Eigenvalues & Eigenvectors

L16

This is a core concept in linear algebra.
It's used to study, among other things:

- Difference equations
- Stochastic processes
- Graphs & networks
- Differential equations

We'll focus on difference equations & stochastic processes as applications. Eigenvalues also play a large role in the SVD.

This is also one of the most subtle concepts in the class.

Running (Hopping?) Example:

[DEMO]

In a population of rabbits,

- 1/4 survive their 1st year
- 1/2 survive their 2nd year
- Max lifespan is 3 years
- 1-year-old rabbits have an average of 13 babies
  2-year-old rabbits have an average of 12 babies

This year there are

16 babies
6 1-year-olds
1 2-year-old

Problem: Describe the long-term behavior of the system, both qualitatively & quantitatively.

Let's give names to the state of the system in year k:

$$\left.\begin{aligned} x_k &= \#\text{babies} \\ y_k &= \#\text{1-year-olds} \\ z_k &= \#\text{2-year-olds} \end{aligned}\right\} \text{in year } k \xrightarrow{\text{vector}} v_k = \begin{pmatrix} x_k \\ y_k \\ z_k \end{pmatrix}$$

The rules say:

$$\boxed{\begin{aligned} x_{k+1} &= \quad 13y_k + 12z_k \\ y_{k+1} &= \tfrac14 x_k \\ z_{k+1} &= \quad \tfrac12 y_k \end{aligned}} \qquad \boxed{\begin{aligned} x_0 &= 16 \\ y_0 &= 6 \\ z_0 &= 1 \end{aligned}}$$

state change — initial state

Write as a matrix equation:

$$v_{k+1} = \underbrace{\begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix}}_{A} v_k \qquad v_0 = \begin{pmatrix} 16 \\ 6 \\ 1 \end{pmatrix}$$

What happens in 100 years?

$$v_{100} = A \cdot v_{99} = A(Av_{98}) = \cdots = A^{100} v_0$$

Great! Are we done? No!

→ Computing Av_k 100 times is a lot of work/flops!

→ We get no qualitative understanding that way!

We want to extract the "32:4:1" ratio directly from the matrix A.

Def: A difference equation is a matrix equation of the form

$$v_{k+1} = A v_k \quad \text{with } v_0 \text{ fixed}$$

where:

- $v_k \in \mathbb{R}^n$ is the state of the system at time k
- $v_0 \in \mathbb{R}^n$ is the initial state
- A is an $n \times n$ matrix called the state change matrix

As above, $\boxed{v_k = A^k v_0}$

So a difference equation models a system that "changes state in a linear way."

Solving a difference equation means computing and describing $v_k = A^k v_0$ as $k \to \infty$.

Running Example: Note that if

$$v_0 = \begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix} \rightsquigarrow v_1 = Av_0 = \begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix}\begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix} = \begin{pmatrix} 64 \\ 8 \\ 2 \end{pmatrix} = 2v_0$$

$$\rightsquigarrow v_2 = Av_1 = A(2v_0) = 2(Av_0) = 2(2v_0) = 4v_0$$

$$\rightsquigarrow v_3 = Av_2 = A(4v_0) = 4(Av_0) = 4(2v_0) = 8v_0$$

$$\vdots$$

$$v_k = Av_{k-1} = \cdots = 2^k v_0$$

In this case, it's easy to compute and describe v_k: the population exactly doubles each year!

!! If $Av = \lambda v$ for some scalar λ, then

$$\boxed{A^k v = \lambda^k v} \quad \text{for all } k$$

Of course, in our running example, $v_0 = (16, 6, 1)$ and $Av_0 \neq \lambda v_0$, so how does this help?

→ Answer: diagonalization (next time)

Def: An eigenvector of a square matrix A is a nonzero vector v such that

$$Av = \lambda v \quad \text{for a scalar } \lambda$$

The scalar λ is the eigenvalue.
We also say that v is a λ-eigenvector.

♫ eigenvector song ♫

Upshot so far: If v is a λ-eigenvector of A then $A^k v = \lambda^k v$ is easy to compute.

Eg:
$$\begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix} \begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix} = \begin{pmatrix} 64 \\ 8 \\ 2 \end{pmatrix} = 2 \begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix}$$

This says $\begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix}$ is an eigenvector with eigenvalue 2

Eg: If $Av = v$ $(v \neq 0)$ then v is a 1-eigenvector: $Av = 1v$.
So the nonzero vectors that A doesn't move are the 1-eigenvectors.

Eg: If $Av = 0$ $(v \neq 0)$ then v is a 0-eigenvector: $Av = 0v$.
→ 0 is a valid eigenvalue, not a valid eigenvector.
So the nonzero vectors in $\mathrm{Nul}(A)$ are the 0-eigenvectors

NB: If we allowed 0 to be an eigenvector then every number would be its eigenvalue: $A0 = \lambda 0$ for any $\lambda \in \mathbb{R}$.

Geometrically, "$Av = \lambda v$" means "v and Av are collinear".

A rotates eigenvectors by $0°$ or $180°$

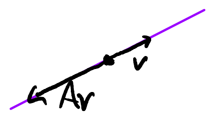

Eg: $A = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$ $\quad A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -x \\ y \end{pmatrix}$: flip over the y-axis.
Where are the eigenvectors?

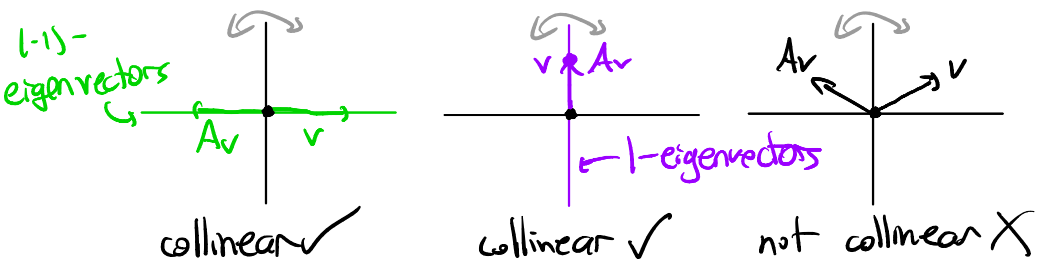

[DEMO]

Eg: $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ is called a shear:

$A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x+y \\ y \end{pmatrix}$ so vectors above the x-axis move right & the vectors below move left.

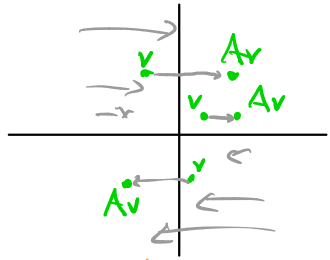

shear
(think "scissors")

Where are the eigenvectors?

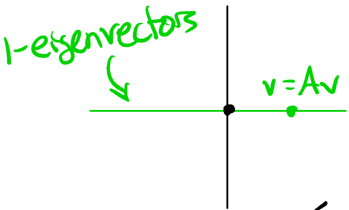

collinear ✓

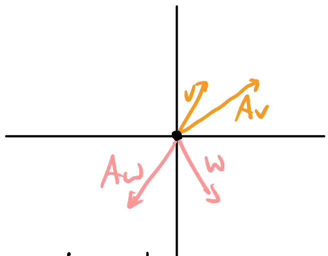

not collinear X

[DEMO]

Eg: $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ $A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -y \\ x \end{pmatrix}$: this is CCW rotation by 90°.

Where are the eigenvectors?

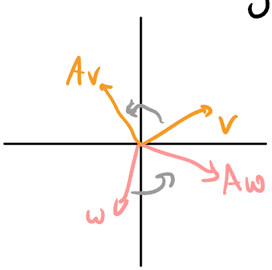

This matrix has no (real) eigenvectors!

[DEMO]

Eigenspaces

Given an eigenvalue λ, how to compute all λ-eigenvectors?

$$\begin{aligned} Av = \lambda v &\iff Av - \lambda v = 0 \\ &\iff Av - \lambda I_n v = 0 \\ &\iff (A - \lambda I_n)v = 0 \\ &\iff v \in \mathrm{Nul}(A - \lambda I_n) \end{aligned}$$

Summary: given λ,

$$\boxed{\text{solving } Av = \lambda v \iff \text{solving } (A - \lambda I_n)v = 0}$$

Def: Let λ be an eigenvalue of an $n \times n$ matrix A. The **λ-eigenspace** of A is

$$\begin{aligned} \mathrm{Nul}(A - \lambda I_n) &= \{\text{all solutions of } Av = \lambda v\} \\ &= \{\text{all } \lambda\text{-eigenvectors and } 0\} \end{aligned}$$

Eg: $A = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}$

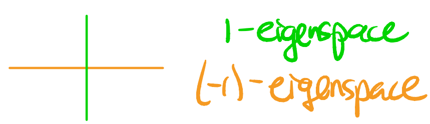

Eg: $A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$

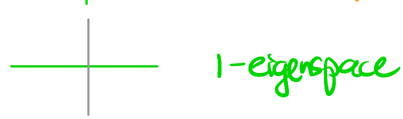

NB: An eigenvalue has infinitely many eigenvectors — a whole (nonzero) subspace worth of them!

NB: If $\mathrm{Nul}(A-\lambda I_n)=\{0\}$ then the only solution of $Av=\lambda v$ is 0 — so λ is not an eigenvalue.

Eg: Find the $\lambda=2$-eigenspace of $A=\begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix}$.

We have to solve $(A-2I_3)x=0$.

$$A-2I_3=\begin{pmatrix} -2 & 13 & 12 \\ 1/4 & -2 & 0 \\ 0 & 1/2 & -2 \end{pmatrix} \xrightarrow{\text{rref}} \begin{pmatrix} 1 & 0 & -32 \\ 0 & 1 & -4 \\ 0 & 0 & 0 \end{pmatrix}$$

$$\xrightarrow{\text{PVF}} x = x_3\begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix}$$

So the 2-eigenspace is

$$\mathrm{Nul}(A-2I_3)=\mathrm{Span}\left\{\begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix}\right\}.$$

Hence all 2-eigenvectors of the rabbit population matrix are scalar multiples of $(32,4,1)$.

[DEMO]

Eg: Find the $\lambda=(-1)$-eigenspace of $A=\begin{pmatrix} -1 & 0 & 0 \\ -1 & 0 & 2 \\ -1 & 1 & 1 \end{pmatrix}$

We have to solve $(A+I_3)x=0$.

$$A+I_3=\begin{pmatrix} 0 & 0 & 0 \\ -1 & 1 & 2 \\ -1 & 1 & 2 \end{pmatrix} \xrightarrow{\text{rref}} \begin{pmatrix} 1 & -1 & -2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

[DEMO]

$$\xrightarrow{\text{PVF}} x = x_2\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} + x_3\begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}$$

So the (-1)-eigenspace is the plane $\mathrm{Span}\left\{\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}\right\}$.

Eg: Find the $\lambda = 0$-eigenspace of $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$

We have to solve $(A - 0I_3)x = 0$, ie, $Ax = 0$.

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \rightsquigarrow \mathrm{Nul}(A) = \mathrm{Span}\left\{ \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix} \right\}$$

This is the 0-eigenspace!

Remember, $Ax = 0x$ means $Ax = 0$, so

$$\boxed{\mathrm{Nul}(A) \text{ is the } 0\text{-eigenspace}}$$

Recall: A square matrix A is invertible $\iff \mathrm{Nul}(A) = \{0\}$.

This just means $Ax = 0x$ has no nonzero solutions.

$$\boxed{A \text{ is invertible} \iff 0 \text{ is not an eigenvalue}}$$

Eg: Find the $\lambda = 3$-eigenspace of $A = 3I_n$.

In this case, $A - 3I_n = 0$! That means

$\mathrm{Nul}(A - 3I_n) = \mathbb{R}^n$.

But of course every nonzero vector is a 3-eigenvector of A:

$$(3I_n)v = 3I_n v = 3v \text{ for any } v \checkmark$$

Eg: Let V be a subspace of $\mathbb{R}^n$. Find the eigenvectors of the projection matrix P_V.

Recall: $x \in V \iff P_V x = 1x$

$x \in V^\perp \iff P_V x = 0x$

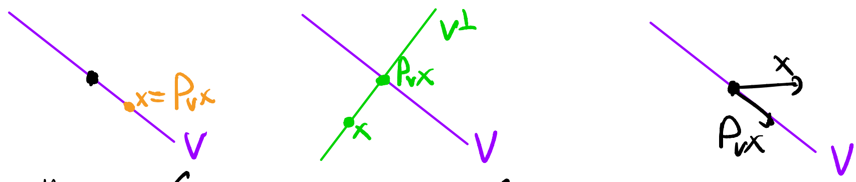

collinear ✓   collinear ✓   not collinear ✗

$$\boxed{\begin{array}{l}\text{The 1-eigenspace of } P_V \text{ is } V\\ \text{The 0-eigenspace of } P_V \text{ is } V^\perp\end{array}}$$

[DEMO]

The Characteristic Polynomial

or: the reason we covered determinants

Given an eigenvalue λ, we know how to find the λ-eigenvectors: $\mathrm{Nul}(A-\lambda I_n)$. How do we find the eigenvalues though?

$$
\begin{aligned}
\lambda \text{ is an eigenvalue of } A \quad & \\
& \iff Av = \lambda v \text{ has a nonzero solution} \\
& \iff (A-\lambda I_n)v = 0 \text{ has a nonzero solution} \\
& \iff \mathrm{Nul}(A-\lambda I_n) \neq \{0\} \\
& \iff A-\lambda I_n \text{ is not invertible} \\
& \iff \boxed{\det(A-\lambda I_n) = 0}
\end{aligned}
$$

This is an equation in λ whose solutions are the eigenvalues!

Eg: Find all eigenvalues of $A = \begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix}$.

We have to solve $\det(A-\lambda I_3) = 0$.

$$\det(A-\lambda I_3) = \det\begin{pmatrix} -\lambda & 13 & 12 \\ 1/4 & -\lambda & 0 \\ 0 & 1/2 & -\lambda \end{pmatrix}$$

$$\overset{\text{expand}}{\underset{\text{Cofactors}}{=}} -\lambda \det\begin{pmatrix} -\lambda & 0 \\ 1/2 & -\lambda \end{pmatrix} + \frac{1}{4}(-1)\det\begin{pmatrix} 13 & 12 \\ 1/2 & -\lambda \end{pmatrix} = \cdots$$

$$= -\lambda^3 + \frac{13}{4}\lambda + \frac{3}{2}$$

So what are the solutions of $p(\lambda) = -\lambda^3 + \frac{13}{4}\lambda + \frac{3}{2} = 0$?

Ask a computer:

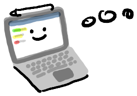

$$-\lambda^3 + \frac{13}{4}\lambda + \frac{3}{2} = -(\lambda-2)\left(\lambda+\frac{1}{2}\right)\left(\lambda+\frac{3}{2}\right)$$

So the eigenvalues are $\lambda = 2, -\frac{1}{2}, -\frac{3}{2}$.

Def: The **characteristic polynomial** of A is

$$p(\lambda) = \det(A - \lambda I_n).$$

The upshot is

$$\boxed{\lambda \text{ is an eigenvalue of } A \iff p(\lambda) = 0}$$