The Geometry of Quadratic Functions

Thomas R. Cameron

April 1, 2026

1 Introduction

In the previous lecture we introduced quadratic programming and observed that the behavior of a quadratic objective function depends strongly on its geometry. In this lecture we study that geometry more carefully. Our main goals are to understand how eigenvalues and eigenvectors describe the geometry of a quadratic, how convexity is determined by the Hessian matrix, and how the level sets of a quadratic explain the location of optimal solutions.

Consider a quadratic function of the form

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T Q\mathbf{x} + \mathbf{c}^T\mathbf{x}$$

where $Q \in \mathbb{R}^{n\times n}$ is symmetric and $\mathbf{c} \in \mathbb{R}^n$. As discussed previously, we may assume that Q is symmetric without loss of generality, since only the symmetric part of Q contributes to the value of $\mathbf{x}^T Q\mathbf{x}$. The matrix Q determines the curvature of the quadratic, while the vector $\mathbf{c}$ shifts the location of the critical point.

2 The Gradient and Hessian

Recall that if

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T Q\mathbf{x} + \mathbf{c}^T\mathbf{x}$$

then the gradient vector is given by

$$\nabla f(\mathbf{x}) = Q\mathbf{x} + \mathbf{c}.$$

The Hessian matrix is the matrix of all mixed second order partial derivatives $\frac{\partial^2 f}{\partial x_i \partial x_j}$. In this case, the Hessian matrix is given by

$$\nabla^2 f(\mathbf{x}) = Q.$$

Since the Hessian matrix of a quadratic function is constant, the second-order behavior is the same everywhere.

A critical point of f is a point $\mathbf{x}^*$ satisfying

$$Q\mathbf{x}^* = -\mathbf{c}.$$

If Q is nonsingular, then there is a unique critical point given by

$$\mathbf{x}^* = -Q^{-1}\mathbf{c}.$$

However, whether this critical point is a minimum, maximum, or saddle depends on the eigenvalues of Q.

3 Eigenvalues and Eigenvectors

Let $Q \in \mathbb{R}^{n\times n}$ be symmetric. A nonzero vector $\mathbf{v}$ is called an eigenvector of Q with eigenvalue λ if

$$Q\mathbf{v} = \lambda\mathbf{v}.$$

Since Q is symmetric, the Spectral Theorem tells us that Q has an orthonormal basis of eigenvectors. Thus, there exist orthonormal vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ and real numbers $\lambda_1, \ldots, \lambda_n$ such that

$$Q\mathbf{v}_i = \lambda_i \mathbf{v}_i$$

for each $i \in \{1, \ldots, n\}$. If we form the orthogonal matrix

$$V = [\mathbf{v}_1 \; \mathbf{v}_2 \; \cdots \; \mathbf{v}_n],$$

and the diagonal matrix

$$D = \operatorname{diag}(\lambda_1, \ldots, \lambda_n),$$

then

$$Q = VDV^T.$$

This decomposition reveals the geometry of the quadratic. For each $\mathbf{x} \in \mathbb{R}^n$, we make the substitution $\mathbf{y} = V^T\mathbf{x}$. Then, the quadratic takes the form

$$\mathbf{x}^T Q\mathbf{x} = \mathbf{y}^T D\mathbf{y} = \sum_{i=1}^{n} \lambda_i y_i^2.$$

Hence, in the eigenvector coordinates, the quadratic has no cross terms. Thus, the eigenvectors determine the principal directions of the graph, while the eigenvalues determine the curvature in those directions.

Let $i \in \{1, \ldots, n\}$. If $\lambda_i > 0$, then the quadratic curves upward in the direction of $\mathbf{v}_i$. If $\lambda_i < 0$, then the quadratic curves downward in the direction of $\mathbf{v}_i$. If $\lambda_i = 0$, then the quadratic is flat in the direction of $\mathbf{v}_i$. Thus the signs of the eigenvalues determine the overall shape of the quadratic.

4 Convexity and Definiteness

Suppose that $\mathbf{x}^*$ is a critical point of $f(\mathbf{x})$, that is, $Q\mathbf{x}^* = -\mathbf{c}$. Then,

$$\begin{aligned} f(\mathbf{x}) &= \frac{1}{2}\mathbf{x}^T Q\mathbf{x} + \mathbf{c}^T\mathbf{x} \\ &= \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^T Q(\mathbf{x} - \mathbf{x}^*) + f(\mathbf{x}^*). \end{aligned}$$

Thus the behavior of f near $\mathbf{x}^*$ is completely determined by the quadratic form associated with Q. Therefore, if $\mathbf{x}^T Q\mathbf{x} > 0$ for all $\mathbf{x} \neq 0$, then $\mathbf{x}^*$ is an absolute minimum. Similarly, if $\mathbf{x}^T Q\mathbf{x} < 0$ for all $\mathbf{x} \neq 0$, then $\mathbf{x}^*$ is an absolute maximum. These cases are described by the *definiteness* of the Hessian matrix Q.

The matrix Q is *positive definite* if

$$\mathbf{x}^T Q\mathbf{x} > 0$$

for all $\mathbf{x} \neq \mathbf{0}$. Equivalently, every eigenvalue of Q is positive. The matrix Q is *positive semidefinite* if

$$\mathbf{x}^T Q\mathbf{x} \geq 0$$

for all $\mathbf{x}$. Equivalently, every eigenvalue of Q is nonnegative. The matrix Q is *negative definite* if

$$\mathbf{x}^T Q \mathbf{x} < 0$$

for all $\mathbf{x} \neq \mathbf{0}$. Equivalently, every eigenvalue of Q is negative. The matrix Q is *negative semidefinite* if

$$\mathbf{x}^T Q \mathbf{x} \leq 0$$

for all $\mathbf{x}$. Equivalently, every eigenvalue of Q is nonpositive. The matrix Q is *indefinite* if it has both positive and negative eigenvalues.

5 Level Sets of Quadratic Functions

The level set of f corresponding to value α is the set of all points satisfying

$$f(\mathbf{x}) = \alpha.$$

If we rewrite f in centered coordinates, then the level sets satisfy

$$\frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^T Q (\mathbf{x} - \mathbf{x}^*) = \alpha - f(\mathbf{x}^*).$$

After diagonalizing Q, this becomes

$$\frac{1}{2}\sum_{i=1}^{n} \lambda_i y_i^2 = \alpha - f(\mathbf{x}^*).$$

Thus the shape of the level sets is determined by the eigenvalues.

- If all eigenvalues are positive, the level sets are ellipsoids.
- If all eigenvalues are negative, then the graph opens downward, and the level sets are again ellipsoids for appropriate level values.
- If Q is indefinite, the level sets are hyperbolas in two dimensions and hyperboloids in higher dimensions.
- If Q is singular, some directions are flat and the level sets may be cylinders, parallel lines, or parabolic regions.

These level sets are especially important in optimization. For a minimization problem with convex quadratic objective, the level sets are nested ellipses or ellipsoids that shrink toward the minimizer. In constrained optimization, the optimal point often occurs where the smallest feasible level set first touches the feasible region.

Example

Consider the quadratic

$$f(x_1, x_2) = x_1^2 + x_1 x_2 + 2x_2^2.$$

Then

$$Q = \begin{bmatrix} 2 & 1 \\ 1 & 4 \end{bmatrix}.$$

Because of the cross term $x_1 x_2$, the level sets are not aligned with the coordinate axes. Instead, they are rotated ellipses. The axes of these ellipses are determined by the eigenvectors of Q.

The characteristic polynomial is

$$\det(Q - \lambda I) = (2 - \lambda)(4 - \lambda) - 1 = \lambda^2 - 6\lambda + 7,$$

so the eigenvalues are

$$\lambda_1 = 3 - \sqrt{2}, \qquad \lambda_2 = 3 + \sqrt{2}.$$

Both are positive, so Q is positive definite and the quadratic is strictly convex.

For $\lambda_1 = 3 - \sqrt{2}$,

$$Q - \lambda_1 I = \begin{bmatrix} -1 + \sqrt{2} & 1 \\ 1 & 1 + \sqrt{2} \end{bmatrix}, \quad \mathbf{v}_1 = \frac{1}{\sqrt{4 - 2\sqrt{2}}} \begin{bmatrix} 1 \\ 1 - \sqrt{2} \end{bmatrix}.$$

For $\lambda_2 = 3 + \sqrt{2}$,

$$Q - \lambda_2 I = \begin{bmatrix} -1 - \sqrt{2} & 1 \\ 1 & 1 - \sqrt{2} \end{bmatrix}, \quad \mathbf{v}_2 = \frac{1}{\sqrt{4 + 2\sqrt{2}}} \begin{bmatrix} 1 \\ 1 + \sqrt{2} \end{bmatrix}.$$

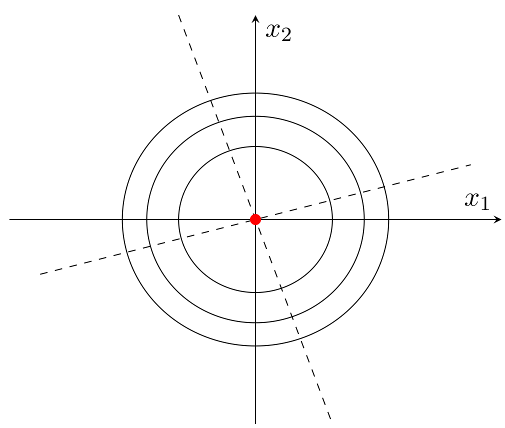

Figure 1: Level sets of $f(x_1, x_2)$. The dashed lines indicate eigenvector directions.

Example

Consider the quadratic

$$f(x_1, x_2) = \frac{1}{2}(x_1 + x_2)^2.$$

Then

$$Q = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

The characteristic polynomial is

$$\det(Q - \lambda I) = \lambda(\lambda - 2),$$

so the eigenvalues are

$$\lambda_1 = 0, \qquad \lambda_2 = 2.$$

Thus Q is positive semidefinite, but not positive definite.

For $\lambda_1 = 0$,

$$Q - \lambda_1 I = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad \mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix}.$$

For $\lambda_2 = 2$,

$$Q - \lambda_2 I = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}, \quad \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

The quadratic has curvature only in the $\mathbf{v}_2$-direction and is flat in the $\mathbf{v}_1$-direction. Hence the level sets are not ellipses. For $\alpha > 0$, the level set

$$f(x_1, x_2) = \alpha$$

is the pair of parallel lines

$$x_1 + x_2 = \pm\sqrt{2\alpha},$$

while for $\alpha = 0$, the level set is the single line

$$x_1 + x_2 = 0.$$

The minimum value is 0, attained at every point on the line

$$x_1 + x_2 = 0.$$

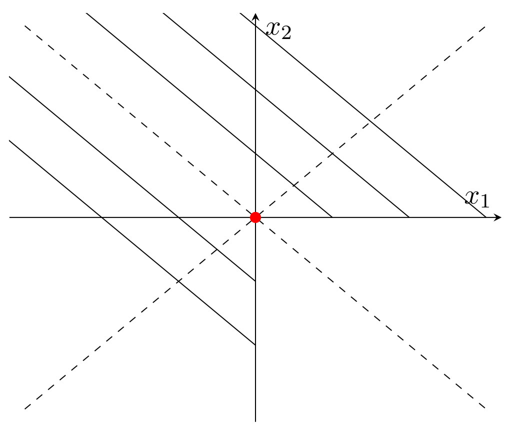

Figure 2: Level sets of $f(x_1, x_2) = \frac{1}{2}(x_1 + x_2)^2$. The dashed lines indicate eigenvector directions.

Example

Consider the quadratic

$$f(x_1, x_2) = x_1^2 - x_2^2.$$

Then

$$Q = \begin{bmatrix} 2 & 0 \\ 0 & -2 \end{bmatrix}.$$

The eigenvalues are

$$\lambda_1 = 2, \qquad \lambda_2 = -2.$$

Since one eigenvalue is positive and the other is negative, Q is indefinite. The origin is a critical point, but it is neither a minimum nor a maximum. Instead, it is a saddle point. Since Q is diagonal, the eigenvector directions coincide with the coordinate axes, and the level sets are hyperbolas aligned with those axes.

For $\lambda_1 = 2$,

$$Q - \lambda_1 I = \begin{bmatrix} 0 & 0 \\ 0 & -4 \end{bmatrix}, \quad \mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

For $\lambda_2 = -2$,

$$Q - \lambda_2 I = \begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{v}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

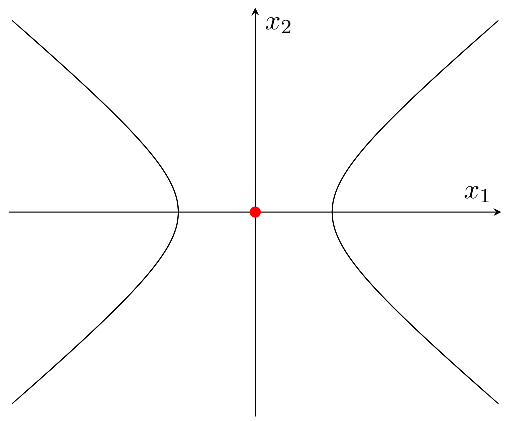

Figure 3: Level sets of $f(x_1, x_2) = x_1^2 - x_2^2$.

6 Summary

The geometry of a quadratic function is governed by the symmetric matrix Q in the term $\frac{1}{2}\mathbf{x}^T Q\mathbf{x}$. The eigenvectors of Q determine the principal directions, and the eigenvalues determine the curvature in those directions. Positive eigenvalues correspond to upward curvature, negative eigenvalues correspond to downward curvature, and zero eigenvalues correspond to flat directions. As a result, the signs of the eigenvalues determine whether the quadratic is convex, concave, or saddle-shaped.

The level sets of a quadratic are translated and rotated versions of the level sets of a diagonal quadratic. Geometrically, the eigenvectors give the directions of the principal axes of the level sets, and the eigenvalues determine how elongated or compressed the level sets are along those axes. When Q is positive definite, the level sets are ellipses or ellipsoids centered at the minimizer. When Q is indefinite, the level sets are hyperbolic. This geometric picture is especially useful in quadratic programming, where the optimal solution can often be understood as the first point at which a level set touches the feasible region.