Synthetic Data Generation & Fine-Tuning Specialist

Author: Santhi Maddipudi

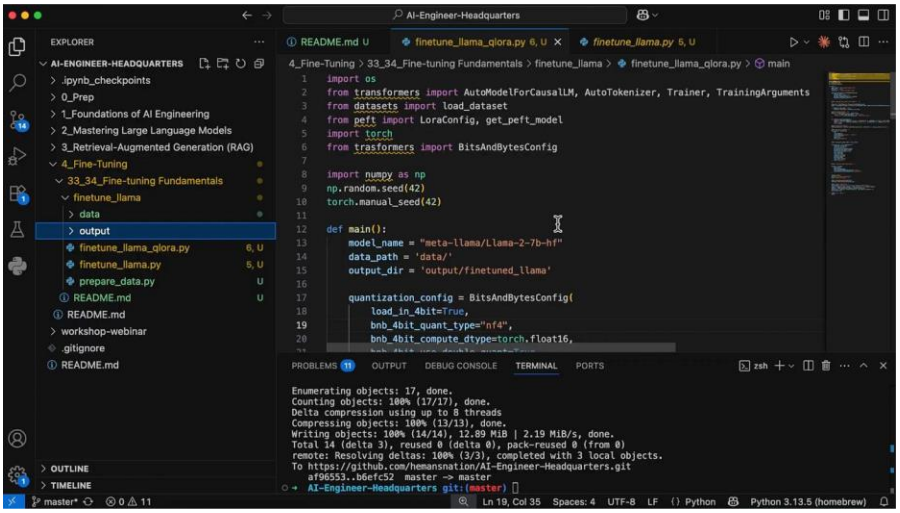

Problem we are solving

GPT-4 costs $30 per million tokens.

For high-volume tasks like generating SQL queries, writing product descriptions, or answering support tickets, these costs destroy budgets.

Meanwhile, GPT-4 is overkill for most tasks.

A small 8B model fine-tuned for one specific job can match GPT-4 quality at 1% of the cost.

Tech stack

Unsloth, Hugging Face TRL, WandB, Ragas, GPT-4 API, Python.

Proof of work :

In progress