





The UChicago Data Science Clinic team partnered with Morningstar to use AI agents to automate fine-tuning open-source frontier models and help build AI on a budget. Building on their previous work with parameter-efficient fine-tuning methods and synthetic data generation, the team built an end-to-end agent framework to handle the entire model development process with minimal human intervention.

To do this, the team wrote all the tools an AI agent needs to train its own models. They:

- created a new method to generate synthetic data from limited real-world examples,
- wrote code to fine-tuning open-source frontier models to excel at specialized tasks,
- designed custom metrics to use reasoning models to evaluate summary quality,
- then compressed the resulting model to make it even cheaper to run.

The team demonstrated that it is possible to build agents that outperform out-of-the-box LLMs across several metrics, including one custom-developed by the team. Recognizing that many evaluation metrics check for letter-by-letter similarity, the team created Deepseek Accuracy, which uses frontier reasoning models to evaluate summarization quality.





The clinic team is hopeful that, as agents' capabilities improve and as their workflows get less linear and more iterative, tools like this will make developers more productive. This approach provided a sustainable path to building models that can compete with cuttingedge open-weights models while controlling for costs related to data acquisition, computational resources, and engineering time.