Teaching LLMs to Reason Like Bayesians: Training on Process Beats Training on Answers

Kabui, Charles

2026-03-11

Read at ToKnow.ai

Google Research published a paper in Nature Communications showing that LLMs struggle with Bayesian inference, the mathematically optimal way to update beliefs as new evidence arrives. In a controlled flight recommendation task where a model must infer a user's preferences over multiple rounds, an optimal Bayesian assistant reached 81% accuracy while off-the-shelf

LLMs performed considerably worse and plateaued after a single interaction instead of improving with more evidence. The fix: fine-tune LLMs not on the correct answers ("oracle teaching") but on the Bayesian assistant's own predictions, including its early wrong guesses. This "Bayesian teaching" approach, a form of distillation, trained the models to maintain uncertainty and update beliefs gradually. The fine-tuned models agreed with optimal Bayesian predictions 80% of the time and, critically, generalized across domains, transferring reasoning learned on synthetic flight data to hotel recommendations and real-world web shopping without retraining.

Any system that personalizes recommendations, adapts to user behavior, or makes sequential decisions under uncertainty benefits from better probabilistic reasoning. Current LLMs default to crude heuristics (like always recommending the cheapest option) instead of tracking individual preferences. Bayesian teaching fixes this with standard supervised fine-tuning, no architectural changes required.

The deeper finding is that teaching models the reasoning process, even when it produces wrong answers early on, beats teaching them correct answers directly. This echoes a pattern emerging across AI: Meta FAIR's multimodal pretraining research similarly found that training on the right structure matters more than brute-forcing the right outputs. Process over product seems to be a general principle for building models that generalize.

Sources:

- Teaching LLMs to Reason Like Bayesians (Google Research Blog)
- Bayesian Teaching Enables Probabilistic Reasoning in Large Language Models (Nature Communications)
- Bayesian Inference (Wikipedia)

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