional capability to compute the function's first derivative as well as its value.

In principle, we might simply search for a zero of the derivative, ignoring the function value information, using a root finder like `rtflsp` or `zbrent` (§§9.2–9.3). It doesn't take long to reject *that* idea: How do we distinguish maxima from minima? Where do we go from initial conditions where the derivatives on one or both of the outer bracketing points indicate that “downhill” is in the direction *out* of the bracketed interval?

We don't want to give up our strategy of maintaining a rigorous bracket on the minimum at all times. The only way to keep such a bracket is to update it using function (not derivative) information, with the central point in the bracketing triplet always that with the lowest function value. Therefore the role of the derivatives can only be to help us choose new trial points within the bracket.

One school of thought is to “use everything you've got”: Compute a polynomial of relatively high order (cubic or above) that agrees with some number of previous function and derivative evaluations. For example, there is a unique cubic that agrees with function and derivative at two points, and one can jump to the interpolated minimum of that cubic (if there is a minimum within the bracket). Suggested by Davidson and others, formulas for this tactic are given in [1].

We like to be more conservative than this. Once superlinear convergence sets in, it hardly matters whether its order is moderately lower or higher. In practical problems that we have met, most function evaluations are spent in getting globally close enough to the minimum for superlinear convergence to commence. So we are more worried about all the funny “stiff” things that high-order polynomials can do (cf. Figure 3.0.1b), and about their sensitivities to roundoff error.

This leads us to use derivative information only as follows: The sign of the derivative at the central point of the bracketing triplet (a, b, c) indicates uniquely whether the next test point should be taken in the interval (a, b) or in the interval (b, c) . The value of this derivative and of the derivative at the second-best-so-far point are extrapolated to zero by the secant method (inverse linear interpolation), which by itself is superlinear of order 1.618. (The golden mean again: see [1], p. 57.) We impose the same sort of restrictions on this new trial point as in Brent's method. If the trial point must be rejected, we *bisect* the interval under scrutiny.

Yes, we are fuddy-duddies when it comes to making flamboyant use of derivative information in one-dimensional minimization. But we have met too many functions whose computed “derivatives” *don't* integrate up to the function value and *don't* accurately point the way to the minimum, usually because of roundoff errors, sometimes because of truncation error in the method of derivative evaluation.

You will see that the following routine is closely modeled on `brent` in the previous section.

```

FUNCTION dbrent(ax,bx,cx,f,df,tol,xmin)
INTEGER ITMAX
REAL dbrent,ax,bx,cx,tol,xmin,df,f,ZEPS
EXTERNAL df,f
PARAMETER (ITMAX=100,ZEPS=1.0e-10)
    Given a function f and its derivative function df, and given a bracketing triplet of abscissas
    ax, bx, cx [such that bx is between ax and cx, and f(bx) is less than both f(ax) and
    f(cx)], this routine isolates the minimum to a fractional precision of about tol using
    a modification of Brent's method that uses derivatives. The abscissa of the minimum is
    returned as xmin, and the minimum function value is returned as dbrent, the returned
    function value.
INTEGER iter
REAL a,b,d,d1,d2,du,dv,dw,dx,e,fu,fv,fw,fx,olde,tol1,tol2,
*   u,u1,u2,v,w,x,xm
    Comments following will point out only differences from the routine brent. Read that
    routine first.
LOGICAL ok1,ok2           Will be used as flags for whether proposed steps are accept-
a=min(ax,cx)              able or not.
b=max(ax,cx)
v=bx
w=v
x=v
e=0.
fx=f(x)
fv=fx
fw=fx
dx=df(x)                  All our housekeeping chores are doubled by the necessity of
dv=dx                     moving derivative values around as well as function val-
dw=dx                     ues.
do 11 iter=1,ITMAX
    xm=0.5*(a+b)
    tol1=tol*abs(x)+ZEPS
    tol2=2.*tol1
    if(abs(x-xm).le.(tol2-.5*(b-a))) goto 3
    if(abs(e).gt.tol1) then
        d1=2.*(b-a)        Initialize these d's to an out-of-bracket value.
        d2=d1
        if(dw.ne.dx) d1=(w-x)*dx/(dx-dw)    Secant method with one point.
        if(dv.ne.dx) d2=(v-x)*dx/(dx-dv)    And the other.
        Which of these two estimates of d shall we take? We will insist that they be within
        the bracket, and on the side pointed to by the derivative at x:
        u1=x+d1
        u2=x+d2
        ok1=((a-u1)*(u1-b).gt.0.).and.(dx*d1.le.0.)
        ok2=((a-u2)*(u2-b).gt.0.).and.(dx*d2.le.0.)
        olde=e              Movement on the step before last.
        e=d
        if(.not.(ok1.or.ok2))then        Take only an acceptable d, and if both
            goto 1                    are acceptable, then take the small-
        else if (ok1.and.ok2)then        est one.
            if(abs(d1).lt.abs(d2))then
                d=d1
            else
                d=d2
            endif
        else if (ok1)then
            d=d1
        else
            d=d2
        endif
        if(abs(d).gt.abs(0.5*olde))goto 1
        u=x+d
        if(u-a.lt.tol2 .or. b-u.lt.tol2) d=sign(tol1,xm-x)
        goto 2

```

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-
 readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMS
 visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

```

endif
1  if(dx.ge.0.) then      Decide which segment by the sign of the derivative.
    e=a-x
  else
    e=b-x
  endif
d=0.5*e      Bisect, not golden section.
2  if(abs(d).ge.tol1) then
    u=x+d
    fu=f(u)
  else
    u=x+sign(tol1,d)
    fu=f(u)
    if(fu.gt.fx)goto 3    If the minimum step in the downhill direction takes us uphill,
                        then we are done.
  endif
du=df(u)     Now all the housekeeping, sigh.
if(fu.le.fx) then
  if(u.ge.x) then
    a=x
  else
    b=x
  endif
  v=w
  fv=fw
  dv=dw
  w=x
  fw=fx
  dw=dx
  x=u
  fx=fu
  dx=du
else
  if(u.lt.x) then
    a=u
  else
    b=u
  endif
  if(fu.le.fw .or. w.eq.x) then
    v=w
    fv=fw
    dv=dw
    w=u
    fw=fu
    dw=du
  else if(fu.le.fv .or. v.eq.x .or. v.eq.w) then
    v=u
    fv=fu
    dv=du
  endif
endif
enddo 11
pause 'dbrent exceeded maximum iterations'
3  xmin=x
   dbrent=fx
   return
END

```

CITED REFERENCES AND FURTHER READING:

- Acton, F.S. 1970, *Numerical Methods That Work*; 1990, corrected edition (Washington: Mathematical Association of America), pp. 55; 454–458. [1]
- Brent, R.P. 1973, *Algorithms for Minimization without Derivatives* (Englewood Cliffs, NJ: Prentice-Hall), p. 78.

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

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 a one-dimensional minimization algorithm 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 [1]. 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, described in §10.8, 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., that 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 a (local, at least) 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 \mathbf{P}_0 , then you can take the other N points to be

$$\mathbf{P}_i = \mathbf{P}_0 + \lambda \mathbf{e}_i \quad (10.4.1)$$

where the \mathbf{e}_i 's are N unit vectors, and where λ is a constant which is your guess of the problem's characteristic length scale. (Or, you could have different λ_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