

(made by David Krueger)

**Don't share without permission!**

Actually I'm just going to put this publicly available on LessWrong.

**If I shared this with you in a slack or email, feel free to edit, else please comment only.**

This is just meant for sharing a few quick thoughts relevant to a few discussions I have had with a few people. I'm not sure everything I write is actually correct, since I haven't used the GPT-3 API myself, and have only spent a few hours on the playground.

#### **How is OpenAI API access limiting?**

- Having full access to something creates a different mindset in practice:
  - You play with it on a whim and discover random unexpected things.
  - You try research ideas more readily.
- Not knowing weights is hard for science of DL / interpretability
- No RL fine-tuning. (probably good to be careful with that one, IMO...)
- No control over fine-tuning hyperparameters.
- API throttling/lack of parallelization
- Iterated retraining (?)
- Multi-head architectures, etc.

Other possible obstacles:

- Lack of expertise: not having much in-house experience with foundation models or related areas of ML (e.g. NLP).
- Lack of culture of such work: if people around you are doing it and talking about it, you are much more likely to as well.

Somewhat related LW post:

<https://www.lesswrong.com/posts/HXxHcRCxR4oHrAsEr/an-update-on-academia-vs-industry-one-year-into-my-faculty>