Eigenvalues & Eigenvectors

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

- Difference equations
- Differential equations
- Markov chains
- Stochastic processes

We will focus on difference equations & differential equations as applications, and we'll also need it to understand the SVD.

It also may be the most subtle set of ideas in the whole class, so pay attention!

Unlike orthogonality, I can motivate eigenvalues with an example right off the bat.

Running Example

In a population of rabbits:

[demo]

- 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, and 1 2-year-old.

Problem: Describe the long-term behavior of this system.

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

x_k = # babies in year k
y_k = # 1-year-olds in year k
z_k = # 2-year-olds in year k

$$v_k = \begin{pmatrix} x_k \\ y_k \\ z_k \end{pmatrix}$$

The rules say:

state change
$$\begin{cases} x_{k+1} = 13y_k + 12z_k \\ y_{k+1} = \frac{1}{4}x_k \\ z_{k+1} = \frac{1}{2}y_k \end{cases}$$

initial state
$$\begin{cases} x_0 = 16 \\ y_0 = 6 \\ z_0 = 1 \end{cases}$$

As a matrix equation,

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

What happens in 100 years?

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

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 $n \times n$ (square) matrix called the state change matrix

So in a difference equation, the state at time $k+1$ is related to the state at time k by a matrix multiplication.

> Solving a difference equation means computing & describing $A^k v_0$ for large values of k.

NB: Difference equations are a very common application. Google's PageRank is a difference equation! (But not in an obvious way.)

NB: Multiplying $A \cdot v_k$ requires n multiplications and $n-1$ additions for each coordinate, so $\approx 2n^2$ flops. If $n = 10{,}000$ and $k = 1{,}000$ this is 100 gigaflops! Plus we get no qualitative understanding of v_k for $k \to \infty$. We need to be more clever.

Observation: If $v_0 = (32, 4, 1)$ instead then

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

So

$$v_2 = A^2 v_0 = A(Av_0) = A(2v_0) = 2Av_0 = 2^2 v_0$$

$$v_3 = A^3 v_0 = A(A^2 v_0) = A(2^2 v_0) = 2^2 A v_0 = 2^3 v_0$$

$$\vdots$$

$$v_k = A^k v_0 = 2^k v_0$$

$$\boxed{\text{If } Av = \lambda v \text{ for a scalar } \lambda, \text{ then } A^k v = \lambda^k v \text{ for all } k}$$

This is easy to compute! And to describe.

Next time: What if $Av \neq (\text{scalar}) \cdot v$? Diagonalization.
(eg. $v_0 = (16, 6, 1)$ above)

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 associated eigenvalue.
We also say v is a λ-eigenvector

♫ eigenvector song ♫

> If v is an eigenvector of A with eigenvalue λ then $A^k v = \lambda^k v$ is easy to compute.

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

So $(32, 4, 1)$ is an eigenvector with eigenvalue 2.

→ This means if you start with 32 babies, 4 1-year rabbits, and 1 2-year rabbit, then the population exactly doubles each year.

Geometrically, an eigenvector of A is a nonzero vector v such that Av lies on the line thru the origin and v.

$$\boxed{A \text{ rotates eigenvectors by } 0° \text{ or } 180°}$$

Eg: $A = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$ $\quad A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -x \\ y \end{pmatrix}$: flip over y-axis.

Where are the eigenvectors?

- $v = \begin{pmatrix} x \\ 0 \end{pmatrix} \rightsquigarrow Av = \begin{pmatrix} -x \\ 0 \end{pmatrix} = -v$

  The (nonzero) vectors on the x-axis are eigenvectors with eigenvalue -1.

  Av & v are on the same line.

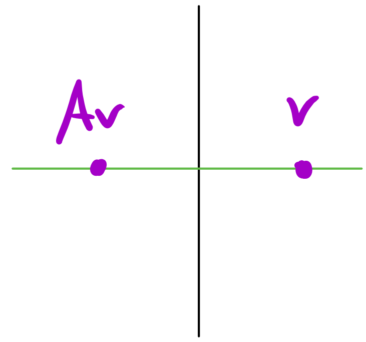

- $v = \begin{pmatrix} 0 \\ y \end{pmatrix} \rightsquigarrow Av = \begin{pmatrix} 0 \\ y \end{pmatrix} = 1 \cdot v$

  The (nonzero) vectors on the y-axis are eigenvectors with eigenvalue $+1$.

  Av & v are on the same line.

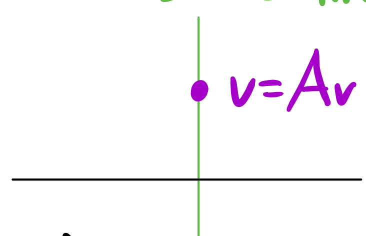

- $v = \begin{pmatrix} x \\ y \end{pmatrix}$ with $x, y \neq 0$

  $\rightsquigarrow Av = \begin{pmatrix} -x \\ y \end{pmatrix}$ is not a multiple of v

  Av & v are on different lines.

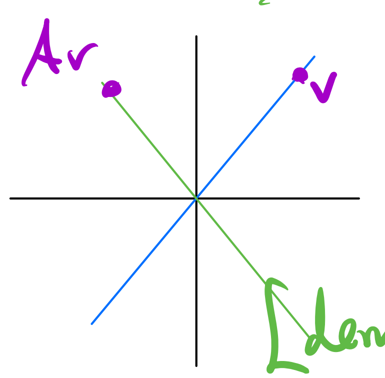

So we've found all eigenvectors [& eigenvalues].

[demo]

Eg: $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ $\quad A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x+y \\ y \end{pmatrix}$: a shear

Where are the eigenvectors?

- $v = \begin{pmatrix} x \\ 0 \end{pmatrix} \rightsquigarrow Av = \begin{pmatrix} x \\ 0 \end{pmatrix} = v$

  The (nonzero) vectors on the x-axis are eigenvectors with eigenvalue 1.

Av & v are on the same line.

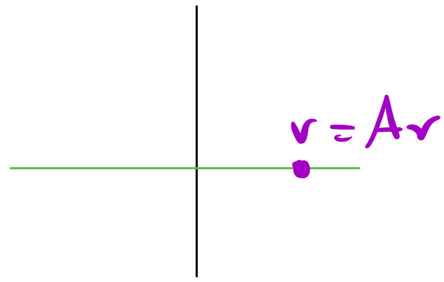

- $v = \begin{pmatrix} x \\ y \end{pmatrix}$ with $x, y \neq 0$:

  $$Av = \begin{pmatrix} x+y \\ y \end{pmatrix}$$

  This is not a multiple of v because $1 = \frac{y}{y} \neq \frac{x+y}{y}$.

Av & v are on different lines.

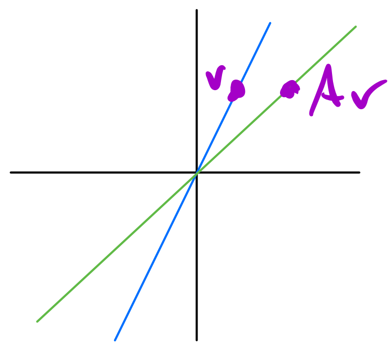

[demo]

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

There are no (real) eigenvectors:

v & Av are never on the same line (unless $v = 0$).

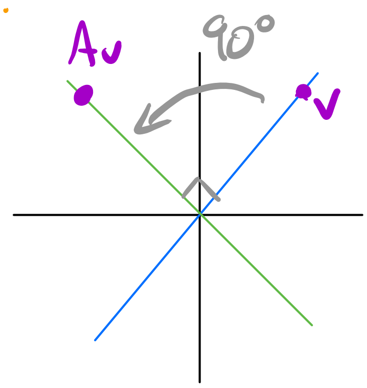

[demo]

Eigenspaces

Given an eigenvalue λ, how do you compute the λ-eigenvectors?

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

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

$$\mathrm{Nul}(A - \lambda I_n) = \{v \in \mathbb{R}^n : Av = \lambda v\}$$
$$= \{\text{all } \lambda\text{-eigenvectors and } 0\}$$

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

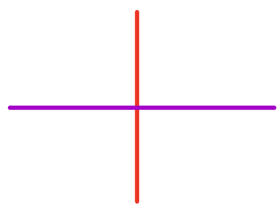

1-eigenspace

(-1)-eigenspace

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

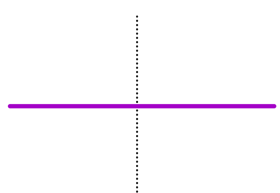

1-eigenspace

Eg: $A = \begin{pmatrix} 0 & 13 & 12 \\ 1/4 & 0 & 0 \\ 0 & 1/2 & 0 \end{pmatrix} \qquad \lambda = 2$

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

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

This line is the 2-eigenspace:
all 2-eigenvectors are multiples of $\begin{pmatrix} 32 \\ 4 \\ 1 \end{pmatrix}$ [demo]

Eg: $A = \begin{pmatrix} -1 & 0 & 0 \\ -1 & 0 & 2 \\ -1 & 1 & 1 \end{pmatrix} \qquad \lambda = -1$

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

$$\xrightarrow{\text{PVF}} \mathrm{Nul}(A - (-1)I_3) = \mathrm{Span}\left\{ \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix} \right\}$$

This plane is the (-1)-eigenspace.
All (-1)-eigenvectors are linear combinations of $\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}$. [demo]

NB: If λ is an eigenvalue then there are infinitely many λ-eigenvectors: the λ-eigenspace is a nonzero subspace.
(This means $A - \lambda I_n$ has a free variable.)

Eg: $A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$ $\lambda = 0 \rightsquigarrow A - \lambda I_3 = A$!

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \overset{\text{rref}}{\rightsquigarrow} \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & 2 \\ 0 & 0 & 0 \end{pmatrix} \overset{\text{PVF}}{\rightsquigarrow} \operatorname{Span}\left\{ \begin{pmatrix} 1 \\ -2 \\ 1 \end{pmatrix} \right\}$$

This line is the 0-eigenspace

NB: 0 is a legal eigenvalue (not an eigenvector) and

$$(\text{0-eigenspace}) = \operatorname{Nul}(A - 0 I_n) = \operatorname{Nul}(A)$$
$$= \{ x \in \mathbb{R}^n : Ax = 0x \}$$

$$\boxed{\text{The 0-eigenspace is the null space}}$$

So if 0 is an eigenvalue of A then $\operatorname{Nul}(A) \neq \{0\}$, so A is not invertible (not FCR).

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

Eg: Let V be a subspace of $\mathbb{R}^n$, P_V the projection matrix. What are the eigenvectors & eigenvalues?

- $P_V b = 1b \iff b = b_V \iff b \in V$

  V is the 1-eigenspace

- $P_V b = 0 = 0b \iff b \in V^\perp$

  $V^\perp$ is the 0-eigenspace [demo]

The Characteristic Polynomial

Given an eigenvalue λ of A, we know how to find all λ-eigenvectors: $\mathrm{Nul}(A-\lambda I_n)$.

How do we find the eigenvalues of A?

Eg: $A = \begin{pmatrix} -1 & 0 & 0 \\ -1 & 0 & 2 \\ -1 & 1 & 1 \end{pmatrix}$ $\lambda = 1$

$$A - 1I_3 = \begin{pmatrix} -2 & 0 & 0 \\ -1 & -1 & 2 \\ -1 & 1 & 1 \end{pmatrix} \xrightarrow{\text{ref}} \begin{pmatrix} -2 & 0 & 0 \\ 0 & -1 & 2 \\ 0 & 0 & 3 \end{pmatrix}$$

This has full column rank: $\mathrm{Nul}(A - 1I_3) = \{0\}$.

This means 1 is not an eigenvalue of A.

Indeed, λ is an eigenvalue of A

$\iff$ $Av = \lambda v$ has a nonzero solution v

$\iff$ $(A - \lambda I_n)v = 0$ has a nonzero solution

$\iff$ $\mathrm{Nul}(A - \lambda I_n) \neq \{0\}$

$\iff$ $A - \lambda I_n$ is not invertible

$\iff$ $\det(A - \lambda I_n) = 0$

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

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

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

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

$$= -\lambda^3 - \frac{1}{4}(-13\lambda - 6) = -\lambda^3 + \frac{13}{4}\lambda + \frac{3}{2}$$

We need to find the zeros (roots) of a cubic polynomial:

$$p(\lambda) = -\lambda^3 + \frac{13}{4}\lambda + \frac{3}{2} = 0$$

Ask a computer:

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

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

Def: The characteristic polynomial of an $n \times n$ matrix A is $p(\lambda) = \det(A - \lambda I_n)$

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