BraintankAI

DISPATCH · LINEAR SUPERPOSITION · 3 MIN

Blend two texts into one input and a language model partly tracks both

Averaging two unrelated passages into one input leaves traces of both in a model's predictions, though pulling them apart takes extra training and still falls short.

8 November 2026 · arXiv preprint, not yet peer reviewed



Photo: Nina Storz

Blend two unrelated passages into a single input, and a language model's guess at the next word lands, roughly, between the guesses it would make for each passage alone. That is the central finding of a 2026 preprint, not yet peer reviewed, by Pavel Tikhonov, Elena Tutubalina and colleagues. The effect is partial and relies mostly on easy-to-guess words. Still, it suggests that a transformer's internal representations can carry two streams of text at once, mixed together below the level of the words themselves.

Transformers, the architecture behind most large language models, are built from heavily non-linear parts. Each pass through one is usually assumed to handle a single stream of meaning. Earlier work by Razzhigaev and colleagues found that the step from one internal layer to the next can often be closely approximated as linear. The authors asked whether that linearity extends to the whole model: does mixing two inputs give the mix of the two outputs? They call this the Superposition Linearity Hypothesis.

They took pairs of equal-length passages from web text and simple stories. At each position they averaged the embeddings, the lists of numbers a model uses to represent each token, or word fragment. The blended sequence went into unmodified models from several families, including Pythia, Qwen2.5 and Llama; the largest reported had 8 billion parameters. These models have vocabularies of at least 50,000 tokens. Even so, the word a model picked for a single passage appeared in the blended output's top 10 about 30 to 40 per cent of the time, and in its top 100 about 60 to 65 per cent of the time. The authors also measured how far the blended output sat from the exact average of the two separate predictions. They compared that gap with the distance between predictions for two unrelated texts. The ratio stayed below one across four models, three distance measures and two passage lengths, ranging from 0.31 to 0.71. Across Pythia's training checkpoints, the additive behaviour was strongest early and weakened as pretraining went on. The authors read this as a property of the architecture rather than something learnt.

A closer look at Qwen2.5-3B tempers the picture. About 65 per cent of positions are classed as predictable, where general word-frequency habits do much of the work; the rest are content positions that depend on the specific passage. With blended input, the model's own choice had a median rank of 6 at predictable positions but 284 at content positions. The authors also tested 200 items from LAMBADA, a benchmark in which a model predicts the final word of a passage. Accuracy fell from 73 per cent unblended to 2.25 per cent blended. In control runs on single passages, the authors scrambled attention, the mechanism by which each token draws on earlier ones. Median ranks then rose into the thousands, which the authors take as evidence that word-frequency habits alone are not enough to keep predictions on track.

The authors then largely restored the property through training, using self-distillation, in which a model learns to imitate a frozen copy of itself. The student saw blended input and was taught to match the teacher's average of the two separate predictions. Training used less than 0.025 per cent of the original pretraining data. For Qwen2.5-3B, the median rank at content positions fell from 284 to 5. At the rank level, the chance that each stream's word made the top five rose from about 30 per cent to more than 60 per cent. Ordinary performance suffered: Pythia-2.8B's LAMBADA accuracy dropped from 0.544 to 0.357.

Getting two separate continuations out of one blended pass is harder, because sampling straight from the mixture hops between the texts. As a proof of concept, the authors combined the large model's blended output with a smaller model reading each passage separately. Their best result came from Llama-3.2-3B guided by a one-billion-parameter model. Two-stream LAMBADA accuracy rose from 0.182 to 0.430, below the 0.540 the small model reached alone. The authors argue that superposed inference could, in theory, double throughput and halve the memory each stream needs to store earlier context.

Much remains open. Most analyses used contexts of 128 tokens or fewer, largely in one language, and only text models were tested. Blends were equal-weight only and involved at most three streams; three streams matched the target less closely than two. Fine-tuning degrades ordinary prediction, and the decoding method still trails a small model working alone. The work shows that traces of two texts survive a single pass through a model. Reading both back out cleanly remains unsolved.

PAPER	*Your Transformer Can Hold Two Thoughts at Once: Evidence of Linear Superposition in LLMs*
AUTHORS	Pavel Tikhonov, Anton Korznikov, Matvey Mikhalchuk, Nikita Dragunov, Temurbek Rahmatullaev, Polina Druzhinina, Anton Razzhigaev, Ivan Oseledets, Elena Tutubalina et al.
VENUE	arXiv preprint — not yet peer reviewed · 24 September 2026
LINK	arxiv.org/abs/2609.29845
INDEX	stacks.braintank.ai