Lecture 4.9. Positive definite and semidefinite forms

April 10, 2020

Let A be a symmetric matrix, and $Q(\mathbf{x}) = \mathbf{x}^T A \mathbf{x}$ the corresponding quadratic form.

Definitions. *Q and A are called positive* semidefinite *if $Q(\mathbf{x}) \geq 0$ for all $\mathbf{x}$. They are called positive* definite *if $Q(\mathbf{x}) > 0$ for all $\mathbf{x} \neq 0$.*

So positive semidefinite means that there are no minuses in the signature, while positive definite means that there are n pluses, where n is the dimension of the space.

We establish some explicit criteria, in terms of A for positive definite matrices.

Theorem 1. *The following conditions are equivalent:*

a) $\mathbf{x}^T A \mathbf{x} > 0$ *for all* $x \neq 0$.

b) All eigenvalues of A satisfy $\lambda_j > 0$.

c) All NW (upper left) minors of A are positive.

d) In the reduction of A to row echelon form, no exchanges are required, and all pivots are positive.

e) $A = R^T R$ *for some non-singular R.*

    Proof.
a)$\longleftrightarrow$b). We have seen from the spectral theorem for symmetric matrices that there is an orthogonal change of the variable $\mathbf{x} = B\mathbf{y}$ which brings our

form to
$$Q(\mathbf{x}) = \lambda_1 y_1^2 + \lambda_2 y_2^2 + \ldots + \lambda_n y_n^2,$$
where λ_j are eigenvalues of A. It is clear that this sum is positive for all $\mathbf{y} \neq 0$ if and only if all λ_j are positive. The condition $\mathbf{y} \neq 0$ is equivalent to $\mathbf{x} \neq 0$ since B is non-singular.

a), b)$\longrightarrow$c). Determinant of a matrix is the product of eigenvalues. So of all eigenvalues are positive, then determinant is also positive. If we restrict a positive definite form to the span of $e_1, \ldots, e_k$, where $k \leq n$, we obtain a positive definite form, so it must have positive determinant. But this determinant is nothing else but the NW minor of A.

c)$\longrightarrow$d). The proof is based on the following formula for the pivots (in the absence of row exchange):
$$d_k = \frac{\det A_k}{\det A_{k-1}}.$$
Here d_k is the k'th pivot, and A_k is the NW $k \times k$ submatrix of A (it consists of the entries in k first rows and columns of A). (When $k = 1$ one has to set $\det A_0 = 1$ in this formula.)

To prove the formula, notice that row operations (no exchanges!) do not change $\det A_k$. Therefore for all $\det A_k$ to be positive, the row echelon form must have no zeros on the main diagonal, and $\det A_k = d_1 \ldots d_k$. Our formula follows.

d)$\longrightarrow$a). Recall the LDU-factorization:
$$A = LDU,$$
where L is lower triangular, with ones on the main diagonal, U is upper triangular, with ones on the main diagonal, and $D = \mathrm{diag}(d_1, \ldots, d_n)$ the diagonal matrix with pivots on the main diagonal. Moreover, if no row exchanges were required, this factorization is unique. Since the matrix A is symmetric, we must have $U = L^T$, so our factorization becomes
$$A = LDL^T.$$
Since all pivots are positive, we can write $D = \sqrt{D}\sqrt{D}$, where $\sqrt{D} = \mathrm{diag}(sqrt\lambda_1, \ldots, \sqrt{\lambda_n})$, (positive square root!). So
$$A = L\sqrt{D}\sqrt{D}L^T = (L\sqrt{D})(L\sqrt{D})^T,$$

since $(\sqrt{D})^T = \sqrt{D}$, and we can set $R = (L\sqrt{D})^T$. This R is non-singular since both L and $\sqrt{D}$ are non-singular.

e)$\longrightarrow$a). We have

$$Q(\mathbf{x}) = \mathbf{x}^T A \mathbf{x} = \mathbf{x}^T R^T R \mathbf{x} = \mathbf{y}^T \mathbf{y} = \|\mathbf{y}\|^2 > 0,$$

where $\mathbf{y} = R\mathbf{x}$. If $\mathbf{x} \neq 0$ then $\mathbf{y} \neq 0$ since R is non-singular.

This completes the proof of the theorem.

Notice that finding eigenvalues is difficult. The simplest way to check that A is positive definite is to use the condition with pivots d). Condition c) involves more computation but it is still a pure arithmetic condition.

Now we state a similar theorem for positive semidefinite matrices. We need one more

Definition. A *principal minor* of A is the determinant of the matrix obtained by removing some rows and columns with the same numbers from A.

Theorem 2. *The following conditions are equivalent:*

a') $\mathbf{x}^T A \mathbf{x} \geq 0$ *for all* $\mathbf{x}$.

b') *All eigenvalues of A are non-negative,* $\lambda_j \geq 0$.

c') *All* **principal** *minors are non-negative.*

e') $A = R^T R$ *for* **some** R *(maybe singular, maybe even not square).*

I highlighted the differences with Theorem 1. First of all, it is no longer sufficient to check only NW minors. For example, in the matrix

$$\begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix},$$

all NW minors are zero, but it is not positive semidefinite: the corresponding quadratic form is $-x_2^2$. But there is one principal minor equal to -1.

Second, there is no analog of condition d). Since some NW minors can be zero, row exchanges can be required. Row exchanges destroy symmetry of the matrix.

(Strang proposes to do "simultaneous row and column exchanges" to preserve symmetry, but this does not help, for example

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

3

Simultaneous row and column exchange does nothing to this matrix, and does not help to bring it to a row echelon form. The associated form

$$2x_1x_2 = (1/2)(x_1 + x_2)^2 - (1/2)(x_1 - x_2)^2$$

is *not* positive semidefinite. Principal minors are $0, 0, -1$. And there is no pivot condition in this case. So this statement of Strang is a mistake).

The proofs of a)$\longleftrightarrow$b), b)$\longrightarrow$c) and a)$\longleftrightarrow$e) are similar to the proofs of the corresponding statements of Theorem 1. The proof that c) implies all other conditions relies on a formula for *all* coefficients of the characteristic polynomial in terms of principal minors, this formula is somewhat complicated, so I omit this part of the proof.

One important application of these theorems is the study of maxima and minima of functions in calculus. For a function of one variable $f(t)$, we have the necessary condition of extremum at the point t_0: derivative $f'(t_0) = 0$. To obtain a sufficient condition (and to tell a maximum from minimum) we use the second derivative: we can write the first two terms of the Taylor formula,

$$f(t) = f(t_0) + f'(t_0)(t - t_0) + (1/2)f''(t_0)(t - t_0)^2 + O((t - t_0)^2),$$

where $O((t - t_0)^2)$ means the term whose ratio to $(t - t_0)^2$ tends to 0 as $t \to t_0$. So, if $f''(t_0) > 0$, the function has a minimum at t_0, and if $f''(t_0) < 0$ it has a maximum. If $f''(t_0) = 0$, the test in not conclusive.

We have a similar test for functions of several variables. Notice that the term $f''(t_0)(t - t_0)^2$ is a quadratic form of one variable $(t - t_0)$.

In several variables, we have the following Taylor formula

$$f(\mathbf{x}) = f(x_0) + f'(x_0)(\mathbf{x} - \mathbf{x}_0) + (1/2)(\mathbf{x} - \mathbf{x}_0)^T f''(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0) + O(\|\mathbf{x} - \mathbf{x}_0\|^2).$$

Here f' is the row vector

$$f' = (\partial f/\partial x_1, \ldots, \partial f/\partial x_n),$$

known as the gradient, and f'' is the matrix made of partial derivatives $B = (b_{i,j})$, where

$$b_{i,j} = \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

This matrix is symmetric.

The necessary condition of extremum is $f'(\mathbf{x}_0) = 0$.

Our function has a minimum at $\mathbf{x}_0$ when B is positive definite, and maximum when it is negative definite (that is $-B$ is positive definite).

For example, when $n = 2$ we obtain a familiar criterion of a minimum of $f(x, y)$ from Calculus:

$$f_{x,x} > 0, \quad f_{x,x} f_{y,y} - f_{x,y}^2 > 0,$$

which is nothing but the condition that the matrix B is positive definite by criterion c) of Theorem 1.

Another application of Theorem 1 is that it described all possible dot products in $\mathbf{R}^n$. Indeed, a dot product was defined as a function which to every two vectors $\mathbf{x}$ and $\mathbf{y}$ assigns a number $(\mathbf{x}, \mathbf{y})$, and has the following properties:

(i) it is linear with respect to $\mathbf{x}$ and with respect to $\mathbf{y}$,

(ii) symmetric, and

(iii) $(\mathbf{x}, \mathbf{x}) > 0$ for all $\mathbf{x} \neq 0$.

We see that each such function is a symmetric (ii) bilinear (i) form, whose associated quadratic form is positive definite (iii).

Thus any dot product on $\mathbf{R}^n$ is given by the formula

$$(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$$

for some positive definite symmetric matrix A. Notice that all theory of dot product was developed from the axioms (i), (ii), (iii) only, so it applied to *every* possible dot product. So for a fixed symmetric positive definite matrix A we can speak of "A-orthogonality", "A- orthogonal projections", etc.

Some of the theory of bilinear and quadratic forms can be generalized to complex case. To obtain a meaningful theory, one defines *Hermitian* forms $B(\mathbf{x}, \mathbf{y})$ which are linear with respect to $\mathbf{y}$ and *anti-linear* with respect to $\mathbf{x}$. The last condition means that

$$B(c_1 \mathbf{x}_1 + c_2 \mathbf{x}_2, \mathbf{y}) = \overline{c_1} B(\mathbf{x}_1, \mathbf{y}) + \overline{c_2} B(\mathbf{x}_2, \mathbf{y}),$$

and have the following symmetry property:

$$B(\mathbf{x}, \mathbf{y}) = \overline{B(\mathbf{y}, \mathbf{x})}.$$

Each Hermitian form is represented by an Hermitian matrix A by the formula

$$B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^* A \mathbf{x}.$$

To each Hermitian form an quadratic form is associated: $Q(\mathbf{x}) = B(\mathbf{x}, \mathbf{x})$, and this quadratic form can be positive definite or positive semidefinite.

In this course we restrict ourselves to the real case when studying quadratic forms.

Geometric interpretation of positive definite quadratic forms.

Let $Q(\mathbf{x}) = \mathbf{x}^T A \mathbf{x}$ be a positive definite form. The set

$$\{\mathbf{x} \in \mathbf{R}^n : Q(\mathbf{x}) = 1\}$$

is called an *ellipsoid*. For example, when $A = I$, $Q(\mathbf{x}) = x_1^2 + \ldots + x_n^2$, then the ellipsoid is the sphere of radius 1. If B is an orthogonal matrix such that $A = B\Lambda B^T$, and we set $\mathbf{y} = B^T \mathbf{x}$, then the equation of the ellipsoid becomes

$$1 = \mathbf{x}^T A \mathbf{x} = x^T B \Lambda B^T \mathbf{x} = \mathbf{y}^T \Lambda \mathbf{y} = \lambda_1 y_1^2 + \ldots + \lambda_n y_n^2.$$

This can be visualized as a distorted unit sphere: if $z_j = \sqrt{\lambda_j} y_j$, then the unit sphere in the coordinates z_j. To obtain it in coordinates y_j it has to be stretched by the factor of $1/\sqrt{\lambda_j}$ in the direction of j-th axis.

The points $y_j = \pm 1/\sqrt{\lambda_j}$, all other coordinates are zero, are the *intercepts* (intersections of our ellipsoid with y_j axes. The original ellipsoid in x_j-coordinates is related to the ellipsoid in y_j coordinates by an orthogonal transformation which can be always taken to be rotation (if determinant of B is -1, just replace one column of B by its negative, to obtain a new B with determinant 1.

The conclusion is that any ellipsoid can be rotated so that its equation takes the standard form

$$\lambda_1 y_1^2 + \ldots + \lambda_n y_n^2 = 1.$$

The intervals between the two intercepts with the same axis are called *principal axes*. The lengths of these principal axes are $1/\sqrt{\lambda_j}$ and the directions $\mathbf{x}_j$ are determined from the equation $B^T \mathbf{x}_j = e_j$, $\mathbf{x}_j = Be_j$, so $\mathbf{x}_j$ is just the jth column of B, that is the j-th eigenvector.

For example, when $n = 2$, an ellipsoid is just an ellipse centered at the origin. The general equation of such an ellipse is

$$ax^2 + 2bxy + cy^2 = 1.$$

the principal axes are directed along the eigenvectors of

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

and the lengths of the principal semi-axes are $1/\sqrt{\lambda}$.

Ellipsoids are frequently encountered in mechanics: for example the ellipsoid of inertia of a rigid body.

Determination of the signs of eigenvalues of a generic symmetric matrix.

Suppose that a symmetric matrix A is generic in the sense that *no raw exchanges are required when bringing it to the row echelon form.* Then we have the LDU-factorization, and as we have seen in the proof of Theorem 1, d)$\longrightarrow$a), in fact

$$A = LDL^T,$$

which means that A is congruent to the diagonal matrix of pivots. Since we also have $A = B\Lambda B^T$, where Λ is the diagonal matrix of eigenvalues, we conclude from the Law of Inertia for quadratic forms that the number of positive (negative) eigenvalues is equal to the number of positive (negative) pivots. This is useful, because finding eigenvalues is hard, while finding pivots is easy.

Moreover, we can use this fact to compute eigenvalues approximately, with any given precision. Indeed, suppose we know that all eigenvalues of a symmetric matrix A are on an interval $[a, b]$. Then we can determine, how many of them are on the left half of this interval, and how many on the right half. For this we consider the matrix

$$A - (a + b)I/2,$$

whose eigenvalues are $\lambda_j - (a+b)/2$ by the Spectral Mapping Theorem. Here λ_j are eigenvalues of A. Using pivots, we can determine how many of these $\lambda_j - (a+b)/2$ are positive and negative. This gives the numbers of eigenvalues

on each half of $[a, b]$. Then one can continue this procedure by dividing each half of $[a, b]$ into two halves again, and so on.

This is one of the practical methods of localizing eigenvalues of a symmetric matrix.