Section 4 – Optimization

In finding roots, we searched for points where a function crossed the horizontal axis, where its value was zero. In optimization, we look for the maximum or minimum values of a function, where $f'(x) = 0$ (recall that if $f''(x) < 0$ at that point it is a maximum, if $f''(x) > 0$ it is a minimum). Additionally, there are often constraints placed on the solution that restrict the domain. A general optimization, or *mathematical programming*, problem can be stated as

Find **x** which minimizes or maximizes *f*(**x**) subject to
$d_i(\vec{x}) \le a_i \quad i = 1, 2, \ldots m$
$e_i(\vec{x}) = b_i \quad i = 1, 2, \ldots p$

where
$\vec{x}$ is the design vector, of order *n*
$f(\vec{x})$ is the objective function
$d_i(\vec{x})$ are the inequality constraints
$e_i(\vec{x})$ are the equality constraints
a_i, b_i are constants

If $f(\vec{x})$ and the constraints are all linear, we have a problem in *linear programming*. If $f(\vec{x})$ is quadratic and the constraints are linear, it is a case of *quadratic programming*. If $f(\vec{x})$ is not linear and not quadratic, and/or the constraints are not all linear, the problem is one of *nonlinear programming*.

If constraints are not included the problem is termed *unconstrained*, if they are it is a problem in *constrained optimization*. The number of *degrees of freedom* is $n - p - m$. Generally, to obtain a solution $p + m \le n$. If $p + m > n$, then the system is *over-constrained*.

One-dimensional problems involve functions of a single variable, and involve climbing hills and valleys to find minima and/or maxima. Multidimensional problems involve two or more independent variables.

The process of finding a maximum is essentially the same as finding a minimum, since a value $\vec{x}^*$ that maximizes $f(\vec{x})$ minimizes $-f(\vec{x})$.

4.1 One-Dimensional Unconstrained Optimization

All of these types of problems can be cast as follows: Find the maximum (or minimum) of a function $f(x)$ of one variable. The difficult part will be to assure ourselves that we have found the global maximum. In general, these problems can be divided into bracketed and open methods.

4.1.1 Golden-Search Method

This is a simple, general purpose search technique for a single variable, similar to the bisection method for finding roots. We start with bounds x_l and x_u that bracket a maximum. To narrow the range, we need to pick an intermediate point to see if a maximum occurs between the bounds, and a fourth point to determine in which part of the interval the maximum occurs (that is, within the lower or upper three points). The selection of these intermediate points is critical to efficiency by minimizing the function evaluations.

With respect to the figure shown, let

$$l_0 = l_1 + l_2$$

and

$$\frac{l_1}{l_0} = \frac{l_2}{l_1} = R = \frac{\sqrt{5}-1}{2} \approx 0.618$$

which is the *Golden Ratio*. With these ratios, the method proceeds as follows:

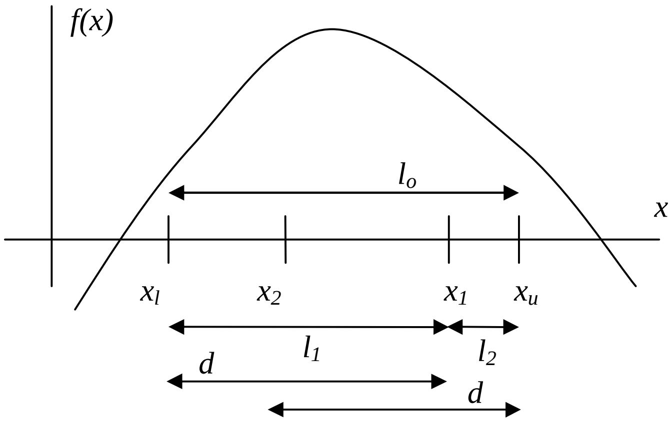

Starting with x_l and x_u, $d = \frac{\sqrt{5}-1}{2}(x_u - x_l)$ and $x_1 = x_l + d$, $x_2 = x_u - d$.

If $f(x_1) > f(x_2)$, then the maximum is within $[x_2, x_u]$, so that for the next step

$$x_{l,new} = x_{2,old}$$
$$x_{2,new} = x_{1,old}$$
$$x_{1,new} = x_l + \frac{\sqrt{5}-1}{2}(x_u - x_l)$$
$$f(x_{2,new}) = f(x_{1,old})$$

If $f(x_2) > f(x_1)$, then the maximum is within $[x_l, x_1]$, so that for the next step

$$x_{u,new} = x_{1,old}$$
$$x_{1,new} = x_{2,old}$$
$$x_{2,new} = x_u - \frac{\sqrt{5}-1}{2}(x_u - x_l)$$
$$f(x_{1,new}) = f(x_{2,old})$$

The size of the interval shrinks by 61.8%., and only one more function evaluation is needed, with every iteration. The relative error in each step is

$$\varepsilon_a = (1-R)\left|\frac{x_u - x_l}{x_{opt}}\right|$$

4.1.2 Quadratic Interpolation

Consider a function $f(x)$ and three points $x_0 < x_1 < x_2$ that bracket an optima of $f(x)$. Fit a parabola through the points $[x_0, f(x_0)]$, $[x_1, f(x_1)]$ and $[x_2, f(x_2)]$ and solve for where the first derivative of that parabola is zero. The maxima of that parabola is at x_3

$$x_3 = \frac{f(x_0)(x_1^2 - x_2^2) + f(x_1)(x_2^2 - x_0^2) + f(x_2)(x_0^2 - x_1^2)}{2f(x_0)(x_1 - x_2) + 2f(x_1)(x_2 - x_0) + 2f(x_2)(x_0 - x_1)}$$

If x_3 is between points x_0 and x_1, then $x_{1,\text{ new}} = x_3$ and $x_{2,\text{ new}} = x_{1,\text{ old}}$ and if x_3 is between points x_1 and x_2 then $x_{0,\text{ new}} = x_{1,\text{ old}}$ and $x_{1,\text{ new}} = x_3$. The method repeats until the relative error

$$\varepsilon_a = \left|\frac{x_3 - x_{3,old}}{x_3}\right|$$

falls below a specified tolerance.

4.1.3 Newton's Method

This is an open method that finds the optimum of $f(x)$ by defining a new function $g(x)=\mathrm{d}f(x)/\mathrm{dx}$ and finding its zeros:

$$x_{i+1} = x_i - \frac{f'(x_i)}{f''(x_i)}$$

This result can also be found by taking the second-order Taylor expansion of $f(x)$ and setting the first derivative equal to zero. This is very similar to the Newton-Raphson method of finding roots, with the advantages of quadratic convergence and the disadvantages of the possibility of divergence and the need of an analytic function. As this method progresses, one needs to check the sign of the second derivative to make sure convergence is to the proper maxima.

4.1.4 Hybrid Methods

Several methods exist that combine the convergence characteristics of bracketing methods when far away from a maxima with the speed and accuracy of an open method when close to the maxima.

4.2 Multidimensional Unconstrained Optimization

Techniques of finding the maxima or minima of multidimensional functions are divided into two general types, those that require derivatives (gradient techniques) or not (direct methods).

4.2.1 Direct Methods

Random Search
This is a brute-force method where the function is evaluated at randomly selected values of the independent variables. If enough points are chosen, the optimum will be found. This technique works on discontinuous and non-differentiable functions and even on difficult undulating functions, and previous attempts can be taken into account to refine the search. Random search techniques are closely related to Monte Carlo methods.

Univariate Method and Pattern Searches
Change one variable at a time to improve the approximation while all the other variables are held constant. This reduces the problem to a series of one-dimensional searches. If we keep track of the general direction of the path, we can find trajectories that shoot directly to the maximum – *pattern searches*.

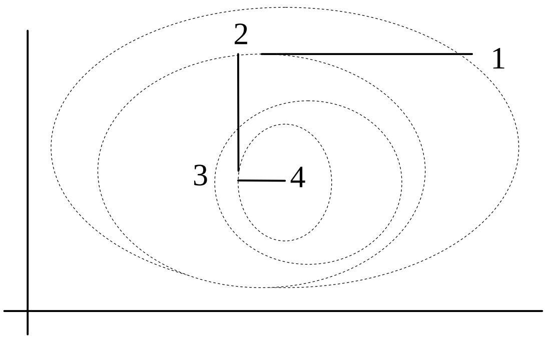

Powell's Method finds two points in a pattern direction by performing two one-dimensional searches in the same direction but with different starting points. The line formed by the two ends is directed towards the maximum along a *conjugate direction*.

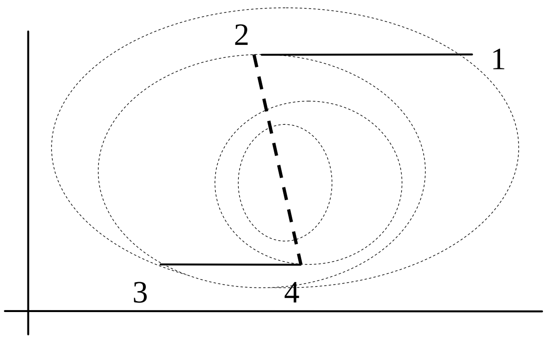

4.2.2 Gradient Methods

In these methods, explicit use of the derivatives is employed to generate algorithms to locate the optima.

The *directional derivative* of a function *f*(*x*,*y*) is $g(x, y) = \frac{\partial f}{\partial x}\cos\theta + \frac{\partial f}{\partial y}\sin\theta$

The *gradient* is in the direction of steepest ascent:

$$\vec{\nabla} f(x, y) = \frac{\partial f}{\partial x}\vec{i} + \frac{\partial f}{\partial y}\vec{j}$$

$$\vec{\nabla} f(\vec{x}) = \begin{bmatrix} \dfrac{\partial f(\vec{x})}{\partial x_1} \\ \vdots \\ \dfrac{\partial f(\vec{x})}{\partial x_n} \end{bmatrix}$$

This direction is the steepest, that is, on the most direct route, to the maximum. When it becomes zero, a local optimum has been reached.

The second derivative, or *Hessian*, tells us if we've reached a local maximum or minimum:

$$|H| = \frac{\partial^2 f}{\partial x^2}\frac{\partial^2 f}{\partial y^2} - \left(\frac{\partial^2 f}{\partial x \partial y}\right)^2 = \det\begin{vmatrix} \dfrac{\partial^2 f}{\partial x^2} & \dfrac{\partial^2 f}{\partial x \partial y} \\ \dfrac{\partial^2 f}{\partial x \partial y} & \dfrac{\partial^2 f}{\partial y^2} \end{vmatrix}$$

Three cases exist:

- If $|H| > 0$ and $\partial^2 f / \partial x^2 > 0$, then $f(x,y)$ is a local minimum in x.
- If $|H| > 0$ and $\partial^2 f / \partial x^2 < 0$, then $f(x,y)$ is a local maximum in x.
- If $|H| < 0$, then $f(x,y)$ is a saddle point.

If the function to be maximized is not accessible, use finite differences to estimate the derivatives:

$$\frac{\partial f}{\partial x} = \frac{f(x+\delta x, y) - f(x-\delta x, y)}{2\delta x}$$

$$\frac{\partial f}{\partial y} = \frac{f(x, y+\delta y) - f(x, y-\delta y)}{2\delta y}$$

$$\frac{\partial^2 f}{\partial x^2} = \frac{f(x+\delta x, y) - 2f(x, y) + f(x-\delta x, y)}{(\delta x)^2}$$

$$\frac{\partial^2 f}{\partial y^2} = \frac{f(x, y+\delta y) - 2f(x, y) + f(x, y-\delta y)}{(\delta y)^2}$$

$$\frac{\partial^2 f}{\partial x \partial y} = \frac{f(x+\delta x, y+\delta y) - f(x+\delta x, y-\delta y) - f(x-\delta x, y+\delta y) + f(x-\delta x, y-\delta y)}{4\delta x \delta y}$$

Note that this is generally not a preferred method, simply due to the number of function evaluations that must be performed.

Steepest Ascent Method

At each step, determine the best direction (gradient) and the best distance in that direction. Following the gradient with an arbitrary step size gives the *method of steepest ascent*. If the function $f(x,y)$ is transformed into a function in h along the gradient using the following substitution:

$$x = x_0 + \frac{\partial f}{\partial x} h$$
$$y = y_0 + \frac{\partial f}{\partial y} h$$
$$f(x, y) \to g(h) = f(x_0 + \frac{\partial f}{\partial x} h, y_0 + \frac{\partial f}{\partial y} h)$$

then one can solve for the maximum step h along that path – *method of optimal steepest ascent*. The method of steepest ascent is linearly convergent and tends to move slowly along long, narrow ridges.

Advanced Gradient Approaches

The *Fletcher-Reeves* conjugate gradient algorithm combines Powell's method of finding conjugate search directions and modifies the steepest ascent method to require that successive gradient search directions.

Newton's Method starts with the second-order Taylor series for $f(\vec{x})$ near $\vec{x} = \vec{x}_i$

$$f(\vec{x}) = f(\vec{x}_i) + \vec{\nabla} f^T(\vec{x}_i)(\vec{x} - \vec{x}_i) - \frac{1}{2}(\vec{x} - \vec{x}_i)^T H_i (\vec{x} - \vec{x}_i)$$

where H_i is the Hessian matrix. At the minimum,

$$\frac{\partial f(\vec{x})}{\partial x_j} = 0 \quad \text{for } j = 1, 2, \dots n$$

so that

$$\vec{\nabla} f = \vec{\nabla} f(\vec{x}_i) + H_i(\vec{x} - \vec{x}_i) = 0$$

If H_i is non-singular, then

$$\vec{x}_{i+1} = \vec{x}_i - H_i^{-1} \vec{\nabla} f$$

which converges quadratically near the optimum, much faster than the method of steepest ascent. Note that analytical derivatives are needed for this method so it is not particularly useful for large numbers of variables. This method may also diverge.

The *Marquadt Method* uses the steepest ascent method when far away from the optimum and Newton's method when near. Modify the diagonal of the Hessian matrix as

$$\vec{H}_i = H_i + \alpha_i I$$

Where α_i is a positive constant and *I* is the identity matrix. At the start of the procedure, α_i is assumed to be large and $\vec{H}_i^{-1} = (1/\alpha_i)I$ which reduces to the method of steepest ascent. As the iterations proceed, $\alpha_i \to 0$ and the method gradually becomes Newton's method.

4.3 Constrained Optimization

We will restrict our discussion of constrained optimization to linear programming (or *linear scheduling*). The basic problem of linear programming is to maximize

$$Z = c_1 x_1 + c_2 x_2 + ... + c_n x_n$$

where x_j is the magnitude of the *j*th *activity* and c_j is the *payoff* for each unit of the *j*th activity undertaken. *Z* is the total payoff. The constraints on the solution can be written in general as

$$a_{i1} x_1 + a_{i2} x_2 + ... + a_{in} x_n \leq b_i$$

where a_{ij} is the amount of the *i*th resource consumed for each unit of the *j*th activity and b_i is the amount of the *i*th resource available. In addition, all activities must be positive, that is, $x_i \geq 0$.

Graphical Solutions

Graphical solutions are limited to two or three dimensions, but are perhaps the most efficient way to solve low-level problems. Plot the constraints as lines (all are linear), and if the problem is properly described they form a *feasible solution space*. Plot the objective function for a particular value of *Z*, and then adjust until the maximum value of Z is found within the feasible solution space.

Constraints that limit the feasible solution space are called *binding*; non-binding constraints do not limit the solution. For any given linear programming problem four possibilities exist:

1. **Unique solution –** the objective function has a maximum at a single point.
2. **Alternate solutions –** if the objective function is parallel to a binding constraint, many solutions may be possible.
3. **No feasible solution –** unsolvable problem, a result of either errors in setting up the problem or over-constraining the solution.
4. **Unbounded problem –** under-constrained system, open-ended solution space.

If a unique solution exists, it always occurs at a point where two or more constraints intersect (*why?*), called *extreme points*. Not all extreme points are feasible, so not all of them have to be examined to find the optimal solution.

The Simplex Method

In order to begin to develop a procedure for solving linear programming problems, define *slack variables* that measure how much of a constrained resource is available. If one slack variable is defined for each resource, the constraints can now be written as equations:

$$a_{i1}x_1 + a_{i2}x_2 + ... + a_{in}x_n + S_i = b_i$$

If $S_i > 0$, there is a surplus of resource *i*, and if $S_i < 0$ we have exceeded the allowable supply of resource *i*.

The system now has more variables than equations, and is *under-specified*. There are *n* structural (original) variables and *m* surplus (slack) variables, resulting in $n + m$ total variables. The difference between the number of variables and the number of equations is the number of variables that must be equal to zero in order to have a unique solution at a feasible extreme point.

For m linear equations with *n* unknowns, set *n-m* variables equal to zero and solve the *m* equations for the remaining unknowns. The *m* variables that are solved for are called *basic variables*; those that are set to zero are called *non-basic variables*. If all the basic variables are non-negative upon solution, the point found is a *basic feasible solution.* The optimum solution will be one of these points.

However, to test all of the extreme points for *m* equations and *n* unknowns one must solve

$$C_m^n = \frac{n!}{m!(n-m)!}$$

systems of *m* equations. A problem consisting of 10 equations and 16 unknowns requires the solution of 8008 sets of 10x10 equations. In addition, many of these points may be infeasible.

The Simplex method avoids these problems by starting with a basic feasible solution (often, all the structural variables are simply set equal to zero), then moves through a sequence of other basic feasible solutions that improve the value of the objective function. Once the optimum is found, the method stops.

First, start at a simple basic solution, one that perhaps sets all of the structural variables to zero. Then increase the value of a non-basic variable so that *Z* increases (called the *entering variable*) and set one of the current basic variables to zero (called the *leaving variable*). The entering variable can be any of the variables that has a negative coefficient, usually the one with the largest negative coefficient is chosen since it leads to the largest increase in *Z.* The leaving variable is chosen from the set of current basic variables by calculating the values at which the constraint lines intersect the constraint corresponding to the leaving variable. Calculate the remaining coefficients by solving the remaining system of equations. The method continues until no negative coefficients remain.

Example: A natural gas refinery receives a fixed amount, 77 m^3, of raw gas per week. The raw gas is processed into two grades, regular and premium. Only one grade of gas can be processed at a time. The facility can make 7 m^3/ton of regular gas and 11 m^3/ton of premium gas. The facility is available for 80 hours per week. Storage capacity is limited to 9 tons of regular and 6 tons of premium gas. Regular takes 10 hours/ton to refine and makes a profit of $150/ton. Premium takes 8 hours/ton to refine and makes a profit of $175/ton. Maximize the profits for this operation.

Maximize $Z = 150x_1 + 175x_2$ (profit)
Subject to
$7x_1 + 11x_2 \le 77$ (raw material constraint, 1)
$10x_1 + 8x_2 \le 80$ (time constraint, 2)
$x_1 \le 9$ (regular storage constraint, 3)
$x_2 \le 6$ (premium storage constraint, 4)
$x_1, x_2 \ge 0$

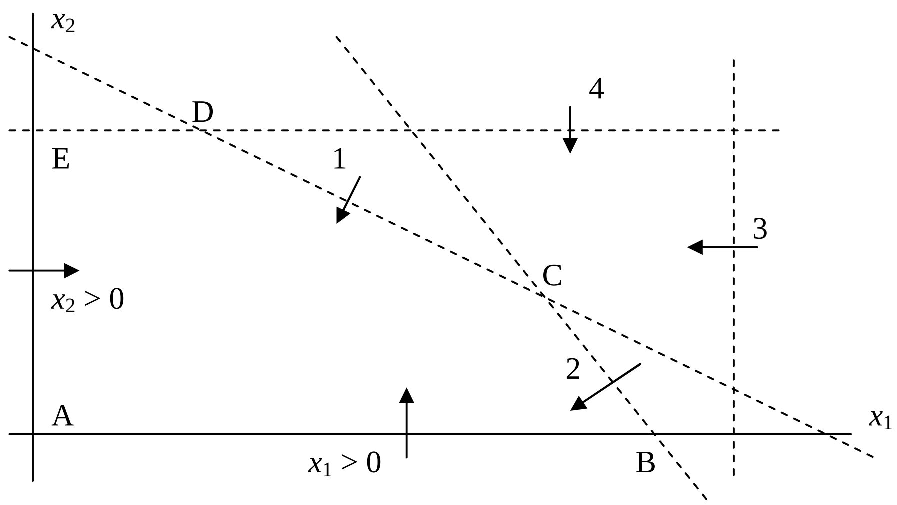

Introduce slack variables for each of the constraints and rewrite these equations in the form:

$Z = 150x_1 + 175x_2 - 0\ S_1 - 0\ S_2 - 0\ S_3 - 0\ S_4 = 0$
$7x_1 + 11x_2 + S_1 = 77$
$10x_1 + 8x_2 + S_2 = 80$
$x_1 + S_3 = 9$
$x_2 + S_4 = 6$

Let $x_1 = x_2 = 0$ and solve the remaining set of equations

$Z = -\ 0\ S_1 - 0\ S_2 - 0\ S_3 - 0\ S_4 = 0$
$S_1 = 77$

$S_2 = 80$
$S_3 = 9$
$S_4 = 6$

to find $Z = 0$, $S_1 = 77$, $S_2 = 80$, $S_3 = 9$, $S_4 = 6$ (point A in the figure above)

Examining the payoff equation, the largest change in Z will come with a change in x_2, this is the entering variable in this step. Calculate the slope of each constraint equation above with respect to x_2:

$x_2 = 77/11 = 7$
$x_2 = 80/8 = 10$
$x_2 = 6/1 = 6$

and choose as the leaving variable the constraint associated with the smallest non-negative value, $S_4 = 0$ (x_1 is still zero) so that the set of equations becomes

$Z = 150x_1 - 0\ S_1 - 0\ S_2 -0\ S_3 -0\ S_4 = 175(6) = \$1050$
$7x_1 + S_1 = 77 – 11(6) = 11$
$10x_1 + S_2 = 80 – 8(6) = 32$
$x_1 + S_3 = 9$
$x_2 = 6$

This is point E ($x_1 = 0$, $x_2 = 6$) in the figure above. Now the largest change in Z will come from a change in x_1, so this is the entering variable in this step. To find the leaving variable, calculate the slopes of the constraint equations with respect to x_1:

$x_1 = 11/7$ smallest non-negative value
$x_1 = 32/10$
$x_1 = 9$

which indicates that S_1 should be the leaving variable. Setting $S_1 = 0$ gives

$Z = 0\ S_1 - 0\ S_2 - 0\ S_3 -0\ S_4 = 150(11/7) + 175(6) = \$1285.71$
$7x_1 + S_1 = 77 – 11(6) = 11$
$S_2 = 80 – 8(6) = -10(11/7) -8(6) = -63.71$
$S_3 = 9 – 11/7 = 7.429$

Which corresponds to point D ($x_1 = 11/7$, $x_2 = 6$) in the figure. The slack variable S_2 cannot be negative, so take that as the next leaving variable $S_2 = 0$. Combined with $S_1 = 0$ gives the system of equations

$Z = 150x_1 + 175x_2 - 0\ S_1 - 0\ S_2 -0\ S_3 -0\ S_4$
$7x_1 + 11x_2 = 77$
$10x_1 + 8x_2 = 80$
$x_1 + S_3 = 9$

$x_2 + S_4 = 6$

Solving this set of equations gives

$x_1 = 4.889$ tons of regular gas
$x_2 = 3.889$ ton of premium gas
$S_3 = 4.111$ excess storage capacity of regular gas
$S_4 = 2.111$ excess storage capacity of premium gas
Z = \$1414 profit
$S_1 = S_2 = 0$, no excess capacity in either raw material or time

Since all the variables are non-negative this must be the optimal solution, corresponding to point C of the graph above.