6

Eigendecomposition

In this chapter we are going to look into the makeup of linear transformations by doing some computations with the matrices that represent them. We will find out that in many cases of interest, we will be able to choose a basis where the matrix that represents the linear transformation is particularly simple. The techniques presented in this chapter are practically simple to define and implement, but their implications pervade all data science. In particular, we will lay the foundations for speaking about Principal Component Analysis (PCA) and Singular Value Decomposition (SVD).

6.1 Eigenstuff

Given a linear transformation

$$\begin{aligned} F\colon \mathbb{R}^n &\to \mathbb{R}^n \\ v &\mapsto Mv \end{aligned}$$

we will say that a non-zero vector v is an *eigenvector* with *eigenvalue* λ if

$$f(v) = Mv = \lambda v.$$

From the point of view of an eigenvector, the effect of the linear transformation is just rescaling.

The set of all eigenvectors with eigenvalue λ along with the zero vector, form a vector space known as the *eigenspace* of λ:

$$E_\lambda = \{v \in \mathbb{R}^n \mid Mv = \lambda v\}$$

This is why instead of being concerned about finding specific eigenvectors, we will rather find entire eigenspaces associated with eigenvalues.

6.2 Finding eigenstuff

Eigenvectors and eigenvalues have a particular fingerprint that make them easy to find and collect. Suppose that $v \in \mathbb{R}^n$ is an eigenvector of f with eigenvalue λ. Then

$$Mv = \lambda v \Rightarrow Mv - \lambda v = \overrightarrow{\mathbf{o}}$$

$$\Rightarrow (M - \lambda\mathbf{I})v = \overrightarrow{\mathbf{o}} \Rightarrow v \in N(M - \lambda\mathbf{I})$$

Consequently, whenever there is a non-zero null-space of the form $N(M - \lambda\mathbf{I})$, there you have the eigenvectors with eigenvalues λ.

6.2.1 Finding eigenvalues

The first step is to find all the possible such eigenvalues. A scalar λ is an eigenvalue if and only if $N(M - \lambda\mathbf{I}) \neq \vec{\mathbf{0}}$, i.e., $\mathbf{dim}(N(M - \lambda\mathbf{I})) > 0$ which is the same to say that $M - \lambda\mathbf{I}$ is not full rank (by Theorem 1). Then all we have to do is to find all the λ's such that $M - \lambda\mathbf{I}$ is not full rank, which in turn is equivalent to

$$\mathbf{det}(M - \lambda\mathbf{I}) = 0.$$

The *characteristic polynomial* of a matrix M is the polynomial in one variable t, denote $Q(t)$, is defined as $\mathbf{det}(M - t\mathbf{I})$.

Proposition 24. *The eigenvalues of f are the roots of the characteristic polynomial of its associated matrix.*

Since the characteristic polynomial of M, $Q(t)$, is a polynomial of degree n, the equation $Q(t) = 0$ will have at most n different real solutions (roots of $Q(t)$).

Note that we do not have to be worried about which matrix representation of f to use. Even though the matrix associated to f will be different in coordinates of different bases, the characteristic polynomial of each matrix representing f will always be the same.

Suppose that $M' = CMC^{-1}$ is the matrix of f in coordinates of another basis, with C the change of basis matrix. Then:

$$\mathbf{det}(M' - t\mathbf{I}) = \mathbf{det}(CMC^{-1} - t\mathbf{I}) =$$

$$= \mathbf{det}(CMC^{-1} - tC\mathbf{I}C^{-1}) = \mathbf{det}(C(M - t\mathbf{I})C^{-1}) =$$

$$= \mathbf{det}(C)\mathbf{det}(M - t\mathbf{I})\mathbf{det}(C^{-1}) = \mathbf{det}(M - t\mathbf{I})$$

If λ is a root of $Q(t)$, then $Q(t)$ can factor as:

$$Q(t) = (t - \lambda)P(t),$$

where $P(t)$ is a degree $n - 1$ polynomial.

For every root λ of $Q(t)$ we can define its *multiplicity* as the maximum k such that $Q(t)$ can factor as:

$$Q(t) = (t - \lambda)^k P(t),$$

where $P(t)$ is a degree $n - k$ polynomial.

There is a trick to compute eigenvalues of lower size matrices that we can state as a proposition:

Proposition 25. *If A is a 2×2 matrix, their eigenvalues λ_1, λ_2 can be computed solving the following system of non-linear equations:*

$$\begin{cases} \lambda_1 + \lambda_2 = \textbf{\textit{trace}}(A) \\ \lambda_1 \lambda_2 = \textbf{\textit{det}}(A) \end{cases}$$

6.2.2 Finding eigenspaces

Computing eigenspaces amounts to computing null-spaces of matrices. If λ is an eigenvalue of f, then we will compute the eigenspace associated to λ as:

$$N(M - \lambda\mathbf{I}) = \{v \in \mathbb{R}^n \mid (M - \lambda\mathbf{I})v = \vec{\mathbf{0}}\}$$

Using Gauss-Jordan elimination we will be able to provide a basis for each eigenspace, thus completely characterizing the eigenvectors associated with each possible eigenvalue.

To summarize the steps for computing eigenstuff:

1. Compute the characteristic polynomial as

$$Q(t) = \mathbf{det}(M - t\mathbf{I})$$

2. Find $\lambda_1, \ldots, \lambda_k$ the roots of $Q(t)$

3. For each λ_i compute a basis for $N(M-\lambda_i\mathbf{I})$ using Gauss-Jordan elimination.

How do eigenvectors of different values relate to one another?

Proposition 26. *If $v_1, \ldots, v_k$ are eigenvectors with pairwise different eigenvalues, they are linearly independent.*

The case $k = 2$ is relatively simple. If v_1, v_2 are eigenvectors (non-zero) with different eigenvalues $\lambda_1 \neq \lambda_2$, they cannot be collinear, since that would lead to a contradiction:

$$v_1 = \lambda v_2 \Rightarrow$$

$$\lambda_1 v_1 = f(v_1) = \lambda f(v_2) = \lambda\lambda_2 v_2 = \lambda_2 v_1$$

We will reason by induction: assuming that we have proved that given any i eigenvectors with different eigenvalues they must be linearly independent, we will prove that the same holds true for $i+1$ eigenvectors. Since we have already proved it for the case of two vectors, it must hold for three, four... and any number of eigenvectors.

Suppose that $v_1, \ldots, v_i, v_{i+1}$ are not linearly independent. Reindexing if necessary, we know that

$$v_{i+1} = \alpha_1 v_1 + \ldots + \alpha_i v_i$$

By applying f we get, on the one hand,

$$f(v_{i+1}) = \lambda_{i+1} v_{i+1} = \lambda_{i+1}\alpha_1 v_1 + \ldots + \lambda_{i+1}\alpha_i v_i$$

and on the other,

$$f(v_{i+1}) = \alpha_1 f(v_1) + \ldots + \alpha_i f(v_i) = \alpha_1\lambda_1 v_1 + \ldots + \alpha_i\lambda_i v_i$$

Given the fact that $v_1, \ldots, v_i$ are linearly independent, their coefficients must be the same in both experssions of $f(v_{i+1})$, therefore

$$\begin{cases} \lambda_{i+1}\alpha_1 = \lambda_1\alpha_1 \\ \quad\ldots \\ \lambda_{i+1}\alpha_i = \lambda_i\alpha_i \end{cases}$$

Since $v_{i+1} \neq \vec{0}$, $\alpha_j \neq 0$ for some $1 \leq j \leq i$, implying that $\lambda_{i+1} = \lambda_j$. But this contradicts the fact that the $v_1, \ldots, v_{i+1}$ have pairwise different eigenvalues.

We can make an even stronger claim in the case where the matrix of the linear transformation is symmetric.

Proposition 27. *If $v_1, \ldots, v_k$ are eigenvectors of a symmetric linear transformation with pairwise different eigenvalues, they are orthogonal.*

Take two eigenvectors v_1, v_2 with different eigenvalues and let's check that their dot product vanishes. The idea is that, because $M^t = M$, then $Mv_1 \cdot v_2 = v_1 \cdot Mv_2$. On the one hand,

$$v_2^t M v_1 = v_2 \cdot \lambda_1 v_1 = \lambda_1 v_1 \cdot v_2$$

By taking the transpose of the previous expression, which should give the same result because we are taking the transpose of a 1×1 matrix,

$$(v_2^t M v_1)^t = v_1^t M^t v_2 = v_1^t M v_2 = v_1 \cdot \lambda_2 v_2 = \lambda_2 v_1 \cdot v_2$$

Therefore $(\lambda_1 - \lambda_2)(v_1 \cdot v_2) = 0$. Since $\lambda_1 \neq \lambda_2$, $v_1 \cdot v_2 = 0$ follows.

6.3 Matrix diagonalization

The matrix of a linear transformation $f : \mathbb{R}^n \longrightarrow \mathbb{R}^n$ would be very simple if expressed in coordinates of a basis $v_1, \ldots, v_n$ of $\mathbb{R}^n$ made exclusively of eigenvectors of f, also known as an *eigenbasis* of f. In fact, the matrix would be diagonal, and diagonal matrices are extremely confortable to work with:

$$\begin{bmatrix} \lambda_1 & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & \lambda_n \end{bmatrix}$$

Although it is not always the case that such an eigenbasis exists for a given f, we will study the conditions where this is possible. We will say that a linear transformation is *diagonalizable* if it has an eigenbasis.

6.3.1 Checking diagonalization

As usual, we are given the matrix M and asked to verify whether it is diagonalizable. Suppose that we have found the eigenvalues $\lambda_1, \ldots, \lambda_k$ and their respective eigenspaces $E_{\lambda_i} = N(M - \lambda_i \mathbf{I})$.

Proposition 28. *M is diagonalizable if and only if*

$$\sum_{i=1}^{k} \boldsymbol{dim}(E_{\lambda_i}) = n.$$

In this situation, if $\mathscr{B}_i$ is a basis of E_{λ_i}, then $\mathscr{B} = \mathscr{B}_1 \cup \ldots \cup \mathscr{B}_k$ is a basis of $\mathbb{R}^n$.

Here are some examples of linear transformations that are diagonalizable:

- Projections
- Symmetries
- Dilation/Contraction

The following matrices, however, are not diagonalizable:

$$A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

$$B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

Can you explain why?

6.3.2 Symmetric case

We saw previously that symmetric matrices enforce stronger conditions on the eigenspaces, in the sense that eigenspaces with distinct eigenvalues turn out to be orthogonal to each other. Symmetric matrices come out often in many applications and – brace yourselves – they will show when we speak about differential calculus with many variables.

Theorem 2 (Spectral theorem for symmetric matrices)**.** *If M is symmetric, the following holds:*

- *The characteristic polynomial of M factorizes as:*

$$Q(t) = \prod_{i=1}^{k} (t - \lambda_i)^{k_i}.$$

- ***dim**$(E_{\lambda_i}) = k_i$ for $1 \leq i \leq k$; in particular, M is diagonalizable.*

- *If $\mathscr{B}_i$ is an orthonormal basis of E_{λ_i}, then*

$$\mathscr{B} = \mathscr{B}_1 \cup \ldots \cup \mathscr{B}_k$$

  is an orthonormal basis of $\mathbb{R}^n$. In particular, the encoder-decoder matrices of $\mathscr{B}$ are orthogonal matrices.

6.3.3 Positive semi-definite case

A particularly interesting case study arises when considering matrices that can be expressed as a matrix product of the form $A^t A$, as is the case of *covariance matrices*, which will be introduced later in chapter 8.

A rank $n \times n$ matrix M is referred to as *positive semi-definite* if it can be expressed as a matrix product $M = A^t A$ where A is row full-rank, i.e., A has $\mathbf{rank}(A)$ rows.

A positive semi-definite matrix is automatically symmetric:

$$M^t = (A^t A)^t = A^t (A^t)^t = A^t A = M.$$

Additionally, all its eigenvalues must be non-negative.

Proposition 29. *For a symmetric matrix M the following statements are equivalent:*

1. *M is positive semi-definite*
2. *M satisfies $v^t M v \geq 0$ for any column vector $v \in \mathbb{R}^n$.*
3. *All the eigenvalues of M are non-negative.*

To see 1) $\Rightarrow$ 2) note the following chain of identities:

$$v^t M v = v^t A^t A v = (Av)^t (Av) = Av \cdot Av = \|Av\|^2 \geq 0$$

To see 2) $\Rightarrow$ 3) suppose that v is an eigenvector of M with eigenvalue λ. Then

$$v^t M v = v^t (\lambda v) = \lambda \|v\|^2 \geq 0,$$

which forces $\lambda \geq 0$. To see 3) $\Rightarrow$ 1) we can arrange the non-negative eigenvalues in a diagonal matrix D, with

$$d_{11} \geq \ldots \geq d_{nn} \geq 0.$$

Since M is symmetric, we can express D as $Q^t M Q$ with Q orthogonal. Since all the entries of D are non-negative, D can be factored using its square root: a diagonal matrix, denoted $\sqrt{D}$, whose entries are just the square roots of those of D. With all this in mind we can express M as

$$M = Q^t \sqrt{D}\sqrt{D}Q = (\sqrt{D}Q)^t(\sqrt{D}Q).$$

To conclude, taking A to be the matrix whose rows are the rows of $\sqrt{D}Q$ corresponding to the non-zero entries of $\sqrt{D}$, we have a decomposition $M = A^t A$ with A having $\mathbf{rank}(A)$ rows.

6.4 TL;DR

Among all the possible bases one can choose to represent linear transformations from $\mathbb{R}^n$ to $\mathbb{R}^n$, also known as "endomorphisms", eigenbases, whenever available, are arguably the best. The only caveat is that this is not always possible. Linear transformations for which this is possible are referred to as "diagonalizable" and there are quite a few interesting examples meeting this criterion, i.e. symmetric endomorphisms.