

Interpreting Large Language Models (LLMs) with Geometry

Target Audience:

The content is targeted toward a non-expert audience, including high school or earlier stage students, undergraduate students in non-CS related majors, educators, professionals in other domains, or the general public. Ideally, the audience has a general interest in how large language models (LLMs) work. No prior experience with programming or AI is required.

Expected Read/Lecture Time:

15 minutes

Brief Description:

This material introduces students to the mathematics behind Large Language Models (LLMs) through the lens of geometry. The lecture slides connect familiar trigonometry concepts (i.e., cosine from SOH-CAH-TOA) to the more advanced idea of **cosine similarity** and **hyperspheres** used in higher dimensional geometry. The lecture slides provide a concise, simple introduction to the higher dimensional mathematics behind LLMs. Students would learn these topics in a Calculus 3 or Linear Algebra course.

Students will:

- Review the basic terminology related to LLMs to provide a shared vocabulary throughout the lecture
- Recall how to calculate **cosine** in right triangles (two-dimensions)
- Learn how cosine extends into **higher-dimensional spaces** where LLMs embed text data as numbers
- Learn how **cosine similarity** measures **semantic similarity**—the closeness in meaning between words, phrases, or ideas
- Understand why geometry is central to how LLMs process, compare, and retrieve information

Optional:

This lecture slide deck includes optional **advanced material**. The advanced material presents state of the art research on what the LLM researchers are actively investigating. It introduces a unit hypersphere (the higher dimensional unit circle from trigonometry), and derives how cosine similarity is related to the standard dot product (also known as the Euclidean standard product). The slides summarize new research on how different dot products can be used to support the **Linear Representation Hypothesis**.

This slide deck blends **clear visuals**, **real-world analogies**, and **accessible math** to demystify how LLMs organize meaning. By the end, students will see that geometry is not just for shapes on paper—it's the hidden compass to interpreting AI's ability to understand and process language.