

Self-Assessment: Answer Key and Review Guide

CS 135: Computational Semantics · Fall 2026

Score one point per question, twenty-one in total (Section G is not scored). Partial credit is a matter for your own judgment — the number matters much less than which sections came out weak.

Each section below is headed by a note on where in CS 135 the material is actually used. That note is the point of this document: the geometric algebra unit does not use linear algebra generally, it uses these specific ideas, and knowing which ones lets you target your review.

Section A — Vectors and Vector Spaces

Where this is used: everywhere; Session 12 in particular

A1. $(4, -11, 13)$.

$3\mathbf{u} = (6, -3, 9)$; $2\mathbf{v} = (2, 8, -4)$; subtract componentwise.

A2. (a).

A vector space must contain the zero vector and be closed under addition and scalar multiplication. (b) fails the first test — $(0,0,0)$ does not satisfy $x+y+z=1$. (c) fails the third — multiplying by -1 leaves the set.

A3. False — or, more precisely, true only relative to a chosen basis.

A vector is a geometric object: a direction with a magnitude. The list $(2, -1, 3)$ is its *coordinates* with respect to a particular basis; choose a different basis and the same vector has a different list. This distinction does almost no work in a computational linear algebra course and is central here. Geometric algebra is built to express relationships between vectors that hold whatever coordinates you use, and Session 12 opens by insisting on the difference. If you answered “true,” this is the single most valuable idea to sit with before Week 7.

Section B — Span, Independence, Basis

*Where this is used: Sessions 13 and 17 — a blade *is* a subspace*

B1. No — they are dependent, since $\mathbf{w} = \mathbf{u} + \mathbf{v}$.

$(1,0,1) + (0,1,1) = (1,1,2)$. Any set in which one member is a combination of the others is dependent.

B2. 1.

All three vectors are multiples of $(1,2)$, so the span is a single line.

B3. 5.

A basis for $\mathbb{R}^n$ always has exactly n elements — that is what dimension means.

B4. S must be linearly **dependent**, and it contains a basis for $\mathbb{R}^4$ as a subset.

Seven vectors cannot be independent in a 4-dimensional space. In Session 13 you will meet the geometric algebra version of this fact: wedging together more vectors than the dimension allows gives exactly zero.

Section C — Inner Products, Norms, Orthogonality, Projection

Where this is used: Sessions 7–9 (embedding similarity) and Session 14 (contraction, projection, rejection)

C1. 5.

$$\sqrt{(9 + 16)} = \sqrt{25}.$$

C2. 0 — the vectors are **orthogonal**.

$$(1)(2) + (2)(-1) + (2)(0) = 2 - 2 + 0.$$

C3. They are perpendicular: the angle between them is 90° .

In Session 14, orthogonality becomes the condition under which contraction returns zero — the algebra's way of saying *there is nothing of this in that*.

C4. $(3, 0)$.

$\text{proj} = ((\mathbf{u} \cdot \mathbf{v}) / \|\mathbf{v}\|^2) \mathbf{v} = (3/1)(1, 0)$. Session 14 gives the geometric algebra form of this same operation, written with a contraction and an inverse.

C5. $1/\sqrt{2} \approx 0.707$, corresponding to 45° .

$\cos \theta = (\mathbf{u} \cdot \mathbf{v}) / (\|\mathbf{u}\| \|\mathbf{v}\|) = 1/(\sqrt{2} \cdot 1)$. This is the similarity measure that every embedding model in Part II relies on, so it is worth being fluent rather than merely familiar with it.

Section D — Matrices as Linear Maps

Where this is used: Sessions 8–9 (embeddings, attention) and Session 11 (tensor product representations)

D1. $(2, 3)$.

D2. The columns are the images of the basis vectors: column j is $T(\mathbf{e}_j)$.

This is the most useful single fact about matrices, and the one most often missed. A matrix is not a grid of numbers; it is a record of where the basis went.

D3. 3×4 .

D4. False. Matrix multiplication does not commute in general.

Hold on to this. When the geometric product turns out to be non-commutative in Session 12, that is not a defect or an oddity — it is the same phenomenon, and it is precisely where the information about *order* is stored. Composition of transformations has never commuted.

Section E — Determinants, Area, and Orientation

Where this is used: Sessions 12–13. This is the most load-bearing section in the assessment.

E1. 5.

$$(2)(3) - (1)(1).$$

E2. The **area** of the parallelogram spanned by the two columns (equivalently, the rows).

E3. -5 . The sign records **orientation** — the handedness or sense of rotation from the first column to the second.

This is the crux. In Session 12 you meet the wedge product $\mathbf{a} \wedge \mathbf{b}$, which generalizes exactly this: an object carrying an area *and* an orientation, with $\mathbf{a} \wedge \mathbf{b} = -(\mathbf{b} \wedge \mathbf{a})$. If the sign of the determinant has always seemed like an accounting convention to you, that impression is worth revising before Week 7, because the whole of geometric algebra is built on taking it seriously as geometry.

E4. The columns are linearly dependent — parallel — so the parallelogram they span is degenerate and has zero area.

The geometric algebra counterpart: $\mathbf{a} \wedge \mathbf{a} = 0$, and more generally the wedge of dependent vectors vanishes.

E5. The **volume** of the parallelepiped spanned by the three columns.

And in n dimensions, the n -dimensional hypervolume. Session 13 gives these objects a name — blades of grade 2, 3, and higher — and lets you compute with them directly rather than only measuring them.

Section F — Rotations and Orthogonal Transformations

Where this is used: Session 15 — rotors

F1. $[[\cos \theta, -\sin \theta], [\sin \theta, \cos \theta]]$.

F2. The transformation is **orthogonal**: $A^T A = I$, equivalently $A^{-1} = A^T$.

Rotations and reflections are exactly the orthogonal transformations. Session 15 reconstructs both without matrices, using rotors — and the Cartan–Dieudonné theorem, that every rotation is a composition of reflections, is what makes that possible.

F3. No — rotations about different axes do not commute; order matters.

Try it with a book: rotate 90° about the vertical axis then 90° toward you, then repeat in the opposite order. The results differ. Rotors compose by multiplication, and their non-commutativity is inherited directly from this fact.

Section G — Optional (not scored)

G1. SVD factors any matrix into a rotation, a scaling along orthogonal axes, and another rotation — revealing the directions along which the map stretches most.

Used conceptually in Session 8 for dimensionality reduction. Taught in the course; not assumed.

G2. A nonzero vector whose direction is unchanged by the matrix — $Av = \lambda v$, with the scalar λ the eigenvalue.

Background only; not required for the geometric algebra unit.

What Your Score Means

18–21. You have what the course needs. Skim Dorst Chapter 1 before Week 7 and you will be fine.

13–17. Normal and workable. Identify the one or two weakest sections and spend two or three hours on them using the guide below, ideally before Week 6.

8–12. Close the gap deliberately rather than hoping. Work through the review material for your weak sections during Weeks 2–5, while Part I is running and the load is lighter, and come to office hours once to check your progress.

Below 8. Please come to office hours in the first two weeks. This is not a judgment about whether you should be here — several strong students in past years started in this range — but it does mean front-loading some work while there is room for it. We will put together a specific plan with you.

Any section at 0–1, regardless of total. Treat that section as the priority, especially if it is Section E. A high total with a weak Section E is a worse position for this course than a middling total with a strong one.

Review Guide

All sources below are freely available or on Moodle. You do not need to work through a whole textbook; target the sections you missed.

Fastest route, if you have three hours

- **3Blue1Brown, Essence of Linear Algebra** (YouTube, ~15 short episodes). Episodes 1–4 cover Sections A and B; episode 5 covers determinants and orientation (Section E) better than any text; episodes 9–10 cover inner products (Section C). This is the single highest-return resource for this course specifically, because it is geometric throughout rather than computational.
- **Macdonald, Linear and Geometric Algebra, Chapters 1–3.** Written by someone who teaches geometric algebra afterwards, so the emphasis lands where we need it.

By section

- **Section A (vectors, vector spaces):** Macdonald Ch. 1; 3Blue1Brown eps. 1–2. Focus on basis-dependence of coordinates.
- **Section B (span, independence, basis):** Macdonald Ch. 2; 3Blue1Brown eps. 2–3.
- **Section C (inner products, projection):** Macdonald Ch. 3; 3Blue1Brown eps. 9–10. Make sure you can compute a projection without looking up the formula.
- **Section D (matrices as maps):** 3Blue1Brown eps. 3–4 and 6. The key idea is columns-as-images-of-basis-vectors.
- **Section E (determinants, area, orientation):** 3Blue1Brown ep. 5 — watch it twice. Then Dorst Ch. 2, §2.1–2.4, which begins the transition to the wedge product. **Prioritize this section above the others.**
- **Section F (rotations, orthogonality):** 3Blue1Brown eps. 3 and 13; Macdonald Ch. 5. Being able to write the 2D rotation matrix from memory is enough.

Optional review session

If enough people report gaps in the same areas, we will run an optional evening review session in Week 4 or 5, before the distributional unit ends and while there is still slack. Indicate interest on the reporting form.

Reminder: submit your six section scores and confidence ratings via the form on the course site by Monday night. It takes two minutes, it is not graded, and it determines how we pace Sessions 12–17.