Eigenvalues and eigenvectors

April 19, 2010

Recap: Matrices and linear transformations of vectors

Last time, we saw that matrices describe linear transformations of vectors:

- The columns of a matrix contain the transformed coordinate axis.
- The original coordinate axis can be stretched, shrunk and rotated
- The image of a rectangle is a parallelogram (or higher-dimensional analog)

A linear transformation

$$\mathbf{M} = \begin{pmatrix} 1 & 1/4 \\ 1/4 & 1 \end{pmatrix}$$

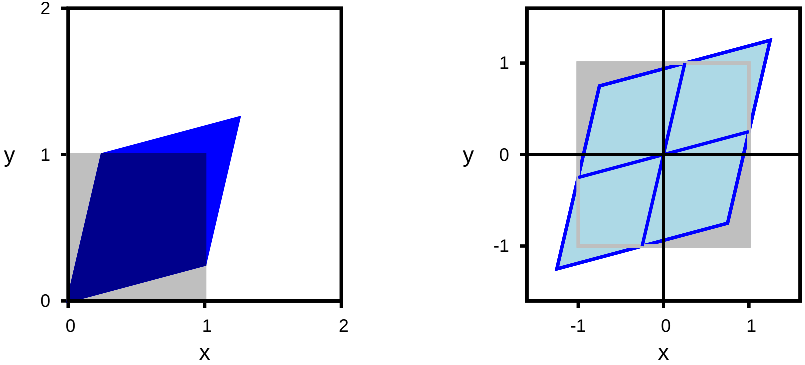

Linear dynamical systems

One of the reasons we are interested in matrices and linear transformation of vectors is that they can be used to describe linear dynamical systems. A **discrete linear dynamical system**

$$
\begin{aligned}
x_1(t+1) &= a_{11}x_1(t) + \ldots a_{1n}x_n(t) \\
&\vdots \\
x_n(t+1) &= a_{n1}x_1(t) + \ldots a_{nn}x_n(t)
\end{aligned}
$$

can be written as

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t)$$

Linear dynamical systems

Similarly, a **continuous linear dynamical system**

$$\frac{dx_1}{dt} = a_{11}x_1(t) + \dots a_{1n}x_n(t)$$
$$\vdots$$
$$\frac{dx_n}{dt} = a_{n1}x_1(t) + \dots a_{nn}x_n(t)$$

can be written as

$$\frac{d\mathbf{x}}{dt} = \mathbf{A}\mathbf{x}(t).$$

In the following, we will first focus on discrete dynamical systems, as these are more directly related to matrix multiplication.

Remarks on linear functions

- A one-dimensional linear function has the form $f(x) = ax + b$. With multiple variables, we have $f(x_i) = a_1x_1 + a_2x_2 + \ldots a_nxn$. The important point is that there are no products of the variables $x_1 \ldots x_n$, nor any powers, roots, sines and cosines or other complicated functions of them.
- Linear functions have the properties that

$$f(x + y) = f(x) + f(y) \text{ and}$$
$$f(cx) = cf(x) \text{ for constant } c.$$

Remarks on linear dynamical systems

- If a dynamical system behaves linearly, cause and effect are directly proportional. Small causes have small effects and large causes have large effects.
- In contrast, in non-linear systems, small causes can have large effects.
- The extreme case of a non-linear system is deterministic chaos, where infinitely small causes can have arbitrarily large effects (the "butterfly effect").
- Linear systems might appear an arbitrary special case. However, their use is justified by Taylor's theorem, which states that, on small scales, any smooth function/system can be approximated by a linear function/system.

Eigenvectors: Vectors that are not rotated

$$\mathbf{M} = \begin{pmatrix} 1 & 1/4 \\ 1/4 & 1 \end{pmatrix}$$

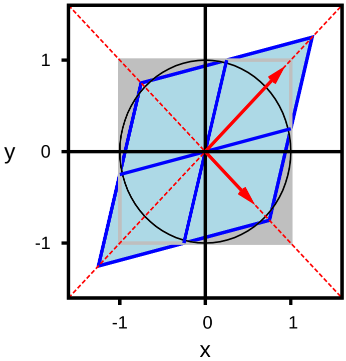

- Vectors that are not rotated under the transformation of a matrix are called **eigenvectors**.
- The factor by which they are stretched (or shrunk) is called **eigenvalue**.
- All multiples of an eigenvector are also eigenvectors.

Examples

$$\mathbf{M} = \begin{pmatrix} 1 & 0.2 \\ 0 & 0.9 \end{pmatrix}$$

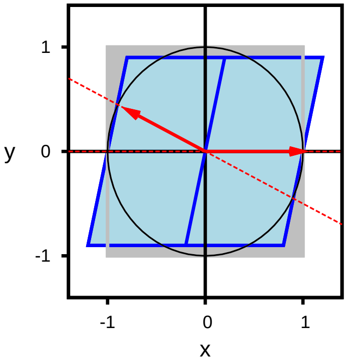

$$\mathbf{M} = \begin{pmatrix} 0.8 & 0 \\ 0 & 1.1 \end{pmatrix}$$

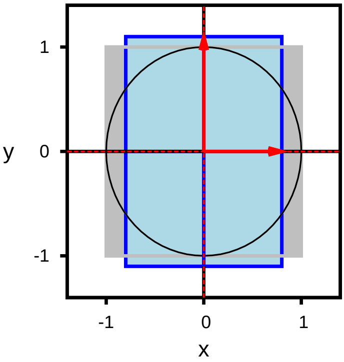

Special cases

$$\mathbf{M} = \begin{pmatrix} 1.2 & 0 \\ 0 & 1.2 \end{pmatrix}$$

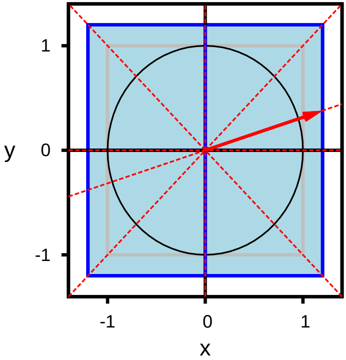

Proportional scaling: Every vector is an eigenvector.

$$\mathbf{M} = \begin{pmatrix} 0.71 & -0.71 \\ 0.71 & 0.71 \end{pmatrix}$$

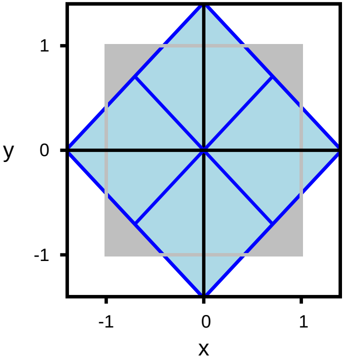

Rotation: No (real) eigenvectors.

The eigenvalue equation

Mathematically, an eigenvector **v** of matrix **M** is defined as a non-zero vector that, when multiplied by the matrix, is only changed by a factor.

$$\mathbf{M}\mathbf{v} = \lambda\mathbf{v}$$

The factor λ is the corresponding eigenvalue.

The characteristic polynomial

It turns out that the eigenvalues can be found without knowledge of the eigenvectors, due to the following calculation involving the identity matrix $\mathbf{I}$:

$$\mathbf{Mv} = \lambda \mathbf{v}$$

$$\mathbf{Mv} - \lambda \mathbf{v} = \mathbf{Mv} - \lambda \mathbf{Iv} = (\mathbf{M} - \lambda \mathbf{I})\mathbf{v} = \mathbf{0}$$

To fulfill this equation (for non-zero $\mathbf{v}$), the matrix $\mathbf{M} - \lambda \mathbf{I}$ must be non-invertible, that is, its determinant must be zero. Thus, the eigenvalues are given by the roots of $\det(\mathbf{M} - \lambda \mathbf{I})$. This determinant is called the **characteristic polynomial**.

The characteristic polynomial of a 2×2 matrix

For a 2×2 matrix $\mathbf{M} = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, we have

$$\det(\mathbf{M} - \lambda \mathbf{I}) = \begin{vmatrix} a - \lambda & b \\ c & d - \lambda \end{vmatrix} =$$
$$(a - \lambda)(d - \lambda) - bc =$$
$$\lambda^2 - \lambda(a + d) + (ad - bc)$$

The eigenvalues of a 2×2 matrix

Solving the previous equation yields the two eigenvalues

$$\lambda_{1/2} = \frac{(a+d) \pm \sqrt{(a+d)^2 - 4(ad-bc)}}{2}$$

Note that $a+d$ is the trace and $ad-bc$ the determinant of $\mathbf{M}$.

If square-root term is zero, the two eigenvalues are identical. For higher-dimensional spaces, eigenvalues may be **repeated** several times.

More about eigenvectors and eigenvalues

- Finding the eigenvalues of a large matrix analytically can be difficult, but there are several useful rules.
- The eigenvalues of a diagonal or triangular matrix are given by the elements along the main diagonal. There are also useful rules for block matrices.
- The sum of the eigenvalues equals the trace of the matrix, and the product of the eigenvalues equals the determinant.

Eigenvectors

Once the eigenvalues are found, the eigenvectors corresponding to each eigenvalue can be determined by solving the system of equations given by the eigenvalue equation $\mathbf{Mv} = \lambda_i \mathbf{v}$.

For a 2×2 matrix $\mathbf{M} = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, we get

$$\mathbf{v} = \begin{pmatrix} b \\ \lambda_i - a \end{pmatrix}.$$

Note that the eigenvectors are determined only up to a constant.

More about eigenvectors

- Different eigenvectors are linearly independent (i.e. they point in different directions; otherwise, they wouldn't be different).
- If there are repeated eigenvalues, a $n \times n$ matrix might have less than n linearly independent eigenvectors. Such a matrix is called *defective*.
- For a non-defective matrix, the eigenvectors form an alternative coordinate system.
- If the matrix is symmetric, the eigenvectors are orthogonal. The new coordinate system is then obtained by a simple rotation of the original one.

Eigenvectors as alternative coordinate system

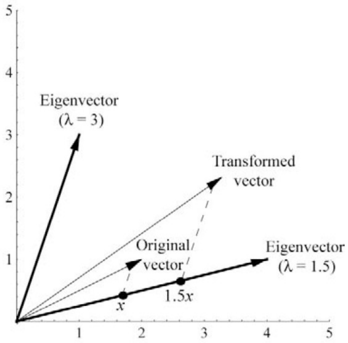

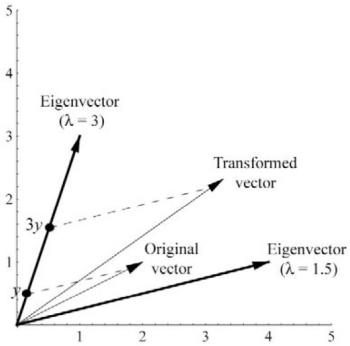

When projected onto a coordinate system given by the eigenvectors, each component of a matrix is multiplied by the respective eigenvalue.

Eigenvectors as alternative coordinate system

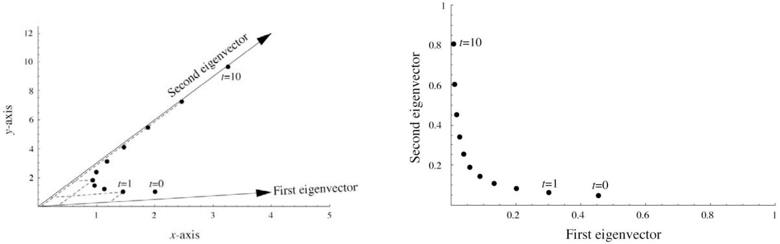

In the transformed coordinate system, the behavior of a linear system is much simpler, as the variables grow or shrink **independently**.

How to do this transformation mathematically?

Let **M** be a $n \times n$ matrix defining a dynamical system $\mathbf{x}(t+1) = \mathbf{M}\mathbf{x}(t)$. Let **A** be the matrix whose columns are the eigenvectors **v** of **M**, and **D** be a diagonal matrix with the corresponding eigenvalues λ:

$$\mathbf{A} = \begin{pmatrix} v_{11} & \dots & v_{1n} \\ \vdots & \vdots & \vdots \\ v_{n1} & \dots & v_{nn} \end{pmatrix}, \quad \mathbf{D} = \begin{pmatrix} \lambda_1 & 0 \dots 0 & & \\ 0 & \lambda_2 & \dots 0 & \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{pmatrix}$$

Diagonalization

With the above definitions,

$$\mathbf{MA} = \mathbf{AD}$$

and

$$\mathbf{M} = \mathbf{ADA}^{-1}$$

This is called a **diagonalization** of **M**. It can only be done if **M** if non-defective, that is, has n linearly independent eigenvectors.

It turns out the rows of $\mathbf{A}^{-1}$ are the *left* eigenvectors of **M** (defined by $\mathbf{u}^T\mathbf{M} = \lambda\mathbf{u}^T$, see below).

Coordinate transformation

With the previous result, we can rewrite our dynamical system as

$$\mathbf{x}(t+1) = \mathbf{M}\mathbf{x}(t) = \mathbf{A}\mathbf{D}\mathbf{A}^{-1}\mathbf{x}(t)$$

This can be interpreted as follows (proofs in Otto and Day, ch. 9):

- $\mathbf{A}^{-1}\mathbf{x}(t)$ performs the transformation into the coordinate system defined by the eigenvectors.
- In this coordinate system, the dynamics of the transformed variable is described by the matrix $\mathbf{D}$ (which is very simple).
- Finally, multiplying with $\mathbf{A}$ transforms the result back into the original coordinate system.

Eigenvectors of symmetric matrices

$$\mathbf{M} = \begin{pmatrix} 1 & 1/4 \\ 1/4 & 1 \end{pmatrix}$$

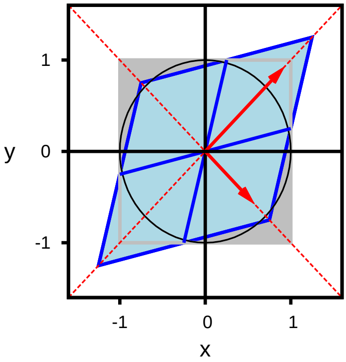

If a matrix **M** is symmetric ($m_{ij} = m_{ji}$), its eigenvectors are orthogonal to each other. The transformation described above it then a simple rotation of the original coordinate system.

Aside: Principal components analysis

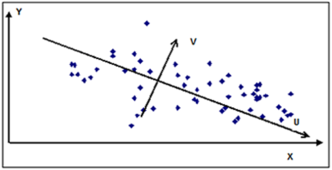

- The principal components of a data set are given by the eigenvectors of the covariance matrix. Their length if proportional to the eigenvalues.
- Along the PCs, the data are uncorrelated.
- The direction of the first PC explains a maximum of variation.
- The second PC explains a maximum of the remaining (uncorrelated) variation, and so on.

Principal components analysis (cont.)

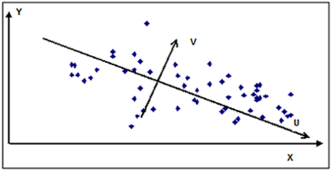

- The first PC can be seen as a form a linear regression that does not assume independent and dependent variables.
- Often, the large PCs can be interpreted in biologically meaningful terms.
- The smaller PCs can be neglected (data reduction).
- As traits are defined arbitrarily, one may as well define them along the PCs.