Linear Transformations: What Do They Do? (continued)

ECE 57000 — September 9, 2026

David I. Inouye

Milestone 1: let stakeholder needs shape your idea

The goal is a credible match between a stakeholder's needs and something you can provide.

Two Starting Paths

- **Stakeholder first (recommended):** identify people, research their work, then explore their needs.
- **Idea first (also okay):** bring a tentative idea to relevant people and let their priorities reshape it.

Your idea is a starting hypothesis. Seek a problem they care about solving.

Reach Out Early

1. Learn enough about their context to ask useful questions.
2. Invite a brief email exchange or call, if they are interested.
3. Ask about needs, pain points, and what would help.
4. Offer a few concise ideas as conversation starters; refine them together as you learn.

Design thinking makes learning from people part of the process

Mode	What You Do
Empathize	Learn about people’s work, experiences, and priorities.
Define	Frame a specific need in its real context.
Ideate	Explore several ways to address that need.
Prototype	Make a small, inexpensive version people can react to.
Test	Observe what helps; revise the idea or the problem definition.

Return to earlier modes as you learn. Stakeholders help shape both the problem and the solution.

A useful first conversation explores their actual experience

Suggested Quickstart

- Prepare around a stakeholder's context; invite a short conversation.
- Ask: "Can you walk me through the last time this was difficult?"
- Follow up: "How do you handle it today? What would make a meaningful difference?"
- Summarize what you heard, check your interpretation, then explore possible help.

Optional Starting Resources

- **Practical tools:** Design Thinking Bootleg cards (PDF) (website) — begin with *Interview Preparation* and *Interview for Empathy* (cards 3–4).
- **Another process map:** Design Council's Double Diamond — Discover → Define → Develop → Deliver; explore and narrow both the problem and possible solutions.

Friday connected dimensionality reduction to linear transformations

Formalize the representation

- One observation: $\boldsymbol{x}_i \in \mathbb{R}^d$
- Reduced features: $\boldsymbol{z}_i = f(\boldsymbol{x}_i) \in \mathbb{R}^k$
- First foundational case: f is linear

Reason through the operator

- $\boldsymbol{y} = A\boldsymbol{x}$: each output is a linear combination
- $Y = XA^T$: one operator transforms every observation
- Shapes expose invalid matrix manipulations

Diagonal matrices scale features independently; the Euclidean norm formalizes vector size and pairwise distance.

The Euclidean norm measures size and distance

For $\boldsymbol{x} \in \mathbb{R}^d$:

$$\|\boldsymbol{x}\|_2 = \sqrt{x_1^2 + x_2^2 + \cdots + x_d^2}$$

For two points $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$:

$$\mathrm{dist}(\boldsymbol{x}_i, \boldsymbol{x}_j) = \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2$$

The subscript 2 identifies the Euclidean norm. Other norms measure size differently.

Orthogonal vectors meet at a right angle

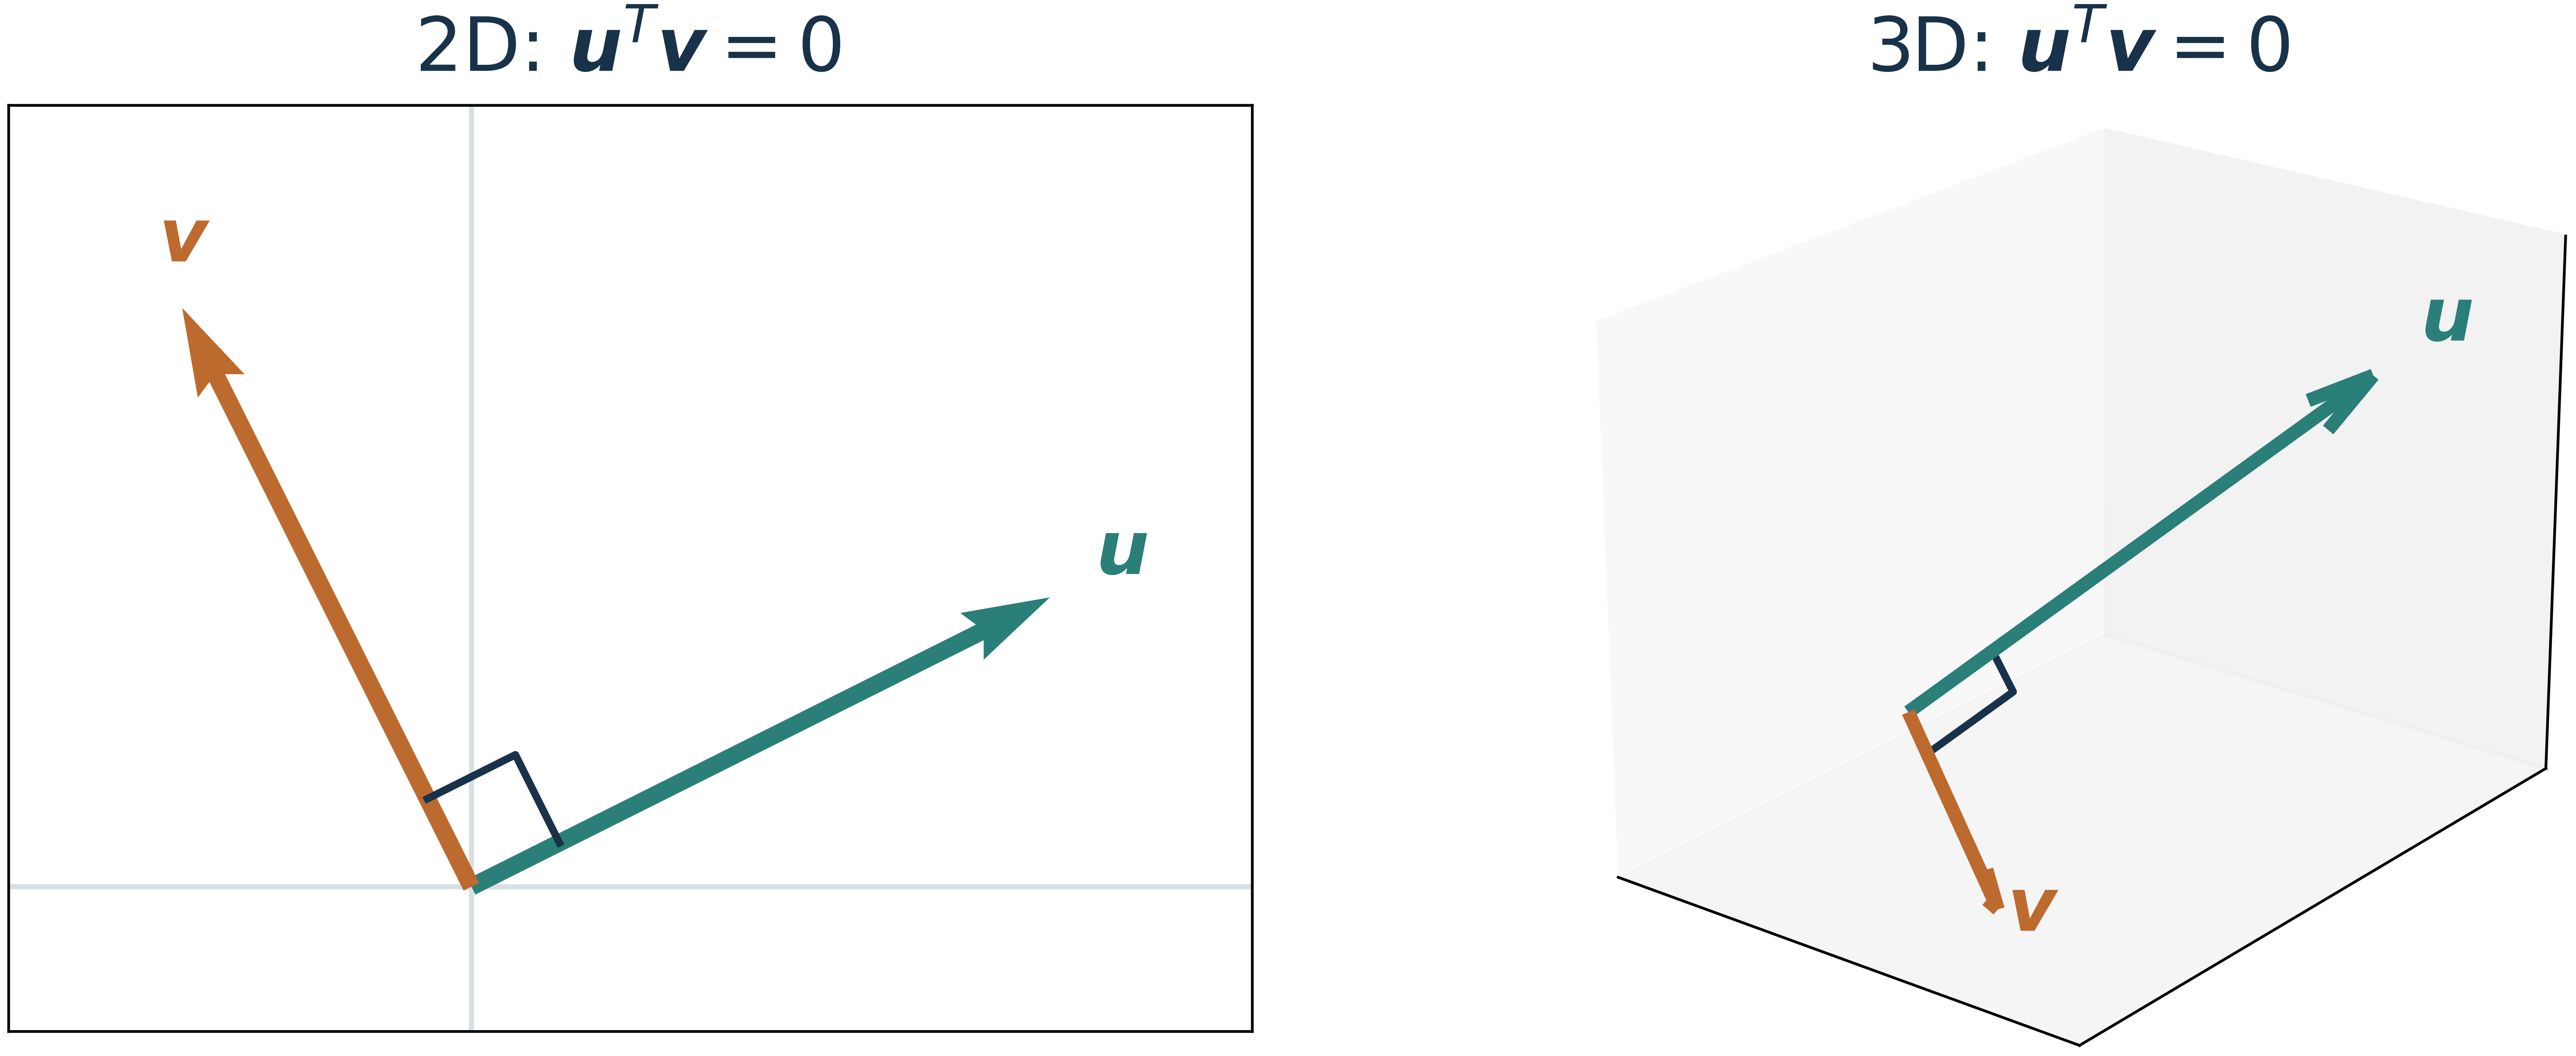

$$\boldsymbol{u} \perp \boldsymbol{v} \quad \Longleftrightarrow \quad \boldsymbol{u}^T\boldsymbol{v} = 0$$

In higher dimensions, we compute orthogonality

Are these dense five-dimensional vectors orthogonal?

$$\boldsymbol{u} = \begin{bmatrix} 1 \\ 2 \\ -1 \\ 3 \\ 4 \end{bmatrix}, \qquad \boldsymbol{v} = \begin{bmatrix} 2 \\ -1 \\ 3 \\ 2 \\ -\frac{3}{4} \end{bmatrix}$$

$$\boldsymbol{u}^T \boldsymbol{v} = 2 - 2 - 3 + 6 - 3 = 0 \qquad \Longrightarrow \qquad \boldsymbol{u} \perp \boldsymbol{v}$$

Orthogonality is defined in any dimension, even when geometry cannot be visualized.

An orthogonal matrix has orthonormal columns

For a square matrix $Q = [\boldsymbol{q}_1 \ \cdots \ \boldsymbol{q}_d] \in \mathbb{R}^{d\times d}$, **orthonormal** means mutually orthogonal columns, each of unit length.

Higher dimensions	Interpretation	1D special case	Interpretation
$I\boldsymbol{x} = \boldsymbol{x}$	The identity matrix leaves a vector unchanged.	$1x = x$	The scalar 1 leaves a number unchanged.
$(Q^TQ)_{ij} = \boldsymbol{q}_i^T\boldsymbol{q}_j$	Multiplication collects all column inner products.	$q\,q = q^2$	There is only one entry to multiply by itself.
$Q^TQ = I$	Inner products are 1 for the same column, 0 for distinct columns.	$q^2 = 1$	Hence $q \in \{-1, 1\}$.
$\boldsymbol{y} = Q\boldsymbol{x}$	Rotates or reflects about the origin, preserving length.	$y = qx$	Keeps or reverses direction while preserving length.

Do orthogonal transformations preserve every pairwise distance?

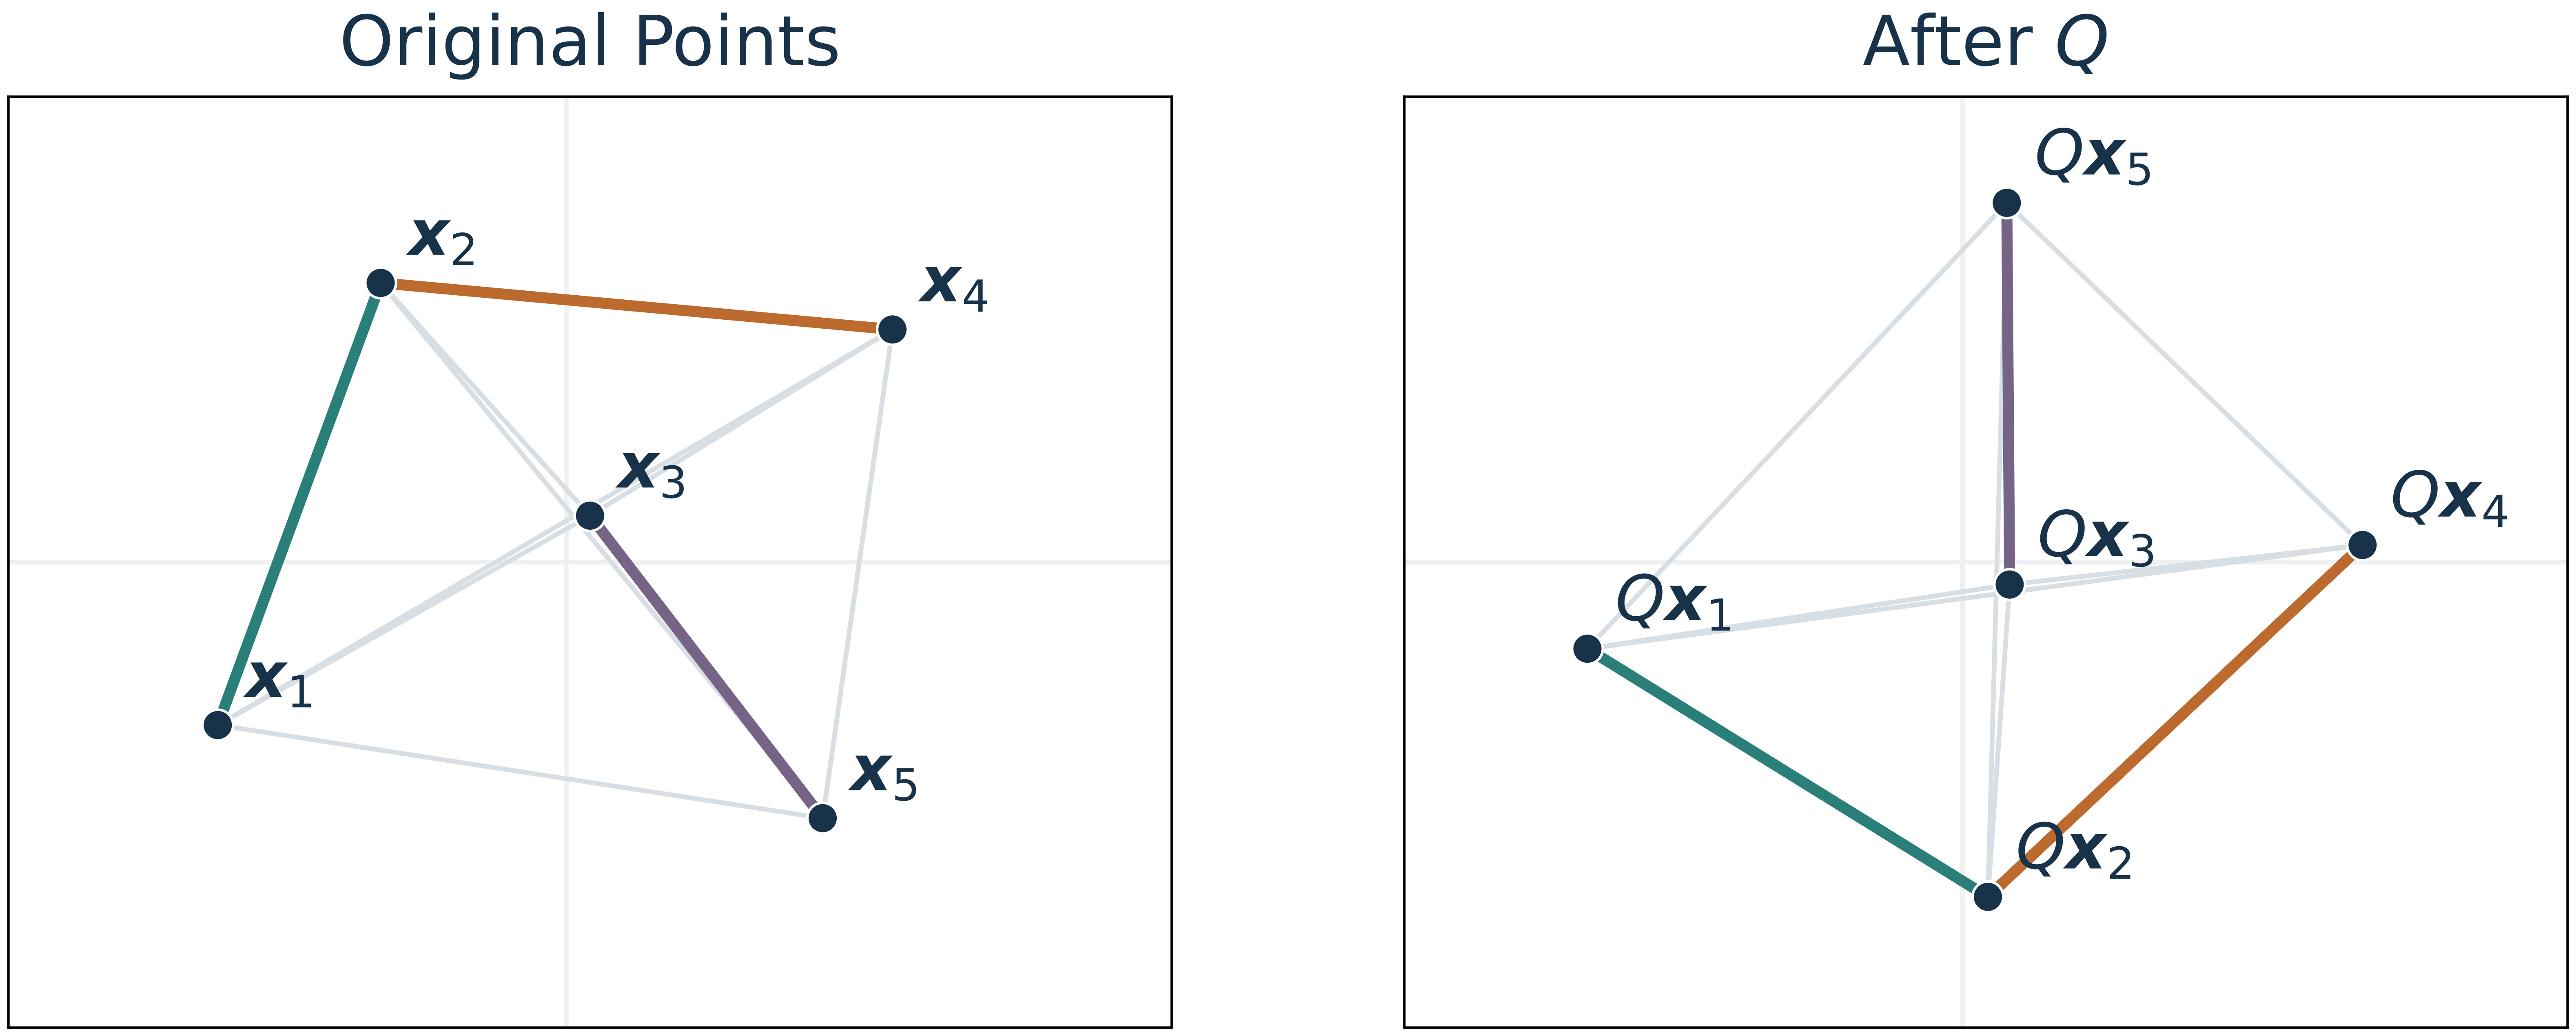

Is $\|Qx_i - Qx_j\|_2$ equal to $\|x_i - x_j\|_2$ for every pair?

Orthogonality preserves all pairwise Euclidean distances

For any $\boldsymbol{x}_i, \boldsymbol{x}_j \in \mathbb{R}^d$ and orthogonal Q:

$$\begin{aligned}
\|Q\boldsymbol{x}_i - Q\boldsymbol{x}_j\|_2 &= \|Q(\boldsymbol{x}_i - \boldsymbol{x}_j)\|_2 && \text{(Distributivity)} \\
&= \sqrt{(\boldsymbol{x}_i - \boldsymbol{x}_j)^T Q^T Q (\boldsymbol{x}_i - \boldsymbol{x}_j)} && \text{(Euclidean norm)} \\
&= \sqrt{(\boldsymbol{x}_i - \boldsymbol{x}_j)^T (\boldsymbol{x}_i - \boldsymbol{x}_j)} && (Q^T Q = I) \\
&= \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2 && \text{(Euclidean norm)}
\end{aligned}$$

Diagonal and orthogonal matrices have simple interpretations

Diagonal: Scale Each Coordinate

$$D = \operatorname{diag}(d_1, \dots, d_d), \qquad y_j = d_j x_j$$

- Each coordinate undergoes ordinary scalar multiplication.
- The sign keeps or reverses its direction.
- The magnitude stretches or shrinks it; zero collapses it.

Orthogonal: Rotate or Reflect

$$\boldsymbol{y} = Q\boldsymbol{x}, \qquad Q^T Q = I$$

- Rotate, reflect, or combine both about the origin.
- Keep lengths, angles, and pairwise distances unchanged.
- Change orientation without stretching or shrinking.

In 1D, sign and magnitude explain what multiplication does

One Dimension

For $a \neq 0$:

$$y = ax = \underbrace{\mathrm{sign}(a)}_{\text{keep or reverse}} \quad \underbrace{|a|}_{\text{scale}} \quad x$$

For example, $-3x = (-1)(3)x$:

- 3: stretch by a factor of three.
- -1: reverse direction.

Higher Dimensions

$$\boldsymbol{y} = A\boldsymbol{x}$$

We know how to **calculate** the output. But for an arbitrary matrix:

- Which directions change, and how?
- How much does each direction stretch or shrink?
- Can any direction collapse entirely?

Can we decompose any A into rotations or reflections and independent scalings so that we can **understand and analyze** what it does?

SVD separates every linear transformation into three understandable steps

Let $U\Sigma V^T$ be the **singular value decomposition (SVD)** of A, i.e., $A = U\Sigma V^T$.

U and V are orthogonal; Σ is diagonal (possibly rectangular).

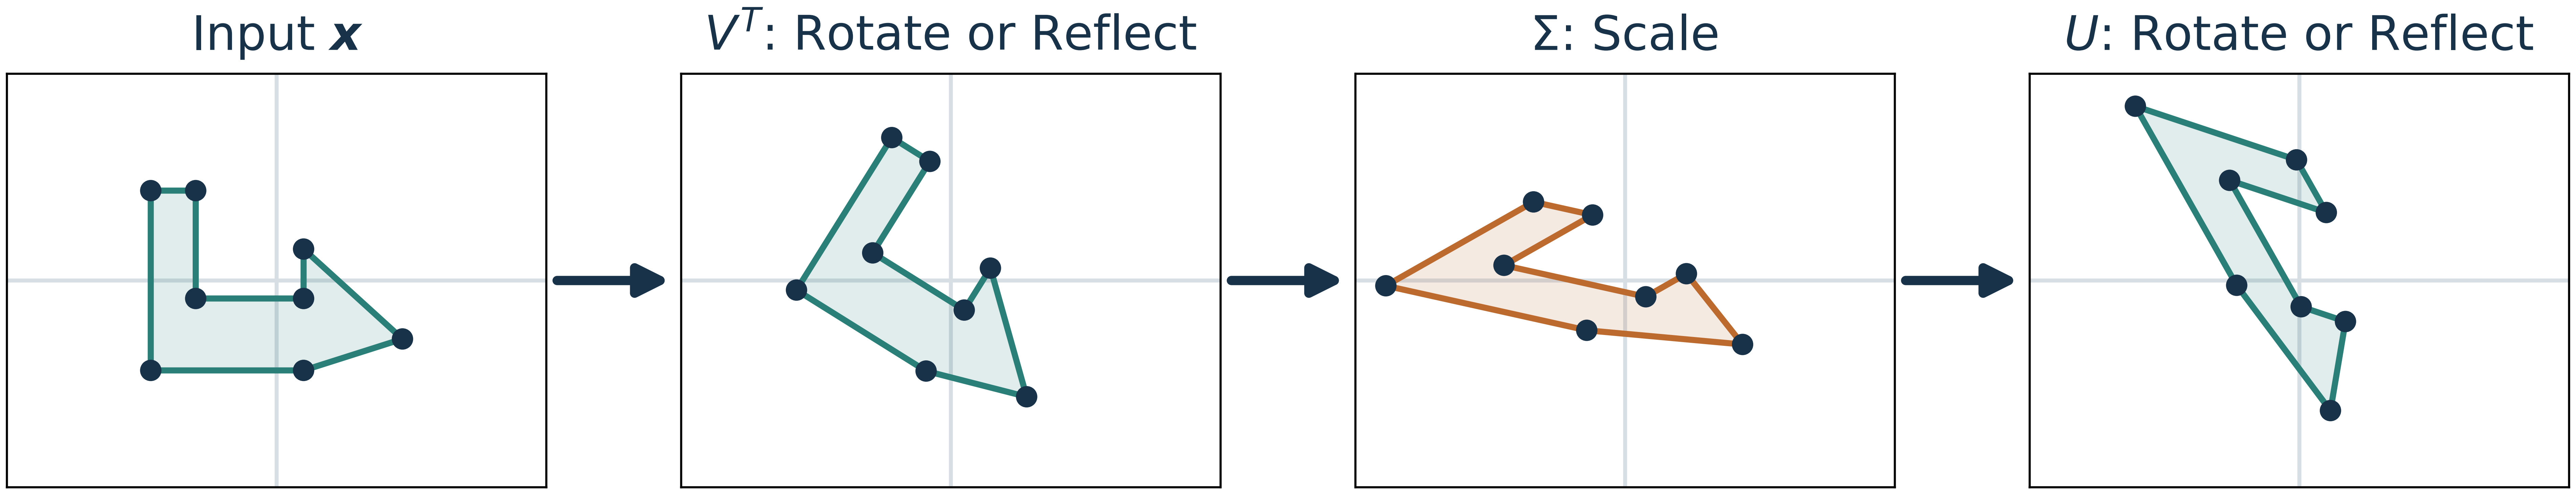

$$\boldsymbol{y} = A\boldsymbol{x} = U\Sigma V^T\boldsymbol{x} = U\Big(\Sigma(V^T\boldsymbol{x})\Big)$$

Each SVD factor has a precise shape and role

For $A \in \mathbb{R}^{m \times d}$, the **full SVD** is:

$$\underbrace{A}_{m\times d} = \underbrace{U}_{m\times m} \quad \underbrace{\Sigma}_{m\times d} \quad \underbrace{V^T}_{d\times d}$$

Factor	Mathematical Property	What It Does
U	Square and orthogonal: $U^T U = I_m$	Rotates or reflects in the output space.
Σ	Same shape as A; $\Sigma_{ij} = 0$ whenever $i \neq j$	Scales directions; zero scalings can collapse them.
V^T	Square and orthogonal: $V^T V = V V^T = I_d$	Rotates or reflects in the input space.

Σ is **rectangular diagonal**: only $\Sigma_{ii} = \sigma_i \geq 0$ may be nonzero, for $i = 1, \ldots, \min(m, d)$.

A and Σ share a shape; the orthogonal factors are square

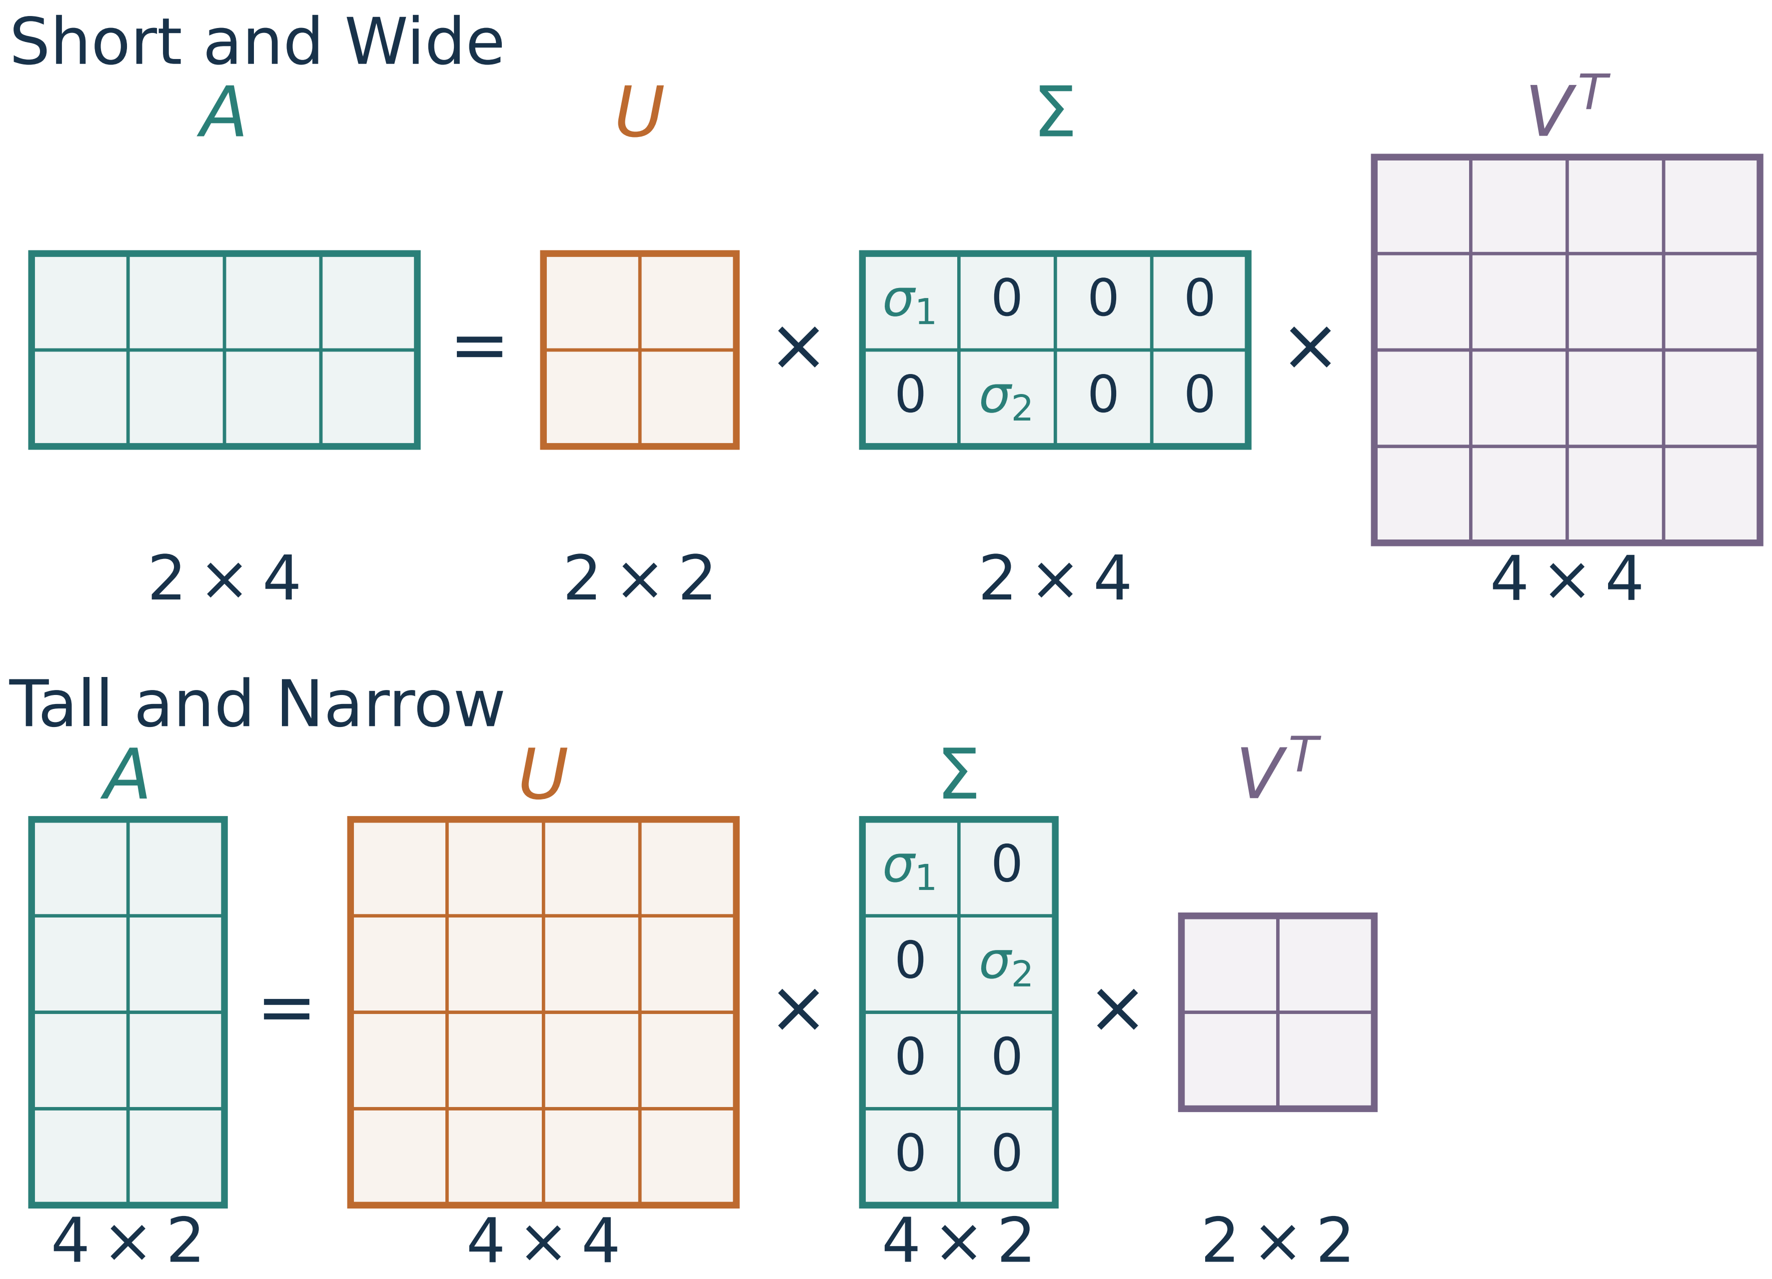

U: square, sized by the **rows** of A.

V^T: square, sized by the **columns** of A.

Singular values tell us how much each direction scales

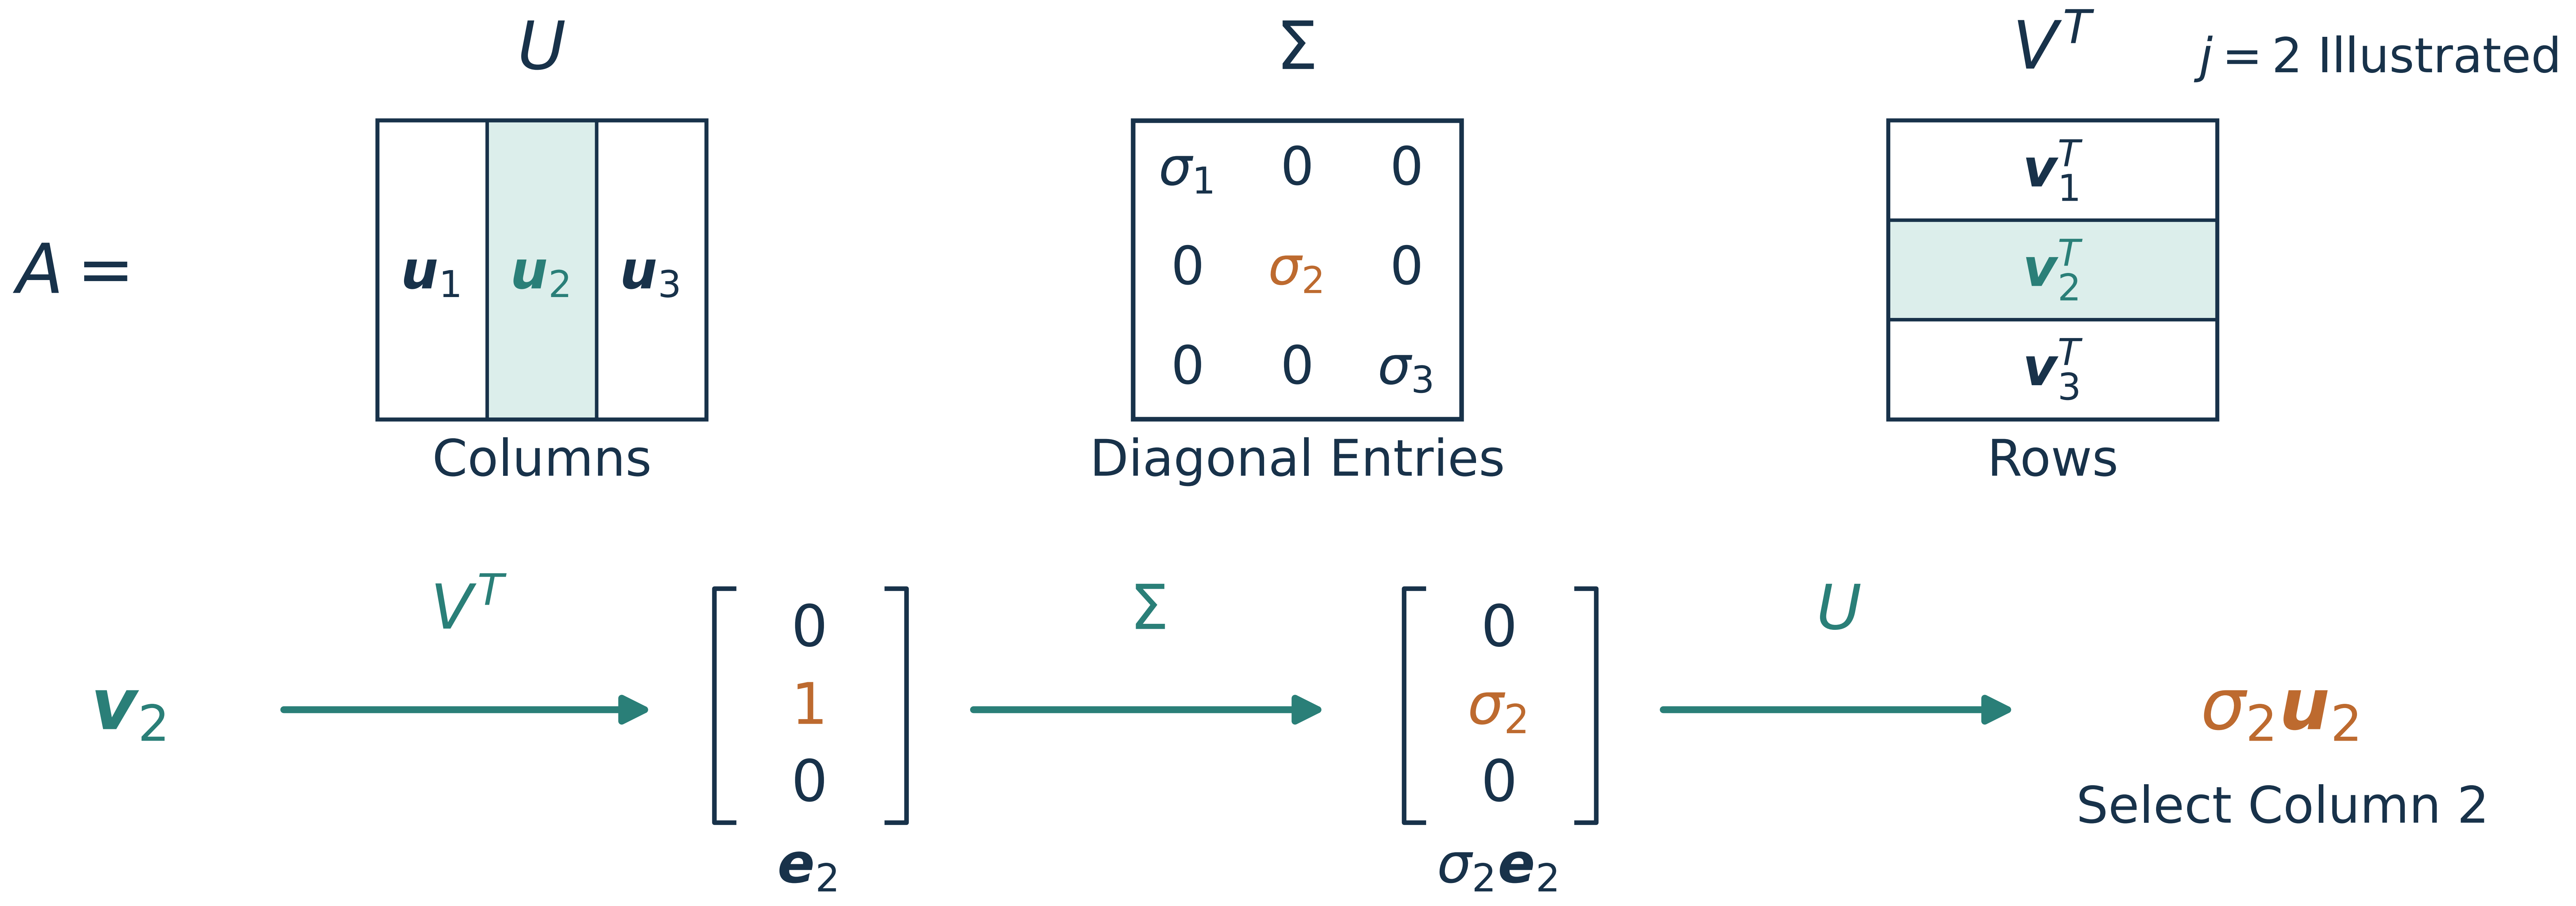

$$V^T \boldsymbol{v}_j = \boldsymbol{e}_j \quad \Longrightarrow \quad A\boldsymbol{v}_j = U(\Sigma \boldsymbol{e}_j) = \sigma_j \boldsymbol{u}_j, \qquad j \leq \min(m, d).$$

The largest singular value gives the maximum stretching

Order the singular values: $\sigma_1 \geq \sigma_2 \geq \cdots \geq 0$.

$$\|A\boldsymbol{x}\|_2 \leq \sigma_1 \|\boldsymbol{x}\|_2, \qquad \|A\|_2 := \max_{\|\boldsymbol{x}\|_2 = 1} \|A\boldsymbol{x}\|_2 = \sigma_1.$$

- V^T and U preserve lengths.
- Each coordinate scaling in Σ is at most σ_1.
- The input $\boldsymbol{v}_1$ attains the largest scaling.

Rank counts the independent directions that survive

$$\operatorname{rank}(A) = \#\{j : \sigma_j > 0\} \le \min(m, d)$$

For example:

$$\Sigma = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \qquad \Longrightarrow \qquad \operatorname{rank}(A) = 2$$

- Two independent input directions produce nonzero output directions.
- The third input direction disappears: $A\boldsymbol{v}_3 = \mathbf{0}$.
- Rotations and reflections change which directions these are, but not the rank.