

Linear Algebra Review Guide

CS 135: Computational Semantics -- From Logic to Geometry · Fall 2026

James Pustejovsky

Companion to the Session 1 Self-Assessment · Posted with the answer key

How to use this guide

The self-assessment was a map, and this is the legend. It is organized by section (A through F) so that you can go directly to whatever came out weakest and skip the rest. For each section it tells you why that material shows up in the course and when, what you need to know (which is narrower than a full linear algebra course), what the common errors on the assessment were, where to review, and how to check whether the review worked.

Two points before you start.

- **Trust the gap, not just the score.** If you scored well on a section but circled a low confidence, you probably know the mechanics and not the picture; read the "why it matters" paragraph and do the self-check, and you are likely done. If you scored low but circled a high confidence, that is the section to spend real time on -- it usually means a specific misconception, and the "common mistakes" list is aimed at exactly that.
- **The window is Weeks 2 through 5.** Sessions 4 through 6 are logic and composition and make no use of linear algebra at all. The embedding sessions (7 through 9) start to lean on Sections C and D, and the geometric algebra unit from Week 7 onward leans on everything, Section E above all. A few hours spread across the next four weeks will close any gap on this list; the same hours spent during Week 8 will not.

Priority order

If more than one section came out low, review in this order. It is ordered by how much the course depends on the section and how soon.

Priority	Section	Reason
1	E -- Determinants, area, orientation	The geometric algebra unit is built on this. The outer product is a determinant that remembers its orientation. Nothing in Sessions 12 and 13 makes sense without it.
2	C -- Inner products, projection	Needed twice: for embedding similarity in Sessions 7 through 9, then again for contraction and projection in Session 14.
3	B -- Span, independence, basis	Blades represent subspaces; a subspace is a span; a blade is zero exactly when its factors are dependent. Sessions 13 and 17.
4	F -- Rotations, orthogonal maps	Rotors (Session 15) are the geometric algebra version of rotation. You need to know what a rotation is before you see it rebuilt.

Priority	Section	Reason
5	D -- Matrices as linear maps	Embedding matrices and attention weights (Sessions 8 and 9); tensor products (Session 11).
6	A -- Vectors and vector spaces	Foundation for everything, but most students have it. Review only if A2 or A3 went wrong.

Section G (SVD, eigenvectors) is not on this list. It is taught in the course, not assumed, and there is nothing to review for it.

What you do not need

The course uses a narrow slice of linear algebra. You do not need fluency with Gaussian elimination as an algorithm, row-reduced echelon form, LU factorization, the rank--nullity theorem in its general form, Cramer's rule, or anything about eigenvalues beyond what Session 16 introduces. If a review resource spends a chapter on solving systems of equations, skip it. The emphasis here is on the geometry: what vectors, subspaces, inner products, determinants, and rotations mean as pictures, and how the algebra tracks the picture.

The resources this guide points to

Four sources are used throughout. All are free or in the library; none requires purchasing anything.

- **3Blue1Brown, Essence of Linear Algebra.** A series of short videos (roughly ten to fifteen minutes each) built entirely around geometric intuition. For this course it is the single best resource, because its point of view -- matrices are transformations, determinants are scale factors, dot products are projections -- is exactly the point of view geometric algebra generalizes. Cited below by chapter number.
- **Strang, Introduction to Linear Algebra, and the MIT OCW 18.06 lectures.** The standard undergraduate treatment. Cited by book section; the OCW lectures follow the book and are useful if you learn better from a lecture.
- **Hefferon, Linear Algebra.** A free, open textbook with full solutions, good for doing exercises. Its "Geometry of Determinants" and "Geometry of Linear Maps" topics are particularly well matched to this course.
- **Immersive Linear Algebra (immersivemath.com).** A free online text with interactive figures. Good for Sections C, E, and F, where being able to drag a vector and watch a projection, area, or rotation change is worth more than a page of text.

Chapter and section numbers below are for the current editions as of this writing. If a number does not match what you have, use the topic name; the topics have not moved.

Section A -- Vectors and Vector Spaces

Where it is used: Throughout; Session 12 in particular.

Suggested review time if this section came out low: 1 to 2 hours. Most students only need A2 and A3.

Why it matters in this course

Session 12 introduces the geometric product, and the first thing it asks of you is to stop thinking of a vector as a list of numbers. The list is a description of the vector relative to a basis; the vector is the arrow. Geometric algebra is built to work with the arrow directly, and much of its payoff comes from expressions that never mention coordinates. Question A3 was on the assessment specifically to see who already draws that distinction.

The subspace test (A2) matters for a similar reason. Blades in geometric algebra represent subspaces, and a subspace has to pass through the origin and be closed under the vector operations. The plane $x + y + z = 1$ is not one, and neither is the positive octant. If you answered (d), it is worth understanding why (a) is the only correct answer.

What you need to know

- Vector addition and scalar multiplication, componentwise, and what they look like as arrows (parallelogram rule, stretching).
- The definition of a vector space, and the three-part test for whether a subset is a subspace: contains the zero vector, closed under addition, closed under scalar multiples.
- The distinction between a vector and its coordinate representation in a basis. The same arrow has different coordinates in different bases.

Common mistakes on the assessment

- A2: choosing (d) because (b) "looks like" (a). The plane $x + y + z = 1$ does not contain the origin, so it fails immediately. Choosing (c) because the set is closed under addition; it is, but it is not closed under multiplication by negative scalars.
- A3: answering "true" without qualification. The intended answer is that this is true only relative to a choice of basis; a vector is an element of the space, and the three numbers are its coordinates.

Where to review

- 3Blue1Brown, Chapter 1 ("Vectors, what even are they?") and Chapter 16 ("Abstract vector spaces"). The second one is the direct answer to A3.
- Strang, Sections 1.1 and 3.1 (vector spaces and subspaces). The subspace examples in 3.1 cover A2 exactly.
- Hefferon, Chapter Two, Section I (Definition of Vector Space), particularly the subspace examples and exercises.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- A1. Let $u = (1, -2, 0)$ and $v = (3, 1, 1)$. Compute $2u + v$.
- A2. Is the set of all (x, y) in $\mathbb{R}^2$ with $xy = 0$ a subspace of $\mathbb{R}^2$? Give a one-line reason.
- A3. The vector $(1, 1)$ in $\mathbb{R}^2$ has coordinates $(1, 1)$ in the standard basis. What are its coordinates in the basis $\{(1, 1), (1, -1)\}$?

Section B -- Span, Independence, Basis

Where it is used: Sessions 13 and 17.

Suggested review time if this section came out low: 2 to 3 hours.

Why it matters in this course

In Session 13 you will meet the outer product of two vectors, $a \wedge b$, which represents the plane spanned by a and b , together with an orientation and an area. It is zero exactly when a and b are linearly dependent. The outer product of three vectors represents the volume they span and is zero exactly when they are dependent. So the concept of independence is not background here: it is the condition for a blade to be nonzero, and "dimension of the span" becomes "grade of the blade".

Session 17 goes further and uses blades to represent subspaces of a semantic space directly. If you are not comfortable with what it means for a set of vectors to span a subspace, or with the fact that a basis for an n -dimensional space has exactly n vectors, Session 17 will be harder than it needs to be.

What you need to know

- The span of a set of vectors, as the set of all linear combinations, and as a picture: a line, a plane, all of space.
- Linear independence, both as a definition (the only combination giving zero is the trivial one) and as a picture (no vector in the set lies in the span of the others).
- Basis: a linearly independent spanning set. Dimension: the number of vectors in any basis. Every basis of $\mathbb{R}^n$ has exactly n vectors.
- The consequences: more than n vectors in $\mathbb{R}^n$ are always dependent; fewer than n never span; a spanning set can be thinned to a basis and an independent set can be extended to one.

Common mistakes on the assessment

- B1: computing something and concluding the vectors are independent without noticing that $w = u + v$. Whenever one vector is the sum of two others, the set is dependent. A quick test in $\mathbb{R}^3$ is the determinant of the matrix with the three vectors as columns; here it is zero.
- B2: answering 3 because there are three vectors. All three lie on the same line through the origin; the span is a line and the dimension is 1.
- B4: concluding only "S spans $\mathbb{R}^4$ " or nothing. Seven vectors in $\mathbb{R}^4$ must be dependent (seven exceeds four). Since S spans, some four of them form a basis. The set is not itself a basis.

Where to review

- 3Blue1Brown, Chapter 2 ("Linear combinations, span, and basis vectors"). Watch this one even if you scored well; the way it draws span is the way blades will be drawn in Session 13.

- Strang, Section 3.4 (Independence, basis, and dimension). This is the definitive statement of everything above, with the counting arguments for B4.
- Hefferon, Chapter Two, Sections II and III (Linear Independence; Basis and Dimension). The exercises with solutions are the best place to practice deciding independence by hand.
- OCW 18.06, Lecture 9 (Independence, basis, and dimension), if you prefer a lecture.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- B1. Are $(1, 2, 3)$ and $(2, 4, 6)$ linearly independent? What is the dimension of their span?
- B2. What is the dimension of $\text{span}\{(1, 0, 0), (0, 1, 0), (1, 1, 0)\}$ in $\mathbb{R}^3$? Describe the span geometrically.
- B3. Five vectors are given in $\mathbb{R}^3$. Without computing anything, what can you say about their independence? Can they span $\mathbb{R}^3$?
- B4. Give a basis for the plane $x + y + z = 0$ in $\mathbb{R}^3$. How many vectors does it need?

Section C -- Inner Products, Norms, Orthogonality, Projection

Where it is used: Sessions 7 through 9 (embedding similarity) and Session 14 (contraction and projection).

Suggested review time if this section came out low: 2 to 3 hours.

Why it matters in this course

This is the section the course uses first. In Sessions 7 through 9, word and sentence embeddings are vectors, and almost every question about them -- how similar are two words, which context does a word best fit, what does an attention weight measure -- is answered by an inner product. Cosine similarity (C5) is an inner product divided by two norms, and you will compute it constantly.

It returns in Session 14 with the geometric product, which splits into a symmetric part (the inner product) and an antisymmetric part (the outer product from Section B). Projection of a vector onto a subspace, which you did in C4 for a line, is done there with a single algebraic expression that works for planes and higher-dimensional subspaces without change. That expression only makes sense if you already know what projection onto a line is doing.

What you need to know

- The dot product, computed componentwise, and its geometric meaning: $u \cdot v = \|u\| \|v\| \cos \theta$.
- The norm $\|u\| = \sqrt{u \cdot u}$, and the unit vector $u / \|u\|$.
- Orthogonality: $u \cdot v = 0$ means the vectors are perpendicular (for nonzero u and v). The sign of the dot product tells you whether the angle is acute, right, or obtuse.

- Projection of u onto v : $\text{proj}_v u = (u \cdot v / v \cdot v) v$. The result is a vector along v , and u minus that projection is orthogonal to v .
- Cosine similarity: $(u \cdot v) / (\|u\| \|v\|)$, which ranges over $[-1, 1]$ and is 1 for parallel, 0 for orthogonal, -1 for opposite vectors.

Common mistakes on the assessment

- C2: computing $u \cdot v = 0$ correctly but not saying what it means. The point of the question is the interpretation: the vectors are orthogonal.
- C4: returning the scalar 3 rather than the vector $(3, 0)$. A projection onto a line is a vector on that line. The scalar 3 is its length (the scalar projection); the question asked for the vector.
- C5: stopping at $1/\sqrt{2}$ without naming the angle, or naming it as 90° . Since $\cos \theta = 1/\sqrt{2} \approx 0.707$, $\theta = 45^\circ$.

Where to review

- 3Blue1Brown, Chapter 9 ("Dot products and duality"). The duality half is optional, but the first half is the cleanest picture of why the dot product measures projection.
- Strang, Sections 1.2 (Lengths and dot products) and 4.2 (Projections). Section 4.2 starts with projection onto a line, which is what you need, before moving to subspaces.
- Immersive Linear Algebra, the chapter on the dot product. Drag the vectors and watch the projection.
- OCW 18.06, Lecture 15 (Projections onto subspaces), first twenty minutes.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- Compute $\|(1, 2, 2)\|$.
- Compute the vector projection of $u = (2, 3)$ onto $v = (1, 1)$. Check that u minus your answer is orthogonal to v .
- What is the cosine similarity of $(1, 0)$ and $(0, 1)$? Of $(1, 0)$ and $(-1, 0)$?
- Two embedding vectors have cosine similarity 0.95. What does this say about the angle between them? What does it not say about their lengths?

Section D -- Matrices as Linear Maps

Where it is used: Sessions 8 and 9 (embeddings, attention) and Session 11 (tensor products).

Suggested review time if this section came out low: 1 to 2 hours.

Why it matters in this course

Everything in Sessions 8 and 9 is matrix multiplication: an embedding matrix takes a one-hot word vector to a dense vector; the query, key, and value matrices in attention are linear maps applied to every token; the attention weights form a matrix that mixes the value vectors. You do

not need to be fast at multiplying matrices by hand, but you need to know what a matrix does -- it sends each basis vector to a column -- and how shapes have to line up (D3).

Session 11 introduces tensor products as one route to compositional semantics, and the argument there compares tensor products with geometric products. The comparison needs D4: that matrix multiplication is not commutative, that AB and BA are different maps. Geometric algebra also has a noncommutative product, but its noncommutativity is controlled and meaningful. That contrast is easier to appreciate once you are used to the matrix case.

What you need to know

- A matrix as a linear map: Ax is a linear combination of the columns of A with coefficients from x . The columns are the images of the basis vectors.
- Shape rules: an $m \times n$ matrix maps $\mathbb{R}^n$ to $\mathbb{R}^m$. The product AB is defined when the column count of A matches the row count of B , and $(m \times n)(n \times k) = m \times k$.
- Composition: AB means "apply B , then A ". This is why $AB \neq BA$ in general.
- Diagonal matrices scale each axis independently; this is the picture behind D1.

Common mistakes on the assessment

- D2: "the coefficients" or "the equations". The columns are the images of the standard basis vectors under T . If you know where the basis goes, you know everything about the map, and the matrix is just that information written down.
- D3: 2×2 (taking the inner dimensions) or 3×4 with the reasoning reversed. The inner dimensions must match and then vanish; the outer ones remain: 3×4 .
- D4: "true". It is false. Rotating and then reflecting is not the same as reflecting and then rotating.

Where to review

- 3Blue1Brown, Chapter 3 ("Linear transformations and matrices") and Chapter 4 ("Matrix multiplication as composition"). Chapter 4 is the picture behind D4.
- Strang, Sections 2.3 and 2.4 (Matrix operations; Rules for matrix operations). Section 2.4 has the noncommutativity examples.
- Hefferon, Chapter Three, Section III (Computing Linear Maps), which develops "the columns are the images of the basis" carefully.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- D1. Let $A = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$. Compute $A(1, 0)$ and $A(0, 1)$. What does A do to the plane?
- D2. A is 4×3 . What is the shape of x if Ax is defined, and what is the shape of Ax ?
- D3. Let R be rotation by 90° counterclockwise and F be reflection across the x -axis, both in $\mathbb{R}^2$. Is $RF = FR$? Test on the vector $(1, 0)$.

Section E -- Determinants, Area, and Orientation

Where it is used: Sessions 12 and 13 build directly on it; used through the end of the course.

Suggested review time if this section came out low: 3 to 4 hours. This is the section to prioritize.

Why it matters in this course

The outer product of two vectors, $a \wedge b$, is the central new object in Session 13. It has a magnitude, which is the area of the parallelogram spanned by a and b , and an orientation, which is the sense (counterclockwise or clockwise) in which a is swept toward b . In two dimensions, $a \wedge b$ is exactly the 2×2 determinant with a and b as columns, multiplied by the unit bivector. The determinant's magnitude is the area (E2); its sign is the orientation (E3); it vanishes when a and b are parallel (E4). Every one of the assessment questions in this section is a fact about the outer product in disguise.

The same holds in three dimensions: $a \wedge b \wedge c$ is the 3×3 determinant times the unit trivector, and its magnitude is a volume (E5). Once you see determinants this way, the geometric algebra unit is a matter of making the determinant into an algebraic object you can add, multiply, and take products with, rather than a number you compute and throw away. If you see determinants only as a formula, the unit will look like it is inventing things.

The one idea students most often lack is that orientation is information. Swapping two columns negates the determinant (E3) and that sign is not a bookkeeping accident; it records which way around the parallelogram you went. Geometric algebra keeps that sign throughout, and it is what lets rotations, reflections, and the difference between "a then b" and "b then a" all live in one algebra. Spend your review time on E3, not E1.

What you need to know

- The 2×2 determinant formula $ad - bc$, and the 3×3 formula by cofactor expansion (you do not need a general $n \times n$ algorithm).
- $|\det|$ of a 2×2 matrix is the area of the parallelogram spanned by its columns; for a 3×3 matrix it is the volume of the parallelepiped. Equivalently, $|\det A|$ is the factor by which A scales all areas (or volumes).
- The sign of the determinant is the orientation: positive means the columns, in order, have the same handedness as the standard basis; negative means the map flips orientation (it includes a reflection).
- $\det = 0$ exactly when the columns are linearly dependent; the parallelogram is flattened to a line (or point) and has no area.
- Swapping two columns negates the determinant. Scaling one column by c scales the determinant by c . The determinant is linear in each column separately (this is the property the outer product inherits).
- $\det(AB) = \det(A) \det(B)$: scale factors compose by multiplying.

Common mistakes on the assessment

- E2: "the size of the matrix" or "how invertible it is". The answer is a specific quantity: the area of the parallelogram whose sides are the two column vectors (equivalently, the area-scaling factor of the map).
- E3: answering 5, or -5 without an interpretation. The determinant becomes -5 . The area is unchanged; what changed is the orientation. Going from a to b sweeps one way around the origin; going from b to a sweeps the other way.
- E4: "the matrix is not invertible" and nothing else. That is true but is not a statement about the columns. The columns are linearly dependent: one is a multiple of the other, so the parallelogram they span is flat.
- E5: "area", by analogy with E2. It is volume: the three columns span a parallelepiped.

Where to review

- 3Blue1Brown, Chapter 6 ("The determinant"). This video is the closest thing to a prerequisite for Session 13 that exists. Watch it twice: once now, once the week before the geometric algebra unit begins. The treatment of negative determinants as orientation flips is exactly the point.
- 3Blue1Brown, Chapter 5 ("Three-dimensional linear transformations") for the volume version.
- Hefferon, Chapter Four, Section I.1 and the Topic "Geometry of Determinants". The Topic is short and treats the determinant as an oriented area from the start, which is the view this course takes.
- Strang, Section 5.1 (The properties of determinants) for the algebraic rules -- especially the sign change under row or column exchange and linearity in each row -- and Section 5.3 for the area and volume interpretation.
- Immersive Linear Algebra, the determinant chapter. The interactive figure that lets you drag a column vector through the other one and watch the sign of the area flip is worth the visit on its own.
- OCW 18.06, Lectures 18 and 19 (Properties of determinants; Determinant formulas and cofactors). Lecture 18 is the important one.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- E1. Compute $\det \begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix}$. Is the map orientation-preserving or orientation-reversing?
- E2. Let $a = (2, 0)$ and $b = (1, 1)$. Compute $\det[a \ b]$ and $\det[b \ a]$. Draw both parallelograms; they are the same shape. What differs?
- E3. The columns of a 2×2 matrix are $(1, 2)$ and $(2, 4)$. Without computing, what is the determinant, and why?
- E4. If $\det A = 3$ and $\det B = -2$, what is $\det(AB)$? Does AB preserve or reverse orientation?
- E5. Scale both columns of a 2×2 matrix by 2. By what factor does the determinant change? By what factor does the area change?

Section F -- Rotations and Orthogonal Transformations

Where it is used: Session 15 (rotors).

Suggested review time if this section came out low: 2 hours.

Why it matters in this course

Session 15 introduces rotors, which are the geometric algebra objects that perform rotations. A rotor is built from two reflections, and applying it to a vector uses a "sandwich" product $R v R$ rather than a matrix multiplication. To appreciate what that buys -- rotors compose by multiplication, they work in any dimension without special cases, and they interpolate cleanly -- you need to know what the matrix version looks like (F1), what property makes a map a rotation (F2), and where matrices become awkward, which is precisely the noncommutativity of 3D rotations (F3).

The link to Section E is direct: a rotation is an orthogonal map with determinant $+1$. An orthogonal map with determinant -1 is a reflection (composed with a rotation). The sign of the determinant, again, is the orientation, and rotors are exactly the orientation-preserving case.

What you need to know

- The 2D rotation matrix $\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$, and why its columns are what they are: they are where $(1, 0)$ and $(0, 1)$ land.
- Orthogonal matrices: $Q^T Q = I$, equivalently the columns are orthonormal, equivalently Q preserves all dot products and hence all lengths and angles.
- $\det Q = \pm 1$ for orthogonal Q . Rotations have $\det +1$; reflections have $\det -1$.
- In 2D, rotations commute (angles add). In 3D, rotations about different axes do not commute.

Common mistakes on the assessment

- F1: sign in the wrong place ($\begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$ is clockwise rotation). Check by applying to $(1, 0)$: counterclockwise rotation by θ should give $(\cos \theta, \sin \theta)$, which means the first column must be $(\cos \theta, \sin \theta)$.
- F2: "it has determinant 1". That is necessary but not sufficient; a shear has determinant 1 and preserves neither lengths nor angles. The characterizing property is orthogonality, $Q^T Q = I$.
- F3: "yes, they commute". They do not. Hold a book, rotate it 90° about a vertical axis then 90° about a horizontal one, and then try the other order; the book ends up in different positions.

Where to review

- 3Blue1Brown, Chapter 3 (revisit the rotation example) and Chapter 5 (three-dimensional transformations). For F3, the video does not treat noncommutativity directly, so also do the physical experiment above.
- Strang, Section 4.4 (Orthonormal bases and Gram--Schmidt), first part only, for the definition and properties of orthogonal matrices. Skip Gram--Schmidt.

- Immersive Linear Algebra, the chapter on rotation; the interactive figures for 3D rotations about different axes demonstrate F3 directly.
- OCW 18.06, Lecture 17 (Orthogonal matrices and Gram--Schmidt), first fifteen minutes.

Self-check

Do these without looking anything up. Answers are in the Appendix.

- F1. Write the matrix for counterclockwise rotation by 90° . Apply it to $(1, 0)$ and to $(0, 1)$.
- F2. Verify that the rotation matrix for angle θ satisfies $R^T R = I$. What is its determinant?
- F3. The matrix $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$ preserves lengths and angles. Is it a rotation? What is its determinant, and what does it do geometrically?
- F4. In 2D, rotate by 30° then by 45° . Then do it in the other order. Same result? Why does this argument fail in 3D?

A suggested schedule for Weeks 2 through 5

This assumes two low sections. Adjust in the obvious way. The plan is deliberately light: an hour or two a week is enough, and the point is to be finished before Session 12, not to be finished this weekend.

Week	Sessions	Review
2	4 and 5 (fragment, quantifiers)	Watch the 3Blue1Brown chapters for your weakest section. Do that section's self-check.
3	6 and 7 (composition, then embeddings begin)	Read the textbook sections for your weakest section; do a handful of exercises from Hefferon with solutions. If Section C is low, finish it this week, since Session 7 uses it.
4	8 and 9 (embeddings, attention)	Second-weakest section: videos and self-check. Redo the original assessment questions for both sections cold.
5	10 and 11 (tensor products)	Watch 3Blue1Brown Chapter 6 (the determinant) again regardless of your Section E score. Come to office hours with anything still unclear.

If three or more sections came out low, or if Section E was low and you are not sure the review is landing, come to office hours in the next two weeks. That is a normal thing to do in this course and the conversation is usually short.

Appendix: Self-check answers

Section A

- A1. $(5, -3, 1)$.
- A2. No. $(1, 0)$ and $(0, 1)$ are both in the set, but their sum $(1, 1)$ is not (since $1 \cdot 1 \neq 0$). It fails closure under addition.

A3. $(1, 0)$: the vector is $1 \cdot (1, 1) + 0 \cdot (1, -1)$. Same arrow, different coordinates.

Section B

- B1. Dependent: $(2, 4, 6) = 2 \cdot (1, 2, 3)$. The span is a line, dimension 1.
- B2. Dimension 2. The third vector is the sum of the first two. The span is the xy -plane (all vectors with third coordinate 0).
- B3. They must be dependent, since $5 > 3$. They may or may not span $\mathbb{R}^3$; that depends on the vectors. (They span $\mathbb{R}^3$ if and only if some three of them are independent.)
- B4. Two vectors, since the plane is two-dimensional. For example, $(1, -1, 0)$ and $(1, 0, -1)$. Any two independent vectors satisfying $x + y + z = 0$ will do.

Section C

- C1. $3 \cdot (\sqrt{1 + 4 + 4}) = \sqrt{9}$.
- C2. $u \cdot v = 5$, $v \cdot v = 2$, so the projection is $(5/2)(1, 1) = (5/2, 5/2)$. Then u minus the projection is $(-1/2, 1/2)$, and $(-1/2, 1/2) \cdot (1, 1) = 0$.
- C3. 0 (orthogonal, 90°) and -1 (opposite, 180°).
- C4. The angle is small: $\cos \theta = 0.95$ gives $\theta \approx 18^\circ$. It says nothing about lengths; cosine similarity is scale-invariant, so one vector could be a hundred times longer than the other.

Section D

- D1. $A(1, 0) = (0, 1)$ and $A(0, 1) = (-1, 0)$. A rotates the plane counterclockwise by 90° .
- D2. x is in $\mathbb{R}^3$ (a 3×1 column) and Ax is in $\mathbb{R}^4$.
- D3. No. $F(1, 0) = (1, 0)$, then R gives $(0, 1)$; so $RF(1, 0) = (0, 1)$. But $R(1, 0) = (0, 1)$, then F gives $(0, -1)$; so $FR(1, 0) = (0, -1)$. Different.

Section E

- E1. $1 \cdot 4 - 2 \cdot 3 = -2$. Orientation-reversing (the map includes a reflection).
- E2. $\det[a \ b] = 2 \cdot 1 - 1 \cdot 0 = 2$; $\det[b \ a] = -2$. Same parallelogram, same area 2. What differs is the orientation: sweeping from a to b goes counterclockwise; from b to a , clockwise.
- E3. 0. The second column is twice the first, so the columns are dependent and the parallelogram is flat.
- E4. $\det(AB) = 3 \cdot (-2) = -6$. AB reverses orientation (one of the two maps did, and the other did not).
- E5. The determinant changes by a factor of 4 (2 for each column, since the determinant is linear in each column separately). The area also changes by 4: doubling both sides of a parallelogram quadruples its area.

Section F

- F1. $[[0, -1], [1, 0]]$. It sends $(1, 0)$ to $(0, 1)$ and $(0, 1)$ to $(-1, 0)$. This is the matrix A from D1 above.

- F2. $R^T R$ has entries $\cos^2 \theta + \sin^2 \theta = 1$ on the diagonal and $\cos \theta \sin \theta - \sin \theta \cos \theta = 0$ off the diagonal. The determinant is $\cos^2 \theta + \sin^2 \theta = 1$.
- F3. It is not a rotation. Its determinant is -1 . It is a reflection across the x-axis: it preserves lengths and angles but reverses orientation.
- F4. Same result: both give rotation by 75° , because 2D rotations about the origin commute (angles add). In 3D, "rotate by 30° " is not well defined without an axis, and rotations about different axes do not commute; the angle-addition argument only works when all rotations share an axis.