Sun Shines Bright Vol.2

by computational mama, Sameera Mudgal and lekkerying

Fine-tuning a Small Language Model

INTRO

An easy dipstick for any new language model I typically ask a few questions. Some of them include (no web search allowed):

- Can you name 5 non white non male science fiction writers?
- What is the carrier bag theory of fiction?
- Who is Gayatri Chakravarty Spivak?
- Can you explain the thematics of Jacinta Kerketta's work?

Though commercial AI models are trained on a vast amount of texts, simple critical questions, popular non-Western feminist authorship, essential understanding of cultural multiplicities remain missing.

This project is a handmade attempt to see if we can use our own personal compute power to build and fine-tune our own small language model - a sort of pastiche of a language model!

The labour of data

There is a cruel irony that creating the dataset for a machine to learn is the most human labour intensive part of the process. And it feels deeply unsettling that as big tech does this at scale in countries with low labour rates and regulations, these data workers are the least represented in the data.

For our dataset we needed a minimum of 300 question and answer pairs. In our case, it was just Sameera and me creating this dataset. So, we came up with a simple solution. We created a set of closed Q&A like **"Q:What is the capital of India? A: The capital of India is New Delhi."** . And open Q&As like **"Q: What's the idea behind the Carrier Bag Theory of Fiction by Ursula Le Guin? A: The Carrier Bag Theory of Fiction, proposed by Ursula K. Le Guin, reimagines the foundational narrative of human technological development by suggesting that the earliest and most."** .

Ideally, we would encourage you to create the dataset as a group - what we like to call - a "community learned model".
Once created, the dataset must be formatted as "jsonl" file for the fine tuning step.

Fine-tuning the Model

When you have a small language model, you can add more of your own data and style by first freezing like 90% of the model and then training only that small 10% that is left. That method is called LoRA.

You need a Mac with Apple silicon for this, MLX only runs on M1 or newer.

Step 1
- Create a dataset and have it in a jsonl format (Check resources)

Step 2
- Create a virtual environment in Python
- Activate the environment
- Install 'MLX' and 'MLX-LM'

```
python3 -m venv femdata
source femdata/bin/activate
pip install mlx mlx-lm
```

- Run the lora (Make sure you are using the path for the data folder correctly.)

```
mlx_lm.lora \
  --model HuggingFaceTB/SmolLM2-360M-Instruct \
  --train \
  --data path/to/datafolder \
  --iters 800 \
  --batch-size 2 \
  --adapter-path ./adapters/smolFem2
```

- Run the model with the adapter that you created

```
mlx_lm.generate \
  --model HuggingFaceTB/SmolLM2-360M-Instruct \
  --adapter-path ./adapters/smolFem2 \
  --prompt "Instruction: What do we learn from
  carrier bag theory of fiction?\nOutput:" \
  --max-tokens 1000
```

- Fuse the model

```
mlx_lm.fuse \
  --model HuggingFaceTB/SmolLM2-360M-Instruct \
  --adapter-path adapters/smolFem2 \
  --save-path ./models/smolFem \
  --de-quantize
```

Learnings

In the end, the compute was too little to steer the model a fully critical direction.

Here are some of the learnings and outcomes from the process:

- It was also apparent that the "closed QnA" didn't affect the model's learning of new facts.
- This fine-tuning can steer the models tone and conversational style but it is a trial and error process.
- It doesn't seem IMPOSSIBLE to try fine tuning a model, it was only a few lines of terminal commands.
- Can we come together as a community with shared compute, larger handmade dataset to create a "better" model?
- Can testing this process as a group, give us greater insight into model creation and the data labour involved?
- Can we use customized small language models for our projects?

Resources

List of the essays for the open Q&A

The Seeds of Suicide: How Monsanto Destroys Farming — Vandana Shiva

Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective — Donna Haraway

The Master's Tools Will Never Dismantle the Master's House — Audre Lorde

Can the Subaltern Speak? — Gayatri Chakravorty Spivak

Programming is Forgetting: Toward a New Hacker Ethic — Allison Parrish

More info at:
https://solararchive.cmama.xyz/2025/11/howtofinetuneamodel/

made by
computational mama, Sameera Mudgal and lekkerying
on 21/09/2026