Is Cosine-Similarity of Embeddings Really About Similarity?

Xinze LI

Mar 23, 2024

Background

This work is from Netflix & Cornell

Leveraging cosine similarity

- **Quantify semantic similarity between high-dimensional objects** by applying cosine-similarity to a learned low-dimensional feature embedding
- Motivation: The norm of the learned embedding-vectors is not as important as the directional alignment between the embedding-vectors.

Advocation

- ***Cosine similarity*** of the learned embeddings can in fact yield arbitrary results.
- Explanation: **Learned embeddings have a degree of freedom** that can render arbitrary cosine-similarities even though their (unnormalized) dot-products are well-defined and unique.

Solution

Solution

- Apply layer normalization
- Avoid the embedding space, which caused the problems outlined above in the first place, and project it back into the original space.