1.4. Coercive Functions and Global Minimizers

At this stage we can find global minimizers for $f(\mathbf{x})$ if $f(\mathbf{x})$ has a critical point and $Hf(\mathbf{x})$ is always positive definite. But what about global minimization for $f(\mathbf{x})$ in the case in which $Hf(\mathbf{x})$ is not known to be always positive definite? This short section is devoted to showing that this question sometimes has a very simple answer. The answer depends on the following theorem from calculus:

(1.4.1) Theorem. *Let D be a closed bounded subset of R^n. If $f(\mathbf{x})$ is a continuous function defined on D, then $f(\mathbf{x})$ has a global maximizer and a global minimizer on D.*[1]

Note that this theorem does not guarantee a global minimum on R^n, or on any other set that is either unbounded or not closed.

Now we isolate the type of function that is easily handled.

(1.4.2) Definition. A continuous function $f(\mathbf{x})$ that is defined on all of R^n is called *coercive* if

$$\lim_{\|\mathbf{x}\| \to \infty} f(\mathbf{x}) = +\infty.$$

This means that for any constant M there must be a positive number R_M such that $f(x) \geq M$ whenever $\|\mathbf{x}\| \geq R_M$. In particular, the values of $f(\mathbf{x})$ cannot remain bounded on a set A in R^n that is not bounded.

[1] For a proof see *Elements of Real Analysis* by R. G. Bartle.

(1.4.3) Examples

(a) Let $f(x, y) = x^2 + y^2 = \|\mathbf{x}\|^2$. Then

$$\lim_{\|\mathbf{x}\|\to\infty} f(\mathbf{x}) = \lim_{\|\mathbf{x}\|\to\infty} \|\mathbf{x}\|^2 = \infty.$$

Thus $f(x, y)$ is coercive.

(b) Let $f(x, y) = x^4 + y^4 - 3xy$. Note that

$$f(x, y) = (x^4 + y^4)\left(1 - \frac{3xy}{x^4 + y^4}\right).$$

If $\|\mathbf{x}\|$ is large, then $3xy/(x^4 + y^4)$ is very small. Hence

$$\lim_{\|(x,y)\|\to\infty} f(x, y) = \lim_{\|(x,y)\|\to\infty} (x^4 + y^4)\cdot(1 - 0) = +\infty.$$

Thus $f(x, y)$ is coercive.

(c) Let $f(x, y, z) =$

$$e^{x^2} + e^{y^2} + e^{z^2} - x^{100} - y^{100} - z^{100}.$$

Then because exponential growth is much faster than the growth of any polynomial, it follows that

$$\lim_{\|(x,y,z)\|\to\infty} f(x, y, z) = \infty.$$

Thus $f(x, y, z)$ is coercive.

(d) Linear functions on R^2 are never coercive. Such functions can be expressed as follows:

$$f(x, y) = ax + by + c,$$

where either $a \neq 0$ or $b \neq 0$. To see that $f(x, y)$ is not coercive, simply observe that $f(x, y)$ is constantly equal to c on the line

$$ax + by = 0.$$

Since this line is unbounded and $f(x, y)$ is not unbounded on this line, the function $f(x, y)$ is not coercive.

(e) If $f(x, y, z) = x^4 + y^4 + z^4 - 3xyz - x^2 - y^2 - z^2$, then as

$$\|(x, y, z)\| = \sqrt{x^2 + y^2 + z^2} \to \infty,$$

the higher degree terms dominate and force $\lim_{\|(x,y,z)\|\to\infty} f(x, y, z) = \infty$. Thus $f(x, y, z)$ is coercive. The following example helps us avoid some misunderstandings.

(f) Let $f(x, y) = x^2 - 2xy + y^2$. Then:

(i) for each fixed y_0, we have $\lim_{|x|\to\infty} f(x, y_0) = \infty$;
(ii) for each fixed x_0, we have $\lim_{|y|\to\infty} f(x_0, y) = \infty$;
(iii) but $f(x, y)$ is *not* coercive.

Properties (i) and (ii) above are more or less clear because in each case the quadratic term dominates. For example, in case (i), we have for a fixed y_0,

$$f(x, y_0) = x^2 - xy_0 + y_0^2.$$

This function of x is a parabola that opens upward. Therefore

$$\lim_{|x|\to\infty} f(x, y_0) = \infty.$$

To see that $f(x, y)$ is not coercive, factor to learn

$$f(x, y) = x^2 - 2xy + y^2 = (x - y)^2.$$

Therefore if $\|(x, y)\|$ goes to ∞ on the line $y = +x$, we see $f(x, y) = (x - x)^2 = 0$ and hence $f(x, y) = 0$ on the unbounded line $y = x$. Therefore, $\lim_{\|(x,y)\|\to\infty} f(x, y) \neq \infty$ so $f(x, y)$ is not coercive.

The point of this last example is very important. For $f(\mathbf{x})$ to be coercive, it is not sufficient that $f(\mathbf{x}) \to \infty$ as each coordinate tends to ∞. Rather $f(\mathbf{x})$ must become infinite along any path for which $\|\mathbf{x}\|$ becomes infinite. Exercise 31 contains a general result concerning functions of the sort discussed in Example (1.4.3)(a), (f) above.

The reason coercive functions are important is that they all have global minimizers.

(1.4.4) Theorem. *Let $f(\mathbf{x})$ be a continuous function defined on all R^n. If $f(\mathbf{x})$ is coercive, then $f(\mathbf{x})$ has at least one global minimizer.*

If, in addition, the first partial derivatives of $f(\mathbf{x})$ exist on all of R^n, then these global minimizers can be found among the critical points of $f(\mathbf{x})$.

PROOF. To prove the first statement, assume $\lim_{\|\mathbf{x}\|\to\infty} f(\mathbf{x}) = +\infty$. This means that if $\|\mathbf{x}\|$ is large, then so is $f(\mathbf{x})$. Accordingly, there is a number $r > 0$ such that if $\|\mathbf{x}\| > r$, then

$$f(\mathbf{x}) > f(\mathbf{0}).$$

Let $\bar{B}(\mathbf{0}, r)$ be the set $\{\mathbf{x}: \|\mathbf{x}\| \leq r\}$. The function $f(\mathbf{x})$ is continuous at each point of the set $\bar{B}(\mathbf{0}, r)$ and the set $\bar{B}(\mathbf{0}, r)$ is closed and bounded. From Theorem (1.4.1), it follows that $f(\mathbf{x})$ takes a minimum value on $\bar{B}(\mathbf{0}, r)$ at a point $\mathbf{x}^*$ in $\bar{B}(\mathbf{0}, r)$. In other words, $x \in \bar{B}(\mathbf{0}, r)$ implies $f(\mathbf{x}^*) \leq f(\mathbf{x})$. In particular, because $\mathbf{0} \in \bar{B}(\mathbf{0}, r)$, we see that

$$f(\mathbf{x}^*) \leq f(\mathbf{0}).$$

On the other hand, if $\mathbf{x} \notin \bar{B}(\mathbf{0}, r)$, then

$$f(\mathbf{x}) > f(\mathbf{0}) \geq f(\mathbf{x}^*).$$

Summarizing, we have seen that $\mathbf{x} \in \bar{B}(\mathbf{0}, r)$ implies $f(\mathbf{x}) \geq f(\mathbf{x}^*)$ and $\mathbf{x} \notin \bar{B}(\mathbf{0}, r)$ implies $f(\mathbf{x}) > f(\mathbf{x}^*)$. This shows $\mathbf{x}^*$ is a global minimizer of $f(\mathbf{x})$ and completes the proof of the first statement of the theorem.

The second statement holds because global minimizers on R^n are critical points by Theorem (1.2.3). This completes the proof.

This theorem sets up a method for trying to minimize coercive functions on R^n. If $f(\mathbf{x})$ is coercive and the first partial derivatives of $f(\mathbf{x})$ exist on R^n, then its minimizers are found among the critical points. Therefore to minimize $f(\mathbf{x})$ on R^n, merely list the critical points $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(p)}$ of $f(\mathbf{x})$. Then choose the critical point $\mathbf{x}^{(i)}$ such that $f(\mathbf{x}^{(i)})$ is less than or equal to the other $f(\mathbf{x}^{(j)})$ for $j = 1, \ldots, p$. Theorem (1.4.4) guarantees that $f(\mathbf{x}^{(i)})$ is a global minimizer of $f(\mathbf{x})$ on R^n.

(1.4.5) Example. Minimize

$$f(x, y) = x^4 - 4xy + y^4$$

on R^2.

To this end, compute

$$\nabla f(x, y) = \begin{pmatrix} 4x^3 - 4y \\ -4x + 4y^3 \end{pmatrix}$$

and

$$Hf(x, y) = \begin{pmatrix} 12x^2 & -4 \\ -4 & 12y^2 \end{pmatrix}.$$

Note that

$$Hf(\tfrac{1}{2}, \tfrac{1}{2}) = \begin{pmatrix} 3 & -4 \\ -4 & 3 \end{pmatrix},$$

which is certainly not positive definite since $\det Hf(\frac{1}{2}, \frac{1}{2}) = 9 - 16 < 0$. Therefore the tests from the last section are not applicable. But all is not lost because $f(x, y)$ is coercive!

To see that $f(x, y)$ is coercive, note that

$$f(x, y) = x^4 + y^4\left(1 - \frac{4xy}{x^4 + y^4}\right).$$

As $\|(x, y)\| = \sqrt{x^2 + y^2} \to \infty$, the term $4xy/(x^4 + y^4) \to 0$. Hence

$$\lim_{\|(x,y)\| \to \infty} f(x, y) = \lim_{\|(x,y)\| \to \infty} (x^4 + y^4)(1 - 0) = +\infty.$$

Thus $f(x, y)$ is coercive. According to the last theorem $f(x, y)$ has a global minimizer at one of the critical points. Setting $\nabla f(x, y) = 0$, we get $y = x^3$, and $x = y^3$. Hence $x = x^9$ and $x(x^8 - 1) = 0$. This produces three critical points

$$(0, 0), (1, 1), (-1, -1).$$

Now

$$f(0, 0) = 0,$$

$$f(1, 1) = 1 - 4 + 1 = -2,$$

$$f(-1, -1) = 1 - 4 + 1 = -2.$$

Therefore $(-1, -1)$ and $(1, 1)$ are both global minimizers of $f(x, y)$.

1.5. Eigenvalues and Positive Definite Matrices

If the eigenvalues of a symmetric matrix are available, then they can easily be used to recognize definite, semidefinite, and indefinite matrices. The goal of this section is to see why this is so. Here is some background from linear algebra.

If A is an $n \times n$-matrix and if $\mathbf{x}$ is a nonzero vector in R^n such that $A\mathbf{x} = \lambda\mathbf{x}$ for some real or complex number λ, then λ is called an *eigenvalue* of A. If λ is an eigenvalue of A, then any nonzero vector $\mathbf{x}$ that satisfies the equation $A\mathbf{x} = \lambda\mathbf{x}$ is called an *eigenvector of A corresponding to λ*. Since λ is an eigenvalue of an $n \times n$-matrix A if and only if the homogeneous system $(A - \lambda I)\mathbf{x} = \mathbf{0}$ of n equations in n unknowns has a nonzero solution $\mathbf{x}$, it follows that the eigenvalues of A are just the roots of the *characteristic equation*

$$\det(A - \lambda I) = 0.$$

Since $\det(A - \lambda I)$ is a polynomial of degree n in λ, the characteristic equation has n real or complex roots if we count multiple roots according to their multiplicities, so an $n \times n$-matrix A has n real or complex eigenvalues counting multiplicities.

Symmetric matrices have the following special properties with respect to eigenvalues and eigenvectors:

(1) All of the eigenvalues of a symmetric matrix are real numbers.
(2) Eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.
(3) If λ is an eigenvalue of multiplicity k for a symmetric matrix A (that is, λ is a root of characteristic equation $\det(A - \lambda I) = 0$, k times), there are k linearly independent eigenvectors corresponding to λ. By applying the Gram–Schmidt Orthogonalization Process, we can always replace these k linearly independent eigenvectors with a set of k mutually orthogonal eigenvectors of unit length. (For another view of this, see Example (7.2.4).)

By combining (2) and (3), we see that if A is an $n \times n$-symmetric matrix, then there are n mutually orthogonal unit eigenvectors $\mathbf{u}^{(1)}, \ldots, \mathbf{u}^{(n)}$ corresponding to the n eigenvalues $\lambda_1, \ldots, \lambda_n$ of A (with repeated eigenvalues listed according to their multiplicity). If P is the $n \times n$-matrix whose ith column is the unit eigenvector $\mathbf{u}^{(i)}$ corresponding to λ_i, and if D is the diagonal matrix with the eigenvalues $\lambda_1, \ldots, \lambda_n$ down the main diagonal, then the following matrix equation holds:

$$AP = PD,$$

because $A\mathbf{u}^{(i)} = \lambda_i \mathbf{u}^{(i)}$ for $i = 1, \ldots, n$. Since the matrix P is orthogonal (that is, its columns are mutually orthogonal unit vectors), P is invertible and the inverse P^{-1} of P is just the transpose P^T of P. It follows that

$$P^T AP = D,$$

that is, the orthogonal matrix P *diagonalizes* A. If $Q_A(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x}$ is the quadratic form associated with the symmetric matrix A and if $\mathbf{x} = P\mathbf{y}$, then

$$Q_A(\mathbf{x}) = \mathbf{x} \cdot A\mathbf{x} = (P\mathbf{y})^T A(P\mathbf{y}) = \mathbf{y}^T P^T AP\mathbf{y}$$
$$= \mathbf{y}^T D\mathbf{y} = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \cdots + \lambda_n y_n^2.$$

Moreover, since P is invertible, $\mathbf{x} \neq 0$ if and only if $\mathbf{y} \neq 0$. Also, if $\mathbf{y}^{(i)}$ is the vector in R^n with the ith component equal to 1 and all other components equal to zero, and if $\mathbf{x}^{(i)} = P\mathbf{y}^{(i)}$, then

$$Q_A(\mathbf{x}^{(i)}) \; \lambda_i$$

for $i = 1, 2, \ldots, n$. These considerations yield the following eigenvalue test for definite, semidefinite, and indefinite matrices.

(1.5.1) Theorem. *If A is a symmetric matrix, then:*

(a) *the matrix A is positive definite (resp. negative definite) if and only if all the eigenvalues of A are positive (resp. negative);*
(b) *the matrix A is positive semidefinite (resp. negative semidefinite) if and only if all of the eigenvalues of A are nonnegative (resp. nonpositive);*
(c) *the matrix A is indefinite if and only if A has at least one positive eigenvalue and at least one negative eigenvalue.*

The following example shows how the eigenvalue criteria in Theorem (1.5.1) can be applied to an optimization problem.

(1.5.2) Example. Let us locate all maximizers, minimizers, and saddle points of

$$f(x_1, x_2, x_3) = x_1^2 + x_2^2 + x_3^2 - 4x_1x_2.$$

The critical points of $f(x_1, x_2, x_3)$ are solutions of the system of equations

$$0 = \frac{\partial f}{\partial x_1} = 2x_1 - 4x_2,$$

$$0 = \frac{\partial f}{\partial x_2} = -4x_1 + 2x_2,$$

$$0 = \frac{\partial f}{\partial x_3} = 2x_3.$$

It is easy to check that (0, 0, 0) is the one and only solution of this system. The Hessian of $f(x_1, x_2, x_3)$ is the constant matrix

$$Hf(x_1, x_2, x_3) = \begin{pmatrix} 2 & -4 & 0 \\ -4 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}.$$

The eigenvalues of the Hessian are the solutions of the characteristic equation

$$0 = \det \begin{pmatrix} 2-\lambda & -4 & 0 \\ -4 & 2-\lambda & 0 \\ 0 & 0 & 2-\lambda \end{pmatrix} = (2-\lambda)[(2-\lambda)^2 - 16]$$
$$= (2-\lambda)[\lambda^2 - 4\lambda - 12].$$

Thus, the eigenvalues are $\lambda = 2, 6, -2$, so Theorem (1.3.6) shows that (0, 0, 0) is a saddle point. Since $f(x_1, x_2, x_3)$ has continuous first partial derivatives everywhere on R^3, it follows from (1.2.3) that $f(x_1, x_2, x_3)$ has no other minimizers, maximizers, or saddle points.

Exercises