

The Eigen-Decomposition: Eigenvalues and Eigenvectors

Hervé Abdi¹

1 Overview

Eigenvectors and *eigenvalues* are numbers and vectors associated to square matrices, and together they provide the *eigen-decomposition* of a matrix which analyzes the structure of this matrix. Even though the eigen-decomposition does not exist for all square matrices, it has a particularly simple expression for a class of matrices often used in multivariate analysis such as correlation, covariance, or cross-product matrices. The eigen-decomposition of this type of matrices is important in statistics because it is used to find the maximum (or minimum) of functions involving these matrices. For example, principal component analysis is obtained from the eigen-decomposition of a covariance matrix and gives the least square estimate of the original data matrix.

Eigenvectors and eigenvalues are also referred to as *characteristic vectors and latent roots* or *characteristic equation* (in German, “*eigen*” means “specific of” or “characteristic of”). The set of eigenvalues of a matrix is also called its *spectrum*.

¹In: Neil Salkind (Ed.) (2007). *Encyclopedia of Measurement and Statistics*. Thousand Oaks (CA): Sage.

Address correspondence to: Hervé Abdi

Program in Cognition and Neurosciences, MS: Gr.4.1,

The University of Texas at Dallas,

Richardson, TX 75083-0688, USA

E-mail: herve@utdallas.edu <http://www.utd.edu/~herve>




Figure 1: Two eigenvectors of a matrix.

2 Notations and definition

There are several ways to define eigenvectors and eigenvalues, the most common approach defines an eigenvector of the matrix \mathbf{A} as a vector \mathbf{u} that satisfies the following equation:

$$\mathbf{A}\mathbf{u} = \lambda\mathbf{u} . \quad (1)$$

when rewritten, the equation becomes:

$$(\mathbf{A} - \lambda\mathbf{I})\mathbf{u} = \mathbf{0} , \quad (2)$$

where λ is a scalar called the *eigenvalue* associated to the *eigenvector*.

In a similar manner, we can also say that a vector \mathbf{u} is an eigenvector of a matrix \mathbf{A} if the length of the vector (but not its direction) is changed when it is multiplied by \mathbf{A} .

For example, the matrix:

$$\mathbf{A} = \begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \quad (3)$$

has the eigenvectors:

$$\mathbf{u}_1 = \begin{bmatrix} 3 \\ 2 \end{bmatrix} \quad \text{with eigenvalue } \lambda_1 = 4 \quad (4)$$

and

$$\mathbf{u}_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix} \quad \text{with eigenvalue } \lambda_2 = -1 \quad (5)$$

We can verify (as illustrated in Figure 1) that only the length of \mathbf{u}_1 and \mathbf{u}_2 is changed when one of these two vectors is multiplied by the matrix \mathbf{A} :

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 3 \\ 2 \end{bmatrix} = 4 \begin{bmatrix} 3 \\ 2 \end{bmatrix} = \begin{bmatrix} 12 \\ 8 \end{bmatrix} \quad (6)$$

and

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = -1 \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}. \quad (7)$$

For most applications we normalize the eigenvectors (*i.e.*, transform them such that their length is equal to one):

$$\mathbf{u}^\top \mathbf{u} = 1. \quad (8)$$

For the previous example we obtain:

$$\mathbf{u}_1 = \begin{bmatrix} .8331 \\ .5547 \end{bmatrix}. \quad (9)$$

We can check that:

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} .8331 \\ .5547 \end{bmatrix} = \begin{bmatrix} 3.3284 \\ 2.2188 \end{bmatrix} = 4 \begin{bmatrix} .8331 \\ .5547 \end{bmatrix} \quad (10)$$

and

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} -.7071 \\ .7071 \end{bmatrix} = \begin{bmatrix} .7071 \\ -.7071 \end{bmatrix} = -1 \begin{bmatrix} -.7071 \\ .7071 \end{bmatrix}. \quad (11)$$

Traditionally, we put together the set of eigenvectors of \mathbf{A} in a matrix denoted \mathbf{U} . Each column of \mathbf{U} is an eigenvector of \mathbf{A} . The eigenvalues are stored in a diagonal matrix (denoted $\mathbf{\Lambda}$), where the diagonal elements gives the eigenvalues (and all the other values are zeros). We can rewrite the first equation as:

$$\mathbf{A}\mathbf{U} = \mathbf{\Lambda}\mathbf{U}; \quad (12)$$

or also as:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1}. \quad (13)$$

For the previous example we obtain:

$$\begin{aligned} \mathbf{A} &= \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1} \\ &= \begin{bmatrix} 3 & -1 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 4 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} 2 & 2 \\ -4 & 6 \end{bmatrix} \\ &= \begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix}. \end{aligned} \quad (14)$$

It is important to note that not all matrices have eigenvalues. For example, the matrix $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$ does not have eigenvalues. Even when a matrix has eigenvalues and eigenvectors, the computation of the eigenvectors and eigenvalues of a matrix requires a large number of computations and is therefore better performed by computers.

2.1 Digression:

An infinity of eigenvectors for one eigenvalue

It is only through a slight abuse of language that we can talk about *the* eigenvector associated with *one* given eigenvalue. Strictly speaking, there is an *infinity* of eigenvectors associated to each eigenvalue of a matrix. Because any scalar multiple of an eigenvector is still an eigenvector, there is, in fact, an (infinite) family of eigenvectors for each eigenvalue, but they are all proportional to each other. For example,

$$\begin{bmatrix} 1 \\ -1 \end{bmatrix} \quad (15)$$

is an eigenvector of the matrix \mathbf{A} :

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix}. \quad (16)$$

Therefore:

$$2 \times \begin{bmatrix} 1 \\ -1 \end{bmatrix} = \begin{bmatrix} 2 \\ -2 \end{bmatrix} \quad (17)$$

is also an eigenvector of \mathbf{A} :

$$\begin{bmatrix} 2 & 3 \\ 2 & 1 \end{bmatrix} \begin{bmatrix} 2 \\ -2 \end{bmatrix} = \begin{bmatrix} -2 \\ 2 \end{bmatrix} = -1 \times 2 \begin{bmatrix} 1 \\ -1 \end{bmatrix}. \quad (18)$$

3 Positive (semi-)definite matrices

A type of matrices used very often in statistics are called *positive semi-definite*. The eigen-decomposition of these matrices always exists, and has a particularly convenient form. A matrix is said to be positive semi-definite when it can be obtained as the product of a matrix by its transpose. This implies that a positive semi-definite matrix is always symmetric. So, formally, the matrix \mathbf{A} is positive semi-definite if it can be obtained as:

$$\mathbf{A} = \mathbf{X}\mathbf{X}^T \quad (19)$$

for a certain matrix \mathbf{X} (containing real numbers). Positive semi-definite matrices of special relevance for multivariate analysis positive semi-definite matrices include correlation matrices, covariance, and, cross-product matrices.

The important properties of a positive semi-definite matrix is that its eigenvalues are always positive or null, and that its eigenvectors are pairwise orthogonal when their eigenvalues are different. The eigenvectors are also composed of real values (these last two properties are a consequence of the symmetry of the matrix, for proofs see, *e.g.*, Strang, 2003; or Abdi & Valentin, 2006). Because eigenvectors corresponding to different eigenvalues are orthogonal, it is possible to store all the eigenvectors in an orthogonal matrix (recall that a matrix is orthogonal when the product of this matrix by its transpose is a diagonal matrix).

This implies the following equality:

$$\mathbf{U}^{-1} = \mathbf{U}^T. \quad (20)$$

We can, therefore, express the positive semi-definite matrix \mathbf{A} as:

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T \quad (21)$$

where $\mathbf{U}^T\mathbf{U} = \mathbf{I}$ are the normalized eigenvectors; if they are not normalized then $\mathbf{U}^T\mathbf{U}$ is a diagonal matrix.

For example, the matrix:

$$\mathbf{A} = \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix} \quad (22)$$

can be decomposed as:

$$\begin{aligned} \mathbf{A} &= \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T \\ &= \begin{bmatrix} \sqrt{\frac{1}{2}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{2}} & -\sqrt{\frac{1}{2}} \end{bmatrix} \begin{bmatrix} 4 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} \sqrt{\frac{1}{2}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{2}} & -\sqrt{\frac{1}{2}} \end{bmatrix} \\ &= \begin{bmatrix} 3 & 1 \\ 1 & 3 \end{bmatrix}, \end{aligned} \quad (23)$$

with

$$\begin{bmatrix} \sqrt{\frac{1}{2}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{2}} & -\sqrt{\frac{1}{2}} \end{bmatrix}^T \begin{bmatrix} \sqrt{\frac{1}{2}} & \sqrt{\frac{1}{2}} \\ \sqrt{\frac{1}{2}} & -\sqrt{\frac{1}{2}} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (24)$$

3.1 Diagonalization

When a matrix is positive semi-definite we can rewrite Equation 21 as

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T \iff \mathbf{\Lambda} = \mathbf{U}^T\mathbf{A}\mathbf{U}. \quad (25)$$

This shows that we can transform the matrix \mathbf{A} into an equivalent *diagonal* matrix. As a consequence, the eigen-decomposition of a positive semi-definite matrix is often referred to as its *diagonalization*.

3.2 Another definition for positive semi-definite matrices

A matrix \mathbf{A} is said to be positive semi-definite if we observe the following relationship for any non-zero vector \mathbf{x} :

$$\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0 \quad \forall \mathbf{x}. \quad (26)$$

(when the relationship is ≤ 0 we say that the matrix is negative semi-definite).

When all the eigenvalues of a symmetric matrix are positive, we say that the matrix is *positive definite*. In that case, Equation 26 becomes:

$$\mathbf{x}^T \mathbf{A} \mathbf{x} > 0 \quad \forall \mathbf{x}. \quad (27)$$

4 Trace, Determinant, etc.

The eigenvalues of a matrix are closely related to three important numbers associated to a square matrix, namely its *trace*, its *determinant* and its *rank*.

4.1 Trace

The trace of a matrix \mathbf{A} is denoted $\text{trace}\{\mathbf{A}\}$ and is equal to the sum of its diagonal elements. For example, with the matrix:

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} \quad (28)$$

we obtain:

$$\text{trace}\{\mathbf{A}\} = 1 + 5 + 9 = 15. \quad (29)$$

The trace of a matrix is also equal to the sum of its eigenvalues:

$$\text{trace}\{\mathbf{A}\} = \sum_{\ell} \lambda_{\ell} = \text{trace}\{\mathbf{\Lambda}\} \quad (30)$$

with $\mathbf{\Lambda}$ being the matrix of the eigenvalues of \mathbf{A} . For the previous example, we have:

$$\mathbf{\Lambda} = \text{diag}\{16.1168, -1.1168, 0\}. \quad (31)$$

We can verify that:

$$\text{trace}\{\mathbf{A}\} = \sum_{\ell} \lambda_{\ell} = 16.1168 + (-1.1168) = 15 \quad (32)$$

4.2 Determinant and rank

Another classic quantity associated to a square matrix is its *determinant*. This concept of determinant, which was originally defined as a combinatoric notion, plays an important rôle in computing the inverse of a matrix and in finding the solution of systems of linear equations (the term *determinant* is used because this quantity determines the existence of a solution in systems of linear equations). The determinant of a matrix is also equal to the product of its eigenvalues. Formally, if $|\mathbf{A}|$ the determinant of \mathbf{A} , we have:

$$|\mathbf{A}| = \prod_{\ell} \lambda_{\ell} \text{ with } \lambda_{\ell} \text{ being the } \ell\text{-th eigenvalue of } \mathbf{A} . \quad (33)$$

For example, the determinant of matrix \mathbf{A} (from the previous section), is equal to:

$$|\mathbf{A}| = 16.1168 \times -1.1168 \times 0 = 0 . \quad (34)$$

Finally, the *rank* of a matrix can be defined as being the number of non-zero eigenvalues of the matrix. For our example:

$$\text{rank}\{\mathbf{A}\} = 2 . \quad (35)$$

For a positive semi-definite matrix, the rank corresponds to the dimensionality of the Euclidean space which can be used to represent the matrix. A matrix whose rank is equal to its dimensions is called a *full rank* matrix. When the rank of a matrix is smaller than its dimensions, the matrix is called *rank-deficient*, *singular*, or *multicolinear*. Only full rank matrices have an inverse.

5 Statistical properties of the eigen-decomposition

The eigen-decomposition is important because it is involved in problems of optimization. For example, in principal component

analysis, we want to analyze an $I \times J$ matrix \mathbf{X} where the rows are observations and the columns are variables describing these observations. The goal of the analysis is to find row *factor scores*, such that these factor scores “explain” as much of the variance of \mathbf{X} as possible, and such that the sets of factor scores are pairwise orthogonal. This amounts to defining the factor score matrix as

$$\mathbf{F} = \mathbf{X}\mathbf{P}, \quad (36)$$

under the constraints that

$$\mathbf{F}^T \mathbf{F} = \mathbf{P}^T \mathbf{X}^T \mathbf{X} \mathbf{P} \quad (37)$$

is a diagonal matrix (*i.e.*, \mathbf{F} is an orthogonal matrix) and that

$$\mathbf{P}^T \mathbf{P} = \mathbf{I} \quad (38)$$

(*i.e.*, \mathbf{P} is an orthonormal matrix). There are several ways of obtaining the solution of this problem. One possible approach is to use the technique of the Lagrangian multipliers where the constraint from Equation 38 is expressed as the multiplication with a diagonal matrix of Lagrangian multipliers denoted $\mathbf{\Lambda}$ in order to give the following expression

$$\mathbf{\Lambda} (\mathbf{P}^T \mathbf{P} - \mathbf{I}) \quad (39)$$

(see Harris, 2001; and Abdi & Valentin, 2006; for details). This amount to defining the following equation

$$\mathcal{L} = \mathbf{F}^T \mathbf{F} - \mathbf{\Lambda} (\mathbf{P}^T \mathbf{P} - \mathbf{I}) = \mathbf{P}^T \mathbf{X}^T \mathbf{X} \mathbf{P} - \mathbf{\Lambda} (\mathbf{P}^T \mathbf{P} - \mathbf{I}). \quad (40)$$

In order to find the values of \mathbf{P} which give the maximum values of \mathcal{L} , we first compute the derivative of \mathcal{L} relative to \mathbf{P} :

$$\frac{\partial \mathcal{L}}{\partial \mathbf{P}} = 2\mathbf{X}^T \mathbf{X} \mathbf{P} - 2\mathbf{\Lambda} \mathbf{P}, \quad (41)$$

and then set this derivative to zero:

$$\mathbf{X}^T \mathbf{X} \mathbf{P} - \mathbf{\Lambda} \mathbf{P} = \mathbf{0} \iff \mathbf{X}^T \mathbf{X} \mathbf{P} = \mathbf{\Lambda} \mathbf{P}. \quad (42)$$

Because Λ is diagonal, this is clearly an eigen-decomposition problem, and this indicates that Λ is the matrix of eigenvalues of the positive semi-definite matrix $\mathbf{X}^T \mathbf{X}$ ordered from the largest to the smallest and that \mathbf{P} is the matrix of eigenvectors of $\mathbf{X}^T \mathbf{X}$ associated to Λ . Finally, we find that the factor matrix has the form

$$\mathbf{F} = \mathbf{P} \Lambda^{\frac{1}{2}} . \quad (43)$$

The variance of the factors scores is equal to the eigenvalues:

$$\mathbf{F}^T \mathbf{F} = \Lambda^{\frac{1}{2}} \mathbf{P}^T \mathbf{P} \Lambda^{\frac{1}{2}} = \Lambda . \quad (44)$$

Taking into account that the sum of the eigenvalues is equal to the trace of $\mathbf{X}^T \mathbf{X}$, this shows that the first factor scores “extract” as much of the variances of the original data as possible, and that the second factor scores extract as much of the variance left unexplained by the first factor, and so on for the remaining factors. Incidentally, the diagonal elements of the matrix $\Lambda^{\frac{1}{2}}$ which are the standard deviations of the factor scores are called the *singular values* of matrix \mathbf{X} (see entry on singular value decomposition).

References

- [1] Abdi, H. (2003). Multivariate analysis. In M. Lewis-Beck, A. Bryman, & T. Futing (Eds): *Encyclopedia for research methods for the social sciences*. Thousand Oaks: Sage.
- [2] Abdi, H., Valentin, D. (2006). *Mathématiques pour les sciences cognitives (Mathematics for cognitive sciences)*. Grenoble: PUG.
- [3] Harris, R.j. (2001). *A primer of multivariate statistics*. Mahwah (NJ): Erlbaum.
- [4] Strang, G. (2003). *Introduction to linear algebra*. Cambridge (MA): Wellesley-Cambridge Press.