



Fine Tune - Class 1 Notes

Why this class?

As models become more advanced every quarter, and sometimes more expensive, it's still difficult for authors to get a model to "write like them." Frustrations boil because you want to save time on editing writing mistakes that a model is making over and over and over again.

The answer is a fine tune.

You can fine tune select models on OpenAI, Mistral, and Google.

The public cannot fine tune Claude models at this time (but if you LOVE Claude, you could fine tune a cheaper model to write just like Claude).

And that is what fine tunes do best: they change the style, tone, and output of an LLM response.

They do NOT:

- help the LLM memorize facts.
- work well with generic prompting.

Your Story Information (scene briefs, beats, instructions)

plus

Consistent Prompting with a Fine Tune

equals

Outputs that need minimal editing.

Additionally, as more and more authors turn to using AI to keep their readers happy, a fine tune is YOUR ticket to writing with AI, but never sounding like you wrote with AI.

The fine tune course will

- Walk you through everything you need to know about making a fine tune data set so that you can make your own.
- Provide easy software to make it a snap for you to organize your data and format it properly into JSONL.
- Give you 9 datasets to get started with so you can modify them or test them as is.

We are going to cover over the next 10 modules:

- Intro to Finetuning
- What IS a Dataset?
- Your First Dataset: Better Outlines
- Consistent Outputs: Scene Briefs
- Write Like Me
- New Advances in Fine Tuning
- Conversational AI dataset
- Genre specific fine tune datasets
- DPO make 4o to write like o1
- How to modify/make a dataset yours in Dinotrainer

Let's learn more about fine tunes.

Recommended Viewing:

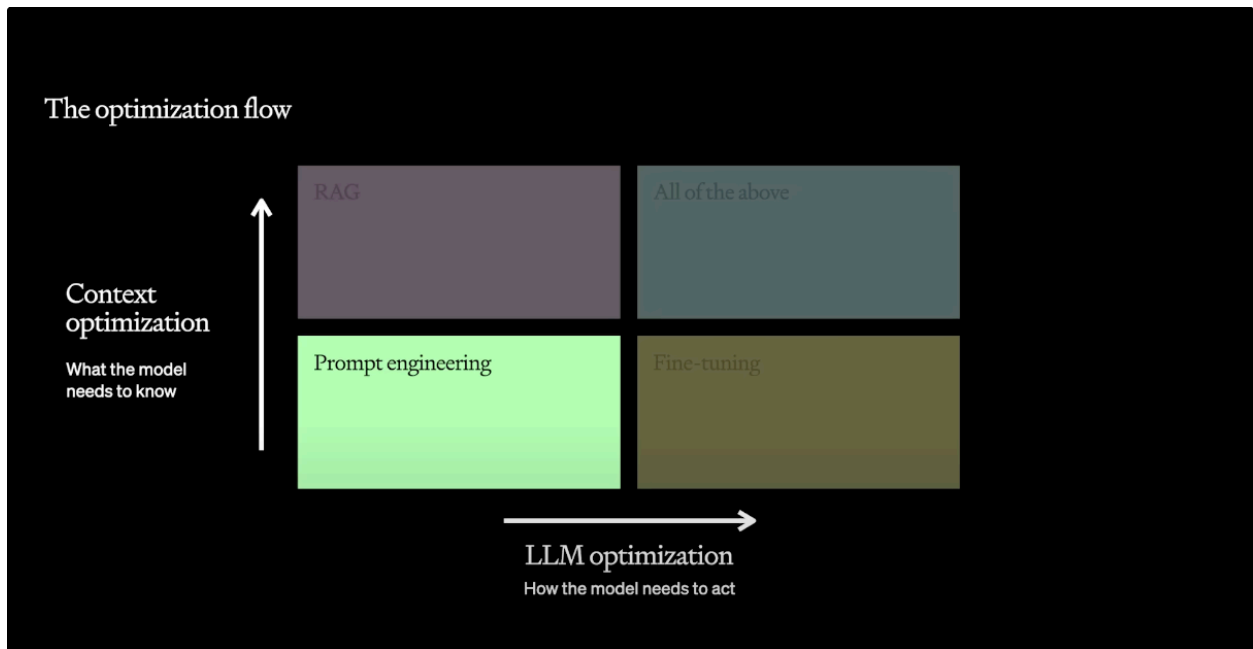
<https://www.youtube.com/watch?v=ahnGLM-RC1Y>

Jump to 31 minutes to see the funny story about training on your own Slack messages....

Too Long, Didn't Watch:

Quadrants of Optimizing LLMs:

This is not easy to do, even by OpenAI's standards!



🔴 YOU MUST HAVE PROMPT ENGINEERING MASTERED BEFORE TRYING TO FINE TUNE!!!

- You can tell me why you like one model over another, such as OpenAI is better at summarizing and Claude, when he wants to write, is more "creative" out of the box.
- You know what a context window is and can explain it to someone else.
- You know the OUTPUT limits of major foundational models.
- You know what role System, User, and Assistant play in OpenAi Playground interface.
- You know how to manipulate Temperature, Top_P, Presence, and Frequency Penalty to get different results.
- You know the limitations and opportunities of LLMs (that they don't think, but you can use chain of thought)
- You can craft specific prompts that give you decent results or the results you wanted.

- When you don't get a result you want, you know how to reroll it or tweak it and try both.

If you do not have the above skills, do not worry. But you need to go get them BEFORE you try to do a fine-tune dataset.

What happens if you don't? You will hate this. All of this. Every bit of this. You will be stuck in a mess of not knowing if you need a fine-tune, don't need a fine-tune, or if it's even working? Most importantly, you will waste money.

Anatomy of a Dataset for GPT3.5 16k

10 Datasets are at least 10 examples of a Prompt (System and User) and a Response (Assistant).

You can use synthetic data, which means a bunch of AI writing that you validate, such as using outputs from a better model to teach a lower model to make better responses. We did this last year feeding GPT 4 Outlines to GPT 3.5 16k. The resulting fine tune wrote a longer and more detailed outline than either baseline.

You can use ONLY human data as well, or a mixture.

Data sets are written in JSONL formatting.

Here are some demos of examples.

For a dataset, you would make multiple versions of these:

System: You are a
fabulous writing
assistant named
FabulousGreg. You
analyze writing and

System: You are a
marketing genius
named
MarkyMarketingMark.
You take a blurb of a

System: You are a
writer's best friend who
always gives good story
ideas mixed with humor
named LoudLizzy.

rewrite it to remove cliches.

User: Rewrite the following: "She was between a rock and a hard place, because she couldn't decide between Brad or Chad. Her stomach was in knots over the two men, and little did she know, her life would never be the same."

Assistant: She found herself torn between two men who loved her: Brad and Chad. The two men could not be more different in how they made her feel, but she didn't want to hurt either one. Somehow, the decision felt final. Too final.

book and write 3 funny short hooks for social media videos.

User: Blurb:

Frankenstein; or, The Modern

Prometheus tells the story of Victor Frankenstein, a young scientist who creates a sapient creature in an unorthodox scientific experiment.

Though *Frankenstein* is very much a product of the Gothic and Romantic literary movements, it also can be thought of as the first true science fiction story, and has been highly influential on subsequent literature and popular culture, spawning a complete genre of horror stories, films, and plays.

Assistant:

Hook 1: You thought he was dead...

User: I need story ideas for my genre of a historical fiction book involving a duck.

Assistant: Sure best friend! I am so glad we are hanging out today! Want to go to the mall later? Here's that story idea:

Premise: A duck is stuck at the Eiffel Tower in Paris during World War II, historical fiction

Outline:

Title: Quacks and Courage: A Duck's Tale From the Eiffel Tower...

Hook 2: Will no one be his friend?

Hook 3: It's alive! And very shocking.

There are now 2 NEW versions of Datasets that we will also be learning in this class:

Conversational (where you can give chains of back and forth messages you would like to be more consistent with the model)

Direct Preference Optimization (this is where you give the model a prompt and then a good output example and a bad output example).

Both of these methods are less than 6 months old.

Limitations of FineTuning

YOU CANNOT USE FINETUNING TO TEACH AN LLM FACTS OR DETAILS ABOUT YOUR BOOK!

You can use fine tuning to demonstrate consistent formatting of responses when the LLM is provided similar types of information.

✅ DO: Make a dataset of a variety of high quality story components for your genre (summary, character list, outlines) and ask for scene briefs. Make the example AI responses high quality, consistently formatted scene briefs you would love to have.

✗ DON'T: Run the same character list over and over again and expect the LLM to know your character.

⚠ RISK: Overfitting the model: Overtraining with a small, unvaried dataset that makes the LLM unable to perform other tasks because it's stuck on a very narrow fine-tuned behavior.



But why can't it learn my characters like it knows public domain characters?

Main training an LLM is billions of pieces of training data (millions run multiple times), and weights are applied to those pieces of data.

A fine-tune doesn't actually add additional tokens to the training data, it rearranges some of them.

It's like Play-doh, there is only so much in the container, but use a tool and you can make different shapes than just your hands.



This is also why your fine-tune is your own.

Your fine-tune makes no permanent change to the core training of the model. You can't add or subtract Play-doh that's in the set.

Each time the LLMs get bigger, we get more colors and bigger containers of Play-doh to play with.



And what's an Overfit? When you try to do too much that you turn all those colors into that terrible brown color. And it's broken. You can't get back the pretty, individual Play-doh colors anymore, or in the case of an LLM, other functionality.

**I was wrong about previous ideas about using weird names in System to add a new star to the vector field... instead the LLM knows those weird names "OutlineMageddon" "MarkyMarketingMark" as the tokens they are.

GPT-3.5 & GPT-4 GPT-3 (Legacy)

Outlinemageddon
MarkyMarketingMark

Clear Show example

Tokens Characters
10 34

Outlinemageddon
MarkyMarketingMark

TEXT TOKEN IDS

GPT-3.5 & GPT-4 GPT-3 (Legacy)

Outlinemageddon
MarkyMarketingMark


Clear Show example

Tokens Characters
10 34

[2729, 3817, 336, 3359, 15357, 198, 9126, 88, 68662, 9126]

TEXT TOKEN IDS

That specific sequence of tokens is part of the rearrangement of other tokens inside the LLM, to rearrange them to be closer or further away from other tokens in the dataset.

 A fine tune is: we rearranged the furniture, and that rearrangement had to involve those specific token sequences that were repeatedly in the training data set.

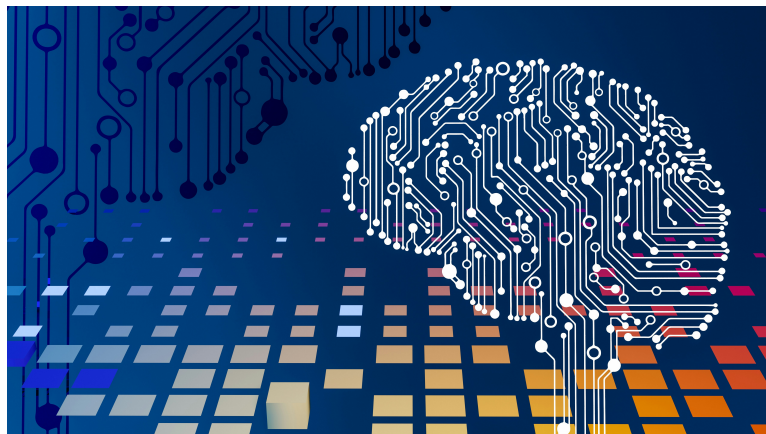
Since those main sequences are gibberish, unlikely to be in a ton of other training, we do help the LLM recognize right away HOW we want the furniture rearranged of some tokens, and not all tokens in the LLM.

But we do not literally write new tokens to the LLM, we only rearrange the tokens inside to change their "address" in relation to other tokens in the vector field.

Sort of like if you imagine these metal bits and screws were "living" in one place on the table, we dropped the magnet (fine tune) and SOME of the metal bits and screws moved directly to the magnet... others stayed put.



A fine tune changes how the existing tokens in the training data set relate to one another, in small "fine tune" amounts. :)



What's Next?

Questions to think about how a fine tune will help you:

What kind of consistent response style do you want from the LLM? (ex: longer, more specific, a particular format, certain tagging)

What kind of prompts (System and User) will you be routinely using in hopes of getting this response? (if you're someone who prompts on the fly, look at old conversations and see if there are any patterns).

Why does the current baseline prompt fall flat? Have you exhausted tricks of prompt engineering like changing hyper-parameters, styles of prompting (chain of thought)?

How will the fine tune be a win?