21-256: Unconstrained extrema

Clive Newstead, Wednesday 11th June 2014

Unconstrained extrema in two dimensions

Local extrema

Intuitively, a *neighbourhood* of a point (a, b) is a set in which there is some positive distance so that you can move in any direction by that distance and still remain in the set. Precisely, a subset U of $\mathbb{R}^2$ is a neighbourhood of (a, b) if, for some $d > 0$, if (x, y) whose distance from (a, b) is $< d$, then (x, y) must lie in U.

Suppose f is a function with domain D and (a, b) is a point in D. We say (a, b) is a...

- ... *local maximum* if $f(a, b) \geqslant f(x, y)$ for all (x, y) in some neighbourhood U of (a, b);
- ... *local minimum* if $f(a, b) \leqslant f(x, y)$ for all (x, y) in some neighbourhood U of (a, b);
- ... *local extremum* if it is a local maximum or a local minimum.

Theorem. Let f be a function of two variables and let (a, b) be a point in its domain. If (a, b) is a local extremum of f and $\nabla f(a, b)$ exists, then $\nabla f(a, b) = \mathbf{0}$.

The converse is not true; a point (a, b) for which $\nabla f(a, b) = \mathbf{0}$ is called a *critical point*. If (a, b) is a critical point but not a local extremum, it is called a *saddle point*.

We can classify critical points using the following theorem:

Theorem (second derivative test). Let f be a function of two variables, let (a, b) be a point in its domain and suppose $\nabla f(a, b)$ exists and is equal to the zero vector. Write

$$D = f_{xx}(a, b) f_{yy}(a, b) - f_{xy}(a, b)^2$$

(a) If $D > 0$ and $f_{xx}(a, b) > 0$ then (a, b) is a local minimum;

(b) If $D > 0$ and $f_{xx}(a, b) < 0$ then (a, b) is a local maximum;

(c) If $D < 0$ then (a, b) is a saddle point;

(d) If $D = 0$ then the test is inconclusive.

This generalises the one-dimensional case: if f is a function of a single variable and $f'(a) = 0$, then the second derivative test states that a is a local minimum if $f''(a) > 0$, a local maximum if $f''(a) < 0$ and the test is inconclusive if $f''(a) = 0$.

Global extrema

The extreme value theorem for functions of a single variable says that a continuous function on a closed and bounded interval $[a, b]$ attains a maximum and minimum value on that interval. We generalise the notion of a closed and bounded interval by introducing *compact* subsets of $\mathbb{R}^2$.

Definition. A subset D of $\mathbb{R}^2$ is...

- ... *closed* if whenever all neighbourhoods of a point (x, y) intersects D, then (x, y) actually lies in D;
- ... *bounded* if there is some $K > 0$ such that the distance between any two points in D is $< K$;
- ... *compact* if it is closed and bounded.

It should be said that the above definition generalises immediately to $\mathbb{R}^n$. (Compact subsets of $\mathbb{R}$ are just finite collections of closed, bounded intervals.)

The compact subsets you encounter will be characterised by being defined using $\leqslant$ signs; their *boundary* is the subset obtained by replacing the $\leqslant$ sign by an $=$ sign. For example, the unit disc $x^2 + y^2 \leqslant 1$ is a compact set, and its boundary is the unit circle $x^2 + y^2 = 1$.

Theorem (extreme value theorem). Let f be a continuous function on a compact set D. Then f attains a maximum and a minimum value on D.

This theorem gives us a feature for finding global maxima and minima of functions defined on compact subsets.

- **Step 1.** Find the value of f at all its critical points and points where ∇f doesn't exist;
- **Step 2.** Find the extreme values of f on the boundary of its domain;
- **Step 3.** The largest value from Steps 1–2 is the global maximum of f; the smallest value is the global minimum.

Step 2 reduces to a one-dimensional problem. As mentioned above, the boundary condition (e.g. $x^2 + y^2 = 1$ if the domain were the subset $x^2 + y^2 \leqslant 1$) allows you to write one variable in terms of the other variable, thus reducing to a single-variable problem which can be solved using techniques from previous calculus courses.

Unconstrained extrema in n dimensions

Definiteness

Let A be an $n \times n$ matrix. The k^{th} *principal minor* of A is the $k \times k$ matrix in the upper left-hand corner of A.

For example, the principal minors of the matrix $\begin{pmatrix} 1 & 2 & 7 & 9 \\ -3 & 0 & 2 & 4 \\ 0 & 1 & -2 & 6 \\ 1 & 1 & 1 & 1 \end{pmatrix}$ are

$$(1)\,, \quad \begin{pmatrix} 1 & 2 \\ -3 & 0 \end{pmatrix}, \quad \begin{pmatrix} 1 & 2 & 7 \\ -3 & 0 & 2 \\ 0 & 1 & -2 \end{pmatrix}, \quad \begin{pmatrix} 1 & 2 & 7 & 9 \\ -3 & 0 & 2 & 4 \\ 0 & 1 & -2 & 6 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

We say a symmetric matrix A is...

- ... *positive definite* (write $A > 0$) if the determinants of its principal minors are all positive;
- ... *negative definite* (write $A < 0$) if the determinants of its principal minors alternate between negative (k odd) and positive (k even);
- ... *indefinite* if it is neither positive definite nor negative definite.

That is, if the signs of the determinants of the principal minors (in increasing order) look like $++++\cdots$ then the matrix is positive definite; if the signs look like $-+-+\cdots$ then the matrix is negative definite.

Note that if the matrix is (positive or negative) definite then the determinants of its even-sized principal minors are always positive. Thus a good check to see if a matrix is indefinite is if its upper left-hand 2×2 matrix has negative determinant.

We will use definiteness of a particular matrix associated with a give function to extract information about its critical points.

Hessian matrix

Let f be a function of n variables $x_1, x_2, \cdots, x_n$. The *Hessian matrix* of f, denoted H_f, is the (variable) matrix whose $(i,j)^{\text{th}}$ component is $\dfrac{\partial^2 f}{\partial x_i \partial x_j}$.

For instance, if f is a function of variables x, y, z then

$$H_f = \begin{pmatrix} f_{xx} & f_{xy} & f_{xz} \\ f_{yx} & f_{yy} & f_{zz} \\ f_{zx} & f_{zy} & f_{zz} \end{pmatrix}$$

Note that if f is sufficiently nice (i.e. second derivatives exist and are continuous) then the order of partial differentiation doesn't matter, so H_f is a symmetric matrix. Thus it makes sense to ask whether the Hessian matrix of a function, evaluated at a point, is positive definite or negative definite.

Theorem (second derivative test). Let f be a function of n variables $x_1, x_2, \cdots, x_n$ and let $(a_1, a_2, \cdots, a_n)$ be a point in the domain of f. Suppose $\boldsymbol{\nabla} f(a_1, \cdots, a_n) = \mathbf{0}$. Then:

- $(a_1, \cdots, a_n)$ is a local minimum if $H_f(a_1, \cdots, a_n) > 0$;
- $(a_1, \cdots, a_n)$ is a local maximum if $H_f(a_1, \cdots, a_n) < 0$;
- The test is inconclusive if $H_f(a_1, \cdots, a_n)$ is indefinite.

This agrees with the tests in lower dimensions:

- If $n = 1$, then $H_f(a) = (f''(a))$, so $H_f(a) > 0$ if and only if $f''(a) > 0$, and $H_f(a) < 0$ if and only if $f''(a) < 0$.
- If $n = 2$, then:
  - $H_f(a, b) > 0$ if and only if $f_{xx}(a, b) > 0$ and $f_{xx}(a, b)f_{yy}(a, b) - f_{xy}(a, b)^2 > 0$;
  - $H_f(a, b) < 0$ if and only if $f_{xx}(a, b) < 0$ and $f_{xx}(a, b)f_{yy}(a, b) - f_{xy}(a, b)^2 > 0$.

  In the 2-variable case the second derivative test gives us the additional condition that if $f_{xx}(a, b)f_{yy}(a, b) - f_{xy}(a, b)^2 < 0$ then (a, b) is a saddle point. In order to generalise this to higher dimensions, we would need to introduce eigenvectors and eigenvalues, which are not within the scope of this course.