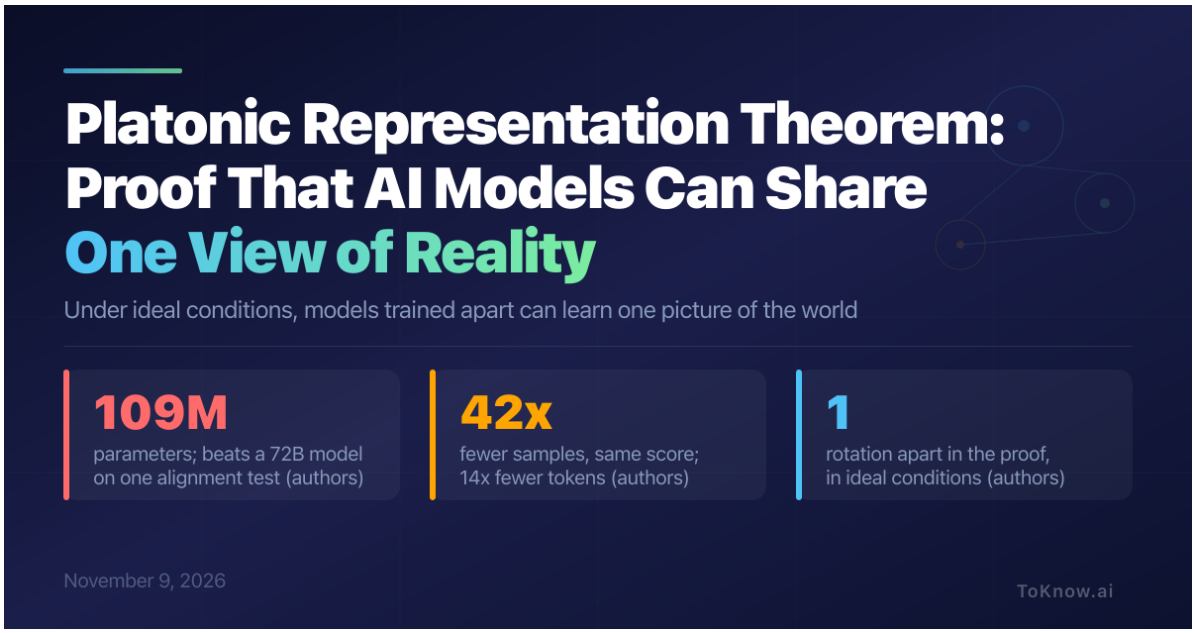
Platonic Representation Theorem: Proof That AI Models Can Share One View of Reality

Kabui, Charles

2026-11-09

Read at ToKnow.ai



Two AI models trained on different data, such as photos and captions, may end up with the same picture of the world. Randall Balestriero and Yann LeCun prove a version of that 2024 MIT idea, the Platonic Representation Hypothesis, for Joint-Embedding Predictive Architectures (JEPAs): models that learn by matching two related inputs in compressed form instead of predicting the next word. When the inputs are good views, such as two captions of one

photo, separate JEPAs learn the same representation (the numbers a model builds for an input) up to one rotation, like one map turned sideways. The proof assumes a finite world, no lost detail, pairs drawn from the same scenes, and perfect training.

The authors argue that next-token prediction, how chatbots learn, is a JEPA with bad views: two parts of one document. On their 0-to-1 alignment test, a 109-million-parameter good-view text JEPA averaged .168 over four runs (.121 to .191), against .154 for Qwen2.5-72B and .151 for Llama-3-70B. The test takes 1,024 Wikipedia image-caption pairs and checks how many of each item's 10 nearest neighbors a text model and the image model DINOv2 share. With 21.6 million samples the JEPA matched next-token training at 900 million (.109 against .111): 42 times fewer samples, 14 times fewer text tokens (word pieces).

Two in three training texts came from an image-captioning model, so they carry some image knowledge. In a smaller run, those texts scored .107 and text-only paraphrases .060. A search on November 9 found no outside rerun.

Read More: AMI Labs, the company Yann LeCun founded to build world models, and the $1.03B it raised

Sources:

- Proving and Improving the Platonic Representation Hypothesis with JEPAs (arXiv)
- Full text with Table 3 and known limitations (arXiv HTML)
- The Platonic Representation Hypothesis, Huh, Cheung, Wang and Isola (arXiv)
- Code and results table (GitHub)
- AI Weekly note on the paper

Disclaimer: *For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation.* ***Read more:*** */terms-of-service*