BraintankAI

DISPATCH · SPECULATIVE DECODING · 4 MIN

# Small drafter speeds up diffusion text models by guessing which blank they fill next

A 0.62M-parameter helper predicts the order in which a diffusion language model commits tokens, cutting the big model's passes without changing its output.

18 November 2026 · arXiv preprint, not yet peer reviewed

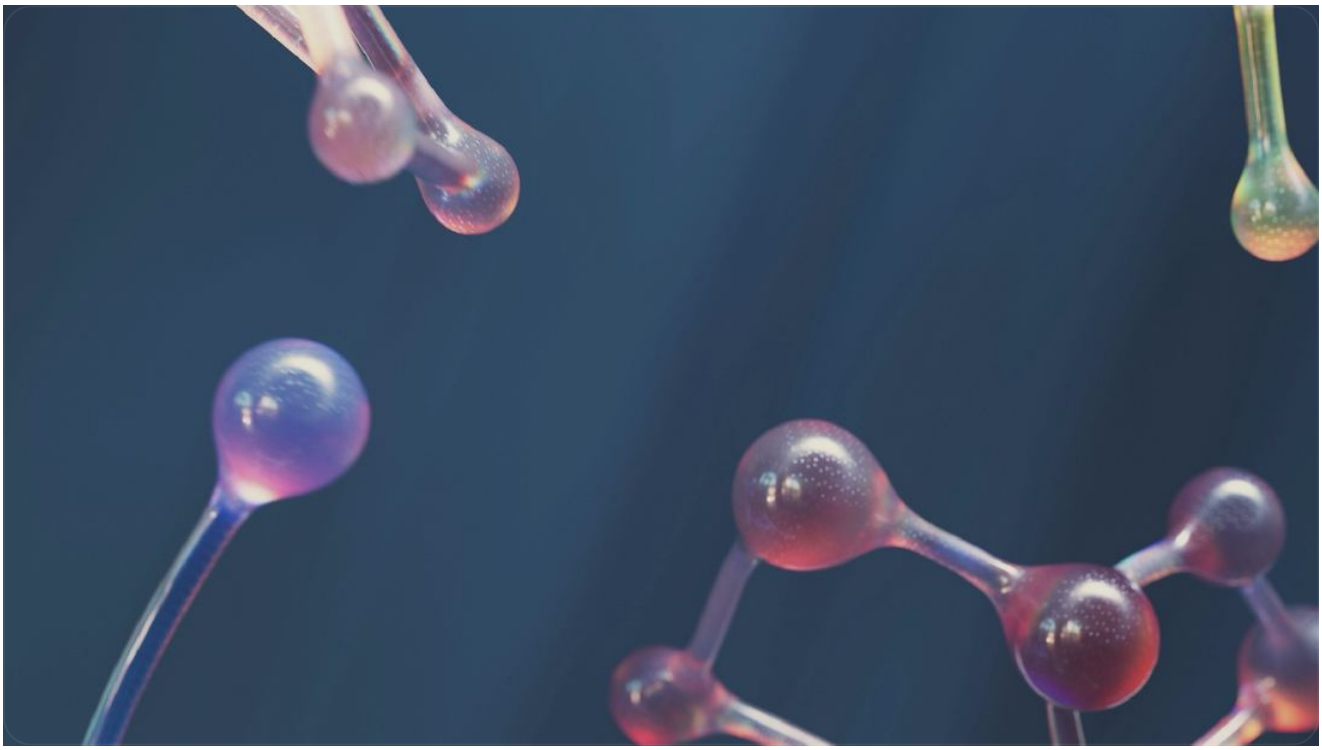


Photo: Google Deepmind

Diffusion language models write text differently from chatbots that produce one word after another. They start from a row of blank, or masked, slots and fill them in over repeated passes of the model. In principle that allows many words at once. In practice, the exact version of the procedure fills only one slot per pass, and the faster shortcuts alter what the model would have written. A 2026 preprint by Eliezer Shlomi, Eliron Rahimi, Avi Mendelson and Chaim Baskin, posted on arXiv and not yet peer reviewed, argues that the missing ingredient for a lossless speed-up is knowing the order in which the slots get filled.

The standard lossless trick for ordinary language models is speculative decoding. A cheap helper guesses several steps ahead, and the big model checks all the guesses in one pass, keeping only those it would have produced anyway. That does not carry over neatly. The models studied fill text in blocks, and within a block every slot attends to every other, so committing one slot changes the

predictions for the remaining slots in that block. A helper therefore has to guess both the word and the position. The authors' diagnostics suggest the word is the easy part. The big model's out-of-date predictions from its previous pass match the eventual token in 92% of first-block commits and at least 90% thereafter. However, the chance that the previous round's most confident slot is the one filled next falls to 3% by block 9. In oracle experiments, being handed the true positions cut model passes on GSM8K by 5.4× to 7.2× as the draft budget rose from 8 to 20. Being handed the true tokens yielded only 3.4–6.5× at a budget of 20.

Their drafter, SPOT, is accordingly tiny and narrow. It is a three-layer transformer with 0.62M parameters. It reads the big model's stale top-32 probabilities, a few summary statistics and a record of which slots the draft has already filled. It sees no text and no internal states. It outputs which slot will commit next, and on held-out runs it picked correctly 75% of the time, against 22% for simply following stale confidence. Tokens are borrowed from the stale predictions. In greedy decoding, the second choice is added when the top one has probability below 0.6. A best-first search assembles a tree of candidate block states. These are appended behind the cached text so a single pass checks them all, and a draft is accepted only if it equals exactly what the big model would have produced. For random sampling, helper and model share one table of Gumbel noise, the random numbers used to draw tokens, which makes each sampled step deterministic. The authors prove that the output distribution then matches unaccelerated sampling at any temperature.

The main tests used SDAR-8B on GSM8K, MBPP, HumanEval and MATH500. SPOT was compared against re-implementations of the Spiffy calibrated draft graph, an SSD-style chain and a BlockSpec-style tree, all run through the authors' own decoder and verifier so that only the drafting differed. At temperature 0, the authors report 3.2–4.7× fewer passes than standard decoding at budget 8 and 4.1–6.8× at budget 20, with 2.5–3.4× throughput at budget 8. Against Spiffy, SPOT needed 16–36% fewer passes in every task and budget. On MBPP at budget 8 it took 56.7 passes, against 83.6 for Spiffy and 269.2 without drafting. Throughput fell as budgets grew, steeply for the chain and tree baselines. At budget 20 the tree became slower than standard decoding on MBPP. At temperature 1, accuracies stayed within 0.04 of unaccelerated runs.

A larger 5.1M-parameter variant, SPOT-Sets, targets dynamic threshold decoding, which fills every slot whose confidence clears a cutoff in one pass. On SDAR-8B at temperature 0, it reached a mean 7.9× throughput over standard decoding, against 5.0× for the threshold decoder alone. That headline uses an acceptance rule that changes the output; the variant that preserves it reached 7.6×.

The authors note several limits. Each draft costs a full block of computation, so budgets must stay small, and verification becomes the bottleneck at wide blocks. Transfer to LLaDA-8B and Dream-7B relied on an approximate cache and was not lossless. Other caveats follow from the setup. The comparison with MRP, the closest learned drafter, uses that system's own reported figures rather than a head-to-head run. The main tables cover four maths and coding benchmarks at a generation length of 256, where base accuracy is sometimes low, such as 0.43 on MBPP. Most results beyond the main tables come from a single seed on one GPU type, and some ablation figures differ from the main table without an explanation in the available text.

braintankAI has covered a theory paper proposing that a lightweight predictor could decide what work a costly AI system performs, though it offered formulas rather than measurements (earlier story). SPOT is a measured instance of that pattern: a model of under a million parameters steering what an 8-billion-parameter model checks.

| | |
|---|---|
| PAPER | *On the SPOT: Enhancing Speculative Decoding for Diffusion Language Models via Commit-Order Prediction* |
| AUTHORS | Eliezer Shlomi, Eliron Rahimi, Avi Mendelson, Chaim Baskin |
| VENUE | arXiv preprint — not yet peer reviewed · 26 October 2026 |
| LINK | arxiv.org/abs/2610.22543 |
| INDEX | stacks.braintank.ai |
| CODE | github.com/Eliezer318/SPOT-dLLM |