CSCI 145 Problem 16: Can a Scientific Probe Survive Retraining?

October 27, 2026

Submission Instructions

Please upload your work to Gradescope by **11:59pm on Wednesday, November 11**.

- Turn in a single PDF of *handwritten* solutions. Typed (including LaTeX) solutions will not be accepted.
- A complete, good-faith attempt earns submission credit. The grader may select any recent assigned problem for a brief conversation, so keep your work and be prepared to explain each solution and answer a related question. Scheduling details will be announced in class.

Problem 16: Can a Scientific Probe Survive Retraining?

A wildlife-monitoring lab retrains its animal-call embedding model each year. For the same call, last year's code $\mathbf{z} \in \mathbb{R}^r$ and this year's code $\widetilde{\mathbf{z}}$ are related by an unknown orthogonal transformation $\widetilde{\mathbf{z}} = \mathbf{R}\mathbf{z}$. Last year, ecologists validated the trait score $\mathbf{a}^\top \mathbf{z}$. To use that score this year, they need a coefficient $\mathbf{b}$ such that $\mathbf{b}^\top \widetilde{\mathbf{z}} = \mathbf{a}^\top \mathbf{z}$ for every call. The lab saved paired codes $(\mathbf{z}_j, \widetilde{\mathbf{z}}_j)$ for reference calls satisfying $\widetilde{\mathbf{z}}_j = \mathbf{R}\mathbf{z}_j$. Let $S = \operatorname{span}(\mathbf{z}_1, \ldots, \mathbf{z}_q)$.

1. Exactly when do the paired reference calls uniquely determine the transferred coefficient $\mathbf{b}$, and what is its value when they do?

Hint: You may use that an orthogonal transformation preserves inner products and that $\mathbb{R}^r = S \oplus S^\perp$.

Takeaway: *An embedding can rotate after retraining even when it keeps the same information. A score from the old model can be transferred uniquely only if its coefficient lies in the span of the saved reference codes. The reference data must cover the directions the score uses.*