An **eigenvector** is a direction the matrix does not turn: $A\mathbf{v} = \lambda\mathbf{v}$. Apply A again and again and you just get $A^n\mathbf{v} = \lambda^n\mathbf{v}$.

Example: $A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$

$$A\begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 3 \\ 3 \end{pmatrix} = 3\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad A\begin{pmatrix} 1 \\ -1 \end{pmatrix} = \begin{pmatrix} 1 \\ -1 \end{pmatrix} = 1\begin{pmatrix} 1 \\ -1 \end{pmatrix}$$

Eigenvalues $\lambda = 3$ and $\lambda = 1$.

Start anywhere, e.g. $\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{2}\begin{pmatrix} 1 \\ 1 \end{pmatrix} + \frac{1}{2}\begin{pmatrix} 1 \\ -1 \end{pmatrix}$:

$$A^5\begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{2}\cdot 3^5\begin{pmatrix} 1 \\ 1 \end{pmatrix} + \frac{1}{2}\begin{pmatrix} 1 \\ -1 \end{pmatrix} = \begin{pmatrix} 122 \\ 121 \end{pmatrix}$$

Almost exactly along $(1,1)$: **the biggest eigenvalue takes over.**

What λ tells you

λ	λ^n as n grows	In real life
$\lambda > 1$	grows without limit	population boom, epidemic
$\lambda = 1$	stays the same	steady state, long-run share
$0 < \lambda < 1$	shrinks to 0	dies out, settles down
$\lambda < 0$	flips sign each step	see-saw, oscillation

Every example in this row uses this one idea: **find the special direction, read off how fast it grows.**

Also inside AI models: training passes numbers through many layers, like multiplying by matrices again and again. If the eigenvalues are too big the numbers explode; too small and they vanish.

Why does the **biggest eigenvalue take over**? Split the starting vector into eigenvector parts. Each part just gets multiplied by its own λ, again and again.

The theory (any 2×2 with two eigenvectors)

Eigenvectors $\mathbf{v}_1, \mathbf{v}_2$ with $|\lambda_1| > |\lambda_2|$. Write the start as

$$\mathbf{x}_0 = c_1\mathbf{v}_1 + c_2\mathbf{v}_2.$$

Each time we apply A, each part is scaled by its λ:

$$A^n\mathbf{x}_0 = c_1\lambda_1^n\mathbf{v}_1 + c_2\lambda_2^n\mathbf{v}_2 = \lambda_1^n\left(c_1\mathbf{v}_1 + c_2\left(\tfrac{\lambda_2}{\lambda_1}\right)^n\mathbf{v}_2\right)$$

Since $\left|\frac{\lambda_2}{\lambda_1}\right| < 1$, the bracket $\to c_1\mathbf{v}_1$. So

$$A^n\mathbf{x}_0 \approx c_1\lambda_1^n\,\mathbf{v}_1 :$$

direction $\to \mathbf{v}_1$, **growth factor** $\to \lambda_1$.

Same thing as diagonalisation: $A^n = PD^nP^{-1}$, and D^n is dominated by λ_1^n.

A simple case: $A = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}$, $\lambda = 3, 1$

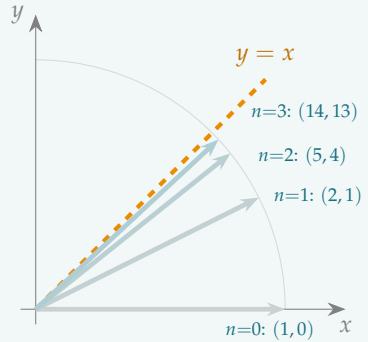

Arrows drawn the same length, so only the direction shows.

$(1,0) = \frac{1}{2}(1,1) + \frac{1}{2}(1,-1)$, so

$$A^n\begin{pmatrix}1\\0\end{pmatrix} = \tfrac{1}{2}3^n\begin{pmatrix}1\\1\end{pmatrix} + \tfrac{1}{2}1^n\begin{pmatrix}1\\-1\end{pmatrix}$$

The 3^n part grows; the 1^n part stays tiny. The arrows swing onto $y = x$: angles $0°$, $26.6°$, $38.7°$, $42.9°$, $44.3°, \cdots \to 45°$.

When a picture is sheared (slanted), almost every arrow in it turns. **The arrows that don't turn are the eigenvectors.**

Before and after the shear

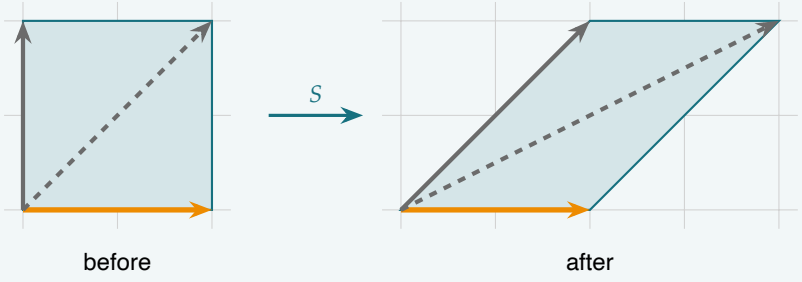

The grey arrows turn. **The orange arrow along the bottom does not**: it keeps its direction and its length.

The maths

$$S = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$ (slides each point sideways by its height)

$$\det(S - \lambda I) = (1-\lambda)^2 = 0 \;\Rightarrow\; \lambda = 1$$

$$(S - I)\mathbf{v} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \mathbf{0} \;\Rightarrow\; y = 0$$

Eigenvector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ with $\lambda = 1$: horizontal arrows are untouched.

Only **one** eigenvector direction here. A shear has no second special direction.

Where you see it: italic fonts and the "slant" tool in photo apps are shears. Anything lying along the eigenvector direction is left exactly as it was.

When something spins, every point moves except the ones on the **axis**. The axis is an **eigenvector with** $\lambda = 1$.

3D: the Earth turning a quarter turn

$$R = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{(rotation about the } z\text{-axis)}$$

$$R \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \qquad \Rightarrow \quad \lambda = 1$$

The North Pole stays put. That direction is the Earth's axis.

Euler (1775): every rotation in 3D has an axis like this, so every 3D rotation has an eigenvector with $\lambda = 1$. Phones and drones use this to describe how they are turning.

2D: a quarter turn has no eigenvector

$$Q = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

$$\det(Q - \lambda I) = \lambda^2 + 1 = 0$$

No real solution: **every arrow in the plane turns**, so there is no real eigenvector.

The eigenvalues are complex: $\lambda = \pm i$.

Complex eigenvalues mean rotation. In our last chapter (coupled differential equations) they are the reason some solutions spiral.

Google's first search ranking (PageRank, 1998): **a page is important if important pages link to it.** The ranking is an **eigenvector with** $\lambda = 1$.

The idea: a random surfer

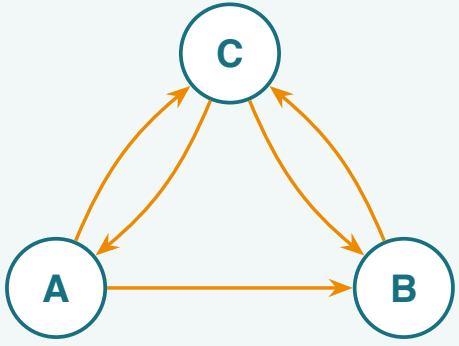

Three web pages. Each arrow is a link.

Someone clicks links at random, forever. On each page they pick one of its links with equal chance.

Where do they spend most of their time? That page ranks first.

The maths: a Markov chain

Column = *from*, row = *to*:

$$T = \begin{pmatrix} 0 & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & 1 & 0 \end{pmatrix} \begin{matrix} \leftarrow A \\ \leftarrow B \\ \leftarrow C \end{matrix}$$

Long-run share of time $\mathbf{s}$: solve $T\mathbf{s} = \mathbf{s}$, i.e. $(T - I)\mathbf{s} = \mathbf{0}$.

$$a = \tfrac{1}{2}c, \qquad b = \tfrac{1}{2}a + \tfrac{1}{2}c, \qquad a + b + c = 1$$

$$\mathbf{s} = \begin{pmatrix} 2/9 \\ 3/9 \\ 4/9 \end{pmatrix} \approx \begin{pmatrix} 0.22 \\ 0.33 \\ 0.44 \end{pmatrix}$$

Ranking: C, then B, then A. $\mathbf{s}$ is the eigenvector of T for $\lambda = 1$.

Real web: billions of pages, so nobody solves the equations. Google multiplied a starting vector by T again and again: $T^n\mathbf{s}_0 \to \mathbf{s}$, because the $\lambda = 1$ eigenvector takes over.

Ecologists track animals by life stage with a matrix. The **largest eigenvalue** is the long-run growth rate; its **eigenvector** is the age mix the population settles into.

The model: juveniles J and adults A

Each year: every adult has 3 young; half the juveniles grow up; half the adults survive.

$$\begin{pmatrix} J_{n+1} \\ A_{n+1} \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & 3 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}}_{L} \begin{pmatrix} J_n \\ A_n \end{pmatrix}$$

$$\det(L - \lambda I) = \lambda^2 - \tfrac{1}{2}\lambda - \tfrac{3}{2} = \tfrac{1}{2}(2\lambda - 3)(\lambda + 1) = 0$$

$$\lambda = \tfrac{3}{2} \quad \text{or} \quad \lambda = -1$$

For $\lambda = \frac{3}{2}$: $-\frac{3}{2}J + 3A = 0$, so $\mathbf{v} = \begin{pmatrix} 2 \\ 1 \end{pmatrix}$.

What happens: start with 10 adults

Year	J	A	J/A	growth
0	0	10	–	–
1	30	5	6.00	$\times 3.50$
2	15	17.5	0.86	$\times 0.93$
3	52.5	16.3	3.23	$\times 2.12$
10	680	350	1.94	$\times 1.47$
20	39 891	19 956	2.00	$\times 1.50$

It wobbles at first, then settles: **2 juveniles per adult** (the eigenvector) and **growth of 50% a year** ($\lambda = \frac{3}{2}$).

Real use: a 1987 study of loggerhead sea turtles used this method and found that protecting older juveniles at sea mattered more than protecting eggs on beaches.

Each Fibonacci number is the sum of the two before: $1, 1, 2, 3, 5, 8, 13, \ldots$ **The golden ratio is an eigenvalue.**

Write the rule as a matrix

$$\begin{pmatrix} F_{n+1} \\ F_n \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}}_{M} \begin{pmatrix} F_n \\ F_{n-1} \end{pmatrix}$$

$$\det(M - \lambda I) = \lambda^2 - \lambda - 1 = 0$$

$$\lambda = \frac{1+\sqrt{5}}{2} \approx 1.618 \quad \text{or} \quad \lambda = \frac{1-\sqrt{5}}{2} \approx -0.618$$

The big one, $\varphi \approx 1.618$, is the **golden ratio**. The small one has $|\lambda| < 1$, so its part dies away.

So each term $\approx 1.618\times$ the last

n	F_n	F_n / F_{n-1}
3	2	2.000
4	3	1.500
5	5	1.667
6	8	1.600
8	21	1.615
10	55	1.618
12	144	1.618

The ratio settles at φ, the largest eigenvalue, just as the population settled at $\lambda = \frac{3}{2}$.

Why it matters: any rule "next term from the last few terms" (loan balances, signal filters, rabbit populations) becomes a matrix, and its largest eigenvalue gives the long-run growth.

Epidemiologists' "R number" (how many people each case infects) is the **largest eigenvalue** of a matrix. $R > 1$: the outbreak grows. $R < 1$: it dies out.

Two groups: children and adults

Column = who is ill, row = who they infect (on average):

$$K = \begin{pmatrix} 1.5 & 0.5 \\ 1 & 1 \end{pmatrix} \begin{matrix} \leftarrow \text{children} \\ \leftarrow \text{adults} \end{matrix}$$

$$\det(K - \lambda I) = \lambda^2 - 2.5\lambda + 1 = (\lambda - 2)(\lambda - 0.5) = 0$$

$$R = \lambda = 2, \qquad \mathbf{v} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

Cases double each generation, settling to equal numbers of children and adults.

Start with one ill child

Generation	Children	Adults	Total
1	1.5	1	2.5
2	2.75	2.5	5.25
3	5.38	5.25	10.6
6	42.7	42.7	85.3

Doubling each time ($\lambda = 2$), with children : adults $\to 1:1$ (the eigenvector).

Vaccinate half of everyone: every entry of K halves, so λ halves to 1. The outbreak stops growing.

Real use: during COVID-19, modellers estimated R this way from contact surveys split by age group, and used it to judge which restrictions mattered most.

Data with many numbers per person is hard to see. The **eigenvectors of its covariance matrix** are the directions where the data spreads out most.

Maths and physics test scores

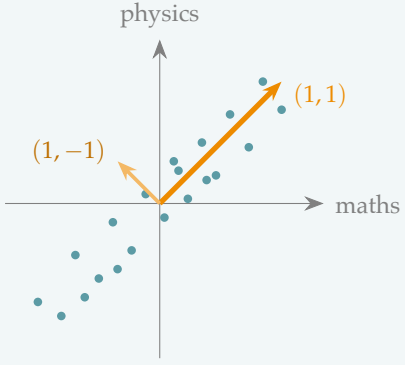

Students who do well in one tend to do well in the other, so the points lie along a slanted band.

The maths

Covariance matrix (scores standardised):

$$C = \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix}$$

$$\det(C - \lambda I) = (1-\lambda)^2 - 0.64 = 0 \;\Rightarrow\; \lambda = 1.8 \text{ or } 0.2$$

$\lambda = 1.8$: eigenvector $(1, 1)$, "good at both".
$\lambda = 0.2$: eigenvector $(1, -1)$, "better at one than the other".

$\frac{1.8}{1.8+0.2} = 90\%$ of the spread lies along $(1, 1)$. One number per student keeps most of the information.

Where it's used: face recognition ("eigenfaces"), compressing images, and drawing maps of how an AI language model groups words: thousands of numbers per word, squashed onto the top two eigenvectors.

Coming in our last chapter: coupled differential equations. There, $\frac{d\mathbf{x}}{dt} = A\mathbf{x}$, and the eigenvalues of A decide the whole shape of the motion.

Both negative

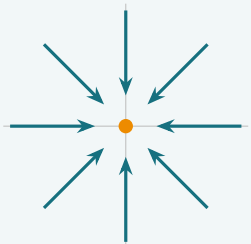

Everything settles to rest.
A swing slowing down.

Opposite signs

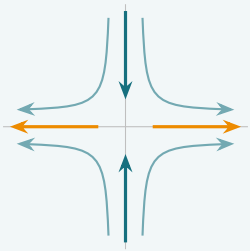

A "saddle": in along one eigenvector, out along the other.

Complex

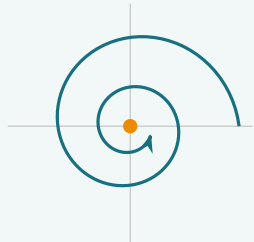

Spirals, like the quarter turn with $\lambda = \pm i$.
Predators and prey.

Engineering: a bridge's natural wobble frequencies come from eigenvalues. London's Millennium Bridge swayed on its opening day in 2000 because crowds' footsteps matched one of them; it closed for nearly two years to be fixed.