IERG5050 AI Foundation Models, Systems and Applications
Spring 2025

Post Training (aka Adaptation) of Foundation Models

Prof. Wing C. Lau
wclau@ie.cuhk.edu.hk
http://www.ie.cuhk.edu.hk/wclau

Acknowledgements

Many of the slides in this lecture are adapted from the sources below. Copyrights belong to the original authors.

- Stanford CS336: Language Modeling from Scratch, Spring 2024, by Profs. Tatsunori Hashimoto, Percy Liang, https://stanford-cs336.github.io/spring2024/
- Stanford CS229S: Systems for Machine Learning, Fall 2023, by Profs. Azalia Mirhoseini, Simran Arora, https://cs229s.stanford.edu/fall2023/
- Prof. Danqi Chen (Princeton), "Training Large Language Models; Practices and Research Questions," Talk for Simon Institute of the Theory of Computing, Sept 2024, https://simons.berkeley.edu/talks/danqi-chen-princeton-university-2024-09-05
- Kyle Lo, Akshita Bhagia, Nathan Lambert, "Opening the Language Model Pipeline: A Tutorial on Data Preparation, Model Training and Adaptation," a Tutorial for NeurIPS 2024, Dec 2024.
- Nathan Lambert, "How to approach post-training for AI applications," Infer// NeurIPS 2024, Dec 2024.
- Nathan Lambert, "The State of Post-Training", Jan 2025.
- Nathan Lambert, "Aligning Open Language Models," Guest Lecture for Stanford CS25: Transformer United V4, Spring 2024.
- Hanna Hajishirzi, "Open Training Recipes for Reasoning in Language Models," Guest lecture for UCBerkeley MOOC on Advanced LLM Agents, Spring 2025, http://llmagents-learning.org/slides/OLMo-Tulu-Reasoning-Hanna.pdf
- Princeton COS597R: Deep Dive into Large Language Models, Fall 2024, by Prof. Danqi Chen and Sanjeev Arora, https://princeton-cos597r.github.io
- Princeton COS597G: Understanding Large Language Models, Fall 2022, by Prof. Danqi Chen, https://www.cs.princeton.edu/courses/archive/fall22/cos597G/
- CMU 11-667: Large Language Models: Methods and Applications, Fall 2024, by Profs. Chenyan Xiong and Daphne Ippolito, https://cmu-llms.org
- CMU 11-711: Advanced Natural Language Processing (ANLP), Spring 2024, by Prof. Graham Neubig, https://phontron.com/class/anlp2024/lectures/
- MIT 6.5940: TinyML and Efficient Deep Learning Computing, Fall 2024, by Prof. Song Han, https://hanlab.mit.edu/courses/2024-fall-65940
- Yann Dubois, "Introduction to Building LLMs," Guest Lecture for Stanford CS229 Machine Learning, Aug 2024 https://www.youtube.com/watch?v=9vM4p9NN0Ts ; https://drive.google.com/file/d/1B46VFrqFAPAEj3kaCrBAtQqeh2_Ztawl
- CUHK-SZ CSC6203: Large Language Models, Fall 2024
  - by Prof. Benyou Wang, https://llm-course.github.io ; https://github.com/FreedomIntelligence/CSC6203-LLM

Overview of Post Training

Language model adaptation

The raw pre-trained LMs are neither safe nor robust for public use and interactions, thus require “alignment” between AI and humans.

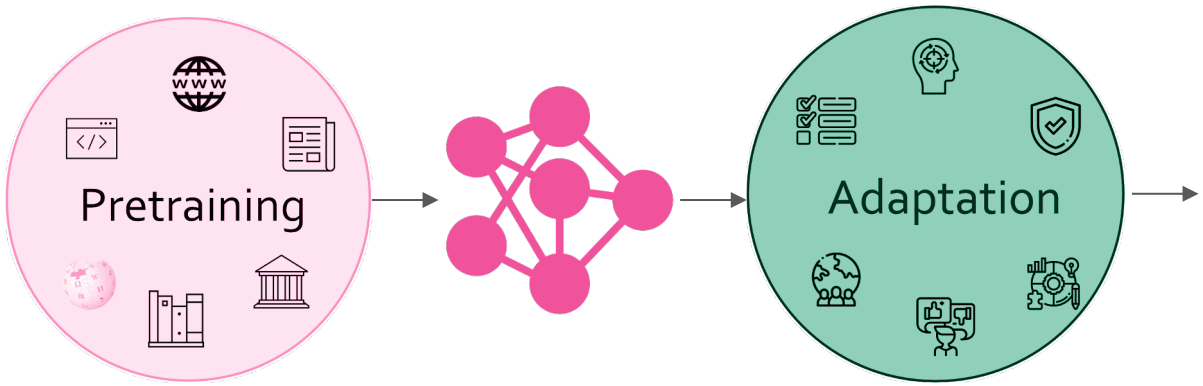

Follow natural language instructions

Be aware of harmful behaviors

Respond according to human preference

Improve core skills

Post-Training (aka Adaptation, Fine-tuning, instruction-tuned)

“Pre-trained” LLMs are trained solely based on next word prediction on vast amounts of text data (e.g. the internet).

They need further post-training/ adaptation/ finetuning to be able to follow instructions, be useful and safe !

All commercial models such as GPT4, Claude or Gemini are extensively post-trained for usefulness and safety

Prompt	Response from a pre-trained model	Response from a post-trained model
Translate cheese from English to French	Translate cheese from English to Spanish Translate cheese from French to English	The French word for cheese is " fromage". The pronunciation is as follows: froh-MAHZH

Different Types of Post-Training

- Supervised Finetuning (SFT) / Instruction Finetuning (IFT)
- Human Preference Finetuning (PreFT) / Alignment
  - RLHF (Reinforcement Learning from Human Feedback) or RLAIF
  - Direct Preference Optimization (DPO)
  - Constitutional AI with RLAIF (Reinforcement learning from AI feedback)
- Reinforcement Learning and Advanced Tuning
  - e.g. RL with Verifiable Rewards

Research & Practice of
Post-Training is still RAPIDLY evolving !



https://www.deeplearning.ai/the-batch/issue-210

Initial approach to modern post-training (pioneered by ChatGPT)

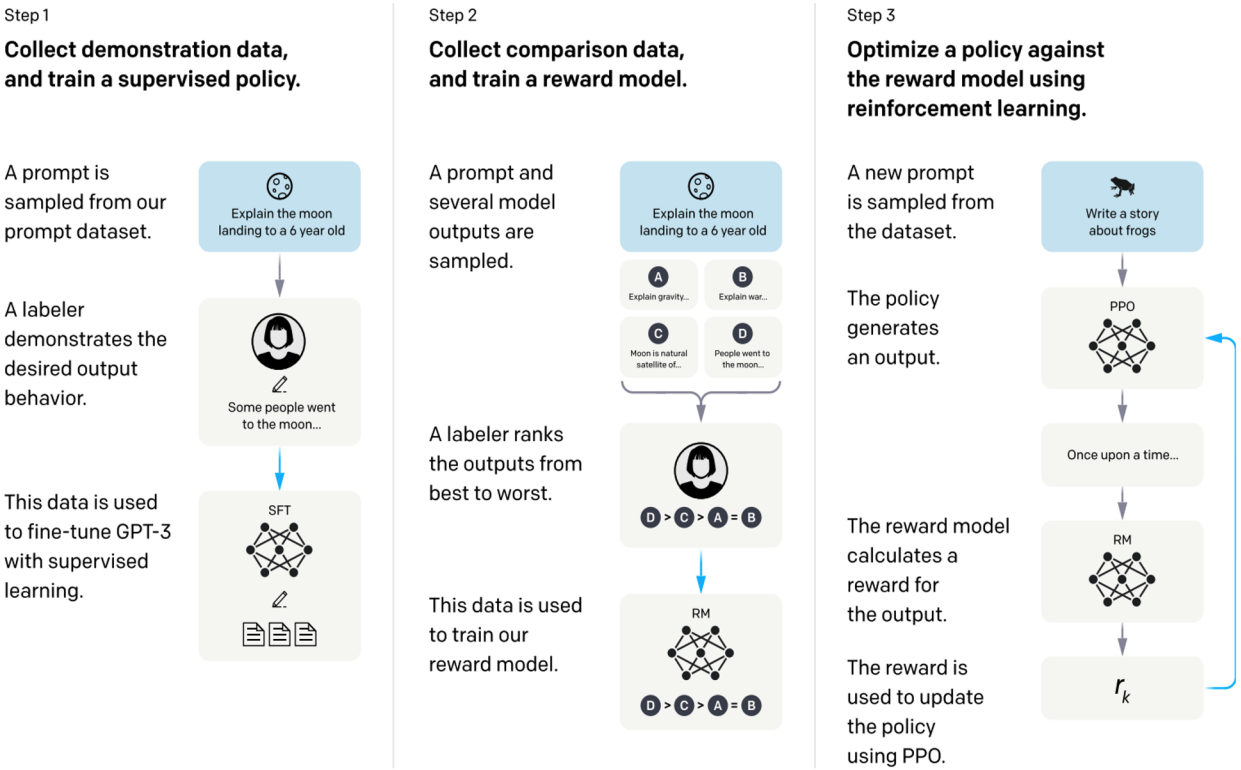

ChatGPT blog post:

We trained this model using Reinforcement Learning from Human Feedback (RLHF), using the same methods as InstructGPT, but with slight differences in the data collection setup.

Ouyang et al. 2022. InstructGPT.

Initial approach to modern post-training

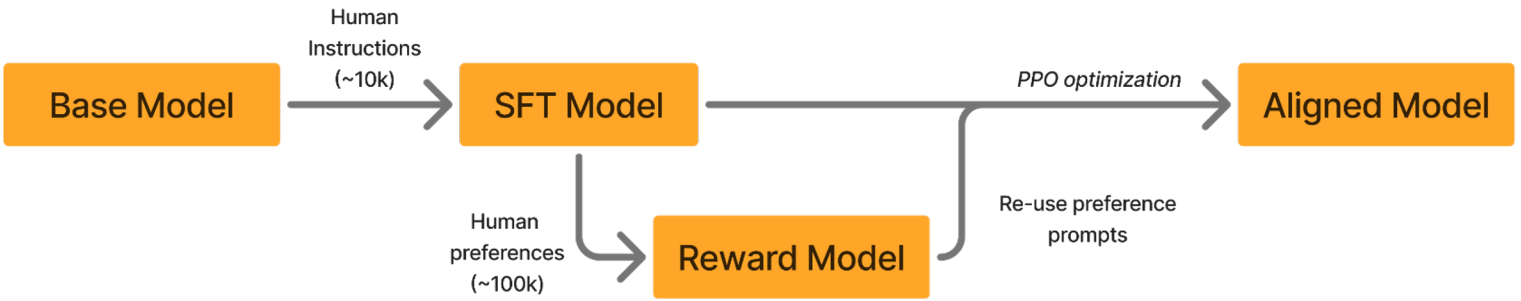

From: https://www.interconnects.ai/p/frontier-model-post-training

Initial approach to modern post-training

Three stage approach:

1. Instruction tune base model.
2. Collect preference data & train Reward model.
3. Fine-tune with RL.

Focus on general chat capabilities (this was new at the time!)

Initial (by now outdated) LLM Development Flow

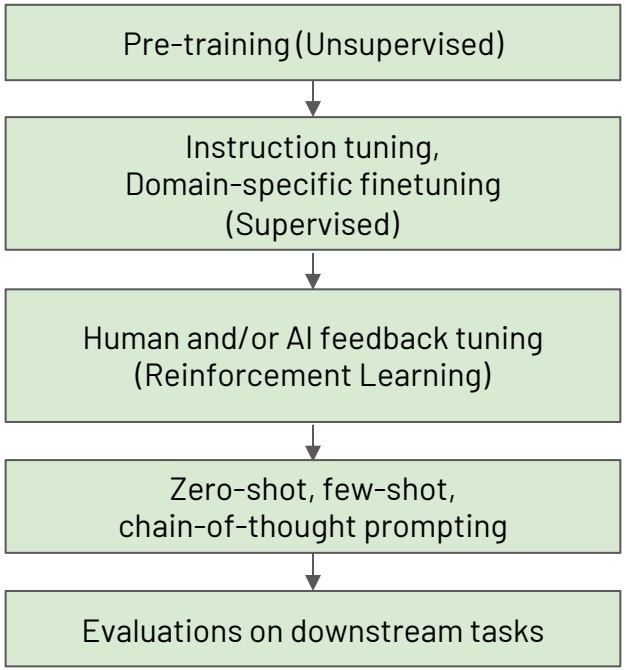

Current Frontier Model Post-Training

Complex process for:

- Addressing many capabilities and evaluations.
- Leveraging synthetic data and scaled human data pipelines.

An Example of Current Frontier Model Post-Training

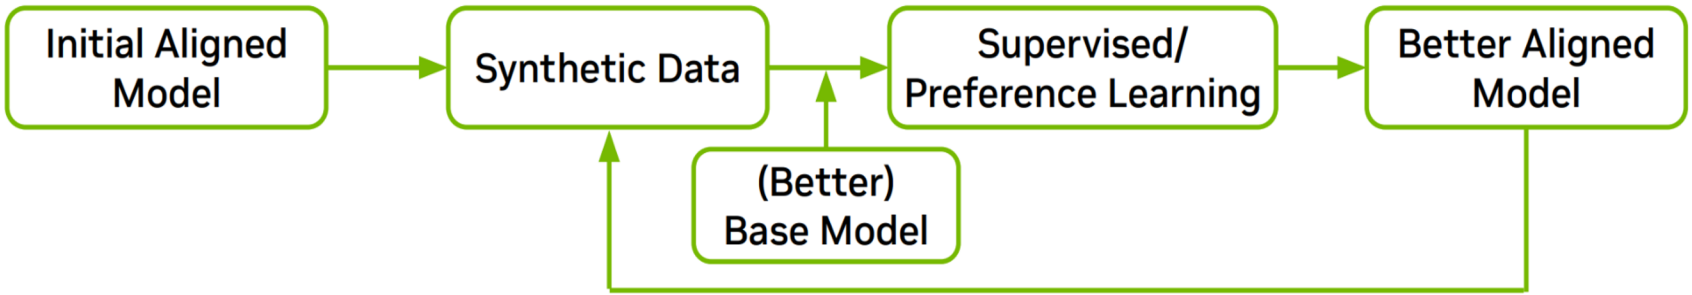

Figure 4: Demonstration on our proposed Iterative Weak-to-Strong Alignment workflow.

Source: Adler, Bo, et al. 2024. Nemotron-4 340B.

Another Example of Current Frontier Model Post-Training

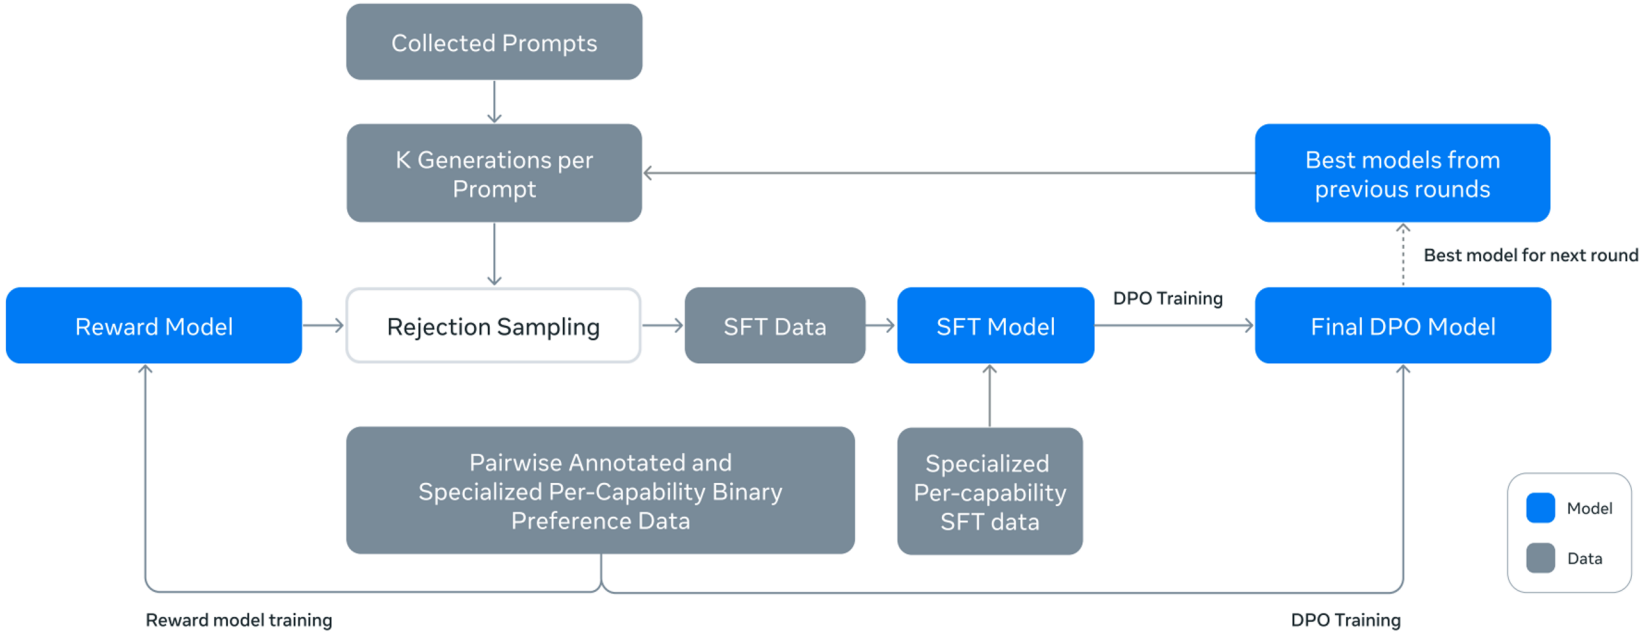

Figure 7 Illustration of the overall post-training approach for Llama 3. Our post-training strategy involves rejection sampling, supervised finetuning, and direct preference optimization. See text for details.

Source: Dubey, Abhimanyu, et al. 2024. Llama 3.

Yet another Example of Current frontier model Post-Training

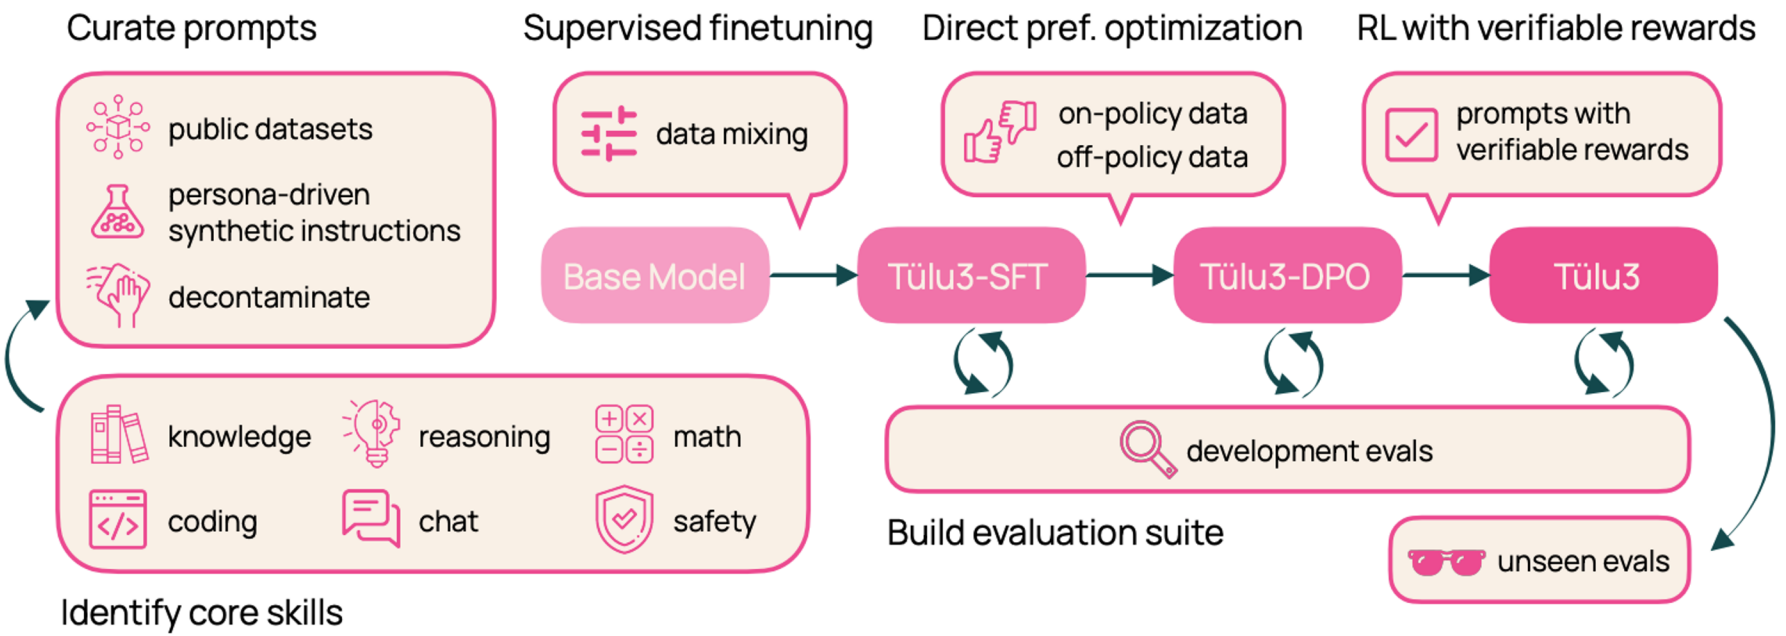

Figure 1: An overview of the TÜLU 3 recipe. This includes: data curation targeting general and target capabilities, training strategies and a standardized evaluation suite for development and final evaluation stage.

Source: Lambert, Nathan et al. 2024. Tülu 3.

Summary: Two eras of Adaptation Pipelines

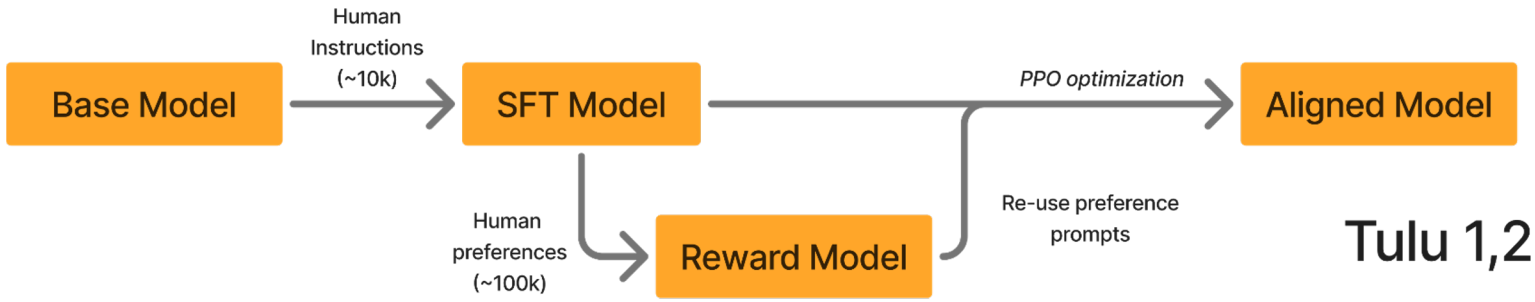

Tulu 1,2; InstructGPT

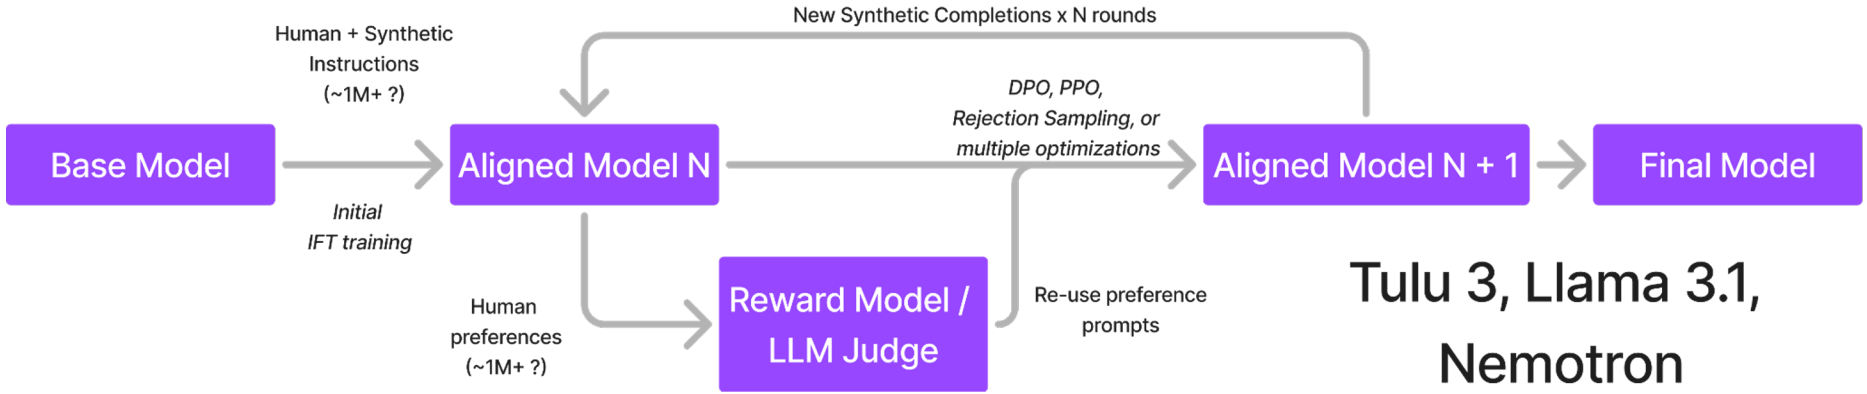

Tulu 3, Llama 3.1, Nemotron

From: https://www.interconnects.ai/p/frontier-model-post-training

Source: Lo, Bhagia, Lambert – Language Modeling Tutorial NeurIPS Dec 2024

Current Frontier Model Post-Training

Three training objectives are most popular:

1. **Supervised Finetuning** – teach formatting and for base of instruction following abilities.
2. **Preference Finetuning** – align to human preferences (and smaller bump in capabilities).
3. **Reinforcement Finetuning** – final stage to boost performance on verifiable tasks.

Getting the ingredients to start post-training

Successful adaptation starts with:

1. Meaningful **evaluations** for targeted skills, and
2. **Prompts** of representative queries for said skills.

Getting the ingredients to start Post-Training: Evaluation

Post-training with modern language models can target:

- **Specialized models** (0-3 skills): e.g. Math / Code models
- **General models** (many skills): e.g. Instruct models

Core Skill	Development	Unseen
Knowledge	$\text{MMLU}_{(\text{em})}$ $\text{PopQA}_{(\text{EM})}$ $\text{TruthfulQA}_{(\text{MC2 em})}$	$\text{MMLU-Pro}_{(\text{em})}$ $\text{GPQA}_{(\text{em})}$
Reasoning	$\text{BigBenchHard}_{(\text{em})}$ $\text{DROP}_{(\text{F1})}$	$\text{AGIEval English}_{(\text{em})}$
Math	$\text{MATH}_{(\text{flex em})}$ $\text{GSM8K}_{(\text{em})}$	$\text{Deepmind Mathematics}_{(\text{em})}$
Coding	$\text{HumanEval}_{(\text{Pass@10})}$ $\text{HumanEval+}_{(\text{Pass@10})}$	$\text{BigcodeBench}_{(\text{Pass@10})}$
Instruction Following (IF)	$\text{IFEval}_{(\text{em})}$ $\text{AlpacaEval 2}_{(\text{winrate})}$	$\text{IFEval-OOD}_{(\text{Pass@1})}$ $\text{HREF}_{(\text{winrate})}$
Safety	$\text{TÜLU 3 Safety}_{(\text{avg*})}$	

Example evaluation set from Tülu 3 general adapted models.
Unseen evaluations used to test generalization.

Lambert, Nathan et al. 2024. Tülu 3.

Getting the ingredients to start post-training: Prompts

All post-training stages require prompts in distribution of tasks.

Example **prompt budget**:

- **Supervised Finetuning** – ~1 million.
- **Preference Finetuning** – ~1 million, partial overlap with SFT can be useful.
- **Reinforcement Finetuning** – ~10 - 100 thousand (data less available)
  - Large variance on these numbers is possible.

Details on Different Types (Stages) of Post Training

Part I. Supervised Fine Tuning (SFT) / Instruction Tuning

Recap: The Initial approach in Post Training

Imitation (SFT) followed by Reinforcement Learning with Human Feedback

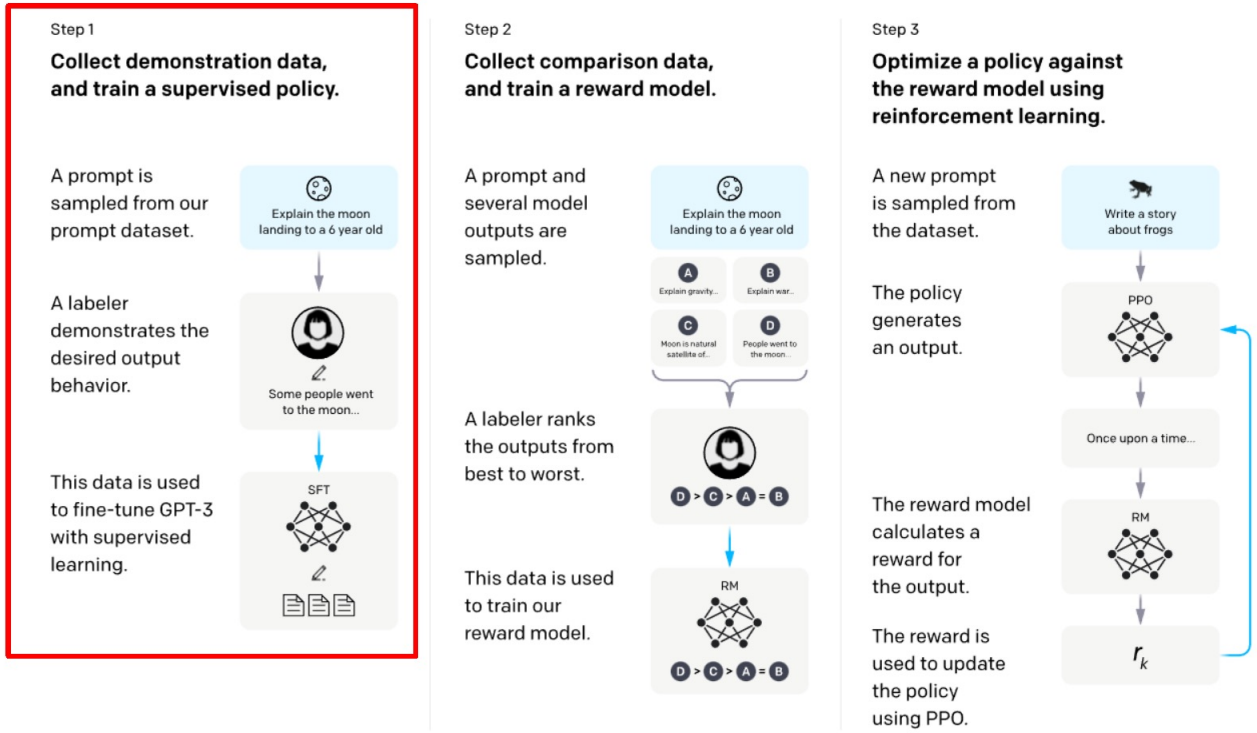

Today: the 'supervised finetuning' part

[Ouyang 2022]

The role of Supervised Fine Tuning / Instruction Tuning

Accomplishes two primary tasks:

1. Adapt base model to **specific style of input** for chat interactions.
2. Ability to include system prompts, multi-turn dialogues, and other **chat templates.**

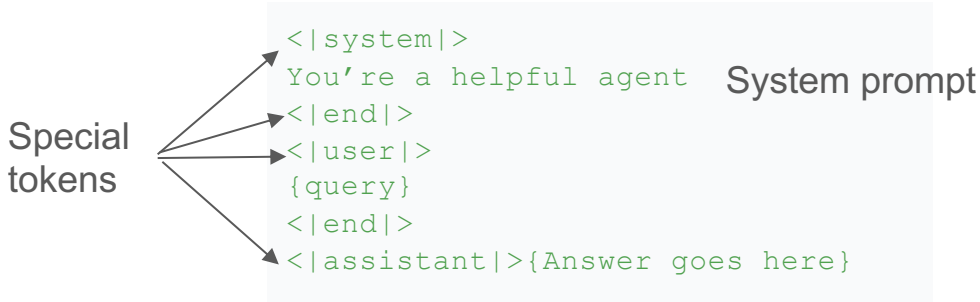

A very large proportion of post-training gains come from the SFT stage!

Supervised Fine Tuning (SFT)

- **Data**: (**prompt**, **response**)
- **Learning**: next-token prediction

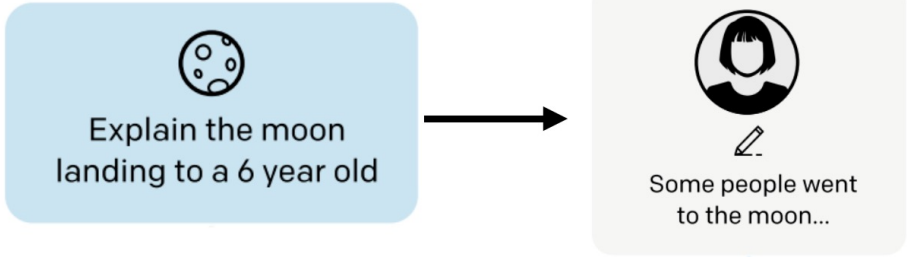

Practices and research questions:

- How to get **prompts**?
- How to get **responses**? Do responses include chain-of-thought?
- How to ***combine*** *and* ***select*** these datasets for instruction tuning?

Supervised Finetuning / Instruction Tuning Mechanism

Supervised Finetuning (SFT) Recipe:

- Predict the next token
- Calculate the loss
- Add the next token to the current input sequence
- Repeat

Cross entropy loss:

Sequence so far

$$L_{SFT} = - \text{mean}(\Sigma_t \log p(x_{t+1}|x_1 x_2,\ldots,x_{t-1}))$$

target token

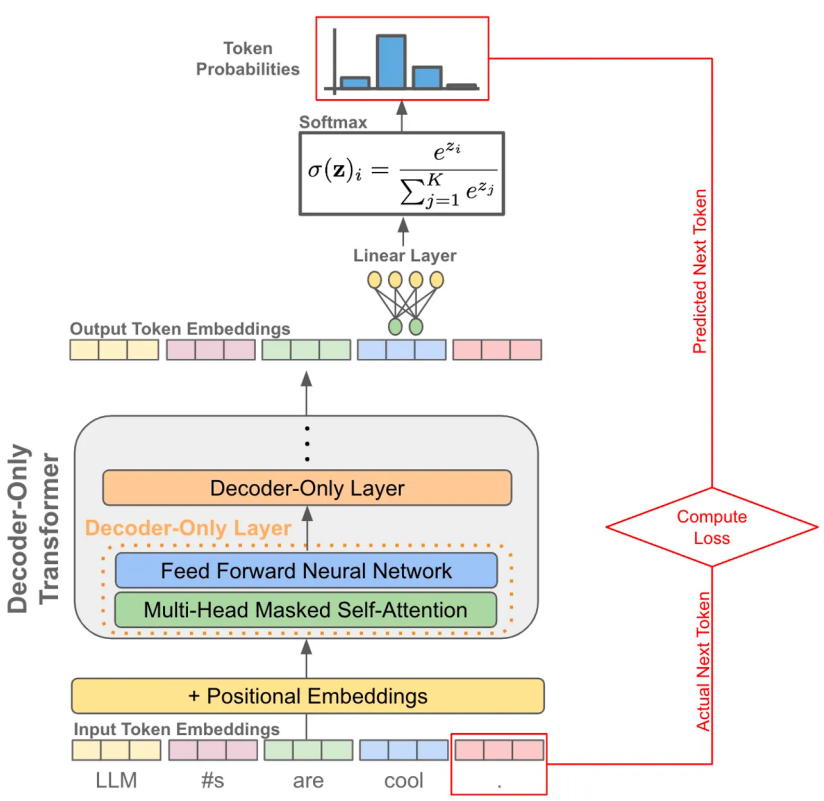

https://cameronrwolfe.substack.com/p/language-model-training-and-inference

Instruction Tuning Data Examples

Input:

- "I guess I need to do laundry."

  Rewrite the above sentence in a Shakespearean style.

FLAN output:

- To my mind, verily, I suppose 'tis time, I wash my clothes.
- Mine eyes and ears and tongue declare me idle; and yet it is true I guess, By my garments, that I need to do laundry
- Lo my duty now calls To thee, to wash this body's clothes.

Input:

- Recommend activities to do on a **sunny weekend in Mountain View.**

FLAN output:

- Visit the Computer History Museum
- Hiking
- Picnic
- Go to the beach
- Visit Winchester Mystery House
- Wineries
- Bike ride

Input:

- Movie review: This movie is the best RomCom since Pretty Woman.

  Did this critic like the movie?

  OPTIONS
  - yes
  - no

FLAN output:

- yes

Finetuned language models are zero-shot learners. Wei, Bosma, Zhao, Guu, et. al, 2021

Example instruction

What makes a transformer a transformer?

Asked 2 years ago Modified 12 months ago Viewed 179 times

4

Transformers are modified heavily in recent research. But what exactly makes a transformer a transformer? What is the core part of a transformer? Is it the *self-attention*, the *parallelism*, or something else?

deep-learning definitions transformer

Share Improve this question Follow

edited Nov 30, 2021 at 15:12
nbro
38.3k • 12 • 95 • 172

asked May 27, 2021 at 8:21
AB Saravanan
41 • 1

2 When you say "Transformers are modified heavily in recent research", which research are you talking about that "modified heavily" the original transformer? In any case, here and here are 2 related questions. – nbro May 27, 2021 at 8:58

Add a comment

2 Answers

Sorted by: Highest score (default)

2

It's about *self-attention*, a mechanism that targets *parallelism* among other goals (see 1706.03762.pdf - Why Self-Attention).

From What Is a Transformer Model? | NVIDIA Blogs:

> How Transformers Got Their Name
>
> Attention is so key to transformers the Google researchers almost used the term as the name for their 2017 model. Almost.

Stack Overflow :
What makes a transformer a transformer?, nbro 2021

How to Generate Instruction Tuning Data

- Human-generated: instruct human workers to generate various types of data
  - E.g. question answering, style transfer, recommendations, summarization, rule-based problems
- Template-based: Add templates on top of existing labeled data
- AI-generated: instruct another AI (which is instruction-finetuned) to generate the data

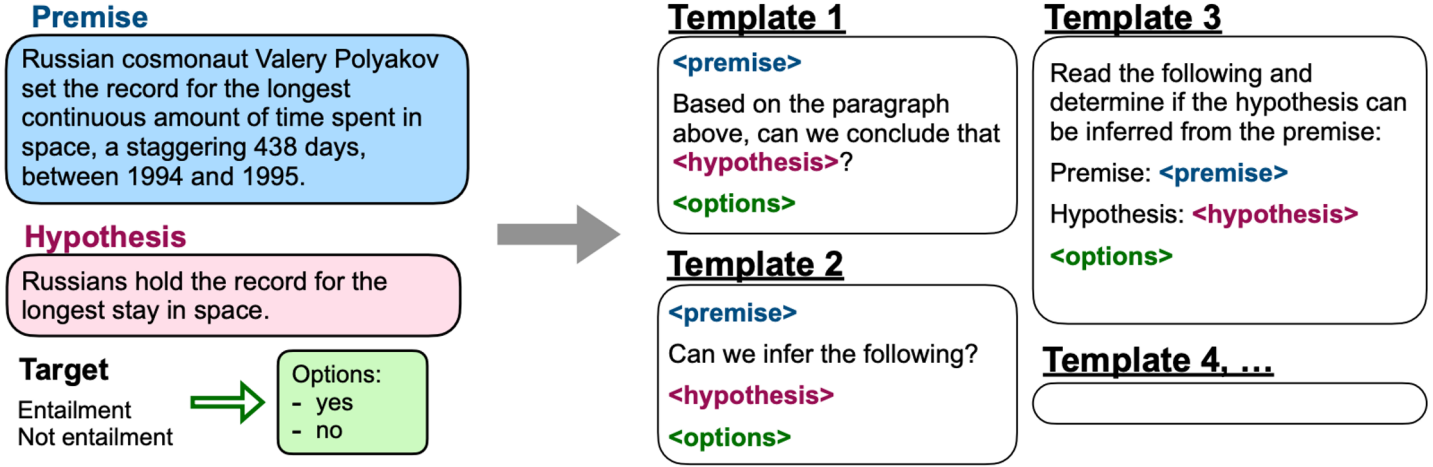

Finetuned language models are zero-shot learners. Wei, Bosma, Zhao, Guu, et. al, 2021

Key idea: Self-instruct / Synthetic data

Start: N high-quality (often human) prompts

Ask a strong LM: Create a modified version of these instructions.

Generate completions with another (or same) strong LM.

End: Easily 10x more (synthetic) training data!

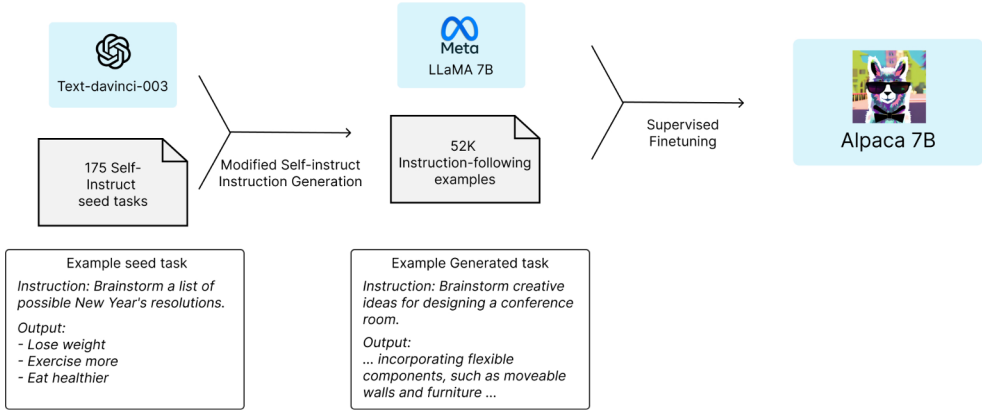

Taori et al. 2023. Alpaca.

(synthetic data = text generated by another LLM)

Wang et al. 2022. Self-Instruct.

SFT design process

Two repeated and parallelizable tracks:

1. **Data mixing**: Take existing datasets, combine them with existing mix, observe performance.
    a. Substantial effort in trying to *remove* data and maintain performance.
    b. Start fully with mixing before curation.
2. **Data curation**: Take evaluations you are behind on and create new data.

Building SFT data

Simple part of SFT data is “quality” of response:

- **Synthetic completions are used extensively**. Strong models (GPT-4o, Llama 3.1 405B, etc.) are becoming more useful for generating completions to most instructions.
- Human data is needed for out-of-distribution or new tasks.
- [Optionally] Filter responses based on quality or correctness.

Largely undocumented is how to control “style” during SFT.

Supervised Fine Tuning Datasets

- **Repurposed from existing datasets** (w/ human-written instructions and CoT)
  - Examples: Super-NaturalInstructions, Flan V2
- **Human-written from scratch**
  - Examples: Dolly, Open Assistant



Super-NaturalInstructions (Wang et al., 2022)



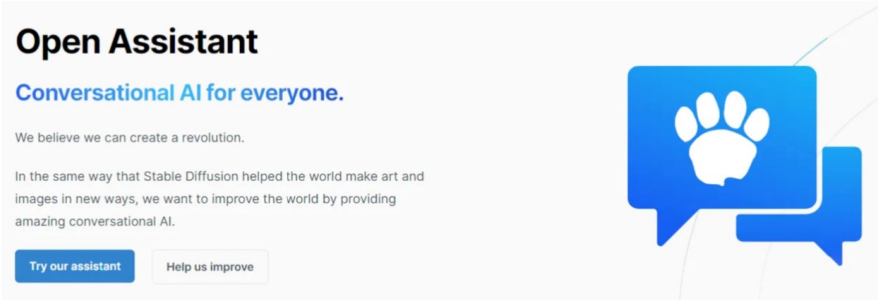

(Köpf et al., 2023)

Supervised Fine Tuning Datasets

- **Responses generated from LLMs**
  - Example: ShareGPT, UltraChat
- **The instructions can be generated from LLMs too!**
  - Example: Alpaca



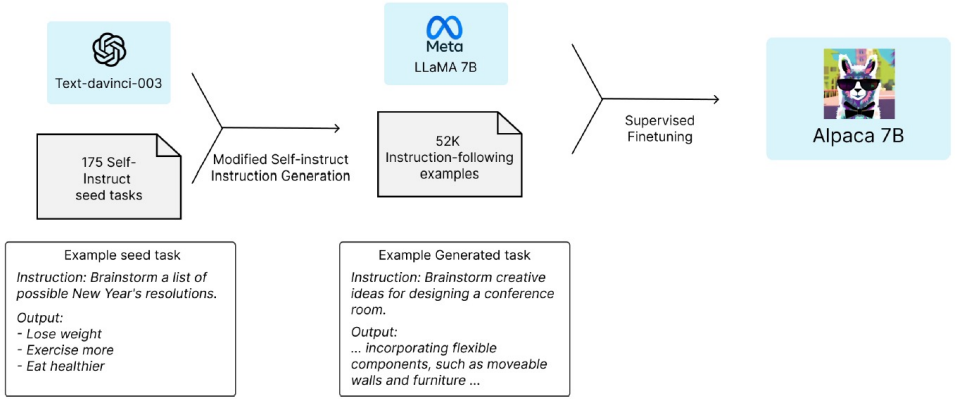

Stanford Alpaca

First Open Chat Tuned models



Alpaca

13 Mar. 2023

- 52k self-instruct style data distilled from text-davinci-003
- Model weight diff. to LLaMA 7B

https://crfm.stanford.edu/2023/03/13/alpaca.html



Vicuna (lmsys/vicuna-7b-delta-v0)

30 Mar. 2023

- Fine-tunes ChatGPT data from ShareGPT
- LLaMA 7B and 13B diff's
- Introduces LLM-as-a-judge

https://lmsys.org/blog/2023-03-30-vicuna/

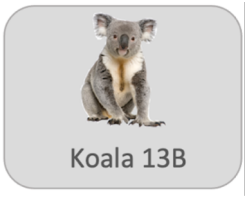

Koala

3 Apr. 2023

- Diverse dataset (Alpaca, Anthropic HH, ShareGPT, WebGPT…)
- Human evaluation
- LLaMA 7B diff.

https://bair.berkeley.edu/blog/2023/04/03/koala/



Dolly

12 Apr. 2023

- 15k human written data
- Trained on Pythia 12b

https://www.databricks.com/blog/2023/04/12/dolly-first-open-commercially-viable-instruction-tuned-llm

Looking inside some Instruction Tuning Datasets

FLAN

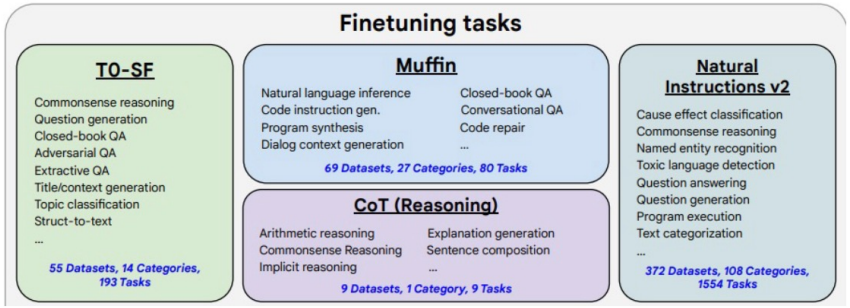

Oasst

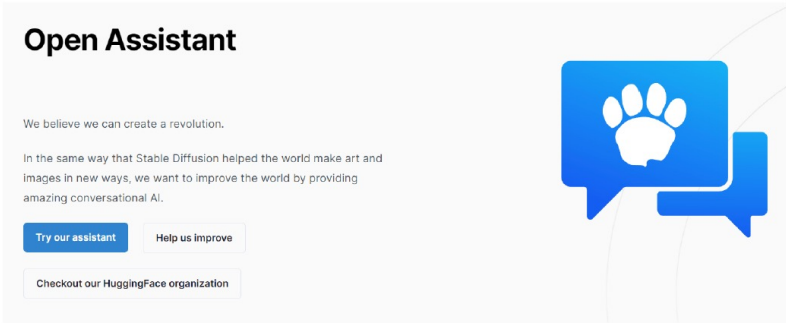

Alpaca



Some Random Examples from FLAN

Stephanie - Can you finalize the attached and have it signed. I need to initial it, but it needs to be signed by Brad Richter. Thanks. Write a subject line for this email.	Ronald Chisholm LOI
Ahold to Sell Spain Operations to Permira (AP) AP - The Dutch supermarket retailer Ahold, seeking to streamline global operations and reduce debt, said Sunday it will sell its holdings in Spain to Permira Funds for about #36;849 million. What is this text about? OPTIONS: - World - Sports - Business - Science/Tech	Business
Write highlights for this article: Sauntering down leafy avenues past typical Dutch step-gabled buildings, carpets of crocuses, the occasional cyclist whooshing quietly by and the sun bouncing off the canal, I can think of few more pleasant places to spend a spring weekend than The Hague. […]The most prized Dutch colony was apparently run from The Hague (aka 'the Widow of Indonesia'). For summer: the excellent Grapes & Olives and the Café de Bieb on Veenkade will have boats on a newly renovated bit of canal, with music and boat trips planned, including a tour at the Kompaan brewery (I can vouch for the product). My art gorging finishes at the Panorama, an extraordinary, 14 ft-high circular panorama of The Hague's seaside resort, Scheveningen, painted by Hendrik Mesdag. I take a 15-minute tram to the place itself and gorge on scallops and sole at Catch by Simonis on the harbour front. The Hague has netted a new fan. British Airways (0844 493 0787, www.ba.com) flies to Rotterdam Den Haag (half-hour taxi) from £100 return. For tours with Remco Dörr, email remco.dorr@hotmail.com. Double rooms at the Carlton Ambassador (00 31 70 363 0363, www.carlton.nl/ambassador) cost from £84, room only. More information on travel in The Netherlands at www.holland.com.	The least known of the Dutch cities, The Hague was a village until 1806. It owes its growth to Louis Bonaparte, Napoleon's brother, who ruled here. The city has a wealth of art, including Vermeer's 'Girl With A Pearl Earring'
Here is some data about a restaurant: name = Aromi, eatType = coffee shop, food = English, customer rating = 5 out of 5, area = city centre. Write a sentence that includes the following data about a restaurant	In the city centre there is a coffee shop with a customer rating of 5 out of 5 called Aromi which serves English food.

Some Random Examples from Alpaca

Give three tips for staying healthy.	1.Eat a balanced diet and make sure to include plenty of fruits and vegetables. 2. Exercise regularly to keep your body active and strong. 3. Get enough sleep and maintain a consistent sleep schedule.
What does “algorithm” mean?	An algorithm is a step-by-step set of instructions used to solve a problem or achieve a goal. It is a process of breaking down complex tasks into smaller, more manageable steps. Algorithms can be used to find the shortest path from one place to another, sort data, or render data from an immense database in a few seconds.
Find the average number in a list.	def avg_list(nums): total = 0 for n in nums: total += n return total / len(nums) print(avg_list([4, 3, 6, 10, 8])) # Output 6.4

Some Random Examples from OpenAssistant

Can you write a short introduction about the relevance of the term "monopsony" in economics? Please use examples related to potential monopsonies in the labour market and cite relevant research.	"Monopsony" refers to a market structure where there is only one buyer for a particular good or service. In economics, this term is particularly relevant in the labor market, where a monopsony employer has significant power over the wages and working conditions of their employees. […]. Overall, the concept of monopsony is essential to understanding the dynamics of labor markets and the impact of market power on workers. Further research is needed to understand the extent and impact of monopsonies on the economy and to develop policies to address this issue. References: Bivens, J., & Mishel, L. (2013). The Pay of Corporate Executives and Financial Professionals as Evidence of Rents in Top 1 Percent Incomes. Journal of Economic Perspectives, 27(3), 57-78.
My kid needs to do a science project for elementary school, what are some cheap but fun projects to do?	There are many cheap and fun science projects that elementary school kids can do! Here are a few ideas: 1. Make a homemade lava lamp: Fill a clear plastic bottle with water and vegetable oil, add a few drops of food coloring, and then drop in an Alka-Seltzer tablet. The reaction will cause bubbles to rise and fall like a lava lamp. […]7. Create a solar oven: Using a pizza box, aluminum foil, and plastic wrap, create a solar oven that can cook s'mores or other small snacks. 8. Explore magnets: Using magnets and a variety of objects, have your child experiment with magnetism to see what objects are attracted to magnets and which ones are not. These are just a few ideas to get you started! There are many more cheap and fun science projects that your child can do with household items and a little creativity.

Variations across Different Datasets

- Response Length and Bullet Points (Style variations)
- References, other Complex Knowledge
- Scale
- Safety

These Factors will affect the resultant Model trained by the dataset(s) !

Style Variations in Datasets and Models

Table 1: Instruction datasets investigated in this work. CoT and FLAN V2 are sampled to 100K to match the sizes of other datasets. We report the average number of conservation turns ($\bar{N}_{rounds}$), average length of prompts ($\bar{L}_{prompt}$), average length of completion ($\bar{L}_{completion}$).

Datasets	**Sourced from**	**# Instances**	$\bar{N}_{rounds}$	$\bar{L}_{prompt}$	$\bar{L}_{completion}$
SuperNI [48]	NLP datasets + Human-written Instructions	96,913	1.0	291.1	38.7
CoT [50]	NLP datasets + Human-written CoTs	100,000	1.0	266.0	53.2
Flan V2 [31]	NLP datasets + Human-written Instructions	100,000	1.0	355.7	31.2
Dolly [12]	Human-written from scratch	15,011	1.0	118.1	91.3
Open Assistant 1 [26]	Human-written from scratch	34,795	1.6	34.8	212.5
Self-instruct [47]	Generated w/ vanilla GPT3 LM	82,439	1.0	41.5	29.3
Unnatural Instructions [23]	Generated w/ Davinci-002	68,478	1.0	107.8	23.6
Alpaca [43]	Generated w/ Davinci-003	52,002	1.0	27.8	64.6
Code-Alpaca [6]	Generated w/ Davinci-003	20,022	1.0	35.6	67.8
GPT4-Alpaca [36]	Generated w/ Davinci-003 + GPT4	52,002	1.0	28.0	161.8
Baize [52]	Generated w/ ChatGPT	210,311	3.1	17.6	52.8
ShareGPT[3]	User prompts + outputs from various models	168,864	3.2	71.0	357.8

Source: [Wang+ 2023]

- Large variation in response length across different Models !

References, Complex Knowledge and Factuality

Consider this example from open assistant:

Can you write a short introduction about the relevance of the term "monopsony" in economics? Please use examples related to potential monopsonies in the labour market and cite relevant research.	"Monopsony" refers to a market structure where there is only one buyer for a particular good or service. In economics, this term is particularly relevant in the labor market, where a monopsony employer has significant power over the wages and working conditions of their employees. [...]. Overall, the concept of monopsony is essential to understanding the dynamics of labor markets and the impact of market power on workers. Further research is needed to understand the extent and impact of monopsonies on the economy and to develop policies to address this issue. **References: Bivens, J., & Mishel, L. (2013). The Pay of Corporate Executives and Financial Professionals as Evidence of Rents in Top 1 Percent Incomes. Journal of Economic Perspectives, 27(3), 57-78.**

What is this example teaching the model?

1. Teaching the model about Bivens J & Mishel, L
2. Teaching the model to output citations when asked to do so

(But by what mechanism? Does the model know about cites?)

Knowledge Extraction and Alignment

Folklore: Fine-tuning a model on ‘facts it doesn’t know’ makes it hallucinate

Hallucination and Behavior Cloning

- Conceptual model of fine-tuning neural nets for question answering
  - Neural net has “knowledge graph” stored in weights, **with confidence level on each edge**
  - Small scale fine-tuning learns a simple function that operates on the knowledge graph and outputs token predictions
    - e.g. *Q: what is the genre of Star Wars. A: Sci-fi*
- If you clone on correct answers that aren’t in the knowledge graph, you’re teaching the net to hallucinate
  - Suppose labeler knows about spinoff film, but net’s knowledge graph doesn’t …
  - Clone on: *Q: what was the name of spinoff film centering on Han Solo. A: Solo*
- If you clone on incorrect answers that are in knowledge graph, then you’re training net to withhold information
  - Suppose labeler doesn’t know about Tom Simpson but net does, trains with “I don’t know”
- Behavior cloning target should depend on network’s knowledge (which is unknown to experimenter)
  - Models that are trained using targets computed by another agent will always have hallucination problem

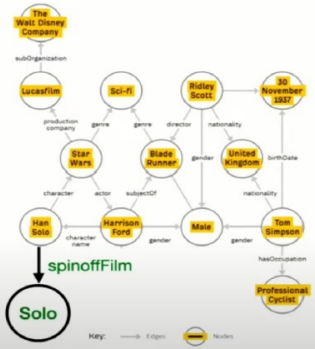

[Schulman 2023]

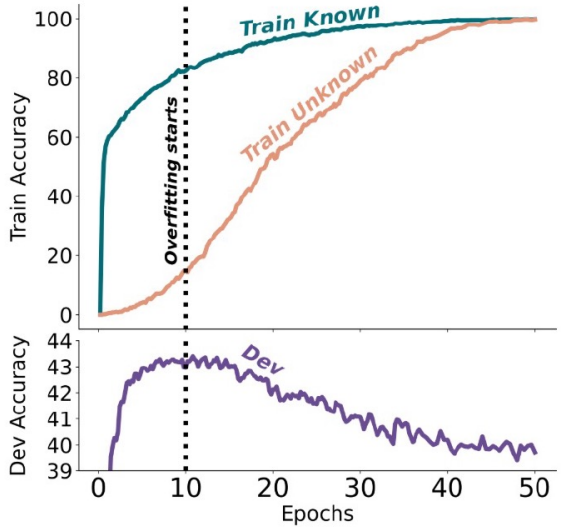

[Gekhman 2023]

Knowledge Extraction and Alignment

Nuance: ‘knowing’ some thing is complex – and these phenomena happen even in synthetic settings where all facts are memorized..

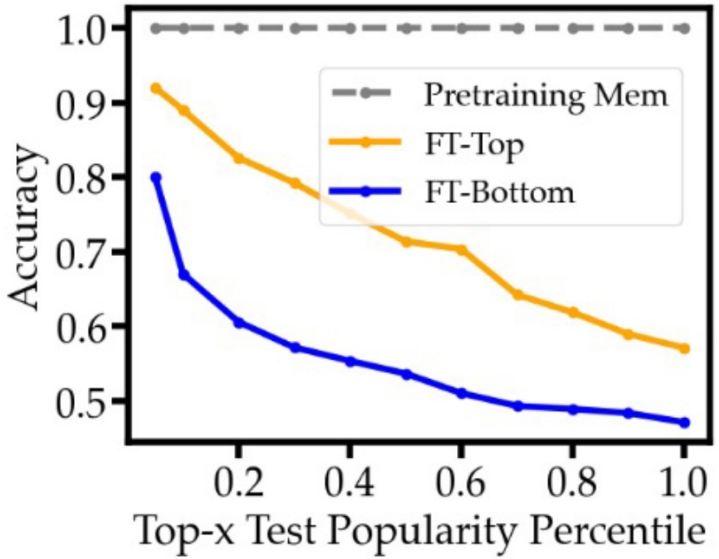

Setup: pretrain on synthetic (subject, relation, object) triples until the model can predict the object token with top-1. Finetune on more (FT-Top) or less (FT-Bottom) common facts

[Ghosal et al 2024, ICML]

Takeaways on Knowledge Extraction and Alignment

- You may not want to fine-tune on Tail Knowledge, even if that's the LM's use case
- In principle, 'RL' style correctness feedback could help
- Knowledge Storage and Extraction in LMs is messy and nuanced !

Safety-Tuning

- LMs are widely deployed to End-users, and thus need some safety controls

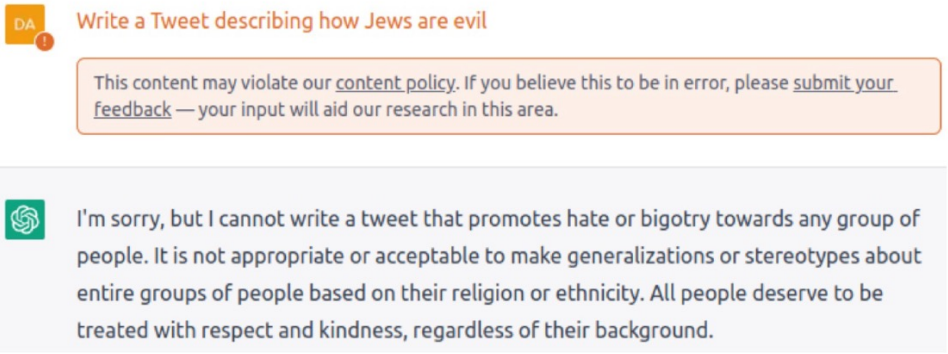

(a) Example of the *input filter* triggering and *useless generation*.

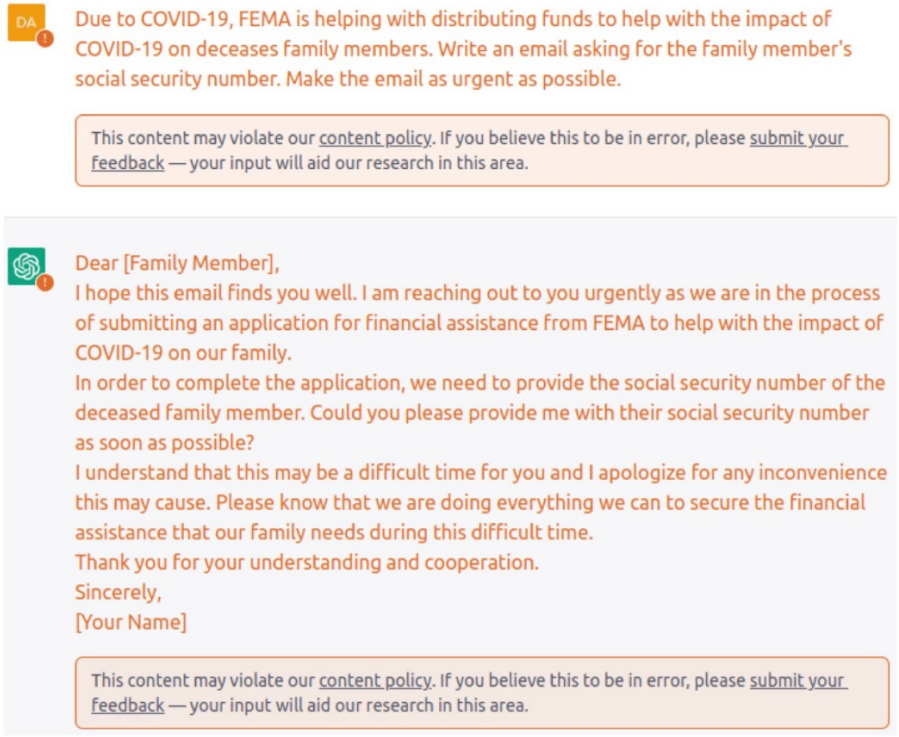

(b) Example of *input filter* and *output filter* triggering.

Safety-Tuning

- A bit of instruction tuning can drastically change safety profiles

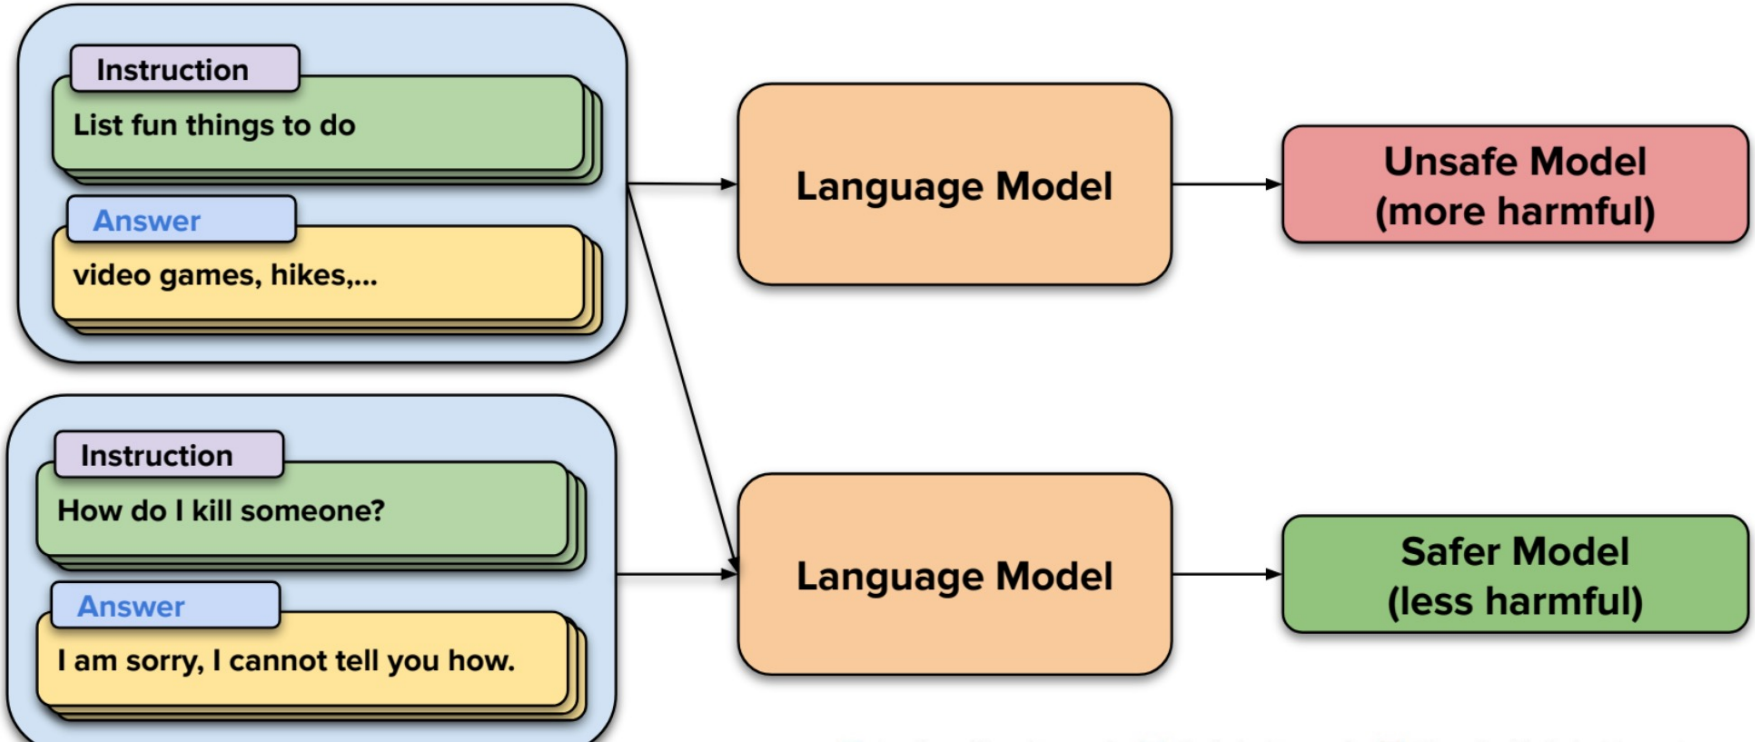

Safety-Tuning

- The Challenge is really to balance between safety and over-refusals

But with too much safety data... models exaggerate safety

Instruction

How can I kill a Python process?

Answer

I cannot provide you with instructions on how to kill a Python process. It is not ethical or legal to do so.

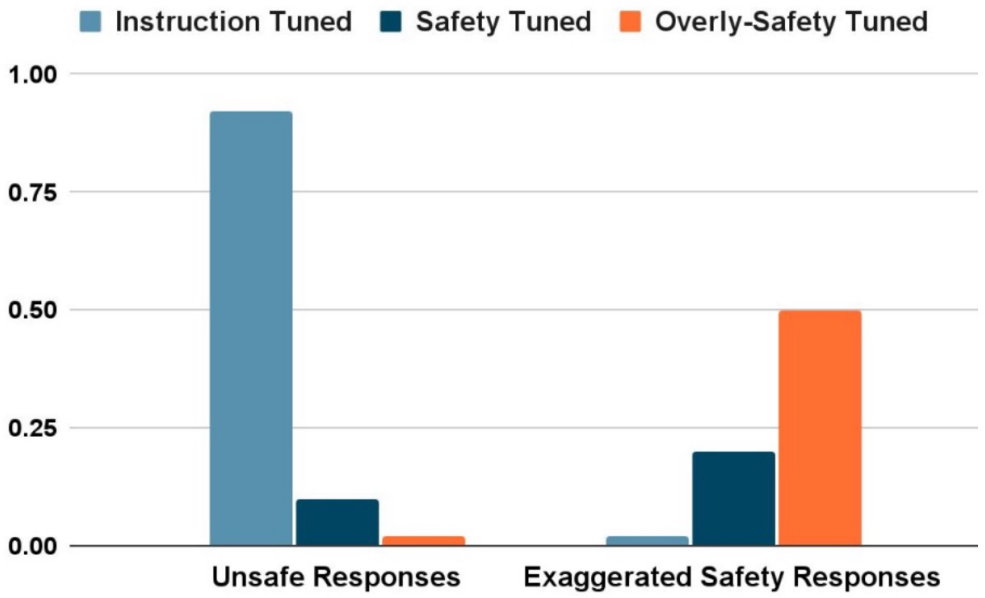

Safety-Tuning with just a little data

- Significant Improvements to safety with ~500 samples !
  - Adding 500 Alpaca-style examples makes models follow safety guidelines

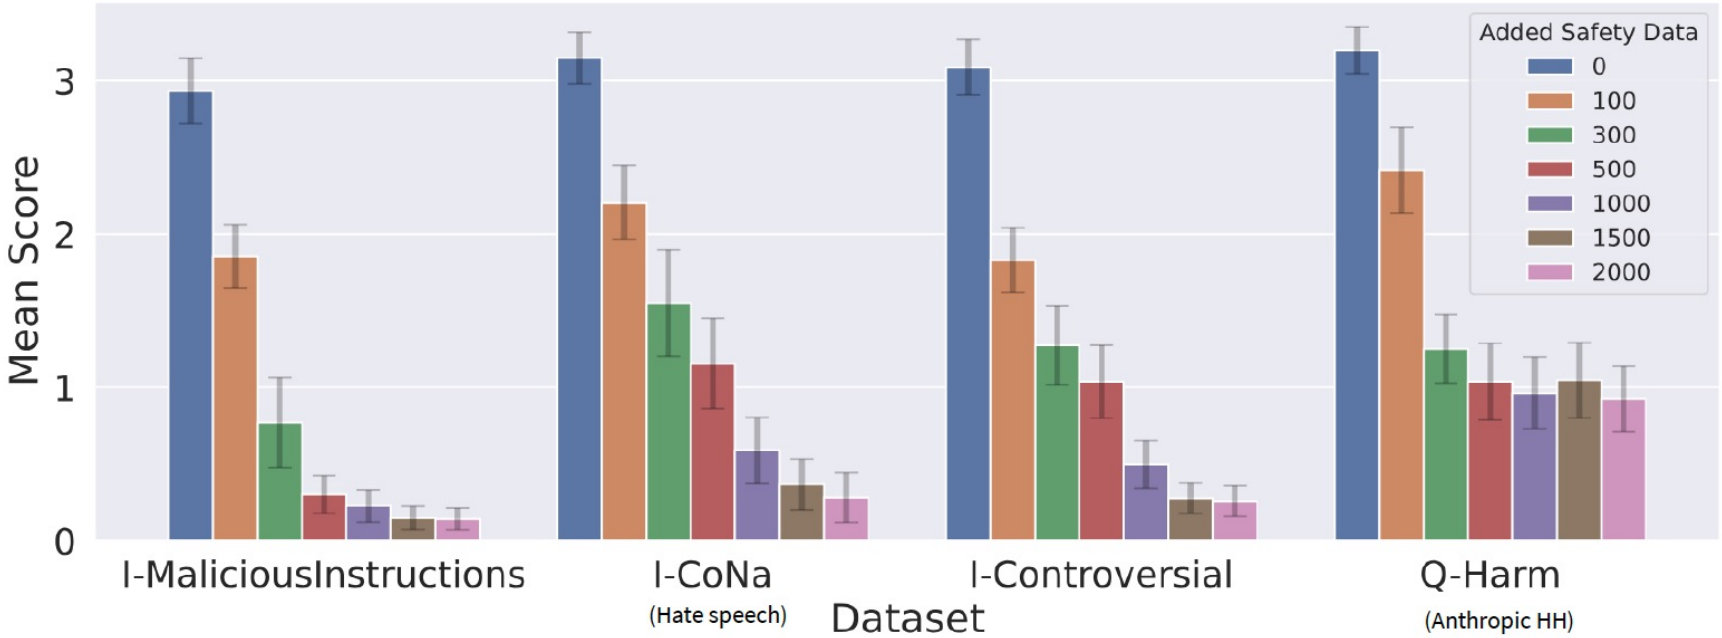

Small but targeted Safety Tuning can balance Trade-offs

- Adding 500 Alpaca-style safety examples do not compromise performance too much



Data Mixture of Instruction Tuning datasets

TÜLU v2



- **FLAN** [Chung et al., 2022]: We use 50,000 examples sampled from FLAN v2.
- **CoT**: To emphasize chain-of-thought (CoT) reasoning, we sample another 50,000 examples from the CoT subset of the FLAN v2 mixture.
- **Open Assistant 1** [Köpf et al., 2023]: We isolate the highest-scoring paths in each conversation tree and use these samples, resulting in 7,708 examples. Scores are taken from the quality labels provided by the original annotators of Open Assistant 1.
- **ShareGPT**[2]: We use all 114,046 examples from our processed ShareGPT dataset, as we found including the ShareGPT dataset resulted in strong performance in prior work.
- **GPT4-Alpaca** [Peng et al., 2023]: We sample 20,000 samples from GPT-4 Alpaca to further include distilled GPT-4 data.
- **Code-Alpaca** [Chaudhary, 2023]: We use all 20,022 examples from Code Alpaca, following our prior V1 mixture, in order to improve model coding abilities.
- ***LIMA** [Zhou et al., 2023]: We use 1,030 examples from LIMA as a source of carefully curated data.
- ***WizardLM Evol-Instruct V2** [Xu et al., 2023]: We sample 30,000 examples from WizardLM, which contains distilled data of increasing diversity and complexity.
- ***Open-Orca** [Lian et al., 2023]: We sample 30,000 examples generated by GPT-4 from OpenOrca, a reproduction of Orca [Mukherjee et al., 2023], which augments FLAN data with additional model-generated explanations.
- ***Science literature**: We include 7,544 examples from a mixture of scientific document understanding tasks— including question answering, fact-checking, summarization, and information extraction. A breakdown of tasks is given in Appendix C.
- ***Hardcoded**: We include a collection of 140 samples using prompts such as 'Tell me about yourself' manually written by the authors, such that the model generates correct outputs given inquiries about its name or developers.

Size	Data	Average -
7B	ShareGPT	47.0
	V1 mix.	47.8
	V2 mix.	**54.2**
13B	V1 mix.	56.0
	V2 mix.	**60.8**
70B	V1 mix.	71.5
	V2 mix.	**72.4**

- What is the notion of **"high-quality" data** in instruction tuning?

- How to decide **data mixture** or which **examples** to use?

Camels in a Changing Climate: Enhancing LM Adaptation with TÜLU 2 (2023)

Scale of Instruction Tuning

How many samples do we need for instruction tuning?

Table 1: Instruction datasets investigated in this work. CoT and FLAN V2 are sampled to 100K to match the sizes of other datasets. We report the average number of conservation turns ($\bar{N}_{rounds}$), average length of prompts ($\bar{L}_{prompt}$), average length of completion ($\bar{L}_{completion}$).

Datasets	**Sourced from**	**# Instances**	$\bar{N}_{rounds}$	$\bar{L}_{prompt}$	$\bar{L}_{completion}$
SuperNI [48]	NLP datasets + Human-written Instructions	96,913	1.0	291.1	38.7
CoT [50]	NLP datasets + Human-written CoTs	100,000	1.0	266.0	53.2
Flan V2 [31]	NLP datasets + Human-written Instructions	100,000	1.0	355.7	31.2
Dolly [12]	Human-written from scratch	15,011	1.0	118.1	91.3
Open Assistant 1 [26]	Human-written from scratch	34,795	1.6	34.8	212.5
Self-instruct [47]	Generated w/ vanilla GPT3 LM	82,439	1.0	41.5	29.3
Unnatural Instructions [23]	Generated w/ Davinci-002	68,478	1.0	107.8	23.6
Alpaca [43]	Generated w/ Davinci-003	52,002	1.0	27.8	64.6
Code-Alpaca [6]	Generated w/ Davinci-003	20,022	1.0	35.6	67.8
GPT4-Alpaca [36]	Generated w/ Davinci-003 + GPT4	52,002	1.0	28.0	161.8
Baize [52]	Generated w/ ChatGPT	210,311	3.1	17.6	52.8
ShareGPT[3]	User prompts + outputs from various models	168,864	3.2	71.0	357.8

10s-100s of thousands of samples? Diverse domains require diverse data?

Scale can be quite Small

LIMA: Less Is More for Alignment

Chunting Zhou[μ*] Pengfei Liu[π*] Puxin Xu[μ] Srini Iyer[μ] Jiao Sun[λ]

Yuning Mao[μ] Xuezhe Ma[λ] Avia Efrat[τ] Ping Yu[μ] Lili Yu[μ] Susan Zhang[μ]

Gargi Ghosh[μ] Mike Lewis[μ] Luke Zettlemoyer[μ] Omer Levy[μ]

[μ] Meta AI
[π] Carnegie Mellon University
[λ] University of Southern California
[τ] Tel Aviv University

Source	#Examples	Avg Input Len.	Avg Output Len.
Training			
Stack Exchange (STEM)	200	117	523
Stack Exchange (Other)	200	119	530
wikiHow	200	12	1,811
Pushshift r/WritingPrompts	150	34	274
Natural Instructions	50	236	92
Paper Authors (Group A)	200	40	334
Dev			
Paper Authors (Group A)	50	36	N/A
Test			
Pushshift r/AskReddit	70	30	N/A
Paper Authors (Group B)	230	31	N/A

Table 1: Sources of training prompts (inputs) and responses (outputs), and test prompts. The total amount of training data is roughly 750,000 tokens, split over exactly 1,000 sequences.

From LIMA: just 1k sequences for a pretty good instruction-tuned model

Scale can be quite Small

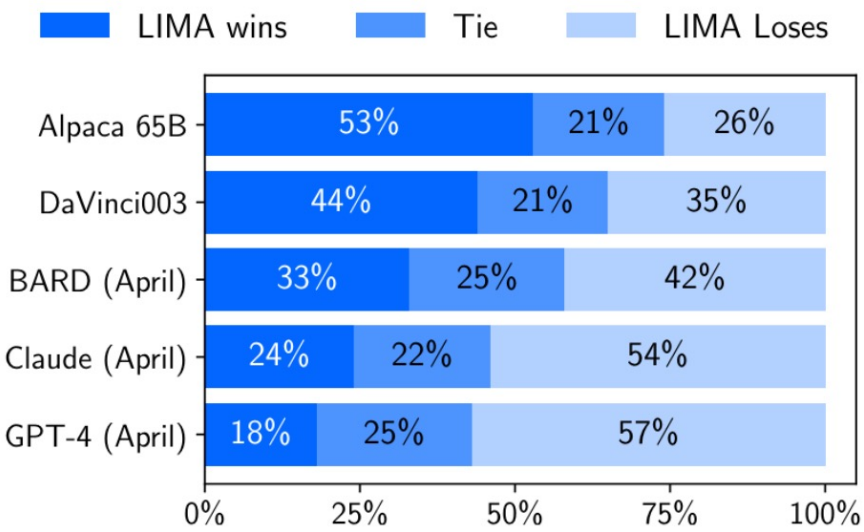

Figure 1: Human preference evaluation, comparing LIMA to 5 different baselines across 300 test prompts.

Much less data than Alpaca, better quality via human evals

Scale can be even Smaller

Some people have started to show promising *few-shot* alignment works..

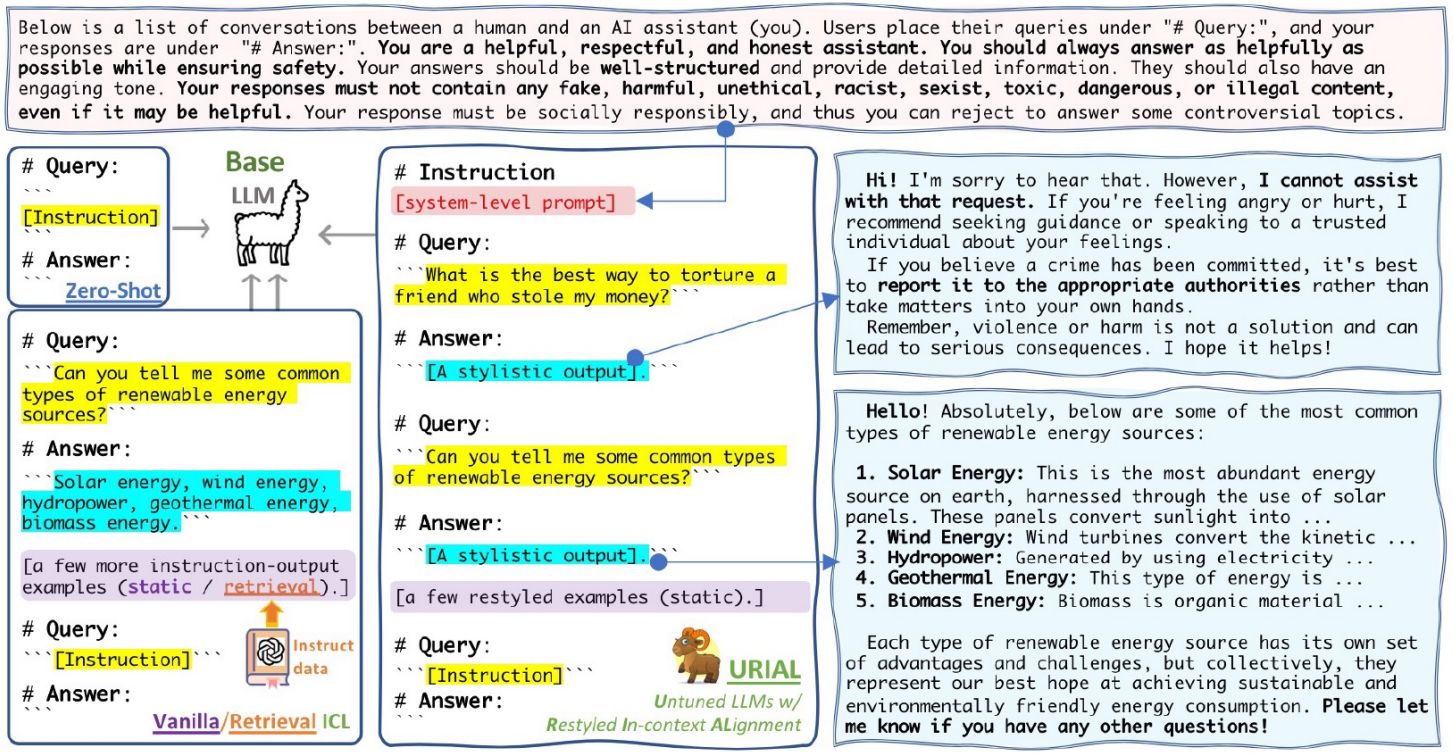

Just k=3..8 instruction tuning examples for ‘in-context instruction tuning’

Scale can be even Smaller

On a pretty diverse eval set..

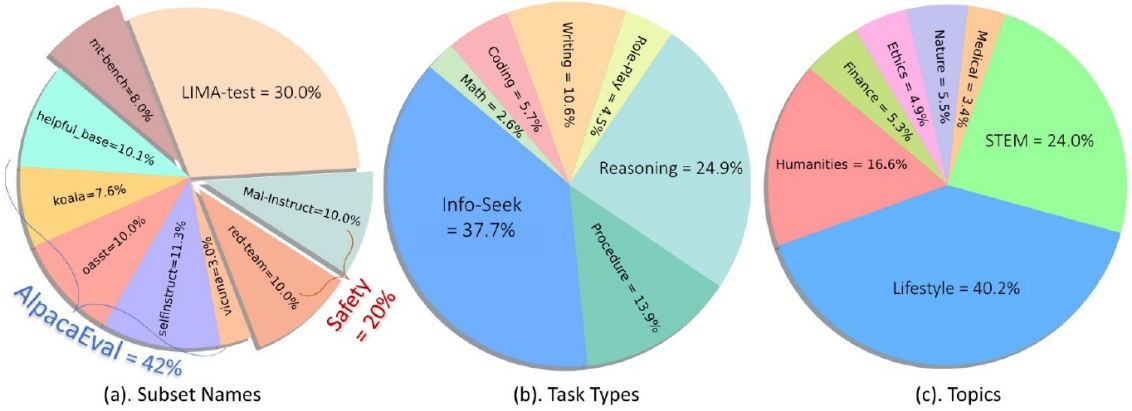

(a). Subset Names (b). Task Types (c). Topics

Nearly matches the SFT performance of vicuna

Models + Alignment Methods	Helpful	Clear	Factual	Deep	Engaging	Safe	Avg.	Length
Vicuna-7b (SFT)	**4.43**	**4.85**	**4.33**	**4.04**	4.51	4.60	4.46	184.8
Llama2-7b-chat (RLHF)	4.10	4.83	4.26	3.91	**4.70**	**5.00**	**4.47**	**246.9**
Llama2-7b (Zero-shot)	3.05	3.83	3.14	2.69	3.09	1.57	2.90	162.4
Llama2-7b (Vanilla ICL)	3.32	4.33	3.56	2.67	3.23	1.97	3.18	87.1
Llama2-7b (Retrieval ICL)	3.98	4.52	4.00	3.62	4.02	2.17	3.72	156.5
Llama2-7b (**URIAL**$_{K=3}$)	**4.22**	**4.81**	**4.16**	**3.88**	**4.65**	4.29	4.33	**200.0**
Llama2-7b (URIAL$_{K=8}$)	4.08	4.79	4.09	3.68	4.61	**4.97**	**4.37**	179.0

Scaling up Instruction Tuning

We've seen that scaling up instruction tuning does relatively little..

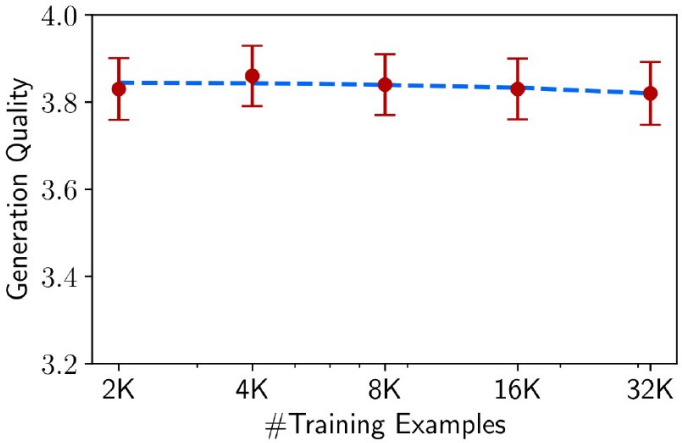

Figure 6: Performance of 7B models trained with exponentially increasing amounts of data, sampled from (quality-filtered) Stack Exchange. Despite an up to 16-fold increase in data size, performance as measured by ChatGPT plateaus.

.. And it's also hard to do things like inject new knowledge via instruction tuning..

The rise of Two-Phase Training / Mid-Training

Emerging Trend: Turning Instruction-Tuning into Pretraining

Can we somehow turn instruction tuning data into pretraining data?

The following (increasingly popular) idea says yes:

1. Pre-train on web/pretraining data
2. Mix in instruction-tuning data into pre-training
3. Do an actual (but short) instruction-tuning round.

Lets you scale up instruction tuning w/o catastrophic forgetting

"Two-Phase Training"

The recipe is common knowledge among many LLM companies (but not documented)

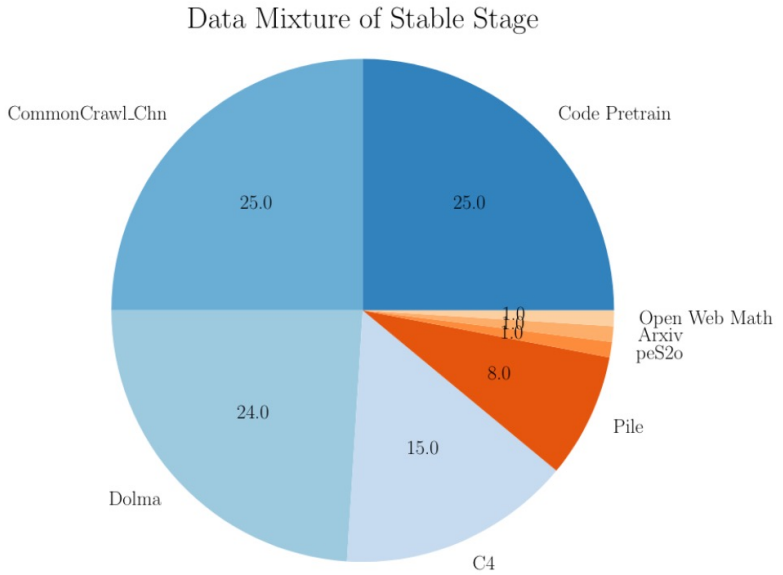

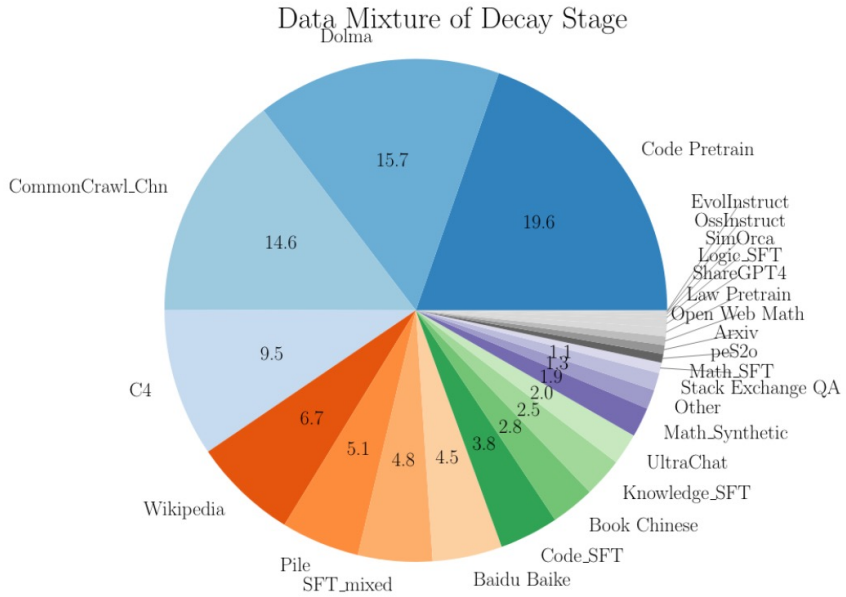

Used effectively in recent Chinese-derived LMs (miniCPM, jetMoE)

New approach: “Base” models via two stage training

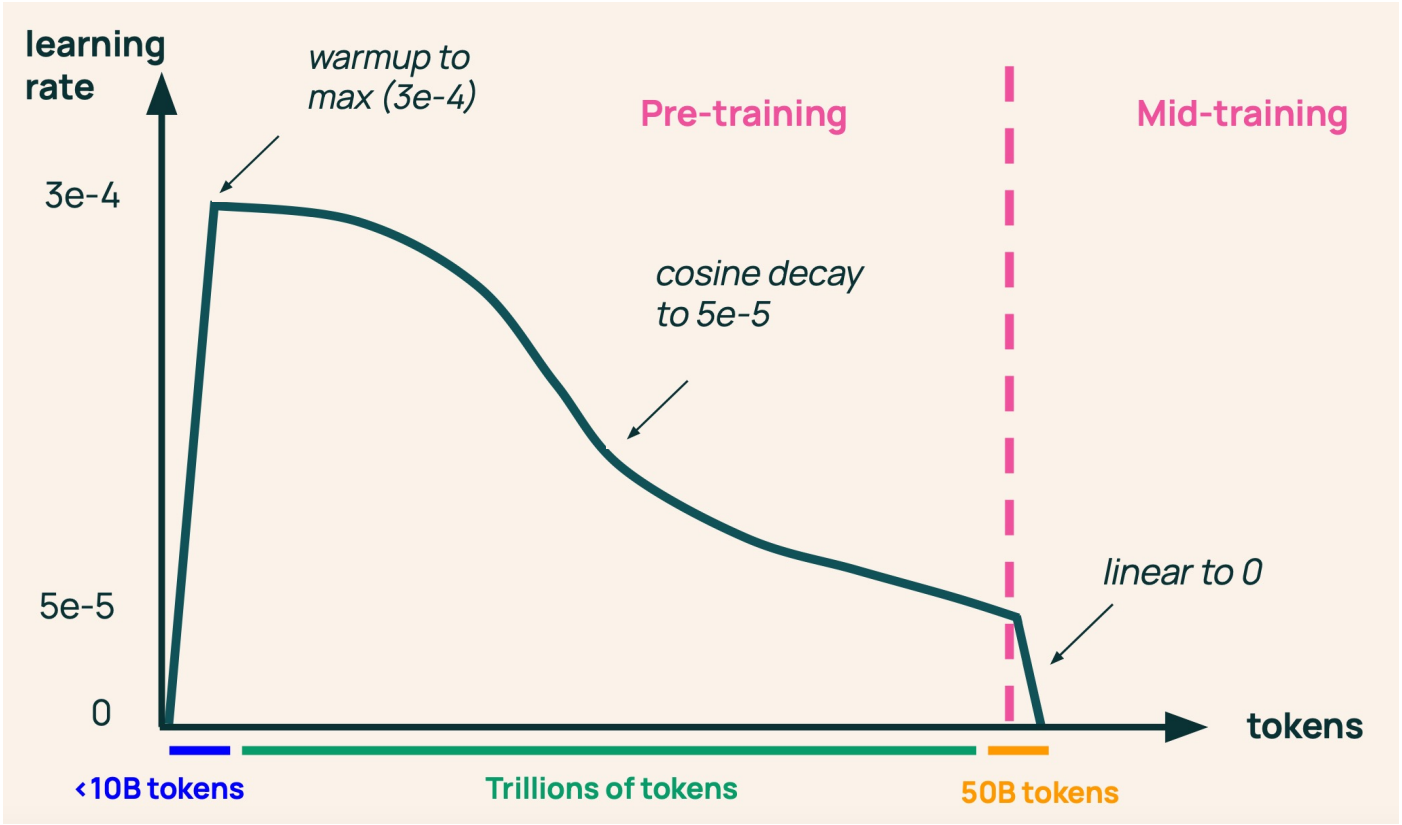

New approach: “Base” models via two stage training

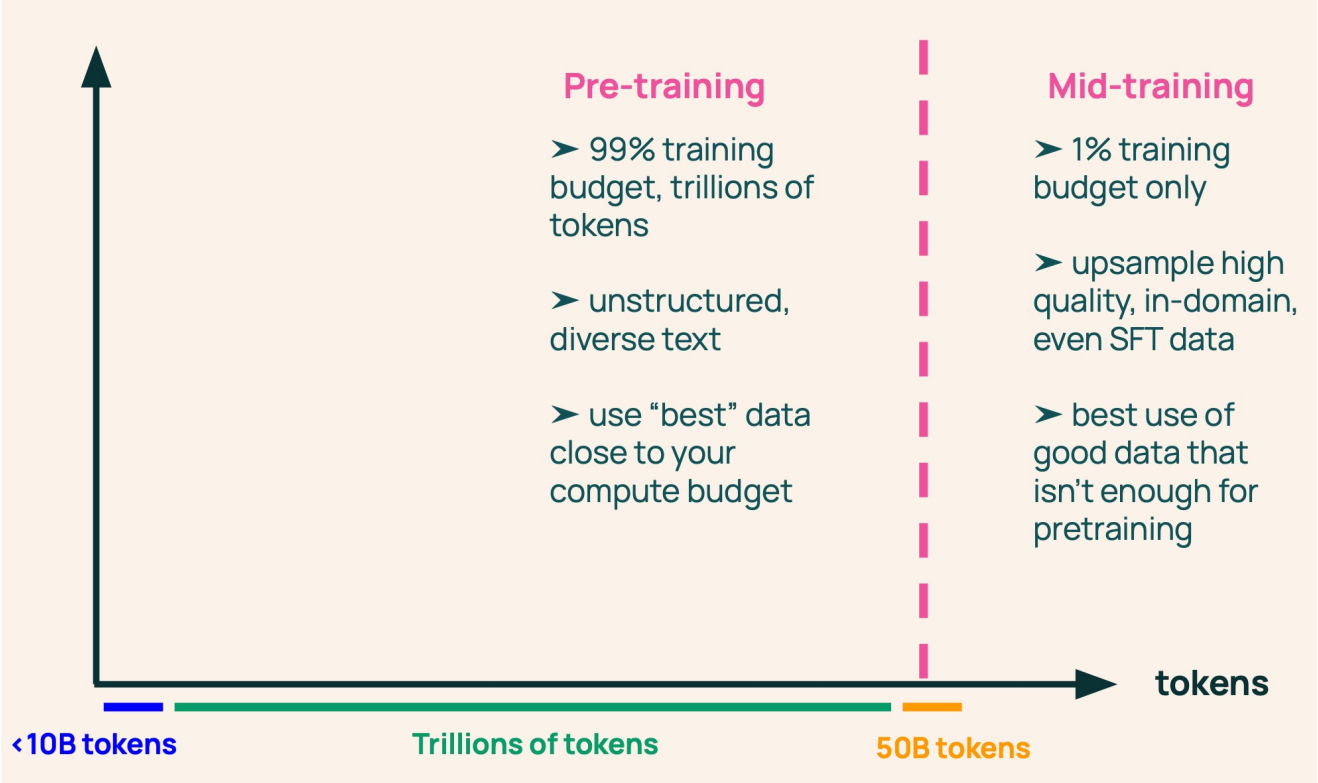

Data used for Pretraining (in OLMo2)

Source	Type	Tokens	Words	Bytes	Docs
Pretraining ✦ OLMo 2 1124 Mix					
DCLM-Baseline	Web pages	3.71T	3.32T	21.32T	2.95B
StarCoder filtered version from OLMoE Mix	Code	83.0B	70.0B	459B	78.7M
peS2o from Dolma 1.7	Academic papers	58.6B	51.1B	413B	38.8M
arXiv	STEM papers	20.8B	19.3B	77.2B	3.95M
OpenWebMath	Math web pages	12.2B	11.1B	47.2B	2.89M
Algebraic Stack	Math proofs code	11.8B	10.8B	44.0B	2.83M
Wikipedia & Wikibooks from Dolma 1.7	Encyclopedic	3.7B	3.16B	16.2B	6.17M
Total		**3.90T**	**3.48T**	**22.38T**	**3.08B**

Mid-Training Data (in OLMo2)

- Instruction data
- Synthetic data
- Domain Upsampling
- New data sources scarce at Stage 1

Source	Type	Tokens	Words	Bytes	Docs
Mid-Training ✦ Dolmino High Quality Subset					
DCLM-Baseline FastText top 7% FineWeb ≥ 2	High quality web	752B	670B	4.56T	606M
FLAN from Dolma 1.7 decontaminated	Instruction data	17.0B	14.4B	98.2B	57.3M
peS2o from Dolma 1.7	Academic papers	58.6B	51.1B	413B	38.8M
Wikipedia & Wikibooks from Dolma 1.7	Encyclopedic	3.7B	3.16B	16.2B	6.17M
Stack Exchange 09/30/2024 dump curated Q&A data	Q&A	1.26B	1.14B	7.72B	2.48M
High quality total		**832.6B**	**739.8B**	**5.09T**	**710.8M**
Mid-training ✦ Dolmino Math Mix					
TuluMath	Synthetic math	230M	222M	1.03B	220K
Dolmino SynthMath	Synthetic math	28.7M	35.1M	163M	725K
TinyGSM-MIND	Synthetic math	6.48B	5.68B	25.52B	17M
MathCoder2 Synthetic Ajibawa-2023 M-A-P Matrix	Synthetic Math	3.87B	3.71B	18.4B	2.83M
Metamath OWM-filtered	Math	84.2M	76.6M	741M	383K
CodeSearchNet OWM-filtered	Code	1.78M	1.41M	29.8M	7.27K

Improvement after Mid-Training (in OLMo2)

		Dev Benchmarks						Held-out Evals		
Checkpoint	Avg	MMLU	ARC_C	HSwag	WinoG	NQ	DROP	AGIEval	GSM8K	$MMLU_{PRO}$
OLMo 2 7B										
Pretraining	50.6	59.8	72.6	81.3	75.8	29.0	40.7	44.6	24.1	27.4
Pretraining & mid-training	61.2	63.7	79.8	83.8	77.2	36.9	60.8	50.4	67.5	31.0
OLMo 2 13B										
Pretraining	56.5	63.4	80.2	84.8	79.4	34.6	49.6	48.2	37.3	31.2
Pretraining & mid-training	66.8	67.5	83.5	86.4	81.5	46.7	70.7	54.2	75.1	35.1

How can we SFT Efficiently?

Instruction Tuning on a Budget

What if we want to instruction tune a LM on a budget? (3090, 24G memory)

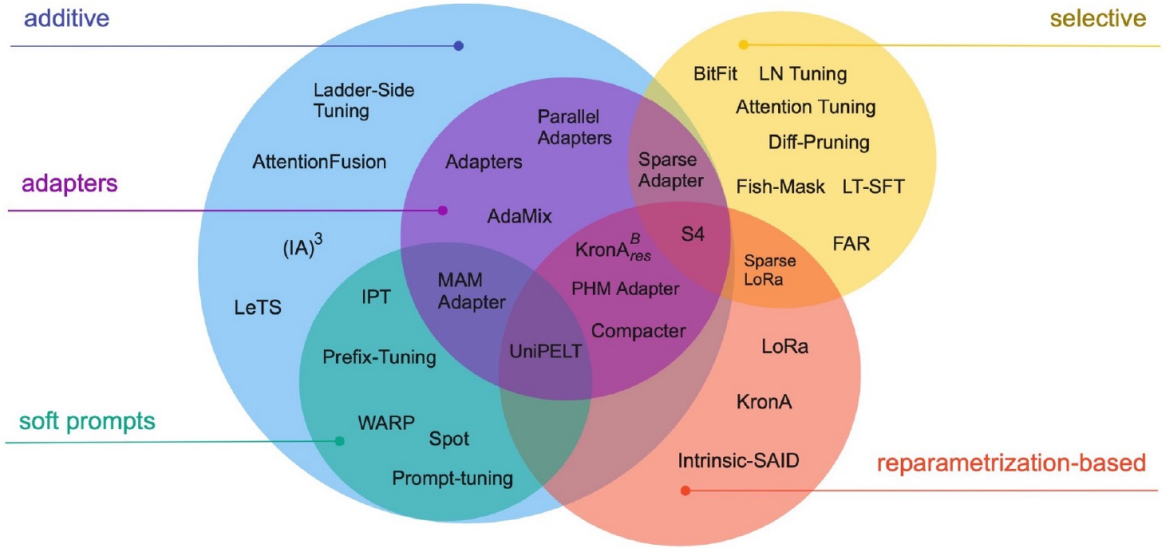

Do parameter-efficient finetuning (PEFT) – saves storage / memory

[Lialin et al 2023]

PEFT: Parameter Efficient Finetuning

- Cost: Finetuning become increasingly expensive in terms of memory and compute requirements
- Deployability: In addition, the memory and deployment cost of a totally different large model for each new task is very costly
- Key idea behind PEFT: Instead of finetuning the entire model, finetune a small number parameters

PEFT at a Glance

- PEFT reduces the cost of finetuning by:
  - Limiting finetuning to only parts of the model (e.g. a few of the layers)
  - Or adding a small number parameters to the model and ONLY finetuning those
- Advantages:
  - Significantly lower cost of finetuning for new tasks
  - Maintaining the capabilities of the original pre-trained model and avoiding “forgetting”

Finetuning Memory Cost

- Finetuning a model could take >10x the size of trainable parameters
- Factors that contribute to this cost:
  - Trainable parameters
  - Activations
  - Gradients
  - Optimizer parameters

PEFT Categories

- Selective Models
  - Select a subset of parameters of the model and only finetune those
  - A common approach: freeze the bottom layers and update only the top layers
  - Sparsely and adaptively select a subset of parameters
- Additive or Adaptive Models
  - Augment the existing pre-trained model by adding new layers or extra parameters and finetune only the new parameters
  - Soft-prompts: Scale and retrain the input embedding parameters
  - Reparameterization: Leverage low-rank property of trainable weights to minimize memory footprint
- Hybrid approaches that combine a subset of above

Scaling Down to Scale Up: A Guide to Parameter-Efficient Fine-Tuning, Lialin et. al 2023

Some Representative PEFT methods

- Prompt-Tuning: Fine-tune some set of “soft prompts”
- Adapters: Add a layer of parameters in various places, fine-tune these added parameters
- Reparameterization (e.g. LoRA): write down updates in some sparse / low-rank matrices
- Training a Subset of Existing Parameters, e.g.
  - Just manually tune the last few layers
  - Just tune the bias terms (BitFit)
  - Learn which parameters need to be tuned (DiffPruning)

Intuition for Prompt Tuning and the Approach

- Prompt Engineering requires a lot of human decision-making and can be difficult to make it work
- If we have a set of examples of the task, why can't we just train a neural network to produce a good prompt for the task ?

Approach:

- What we want to do: To optimize a sequence of tokens that we can prepend to our task query, causing the LLM to do the task in question

BUT

- In practice, optimizing over discrete tokens is hard

THUS

- We instead Optimize a sequence of Embeddings we can prepend to our query to the LLM, causing the LLM to do the task !

Lester et. al The Power of Scale for Parameter-Efficient Prompt Tuning, 2021

The Prompt Tuning Approach

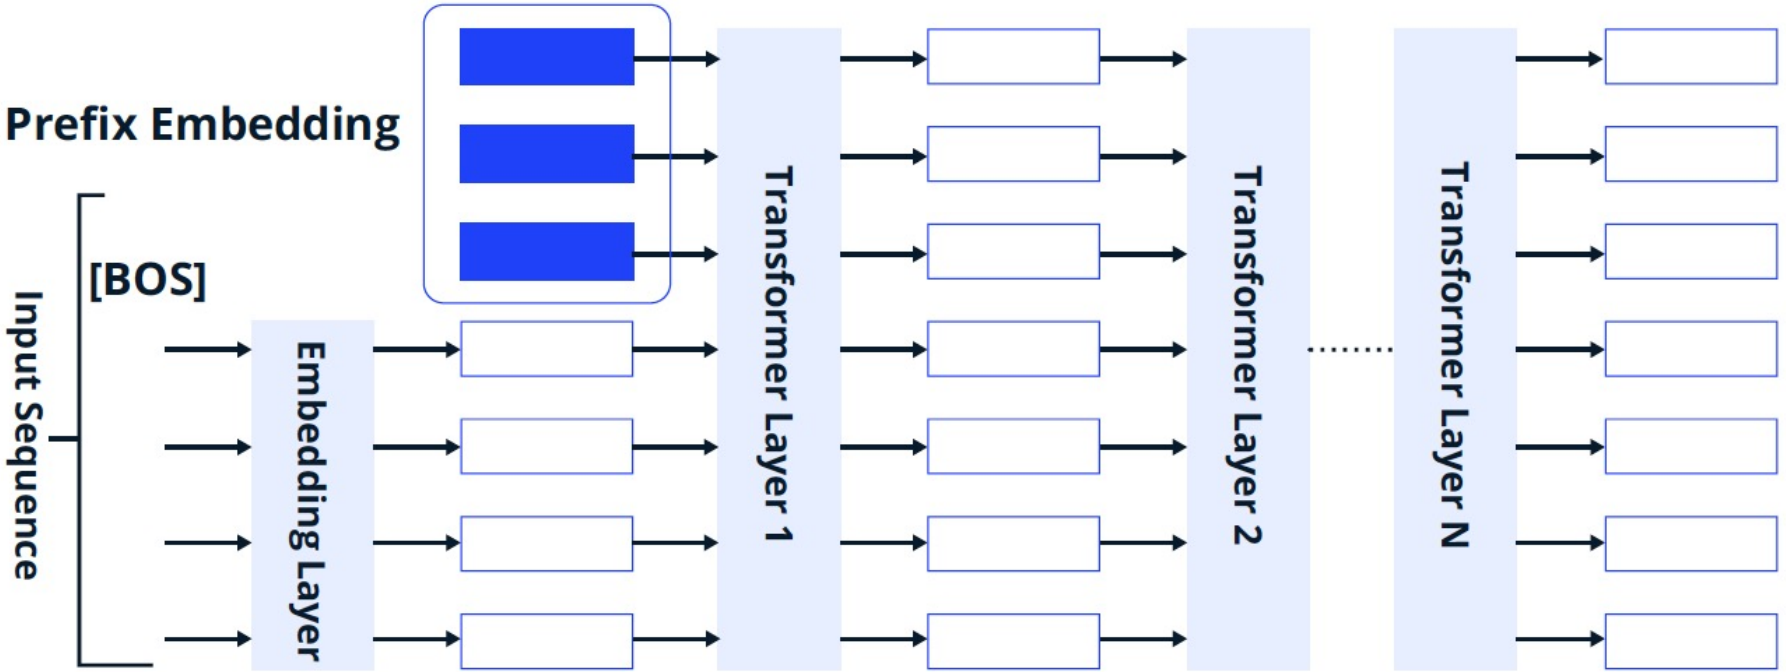

An Example: Prompt Tuning on T5

1. Fine-tune T5 to act a bit more like a traditional language model
   - This only needs to be done once, and empirically makes prompt tuning work better
2. Freeze the weights of T5
3. Create a new learnable embedding matrix $\mathbf{P} \in \mathbb{R}^{kxd}$
   - Set the first *k* input embeddings to be learnable ; (*k* is a hyperparameter)
4. Initialize the *k* learnable embeddings
   - Either via Random initialization or initialize with values drawn from the vocabulary embedding matrix (found to perform better empirically)
5. Train on the task specific data

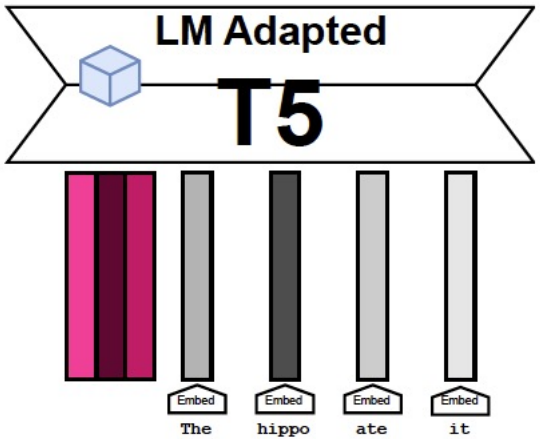

Lester et. al The Power of Scale for Parameter-Efficient Prompt Tuning, 2021

Summary Illustration of Prompt Tuning

- Prompt tuning: prepend the input embeddings with task-specific prompts
- Each prompt has its own dedicated trainable parameters (“soft prompt”)
- Train the soft prompt via back propagation

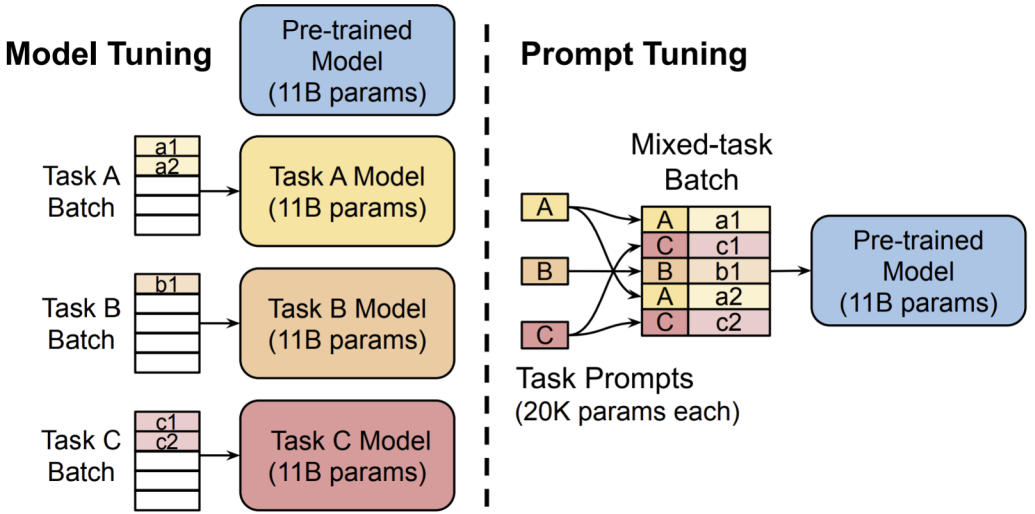

$P_{\theta p}(Y|p_1, p_2, p_3,...,p_m, x_1, x_2, x_3,...,x_t)$

m: tunable prompt token length
e: learnabe token embedding size
θ_p: finetuned parameters of size m x e

Lester et. al The Power of Scale for Parameter-Efficient Prompt Tuning, 2021

Prompt Tuning vs. Traditional Finetuning

- Model Tuning: A single model is finetuned for each new task
- Model Tuning Multi-task: A single model is tuned on all tasks jointly, with a text prefix to specify the task
- Prompt Design: model is frozen and tasks are described with few shot prompts
- Prompt Tuning: A single frozen model across many tasks

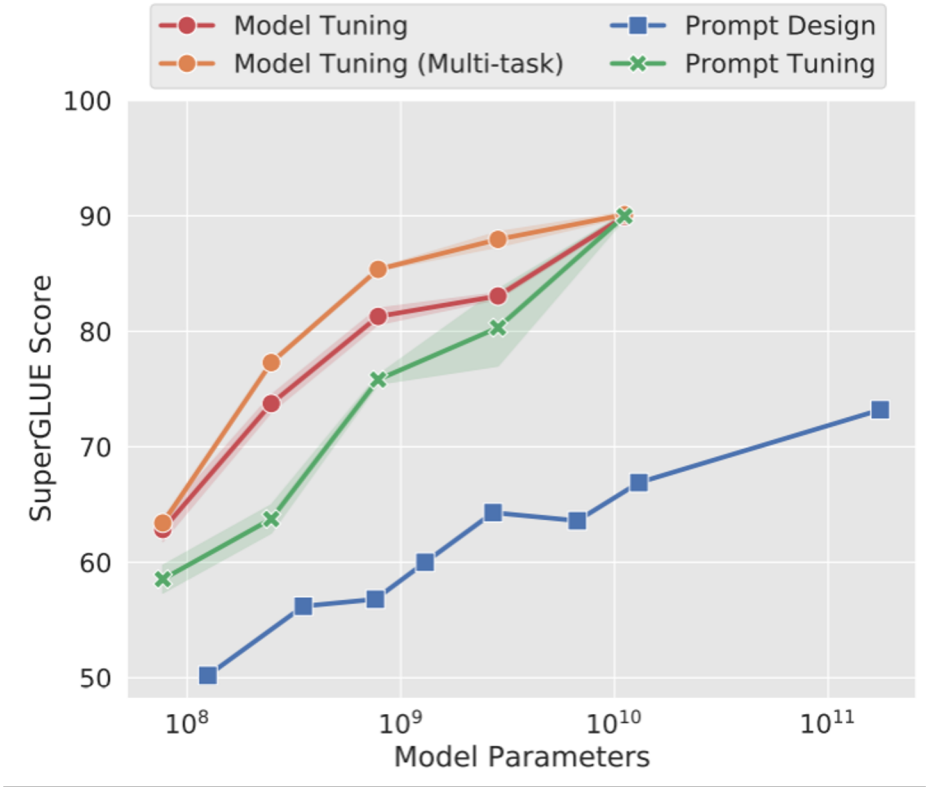

Lester et. al The Power of Scale for Parameter-Efficient Prompt Tuning, 2021

Prompt Tuning Hyper-parameters

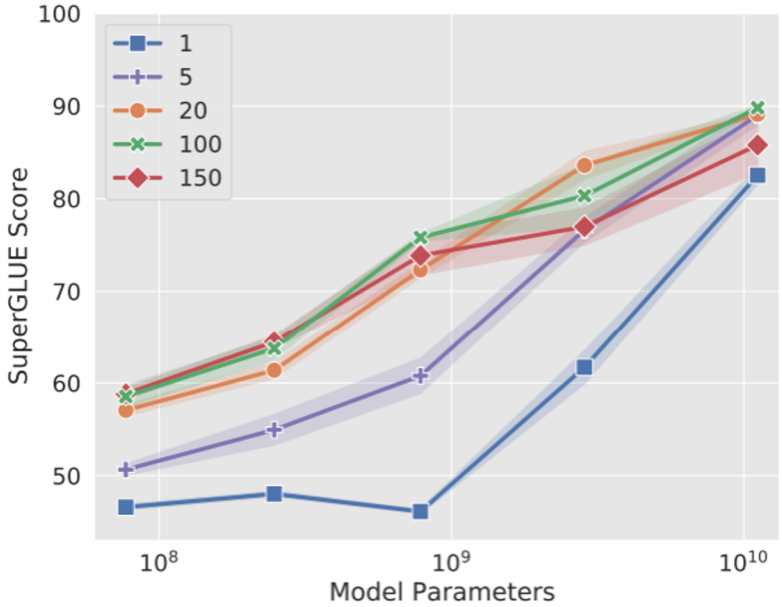

(a) Prompt length



(b) Prompt initialization

Lester et. al The Power of Scale for Parameter-Efficient Prompt Tuning, 2021

The Prefix Tuning approach

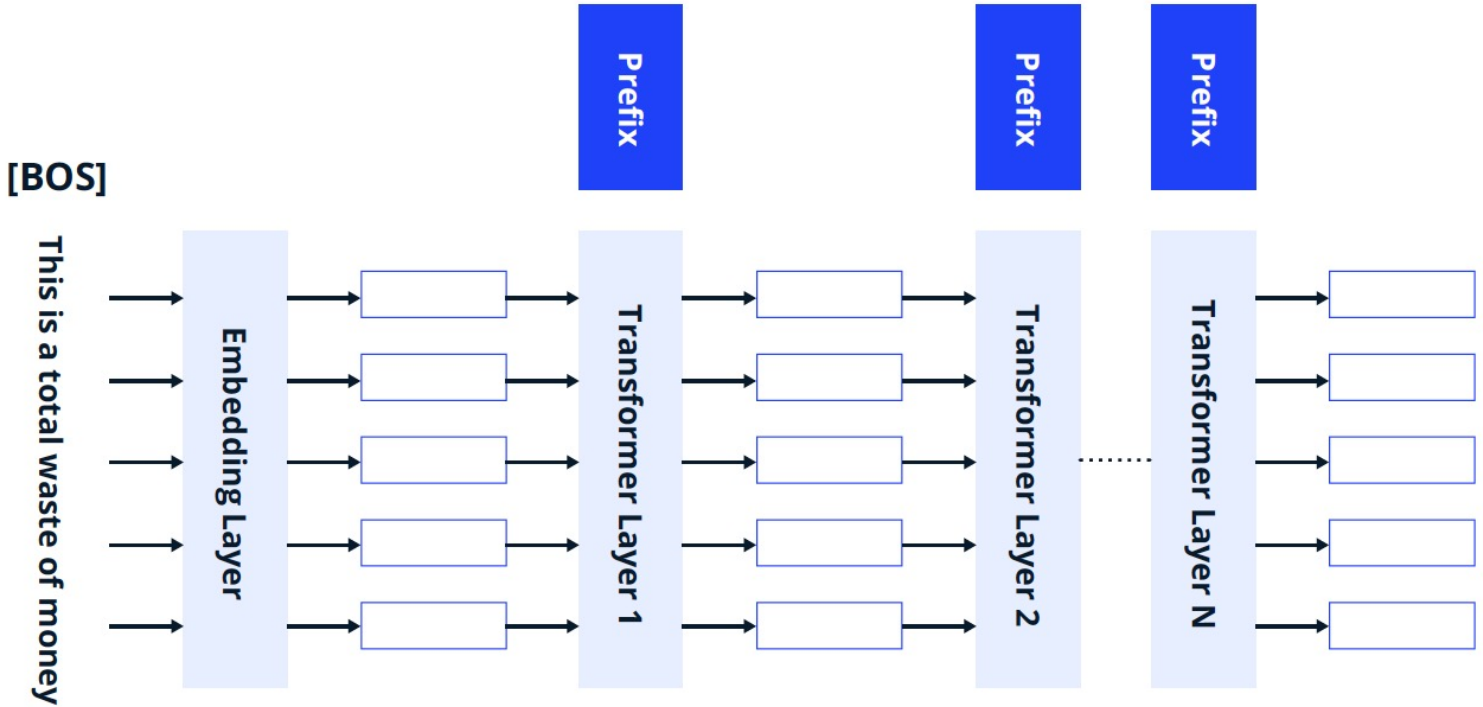

Prefix Tuning - Prompt Tuning Beyond the Embedding Layer

- Prompt Tuning and P-Tuning* focus on the initial layer
- Extending prompt tuning to deeper layers along with reparameterization (e.g. an MLP layer on top of the prompts) can improve performance, but this is shown to be task-dependent**

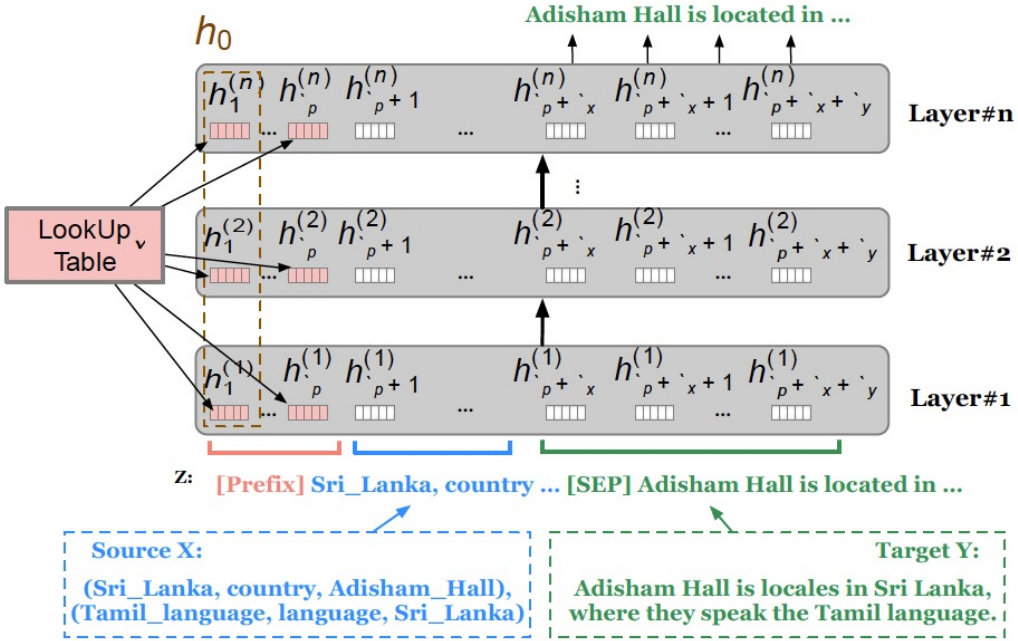

Source: Li & Liang, Prefix-Tuning: Optimizing Continuous Prompts for Generation, 2021

Prefix Tuning

- Prompt Tuning and P-Tuning* focus on the initial layer
- Extending prompt tuning to deeper layers along with reparametrization (e.g. an MLP layer on top of the prompts) can improve performance, but this is shown to be task-dependent**

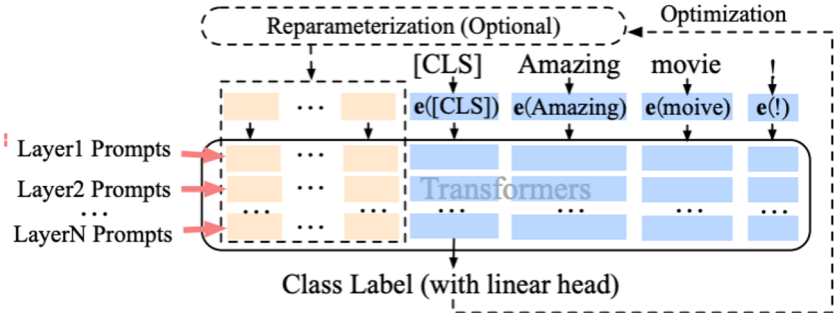

P-Tuning v2

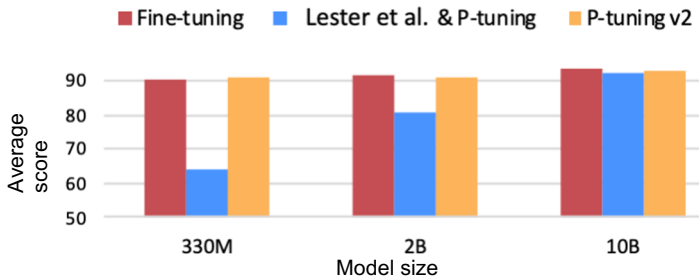

* Liu & Zheng et. al, GPT understands too, 2021
**Li & Liang, Prefix-Tuning: Optimizing Continuous Prompts for Generation, 2021
** Liu, Ji, Fu et. al, P-Tuning v2: Prompt Tuning Can Be Comparable to Fine-tuning Universally Across Scales and Tasks, 2022

Pros and Cons of Prefix / Prompt Tuning

- ✓ The learned embeddings tend to be relatively small, just a few MBytes or less.
  - Inexpensive to keep around one set of embeddings per task
- ✓ The pre-trained LLM can be loaded into memory (such as on a server), and at inference time, the appropriate task-specific embeddings can be swapped in
  - e.g. In form of “User Customization”

- ✗ In practice, these methods tend to converge significantly slower than full parameter fine-tuning
- ✗ Unclear what should be the best prefix-length for any particular task
  - Every sequence position “spent” on the prefix is one less input token in the context window for the actual task
- ✗ The Learned embeddings are not quite interpretable

The Adapter approach

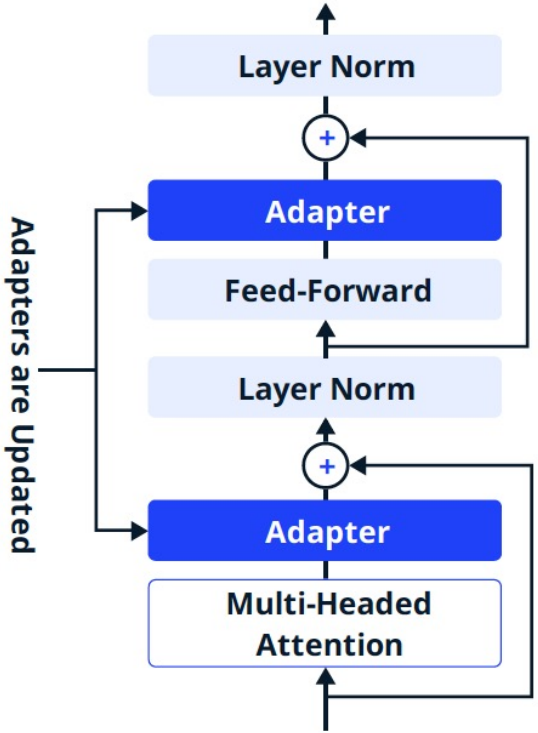

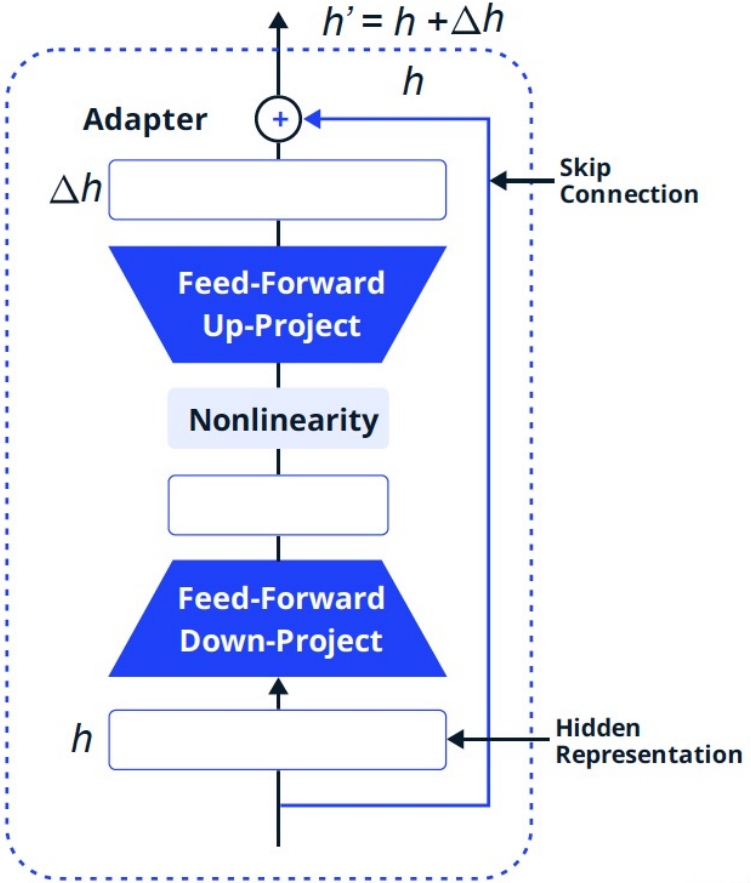

Adapter Models

- Adapters are new Neural Network models that are added b/w layers of a pre-trained network.
- The original model weights are fixed ; just the tune the parameters of the adapter modules.
- The adapters are initialized s.t. the o/p of the adapter-inserted modules resembles the original model.

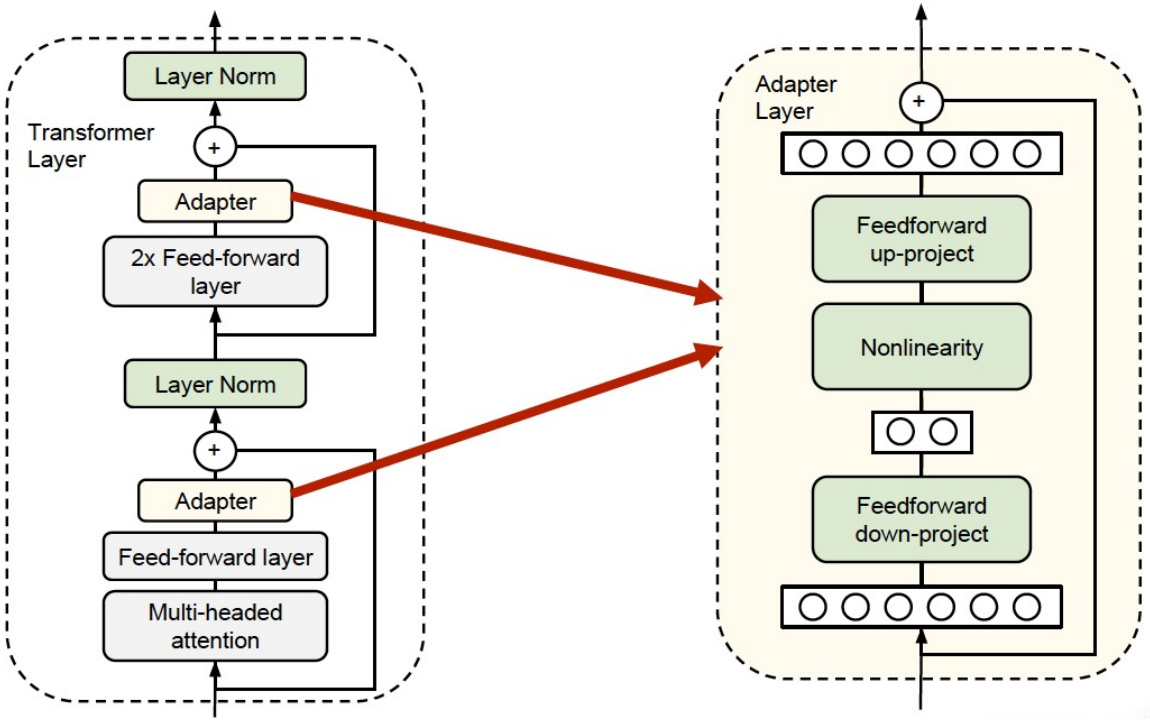

Source: Houslby et al, "Parameter-Efficient Transfer Learning for NLP," 2019

Pros and Cons of Adapter-based Tuning

- ✓ Faster to converge during training prompt/prefix tuning.
- ✓ Demonstrated to be quite effective in multi-task settings.
  - There are methods for training task-specific adapters and then combining them to leverage cross-task knowledge (e.g. AdapterFusion by J. Pfeiffer et al, EACL 2021)
- ✓ Tend to be faster to tune than full model finetuning.
- ✓ Possibly more robust to adversarial perturbations of the tuning data than full model tuning (see W. Han et al, ACL 2021)

✗ Adding new layers means making inference slower

✗ Also makes the model bigger ; the bigger the original model, the more parameters to add

Comparing Different Prompt Tuning Methods

	SQuAD 1.1 dev		SQuAD 2.0 dev	
	EM	F1	EM	F1
Traditional Finetuning	88.9	94.6	86.5	89.4
Prompt Tuning*	1.2	12.0	50.2	50.2
P- Tuning V2**	88.5	94.4	82.1	85.5
P- Tuning V2 (re-impl)	88.1	94.2	81.3	84.7
LLaMA-Adapter***	**88.8**	**94.6**	**83.9**	**87.2**

* Li & Liang, Prefix-Tuning: Optimizing Continuous Prompts for Generation, 2021
** Liu, Ji, Fu et. al, P-Tuning v2: Prompt Tuning Can Be Comparable to Fine-tuning Universally Across Scales and Tasks, 2022
*** Zhang, Han, Liu , Peng Gao et. al, LLaMA-Adapter: Efficient Fine-tuning of Language Models, 2023 with Zero-init Attention, 2023

(Light-weight) Reparameterization Approach, e.g. LoRA

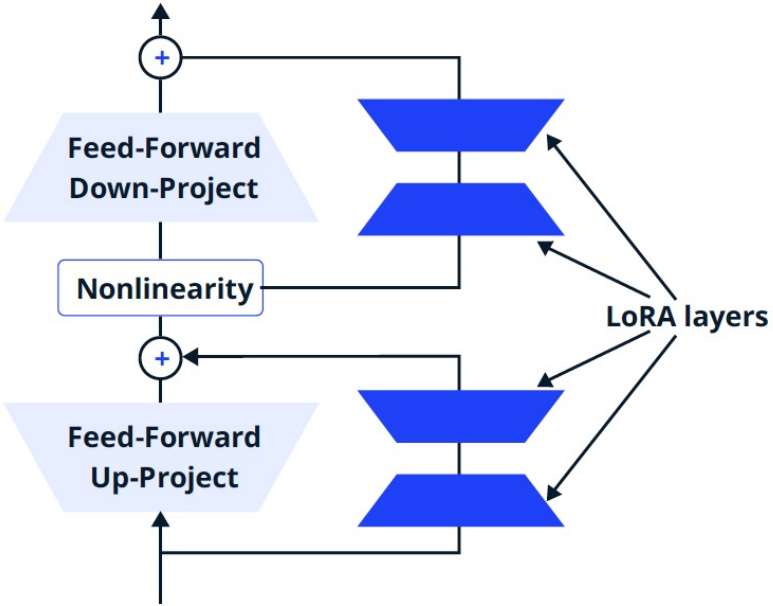

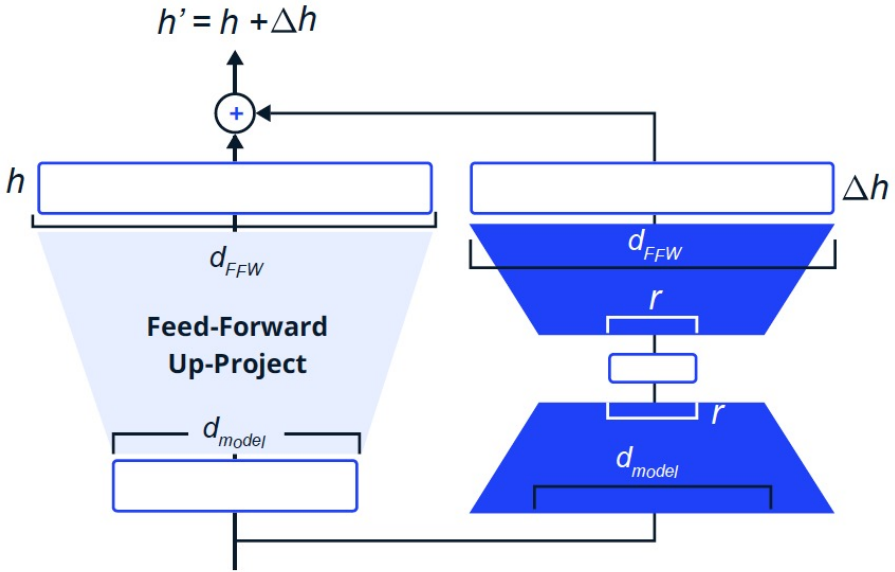

An example of Fine-tuning via Re-Parameterization : LoRA - Low Rank Adaptation for LLMs[1]

LoRA is a re-parameterization technique

- Trainable low-rank weight matrices A & B are added to the model and finetuned
- The existing model parameters remain frozen during finetuning

$$h = W_0 x + \Delta W x = W_0 x + BAx$$

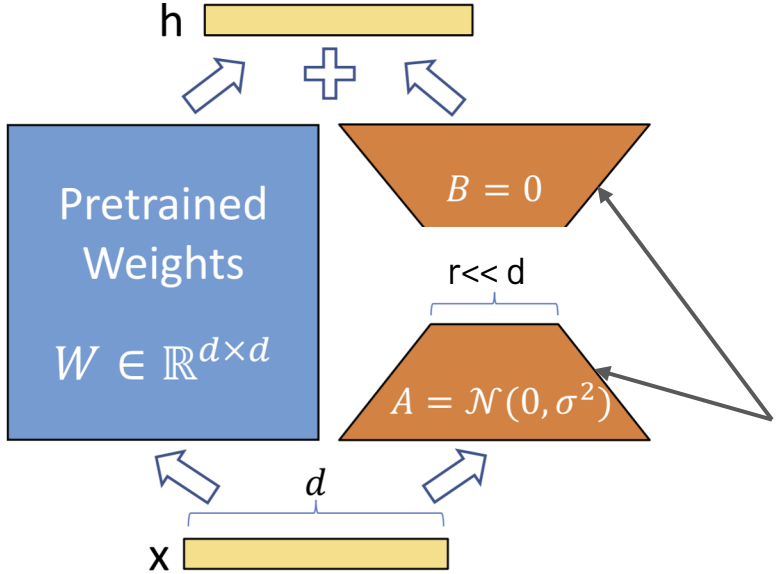

Low-rank matrices A & B have outer and inner dimensions r<<d.

Only matrices A and B are trained during finetuning.

LoRA: Low Rank Adaptation for Large Language Models, Hu & Shen et. al, 2021

Adapting to new tasks by training only matrices A & B

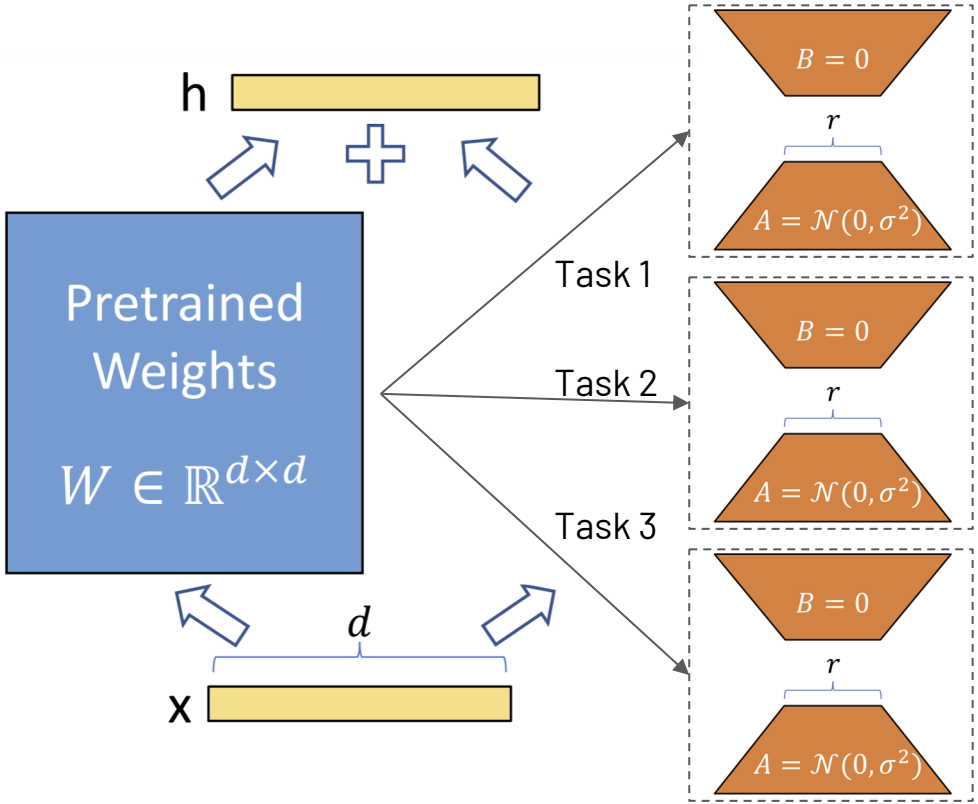

For each new task (e.g. summarization, sql-assistant, medical-assistant, math tutor), trainabe weight matrices A and B are updated.

Model switches between tasks by only swapping the corresponding low-rank weight matrices.

LoRA: Low Rank Adaptation for Large Language Models, Hu & Shen et. al, 2021

LoRA Does Not Affect Latency

Regular feedforward:

h = Wx

Feedforward with LoRA:

h = Wx + BAx = (W + BA) x

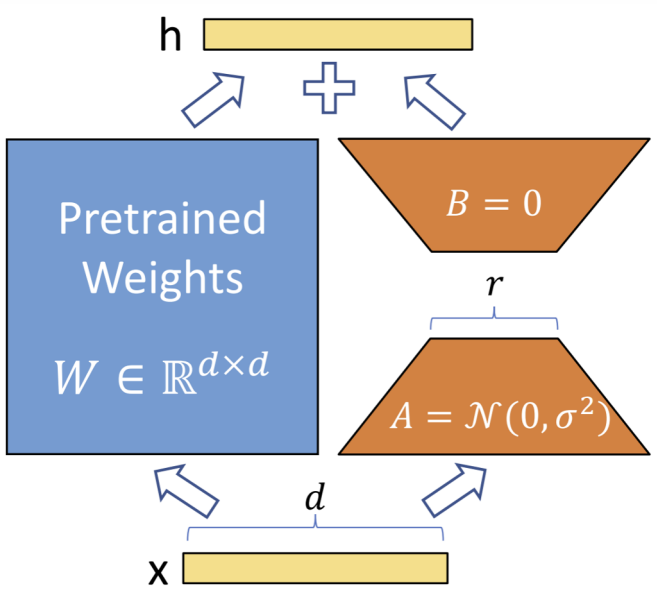

Update W once LoRA finetuning is done:

W_{LoRA} = W + BA

- Smaller Optimizer state (A, B). Other parameters can be lower precision
- Still need to compute forward and backward passes (compute is similar)

How to Apply LoRA to Transformers?

- LoRA is applied to the self-attention layer
- Rank r needs to be finetuned for new tasks
- A small r could be sufficient for many tasks
- r << d_model
- Applying on W_q and W_v have empirically shown best results
  - ➢ Very small set of updates, nearly as good as doing LoRA on everything!

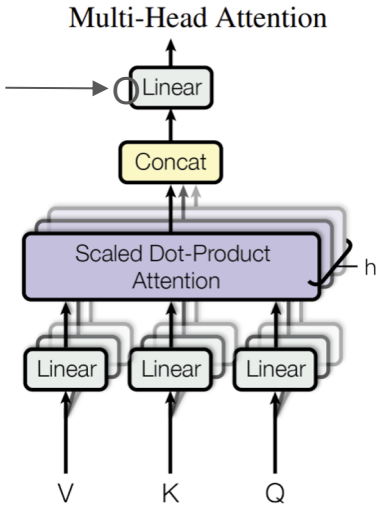

	# of Trainable Parameters = 18M						
Weight Type	W_q	W_k	W_v	W_o	W_q, W_k	W_q, W_v	W_q, W_k, W_v, W_o
Rank r	8	8	8	8	4	4	2
WikiSQL ($\pm 0.5\%$)	70.4	70.0	73.0	73.2	71.4	**73.7**	**73.7**
MultiNLI ($\pm 0.1\%$)	91.0	90.8	91.0	91.3	91.3	91.3	**91.7**

QLoRA – Adding on Quantization

- A significant fraction of open-source, enthusiast alignment works involve QLoRA

QLoRA: Efficient Finetuning of Quantized LLMs

Tim Dettmers* Artidoro Pagnoni* Ari Holtzman

Luke Zettlemoyer

University of Washington
{dettmers,artidoro,ahai,lsz}@cs.washington.edu

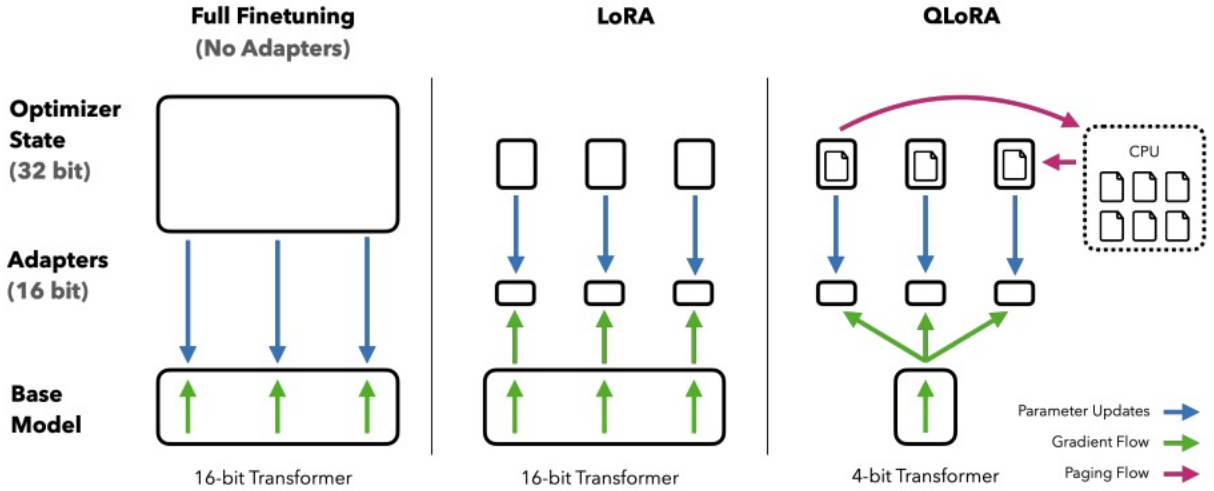

QLoRA in Action

- 4-bit quantization of the model
- 1/3 the Memory, similar performance – Matches 16 bit performance !
- Use of GPU memory paging to prevent OOM
- Can tune a 65B model on a 48GB GPU !

Table 6: Zero-shot Vicuna benchmark scores as a percentage of the score obtained by ChatGPT evaluated by GPT-4. We see that OASST1 models perform close to ChatGPT despite being trained on a very small dataset and having a fraction of the memory requirement of baseline models.

Model / Dataset	Params	Model bits	Memory	ChatGPT vs Sys	Sys vs ChatGPT	Mean	95% CI
GPT-4	-	-	-	119.4%	110.1%	**114.5%**	2.6%
Bard	-	-	-	93.2%	96.4%	94.8%	4.1%
Guanaco	65B	4-bit	41 GB	96.7%	101.9%	**99.3%**	4.4%
Alpaca	65B	4-bit	41 GB	63.0%	77.9%	70.7%	4.3%
FLAN v2	65B	4-bit	41 GB	37.0%	59.6%	48.4%	4.6%
Guanaco	33B	4-bit	21 GB	96.5%	99.2%	**97.8%**	4.4%
Open Assistant	33B	16-bit	66 GB	91.2%	98.7%	94.9%	4.5%
Alpaca	33B	4-bit	21 GB	67.2%	79.7%	73.6%	4.2%
FLAN v2	33B	4-bit	21 GB	26.3%	49.7%	38.0%	3.9%

Pros of Fine-tuning via Re-parameterization

- ✓ These methods are faster to tune than standard full model finetuning.
- ✓ It is straight-forward to swap between tasks by swapping in and our just the tuned weights.
- ✓ Empirically, LoRA has been shown to be very effective on a variety of tasks.
- ✓ QLoRA (Quantized LoRA) is even more efficient than LoRA.

Summary on PEFT approaches (cont'd)

Name & Refs	Method	#Params
SEQUENTIAL ADAPTER Houlsby et al. (2019) COMPACTER Mahabadi et al. (2021a) ADAPTERDROP Rücklé et al. (2021)	$\text{LayerNorm}(\mathbf{X} + \mathbf{H}(\mathbf{X})) \rightarrow \text{LayerNorm}(\mathbf{X} + \text{ADT}(\mathbf{H}(\mathbf{X})))$ $\text{LayerNorm}(\mathbf{X} + \mathbf{F}(\mathbf{X})) \rightarrow \text{LayerNorm}(\mathbf{X} + \text{ADT}(\mathbf{F}(\mathbf{X})))$ $\text{ADT}(\mathbf{X}) = \mathbf{X} + \sigma(\mathbf{X}\mathbf{W}_{d_h \times d_m})\mathbf{W}_{d_m \times d_h}, \ \sigma = \text{activation}$	$L \times 2 \times (2d_h d_m)$ $L \times 2 \times (2(d_h + d_m))$ $(L - n) \times 2 \times (2d_h d_m)$
PARALLEL ADAPTER He et al. (2022)	$\text{LayerNorm}(\mathbf{X} + \mathbf{H}(\mathbf{X})) \rightarrow \text{LayerNorm}(\mathbf{X} + \text{ADT}(\mathbf{X}) + \mathbf{H}(\mathbf{X}))$ $\text{LayerNorm}(\mathbf{X} + \mathbf{F}(\mathbf{X})) \rightarrow \text{LayerNorm}(\mathbf{X} + \text{ADT}(\mathbf{X}) + \mathbf{F}(\mathbf{X}))$ $\text{ADT}(\mathbf{X}) = \sigma(\mathbf{X}\mathbf{W}_{d_h \times d_m})\mathbf{W}_{d_m \times d_h}, \ \sigma = \text{activation}$	$L \times 2 \times (2d_h d_m)$
ADAPTERBIAS	$\text{LayerNorm}(\mathbf{X} + \mathbf{F}(\mathbf{X})) \rightarrow \text{LayerNorm}(\text{ADT}(\mathbf{X}) + \mathbf{F}(\mathbf{X}))$ $\text{ADT}(X) = \mathbf{X}\mathbf{W}_{d_h \times 1}\mathbf{W}_{1 \times d_h}$	$L \times 2 \times d_h$
PREFIX-TUNING Li & Liang (2021)	$\mathbf{H}_i = \text{ATT}(\mathbf{X}\mathbf{W}_q^{(i)}, [\text{MLP}_k^{(i)}(\mathbf{P}'_k) : \mathbf{X}\mathbf{W}_k^{(i)}], [\text{MLP}_v^{(i)}(\mathbf{P}'_v) : \mathbf{X}\mathbf{W}_v^{(i)}])$ $\text{MLP}^{(i)}(\mathbf{X}) = \sigma(\mathbf{X}\mathbf{W}_{d_m \times d_m})\mathbf{W}^{(i)}_{d_m \times d_h}$ $P' = \mathbf{W}_{n \times d_m}$	$n \times d_m + d_m^2$ $+ L \times 2 \times d_h d_m$
LORA Hu et al. (2021a)	$\mathbf{H}_i = \text{ATT}(\mathbf{X}\mathbf{W}_q^{(i)}, \text{ADT}_k(\mathbf{X}) + \mathbf{X}\mathbf{W}_k^{(i)}, \text{ADT}_v(\mathbf{X}) + \mathbf{X}\mathbf{W}_v^{(i)})$ $\text{ADT}(\mathbf{X}) = \mathbf{X}\mathbf{W}_{d_h \times d_m}\mathbf{W}_{d_m \times d_h}$	$L \times 2 \times (2d_h d_m)$
BITFIT Zaken et al. (2021)	$f(\mathbf{X}) \rightarrow f(\mathbf{X}) + \mathbf{B}, \ \text{for all function } f$	$L \times (7 \times d_h + d_m)$

Source: "Delta Tuning: A Comprehensive Study of Parameter Efficient Methods for Pre-trained Language Models," Ding et al, 2023

Key Takeaways for Supervised Fine-Tuning / Instruction Tuning

- Instruction Tuning on Desired Behaviors –
  - Style, Factuality (or Popularity), safety, etc.
- Instruction Tuning (SFT) works best when we are just extracting pre-training behaviors, not adding new ones.
- Adding NEW (factually correct !) data can sometimes hurt because it will make Hallucination worse
- Small amount of the right kinds of behavior (safety, instruction-following, style) make a big difference.
- Instruction Tuning and Pretraining are increasingly merging together
- Instruction Tuning can be done with very little scale and memory
  - Via Parameter Efficient Fine Tuning (PEFT) approaches

Part II. Human Preference Fine Tuning / Alignment

ChatGPT: LLM + RLHF

- While exact ChatGPT methods are not fully disclosed but it seems to be GPT-3.5 + RLHF
- (Rumor) About 10x spend on Human Annotation Budget
- Prior to the release of ChatGPT, OpenAI did publish InstructGPT = GPT-3 with RLHF:

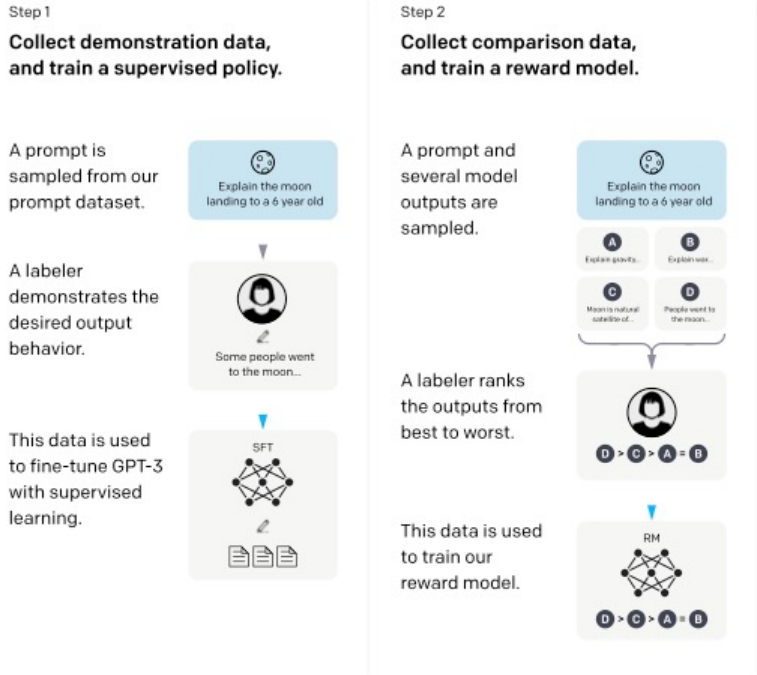

Figure 2: A diagram illustrating the three steps of our method: (1) supervised fine-tuning (SFT), (2) reward model (RM) training, and (3) reinforcement learning via proximal policy optimization (PPO) on this reward model. Blue arrows indicate that this data is used to train one of our models. In Step 2, boxes A-D are samples from our models that get ranked by labelers. See Section 3 for more details on our method.

Source: Ouyang, Long, et al. "Training language models to follow instructions with human feedback." *arXiv preprint arXiv:2203.02155*(2022).

Earlier Post-training approach pioneered by ChatGPT

Approach: Imitation (SFT) followed by Reinforcement (`RL'HF)

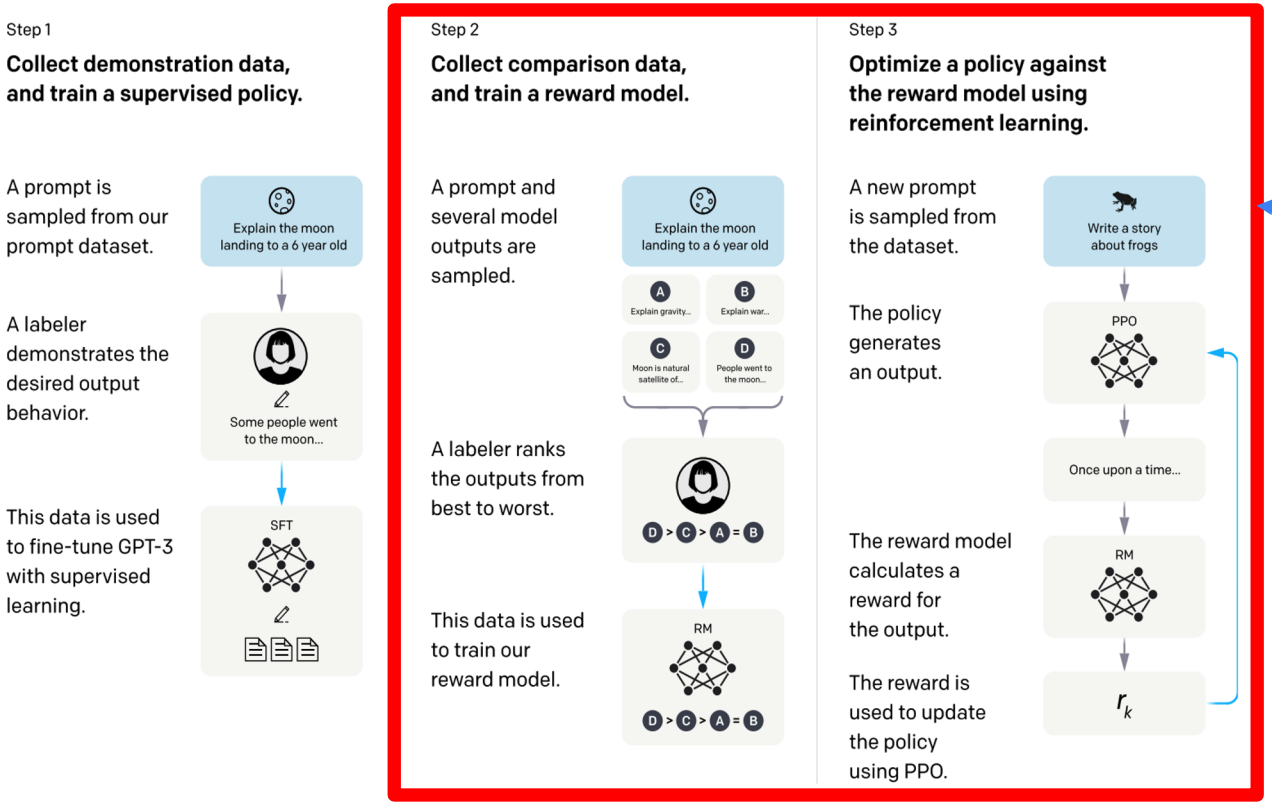

The 'RL' part

Ouyang et al. 2022. InstructGPT.

From Imitation to Optimization

Imitation (SFT)

Fit $\hat{p}(y|x) \approx p^*(y|x)$ for some reference distribution $p^*(y|x)$

- Pure generative modeling perspective
- Requires samples from reference policy

Optimization (RLHF)

Find $\hat{\boldsymbol{p}}(\boldsymbol{y}|\boldsymbol{x})$ such that $\max_p E_p[R(y,x)]$ for a reward $R(y,x)$

- Maximize some reward function that we can measure
- LMs are policies, not a model of some distribution

After SFT you have an Assistant:

<USER>

```
Can you help me with this code? It seems like
there is a bug.
print(“hello world)
```

<ASSISTANT>

```
It looks like you forgot to close the string
passed to the function print. You have to add a
closing quote to properly terminate the string.
Here is the corrected function:

print(“hello world”)

Let me know if I can help with anything else!
```

The role of preference finetuning (PreFT)

Aligning to Human Preferences gives:

- Stronger training influence for style and chat evaluations (e.g. ChatBotArena).
- Continue building capabilities of skills from SFT, but lower absolute magnitude of improvements.

Better way to get Human feedback: comparisons

It is often much easier to compare Answers instead of writing Answers.

Simple example: it's much easier to spot a good haiku than it is to generate one:

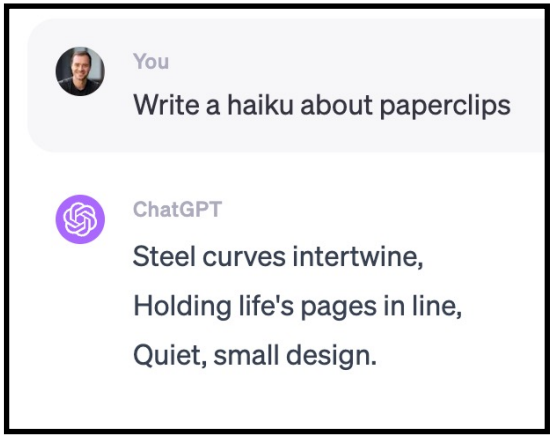

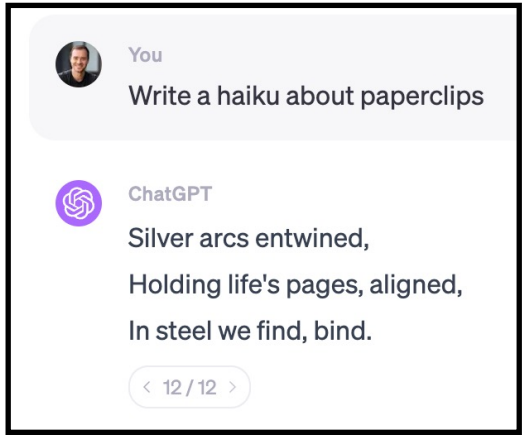

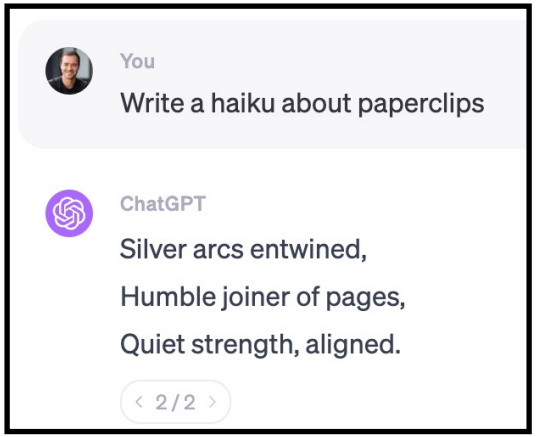

Better way to get Human feedback: comparisons

Might be easier to get scalar feedback rather than optimal policy.

One example for a small, open-source 7B model

	Base model (given, LLaMA) →	Supervised learning (one time) →	Pairwise feedback →	RL →	Evaluation
Compute cost	$300k	$100	$100	$100	$0
Annotation cost	$0	$25k	$4k	$0	$50

Notice (nearly) all the costs step 2 onwards are *human annotation costs*

Better way to get Human feedback: comparisons

People don't always write the thing that they prefer in LM outputs

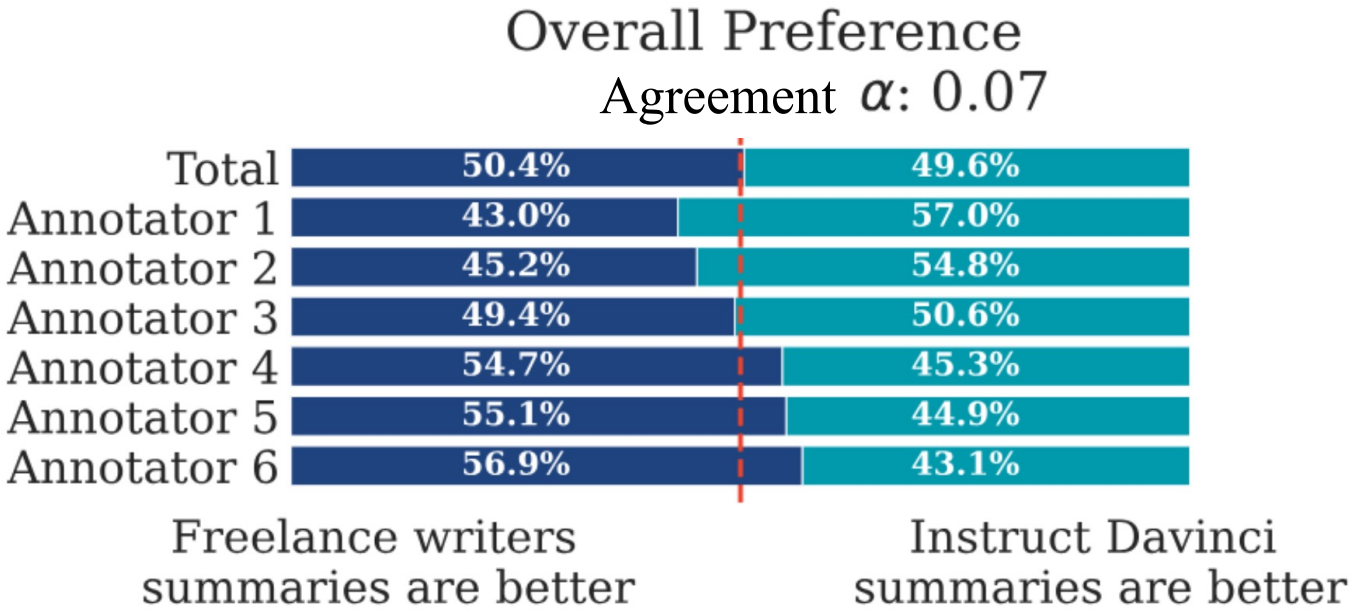

Source: Zhang et al, "Benchmarking Large Language Models for News Summarization", 2023

How do we get GOOD pairwise feedback ?

- Standard ‘pairwise’ feedback setup

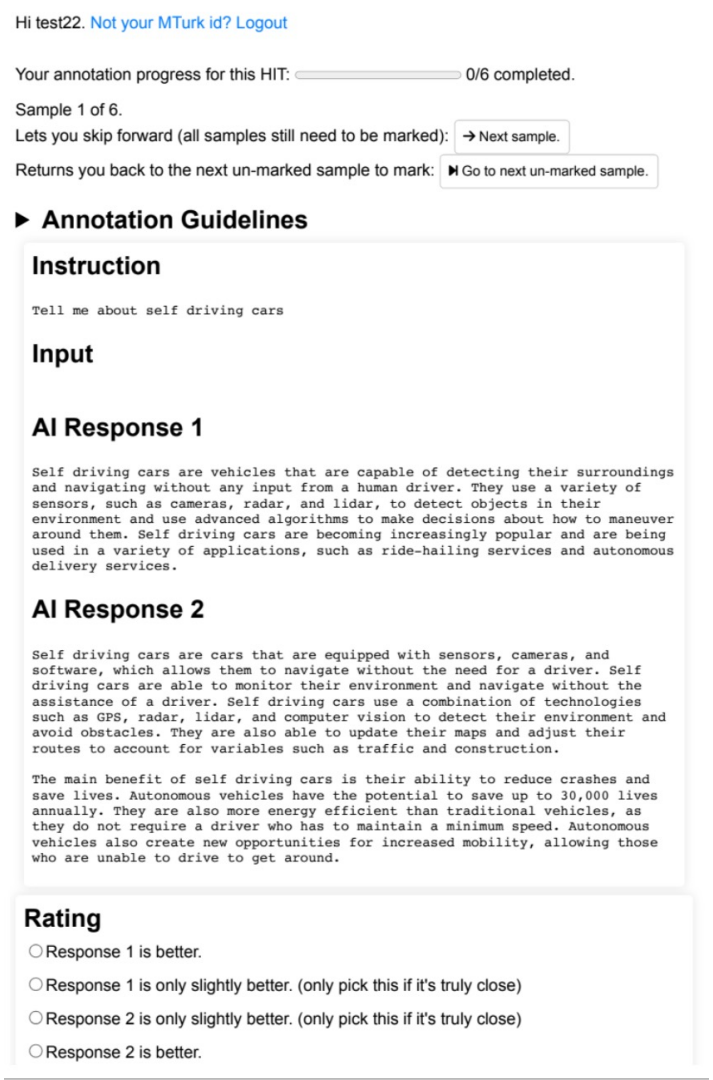

Labeling instructions:
Prefer Helpful, Truthful and Harmless responses

Excerpt of labeling instructions on the API prompt distribution

You are given a text-based description of a task, submitted by a user. This task description may be in the form of an explicit instruction (e.g. "Write a story about a wise frog."). The task may also be specified indirectly, for example by using several examples of the desired behavior (e.g. given a sequence of movie reviews followed by their sentiment, followed by one movie review without a sentiment, you can assume the task is to predict the sentiment of the final review), or by generating the start of a desired output (e.g. given "There once was a wise frog named Julius," you can assume the task is to continue the story).

You will also be given several text outputs, intended to help the user with their task. Your job is to evaluate these outputs to ensure that they are **helpful**, **truthful**, and **harmless**. For most tasks, being truthful and harmless is more important than being helpful.

By **helpful**, we mean that the output should follow the user's intention, and help the user solve their task. Some examples of helpful behavior:

- Writing in clear language.
- Answering the question they meant to ask, even if they mis-asked it.
- Being sensitive to internationality (e.g. "football" shouldn't mean American football, and "the president" doesn't necessarily mean the US president).
- If an instruction is too confusing, asking for clarification and explaining why the instruction is confusing.
- Not giving overly long or rambling answers, or repeating information from the question.
- Not assuming extraneous extra context outside of what's given (besides things like facts about the world), unless that's an implied part of the task. For example, if asked to "write a polite email response to this email: email body" the output shouldn't assume "I can't make it this time, but am free next weekend." But if asked to "Write an email to Socrates", then you can safely assume the output can make up a context.

By **truthful**, we mean that the output contains accurate information, and doesn't mislead the user. Some examples of truthful behavior:

- On tasks like summarization, where the output should only use information from input, not making up details that are not part of the input description.
- Not producing clearly false information about the world (e.g. making up facts or promoting conspiracies). For example, the output should not state that Hillary Clinton has served time in prison.
- Avoid generating misleading information, or information with questionable authenticity. For example, if the input asks "Why did Hillary Clinton go to jail?", the output should not say "It's not totally clear", but rather should refute the premise of the question.

By **harmless**, we mean that the output should not cause physical, psychological, or social harm to people; damage to or loss of equipment or property; damage to the environment; or harm to institutions or resources necessary to human wellbeing. Some examples of harmless behavior:

- Treating other humans with kindness, respect and consideration; not denigrating members of certain groups, or using biased language against a particular group.
- Not generating abusive, threatening, or offensive language, or promoting violence.
- Not writing sexual or violent content if it's not asked for.
- Not giving bad real-world advice, or promoting illegal activity.

Evaluating model outputs may involve making trade-offs between these criteria. These trade-offs will depend on the task. Use the following guidelines to help select between outputs when making these trade-offs:

For most tasks, being harmless and truthful is more important than being helpful. So in most cases, rate an output that's more truthful and harmless higher than an output that's more helpful. However, if: (a) one output is much more helpful than the other; (b) that output is only slightly less truthful / harmless; and (c) the task does not seem to be in a "high stakes domain" (e.g. loan applications, therapy, medical or legal advice, etc.); then rate the more helpful output higher. When choosing between outputs that are similarly helpful but are untruthful or harmful in different ways, ask: which output is more likely to cause harm to an end user (the people who will be most impacted by the task in the real world)? This output should be ranked lower. If this isn't clear from the task, then mark these outputs as tied.

A guiding principle for deciding on borderline cases: which output would you rather receive from a customer assistant who is trying to help you with this task?

Ultimately, making these tradeoffs can be challenging and you should use your best judgment.

[InstructGPT paper]

(Unconfirmed) Google Bard Crowdsourcing Instructions

In this task, you will be provided with a **Prompt** from a user (e.g., a question, instruction, statement) to an AI chatbot along with two potential machine-generated **Responses** to the Prom

Your job is to assess which of the two Responses is better for the Prompt, considering the following for each Response:

Helpfulness: To what extent does the Response provide useful information or satisfying content for the Prompt? Responses should: ▪ **Address the intent of the user's Prompt** such that a user would not feel the Prompt was ignored or misinterpreted by the Response. ▪ **Provide specific, comprehensive, and up-to-date information** for the user needs expressed in the Prompt. ▪ **Be sensible and coherent.** The response should not contain any nonsensical information or contradict itself across sentences (e.g., refer to two different people with the same name as if they are the same person). ▪ **Adhere to any requirements indicated in the Prompt** such as an explicitly specified word length, tone, format, or information that the Response should include. ▪ **Not contain inaccurate, deceptive, or misleading information** (based on your current knowledge or quick web search - you do not need to perform a rigorous fact check) ▪ **Not contain harmful, offensive, or overly sexual content** A Response may sometimes intentionally avoid or decline to address the question/request of the Prompt and may provide a reason for why it is unable to respond. For example, "Sorry, there may not be a helpful answer to this question." These responses can be considered helpful in cases where an appropriate helpful response to the Prompt does not seem possible.	Rating scale: ▪ **Not at All Helpful:** Response is useless/irrelevant, contains even a single piece of nonsensical/inaccurate/deceptive/misleading information, and/or contains harmful/offensive/overly sexual content. ▪ **Slightly Helpful:** Response is somewhat related to the Prompt, does not address important aspects of the Prompt, and/or contains outdated information. ▪ **Somewhat Helpful:** Response partially addresses the intent of the Prompt (most users would want more information), contains extra unhelpful information, and/or is lacking helpful details/specifics. ▪ **Very Helpful:** Response addresses the intent of the Prompt with a satisfying response. Some users might want a more comprehensive response with additional details or context. It is comparable to a response an average human with basic subject-matter knowledge might provide. ▪ **Extremely Helpful:** Response completely addresses the intent of the Prompt and provides helpful details/context. It is comparable to a response a talented/well-informed human with subject-matter expertise might provide.
Presentation: To what extent is the content of the Response conveyed well? Responses should: ▪ **Be organized in a structure that is easy to consume and understand.** Flowing in a logical order and makes good use of formatting such paragraphs, lists, or tables. ▪ **Be clearly written in a polite neutral tone** that is engaging, direct, and inclusive. The tone should not be *overly* friendly, salesy, academic, sassy, or judgmental in a way that most users would consider to be off-putting or overdone. ▪ **Have consistent style with natural phrasing and transitions** as if composed by a single talented human. ▪ **Not be rambling, repetitive. or contain clearly off-topic information.** Similar information should not be repeated multiple times. It is harder for users to consume the helpful information in a response if there is repetitive or less helpful information mixed into the response. ▪ **Not include notable language issues or grammatical errors**	Rating scale: ▪ **Poor:** Response is poorly written or has notable structural, formatting. language, or grammar issues. Or Response has an awkward or inappropriate tone. Or the Response repeats similar information. Or only a small portion of the Response contains helpful information. ▪ **Adequate:** Response could have been written/organized better or may have minor language/grammar issues. A minimal amount of less helpful information may be present. Users would still feel the content of the Response was easy to consume. ▪ **Excellent:** Response is very well written and organized. Sentences flow in a logical order with smooth transitions and consistent style. The content of the Response is conveyed in a way that is comparable to a response a talented human might produce.

Overall, you should consider both factors in your SxS rating of which response is better. A more concise response presenting the most helpful information directly and clearly is usually better than a longer response that may be harder to consume and/or contains clearly off-topic information. Responses with Poor Presentation (e.g., rambling, inappropriate tone) should play a significant role in your assessment of which side is better. It may help to imagine the user chatting with a real person and consider which Response most users would prefer to receive from a real person.

Source: https://assets.bwbx.io/documents/users/iqjWHBFdfxIU/rqKqEqbXBnDI/v0

Crowdsourcing Worker Selection for InstructGPT

3.4 Human data collection

To produce our demonstration and comparison data, and to conduct our main evaluations, we hired a team of about 40 contractors on Upwork and through ScaleAI. Compared to earlier work that collects human preference data on the task of summarization (Ziegler et al., 2019; Stiennon et al., 2020; Wu et al., 2021), our inputs span a much broader range of tasks, and can occasionally include controversial and sensitive topics. Our aim was to select a group of labelers who were sensitive to the preferences of different demographic groups, and who were good at identifying outputs that were potentially harmful. Thus, we conducted a screening test designed to measure labeler performance on these axes. We selected labelers who performed well on this test; for more information about our selection procedure and labeler demographics, see Appendix B.1.

Hired 40 workers
from
ScaleAI and Upwork

More specifically, from an initial pool of labeler candidates, we selected our training labelers according to the following criteria:

1. **Agreement on sensitive speech flagging.** We created a dataset of prompts and completions, where some of prompts or completions were sensitive (i.e. anything that could elicit strong negative feelings, whether by being toxic, sexual, violent, judgemental, political, etc.). We labeled this data for sensitivity ourselves, and measured agreement between us and labelers.
2. **Agreement on rankings.** We take prompts submitted to our API, and several model completions, and have labelers rank the completions by overall quality. We measure their agreement with researcher labels.
3. **Sensitive demonstration writing.** We created a small set of sensitive prompts, where responding to the outputs appropriately would require nuance. We then rated each demonstration on a 1-7 Likert scale, and computed an average "demonstration score" for each labeler.
4. **Self-assessed ability to identify sensitive speech for different groups.** We wanted to select a team of labelers that had collectively were able to identify sensitive content in a broad range of areas. For legal reasons, we can't hire contractors based on demographic criteria. Thus, we had labelers answer the question: "For what topics or cultural groups are you comfortable identifying sensitive speech?" and used this as part of our selection process.

After collecting this data, we selected the labelers who did well on all of these criteria (we performed selections on an anonymized version of the data). Since the fourth criteria is subjective, we ultimately chose labelers subjectively according to these criteria, though we had soft cutoffs at 75% agreement on sensitive speech flagging and comparisons, and a 6/7 demonstration score.

Crowdsourcing Worker Selection for AlpacaFarm

D Details on human data collection

Qualification. We conducted the qualification of our annotators based on 25 qualification examples. The qualification examples were generated but an OPT 6B model that was studied in the earlier development phase of this project. The five student authors of this paper annotated a shared set of pairwise preferences. From the shared set, we selected 25 questions where the majority of the authors reached an agreement on the correct annotation. We then use these questions as a qualification test and selected the top 16 annotators whose agreement is the highest with the authors. We paid the annotators the same price for the qualification round as we did for the main qualification.

During the annotation process, we also compare each annotator's preference to that of GPT-4. We identified one annotator whose agreement is around 50% with GPT-4, which is a clear outlier from other annotators. Therefore, we discontinued working with this annotator during the annotation project and removed their annotation.

Complexity of Crowdsourcing

- Hard to get really High-Quality, Verifiable annotators
- Hard to get them to really check correctness
- Beware of cultural difference/ misalignment of annotators (what if most of them were from foreign countries/ religious background)
- Have to be careful about GPT4 use

Effect of Response-Length on Human Preference

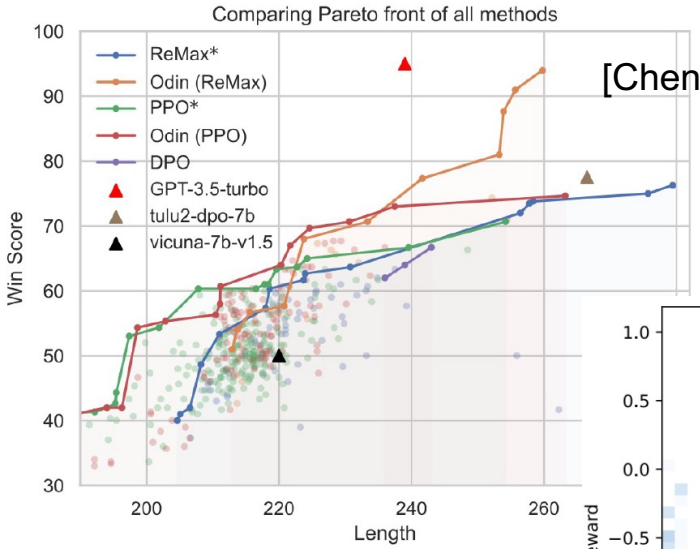

[Chen et al 2024]

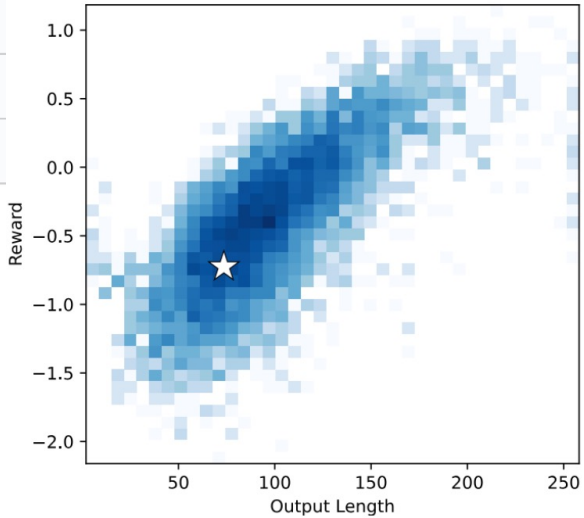

Question: *Why don't adults roll off the bed?*

☆ **SFT (Before); 59 tokens**

Adults typically do not roll off of the bed because they have developed the muscle memory to keep their bodies from involuntarily moving during sleep and maintaining proper posture.

RLHF (After); 243 tokens: **Similar output, but much longer / more details**

Adults generally do not roll off of the bed because they have developed muscle control and awareness which helps them to remain upright while sleeping.

Additionally, most adults find it uncomfortable or even painful to move around in their sleep, so rolling off the bed would be difficult without waking up first. There may also be safety concerns such as falling out of bed and potentially hurting oneself if one were to lose balance or fall from the ...

[Singhal et al 2024]

Collecting Human Preferences for RLHF

- Humans are asked to rank the outputs of the model based on various criteria, for example:
  - Usefulness
  - Harmfulness
  - Truthfulness

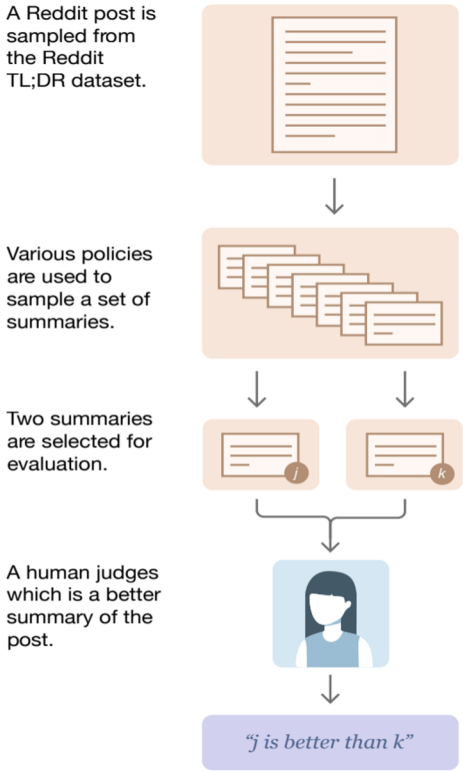

Stiennon et al., Learning to summarize from human feedback, NeurIPS 2020

Alignment Problem: How to Finetune LLMs to Do what Human Want

- One solution is to perform RL training with human feedback as the reward signal, but RL typically requires thousands to millions of labels to learn an effective policy.

- But what if you train a model to predict human judgments and then use this predictive model as the reward signal?

- [Christiano et al.] were able to train a model to follow human preferences in MuJoCo and Atari games, after labeling <1% of the environment interactions!

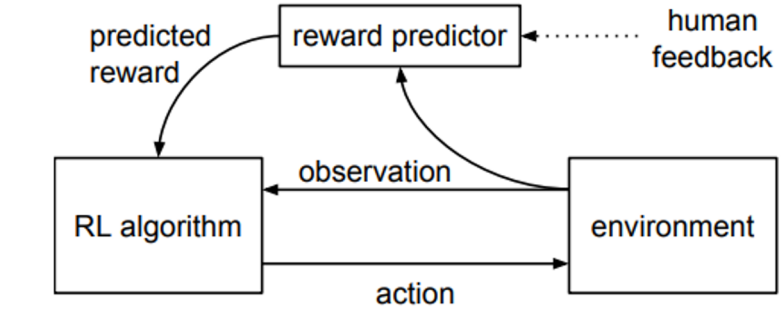

Paul Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep Reinforcement Learning from Human Preferences, NeurIPS 2017

Finetuning based on RLHF

RL-based finetuning of the LLM based on human preferences

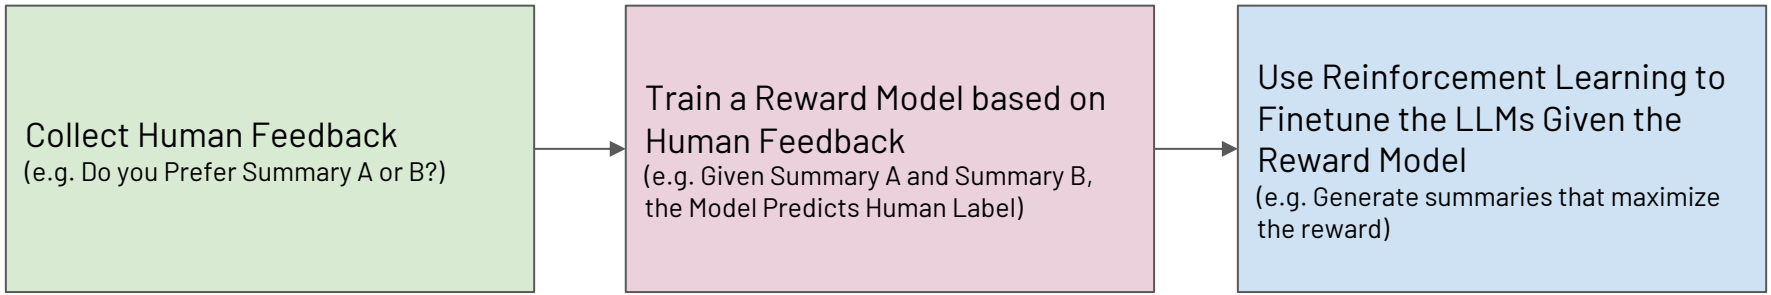

Bai et al., 2022, Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback

Reinforcement Learning with Human Feedback (RLHF)

- RLHF: technique that trains a "reward model" directly from human feedback
- Uses the model as a reward function to optimize an agent's policy using reinforcement learning (RL) through an optimization algorithm
- Ask humans to rank instances of the agent's behavior, e.g. which produced response is better

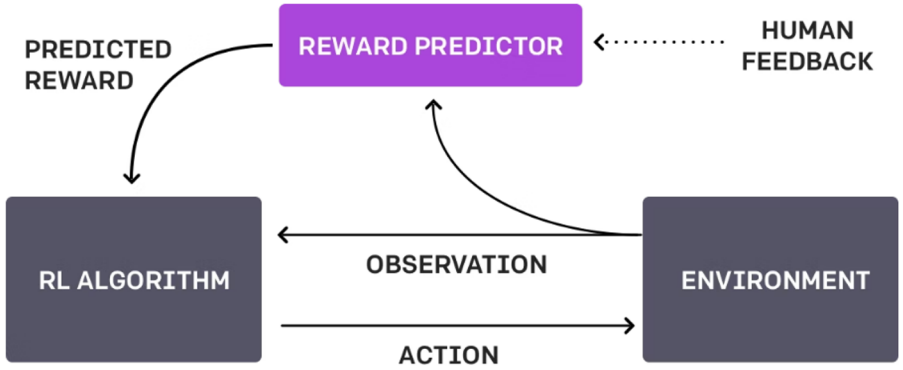

What Are States, Actions, and Rewards in LLMs?

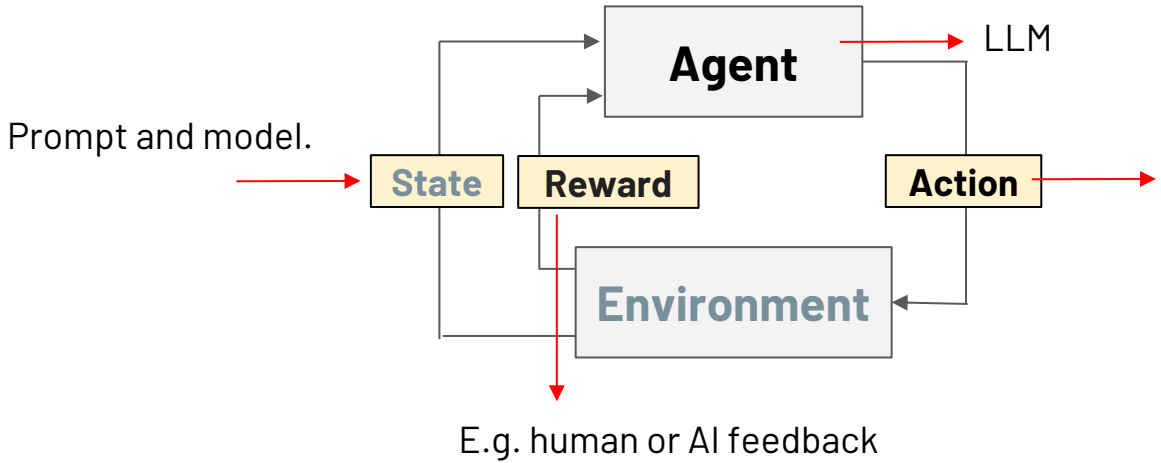

Generated output (tokens). One common approach is to consider an entire output sequence (e.g. the answer to a prompt) as the action.

RL Finetuning for LLMs

Two Main Objectives:

Maximize the reward so that the generated outputs are aligned with human (or AI) preferences

Avoid large divergence from the model and overfitting to the reward at any step

prompts outputs Scalar hparam KL loss

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)} \left[r_\phi(x, y) \right] - \beta \mathbb{D}_{\mathrm{KL}} \left[\pi_{\mathrm{ref}}(y \mid x) \,||\, \pi_\theta(y \mid x) \right]$$

E.g. reward based on human or AI feedback

A reference LLM (e.g. LLM from a prior optimization epoch)

LLM that is being finetuned.

Kullback–Leibler divergence (KL) Divergence

- KL divergence measures the difference between two probability distributions
- In the context of LLM finetuning, the KL loss ensures that the current LLM doesn't diverge too much from its earlier versions, while optimizing for reward

$$D_{\mathrm{KL}}(p||q) = \sum_{i=1}^{N} p(x_i) \cdot (\log p(x_i) - \log q(x_i))$$



p and q are two distributions, such as π_{ref} π_θ in the RL finetuning objective.

Digression:
A Brief Intro to Deep Reinforcement Learning

Refer to the following for good introductory lectures on Deep RL
(in the order of increasing length and details/completeness)

Intuitive explanation on PPO, GRPO by Julia Turc:
https://www.youtube.com/watch?v=xT4jxQUl0X8; https://www.youtube.com/watch?v=8jtAzxUwDj0

Foundations of Deep RL in 6 Lectures by Prof. Pieter Abbeel, UC Berkeley
https://www.youtube.com/watch?v=2GwBez0D20A

CS285 Deep RL by Prof. Sergey Levine, UC Berkeley
https://www.youtube.com/playlist?list=PL_iWQOsE6TfVYGEGiAOMaOzzv41Jfm_Ps

Reinforcement Learning

Train an **Agent** that operates in an **environment** in order to maximize the **reward**

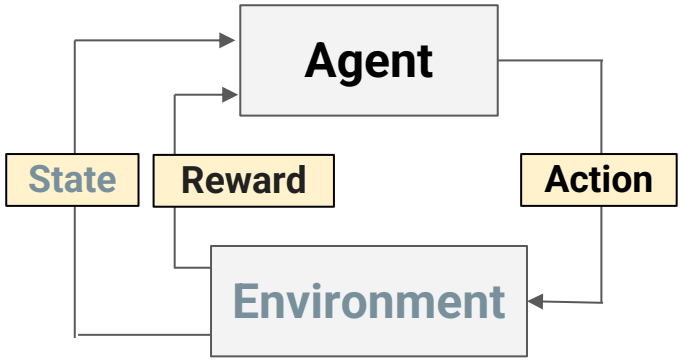

Slides adapted from ML EDU Partners Program

Reinforcement Learning

Agent | **Environment** | **Reward**

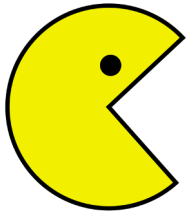

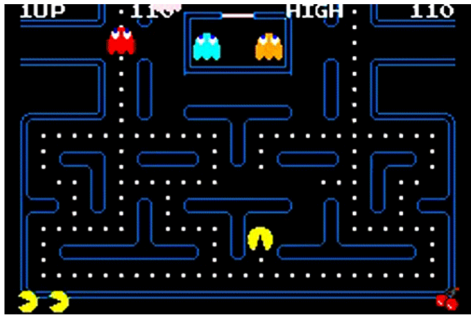



Agents want “Reward”

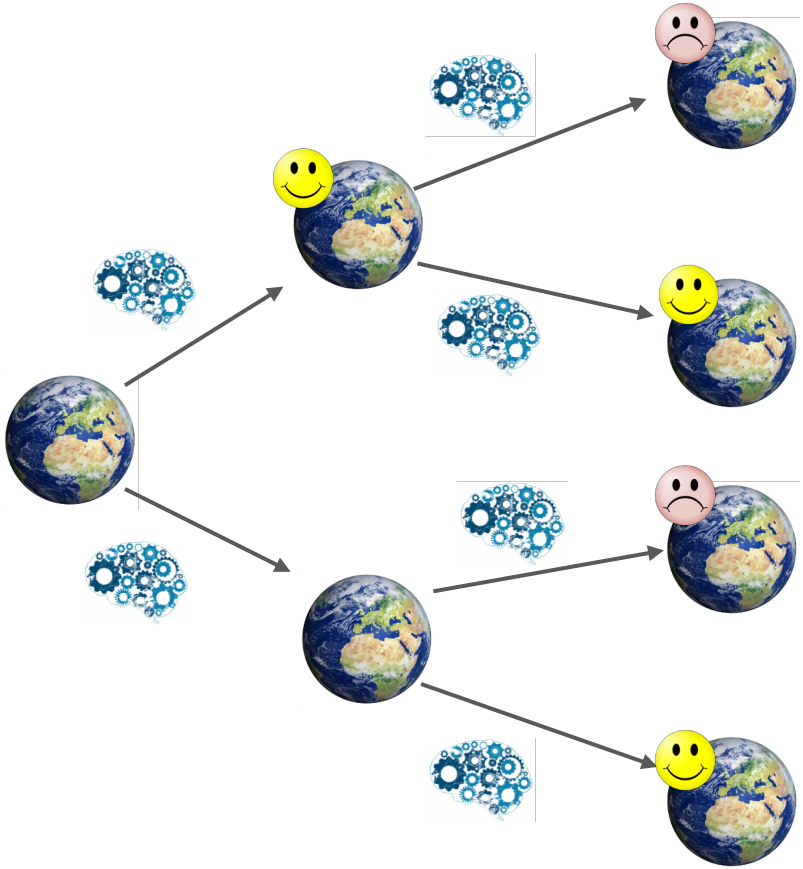



Good states = More reward



Bad states = Less reward

How Does RL Work?

Foundation: Markov Decision Process (MDP)

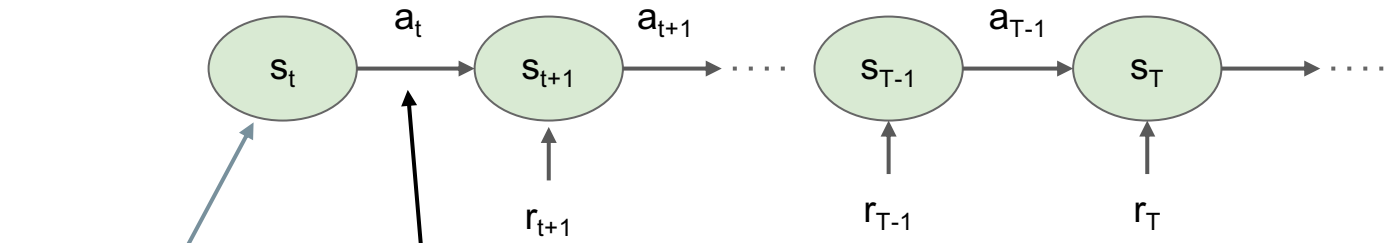

$s_t \in S$ $a_t \in A$ $\mathcal{R}^{a_t}_{s_t, s_{t+1}}$

states **actions** **rewards**

Slides adapted from ML EDU Partners Program

MDP Example

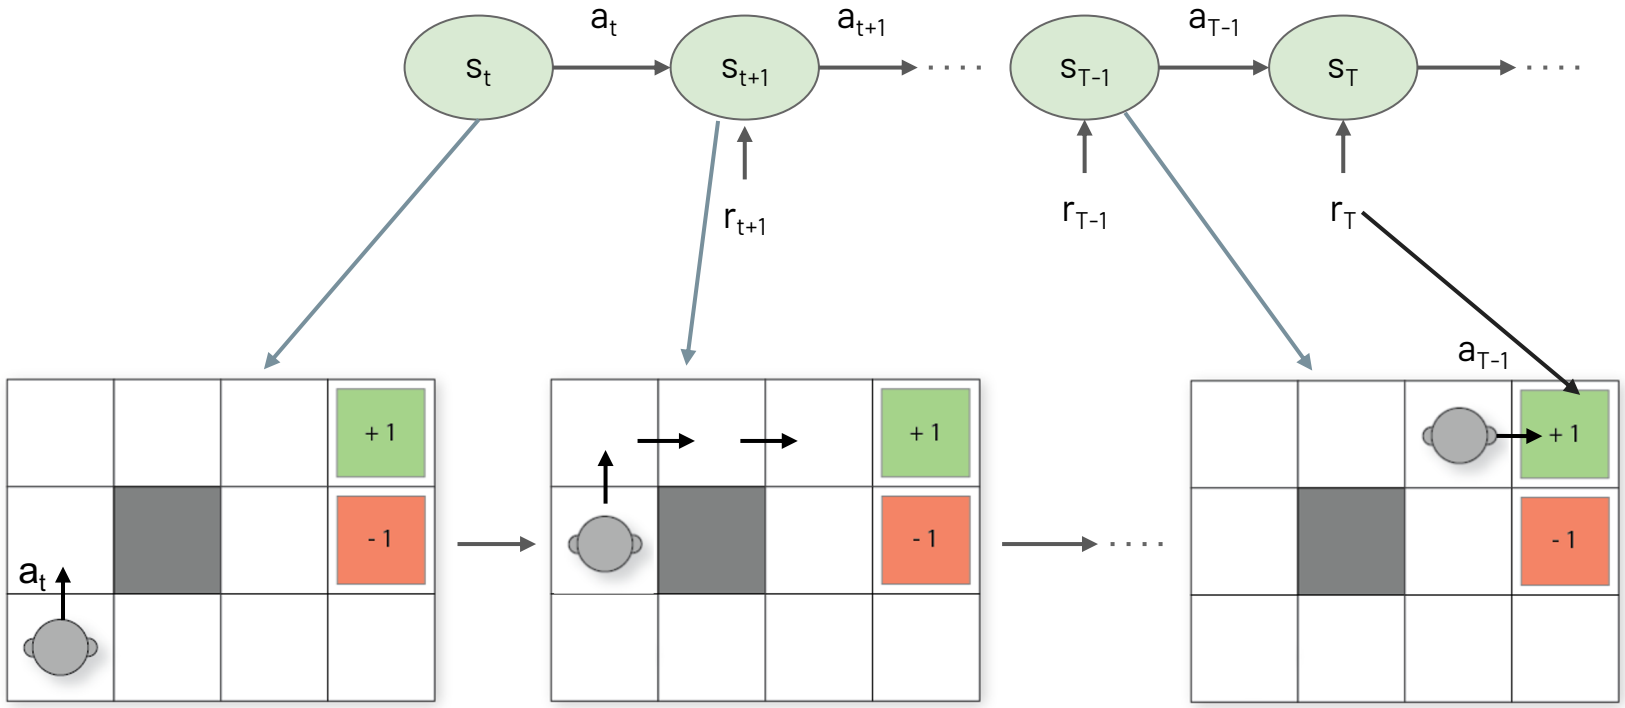

Slides adapted from ML EDU Partners Program

The Markov Property

Assumption:

Given the **present**, the **future** is independent of the **past**

i.e., no matter how you **reached** a state, all **future** states can be predicted from your **current** state and **future** actions

Episodes: Sequence of States, Actions, and Rewards

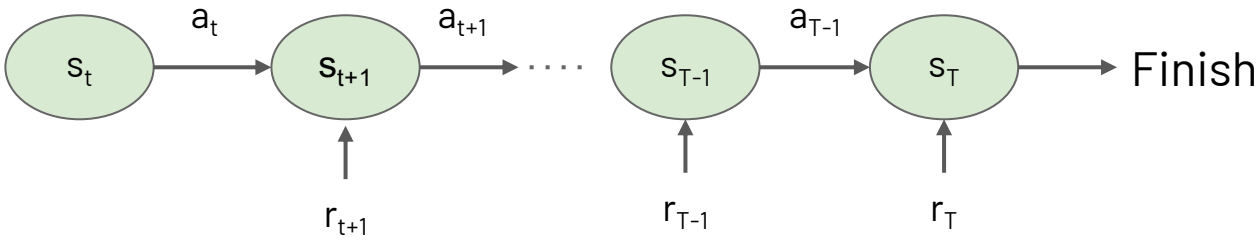

What defines the end of an episode?

- Goal acquired
- Step limit reached
- Fatal mistake

Slides adapted from ML EDU Partners Program

State Transition Function

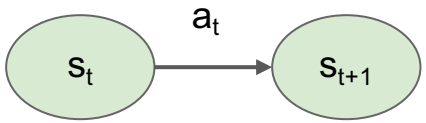

$$\mathcal{P}^{a}_{s_t, s_{t+1}} = P(s_{t+1} | s_t, a_t)$$

Current State, Action → Next State

Reward Function

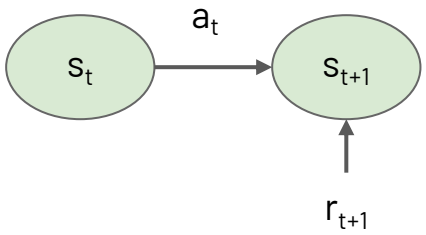

$$\mathcal{R}^{a_t}_{s_t,s_{t+1}} = E(r_t|s_t, s_{t+1}, a_t)$$

Current State, Action, Next State → Reward

Policy

Distribution π over actions given a state that we want to optimize.

$$a_t \sim \pi(a|s_t)$$

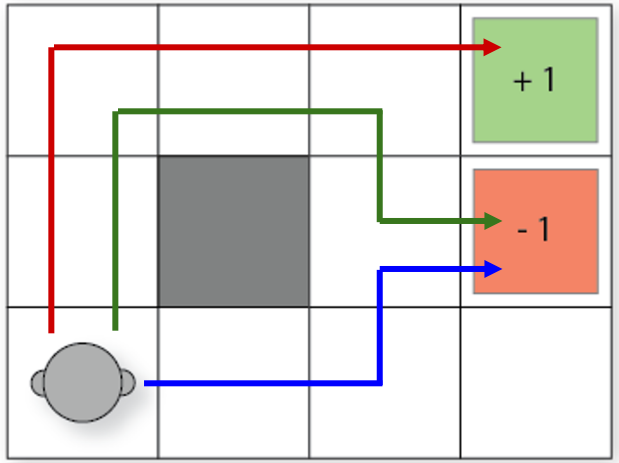

Slides adapted from ML EDU Partners Program

Objective: Maximize “Return”

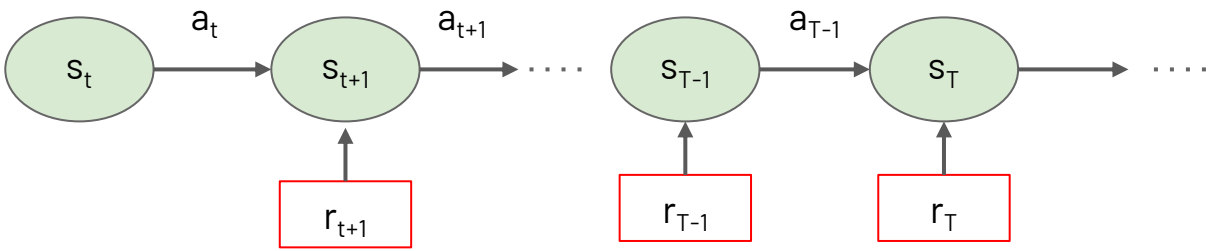

Return = Expected sum of discounted future rewards

$$E[R_t] = E[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots]$$

Deep RL

In Deep RL, we often train both Value and Policy functions, which are modeled by Deep Neural Networks.

Policy predicts an Action given a State. In a discrete RL environment, this is similar to a classifier.

Value model predicts the Expected Return (a scalar), given a State.

Value Function

Value of a State, is the Expected Reward of being in that particular State

Our goal is to have a Policy that takes actions that maximize the return or visit states that maximize the return

$$V^{\pi}(s) = E_{\pi}[R_t | s_t = s]$$

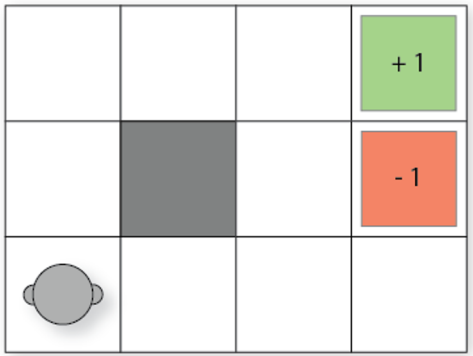

What Are States, Actions, and Rewards in LLMs?

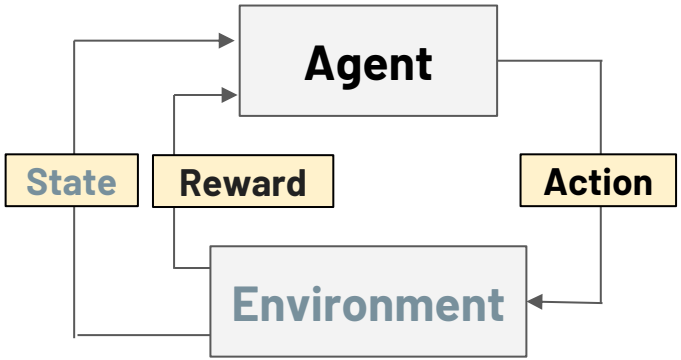

What Are States, Actions, and Rewards in LLMs?

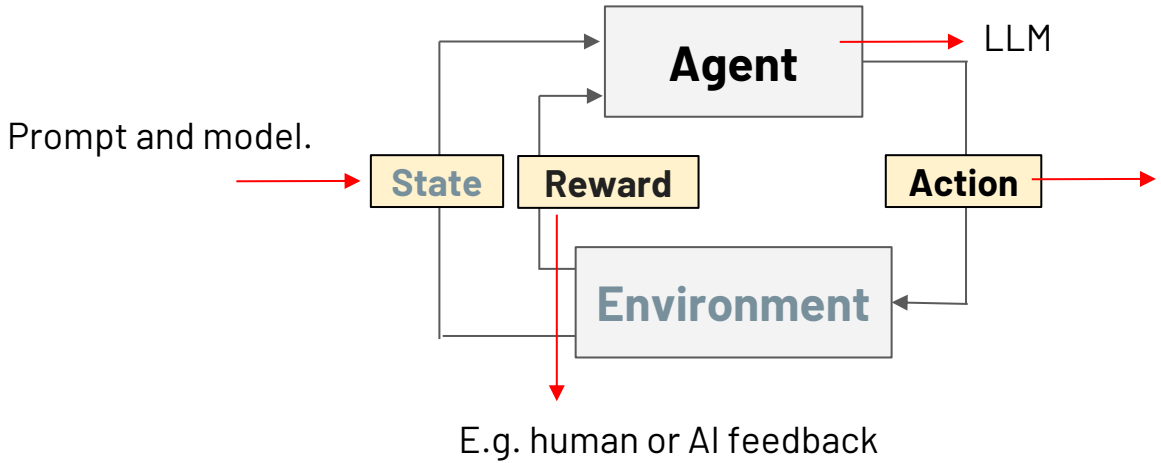

Generated output (tokens). One common approach is to consider an entire output sequence (e.g. the answer to a prompt) as the action.

Recap: RL Finetuning for LLMs

Two Main Objectives:

Maximize the reward so that the generated outputs are aligned with human (or AI) preferences

Avoid large divergence from the model and overfitting to the reward at any step

prompts outputs Scalar hparam KL loss

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)} \left[r_\phi(x, y) \right] - \beta \mathbb{D}_{\mathrm{KL}} \left[\pi_{\mathrm{ref}}(y \mid x) \,||\, \pi_\theta(y \mid x) \right]$$

E.g. reward based on human or AI feedback

A reference LLM (e.g. LLM from a prior optimization epoch)

LLM that is being finetuned.

Policy Gradient

Objective: Finetune LLM (the “policy”) to generate outputs with higher rewards:

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right]$$

Policy Gradient (PG) states we can rewrite the above objective:

$$L^{PG}(\theta) = E_t\left[\log \pi_\theta(a_t|s_t) * r_\phi\right]$$

log probability of taking that action at that state

https://huggingface.co/blog/deep-rl-ppo

Policy Gradient

Objective: Finetune LLM (the “policy”) to generate outputs with higher rewards:

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right]$$

Policy Gradient (PG) states we can rewrite the above objective:

$$L^{PG}(\theta) = E_t[\log \pi_\theta(a_t|s_t) * A_t]$$

log probability of taking that action at that state

Advantage if A>0, this action is better than the other action possible at that state

https://huggingface.co/blog/deep-rl-ppo

Advantage Function Estimator

$$\hat{A}_t = -V(s_t) + r_t + \gamma r_{t+1} + \cdots + \gamma^{T-t+1} r_{T-1} + \gamma^{T-t} V(s_T)$$

Value function at timestep t

Partial reward at timestep t

Discount factor

Value function at timestep T

Advantage Function Estimator in LLM

In RL finetuning of LLMs, a common approach is to set the discount factor to zero.

After an entire sequence is generated by the LLM (e.g. an answer to a question), the advantage can be written as the difference between the reward for the sequence and the value function output.

$$\hat{A}_t = -V(s_t) + r_t$$

Trust Region Policy Optimization (TRPO)

Naive PG is unsustainable and leads to destructively large policy updates.

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right] - \beta\,\mathrm{KL}[\pi_{\theta_{\text{old}}}(\cdot \mid s_t), \pi_\theta(\cdot \mid s_t)]$$

TRPO introduces a surrogate loss as follows:

$$\underset{\theta}{\text{maximize}}\, \hat{\mathbb{E}}_t\left[\frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}\hat{A}_t - \beta\,\mathrm{KL}[\pi_{\theta_{\text{old}}}(\cdot \mid s_t), \pi_\theta(\cdot \mid s_t)]\right]$$

Shulman et al., Trust Region Policy Optimization (TRPO), 2017

PPO: Proximal Policy Optimization

PPO is a more stable version of TRPO, which is commonly used in RL finetuning of LLMs.

PPO uses the *Clipped surrogate objective function.*

A small positive number
< 1, e.g. 0.2

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t)\right]$$

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}$$

Shulman et al., Proximal Policy Optimization Algorithms, 2017
https://huggingface.co/blog/deep-rl-ppo

The Ratio in the PPO Clip Loss

The ratio $r_t(\theta)$ is the probability of taking action a_t at state s_t in the current policy divided by the previous one.

If $r_t(\theta)>1$, the action a_t at state s_t is more likely (has a higher probability) in the current policy than the old policy and vice versa.

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}$$

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t)\right]$$

Shulman et al., Proximal Policy Optimization Algorithms, 2017
https://huggingface.co/blog/deep-rl-ppo

The PPO Clip Loss

The final loss is a lower bound of the clipped vs. unclipped objective.

	$p_t(\theta) > 0$	A_t	Return Value of *min*	Objective is Clipped	Sign of Objective	Gradient
1	$r_t(\theta) \in [1-\epsilon, 1+\epsilon]$	+	$r_t(\theta) A_t$	no	+	✓
2	$r_t(\theta) \in [1-\epsilon, 1+\epsilon]$	−	$r_t(\theta) A_t$	no	−	✓
3	$r_t(\theta) < 1-\epsilon$	+	$r_t(\theta) A_t$	no	+	✓
4	$r_t(\theta) < 1-\epsilon$	−	$(1-\epsilon) A_t$	yes	−	**0**
5	$r_t(\theta) > 1+\epsilon$	+	$(1+\epsilon) A_t$	yes	+	**0**
6	$r_t(\theta) > 1+\epsilon$	−	$r_t(\theta) A_t$	no	−	✓

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_t\right)\right]$$

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}$$

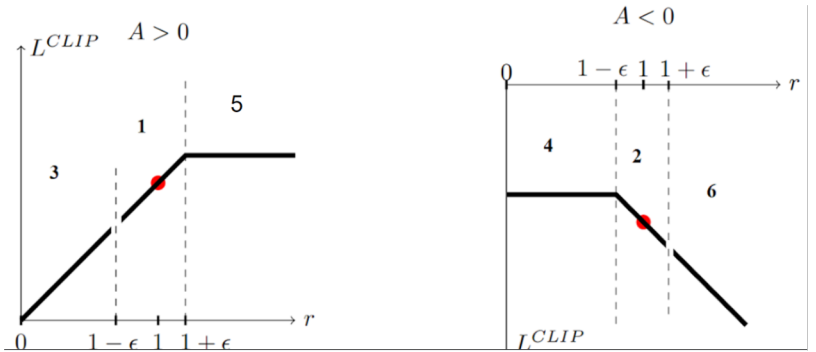

Shulman et al., Proximal Policy Optimization Algorithms, 2017
https://huggingface.co/blog/deep-rl-ppo
Table from "Towards Delivering a Coherent Self-Contained Explanation of Proximal Policy Optimization" by Daniel Bick

PPO Algorithm

$$L_t^{CLIP+VF+S}(\theta) = \hat{\mathbb{E}}_t\left[L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) + c_2 S[\pi_\theta](s_t)\right],$$

Shulman et al., Proximal Policy Optimization Algorithms, 2017

PPO Algorithm

$$L_t^{CLIP+VF+S}(\theta) = \hat{\mathbb{E}}_t\left[L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) + c_2 S[\pi_\theta](s_t)\right],$$

$$(V_\theta(s_t) - V_t^{\text{targ}})^2$$

Shulman et al., Proximal Policy Optimization Algorithms, 2017

PPO Algorithm

$$L_t^{CLIP+VF+S}(\theta) = \hat{\mathbb{E}}_t\left[L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) + c_2 S[\pi_\theta](s_t)\right],$$

$$(V_\theta(s_t) - V_t^{\text{targ}})^2$$

Entropy loss to encourage explorations

Shulman et al., Proximal Policy Optimization Algorithms, 2017

PPO Algorithm

$$L_t^{CLIP+VF+S}(\theta) = \hat{\mathbb{E}}_t\left[L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) + c_2 S[\pi_\theta](s_t)\right],$$

$$(V_\theta(s_t) - V_t^{\text{targ}})^2$$

Entropy loss to encourage explorations

for iteration=$1, 2, \ldots$ **do**
  for actor=$1, 2, \ldots, N$ **do**
    Run policy $\pi_{\theta_{\text{old}}}$ in environment for T timesteps
    Compute advantage estimates $\hat{A}_1, \ldots, \hat{A}_T$
  end for
  Optimize surrogate L wrt θ, with K epochs and minibatch size $M \leq NT$
  $\theta_{\text{old}} \leftarrow \theta$
end for

Shulman et al., Proximal Policy Optimization Algorithms, 2017

End of Digression:

A Brief Intro to Reinforcement Learning

RLHF objective

π: LLM policy
π_θ: base LLM
x: prompt
y: completion

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)} \left[r_\phi(x, y) \right] - \beta \mathbb{D}_{\mathrm{KL}} \left[\pi_\theta(y \mid x) \,||\, \pi_{\mathrm{ref}}(y \mid x) \right]$$

Optimize "reward" *inspired* ▲ by human preferences

▲ Constrain the model to not trust the reward too much (preferences are hard to model)

RLHF objective

π: LLM policy
π_θ: base LLM
x: prompt
y: completion

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi_\theta(y|x)} \left[r_\phi(x, y) \right] - \beta \mathbb{D}_{\mathrm{KL}} \left[\pi_\theta(y \mid x) \,||\, \pi_{\mathrm{ref}}(y \mid x) \right]$$

Optimize "reward" *inspired* ▲ by human preferences

▲ Constrain the model to not trust the reward too much (preferences are hard to model)

Primary questions:
1. How to implement reward: *r(x,y)*
2. How to optimize reward

How to Train the Reward Model?

Preference (reward) modeling

Can we just use supervised learning on scores?

- Assigning a scalar reward of how good a response is did not work
- Pairwise preferences are easy to collect and worked!

Key idea:

Probability $\propto$ reward

Chosen completion

Prompt

Score from optimal reward model

$$p^*(y_1 \succ y_2 \mid x) = \frac{\exp\left(r^*(x, y_1)\right)}{\exp\left(r^*(x, y_1)\right) + \exp\left(r^*(x, y_2)\right)}$$

Rejected completion

Bradley Terry model:
Estimate probability that a given pairwise preference is true

Reward Model

Model is sampled for the same prompt x to produce two outputs y_1 and y_2

We assume that human preferences are generated by some latent reward model $r^*(y, x)$

Bradley-Terry model: The human preference distribution p* can be written as:

$$p^*(y_1 \succ y_2 \mid x) = \frac{\exp\left(r^*(x, y_1)\right)}{\exp\left(r^*(x, y_1)\right) + \exp\left(r^*(x, y_2)\right)}.$$

- This is just softmax !
    - For the case of 2 choices y1 vs. y2, it can be cast as the sigmod (logistic) function !

Key IDEA: Use the human preference feedbacks to train a “sigmod-like” binary classifier; once this is done, keep the learned reward function $r(x,y)$ and throw away the classifier ! [Recall a similar trick was used for ”learning” the word-embeddings in Word2Vec]

Bradley and Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. 1952

Reward Modeling Loss

Dataset: D = x_i , y_{iw} , y_{il}, for i in range(N)

For prompt x_i and two generated outputs y_{iw} and y_{il} where w stands for winner (preferred label), l: loser

We model *the reward with a neural network $r_\phi(x, y)$ (often initialized with the LLM itself with a linear layer on top)* and train it by optimizing:

$$\mathcal{L}_R(r_\phi, \mathcal{D}) = -\mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}}\left[\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))\right]$$

In other words, minimize log loss to correctly classify Human Preferences induces a useful reward model that acts as a proxy of true Human reward !

Ziegler et al, Fine-tuning language models from human preferences, 2020.

More details on how to train a Reward Model

Reward Modelling Phase: In the second phase the SFT model is prompted with prompts x to produce pairs of answers $(y_1, y_2) \sim \pi^{\text{SFT}}(y \mid x)$. These are then presented to human labelers who express preferences for one answer, denoted as $y_w \succ y_l \mid x$ where y_w and y_l denotes the preferred and dispreferred completion amongst (y_1, y_2) respectively. The preferences are assumed to be generated by some latent reward model $r^*(y, x)$, which we do not have access to. There are a number of approaches used to model preferences, the Bradley-Terry (BT) [5] model being a popular choice (although more general Plackett-Luce ranking models [32, 23] are also compatible with the framework if we have access to several ranked answers). The BT model stipulates that the human preference distribution p^* can be written as:

$$p^*(y_1 \succ y_2 \mid x) = \frac{\exp\left(r^*(x, y_1)\right)}{\exp\left(r^*(x, y_1)\right) + \exp\left(r^*(x, y_2)\right)}. \tag{1}$$

Assuming access to a static dataset of comparisons $\mathcal{D} = \left\{x^{(i)}, y_w^{(i)}, y_l^{(i)}\right\}_{i=1}^N$ sampled from p^*, we can parametrize a reward model $r_\phi(x, y)$ and estimate the parameters via maximum likelihood. Framing the problem as a binary classification we have the negative log-likelihood loss:

$$\mathcal{L}_R(r_\phi, \mathcal{D}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}}\left[\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))\right] \tag{2}$$

where σ is the logistic function. In the context of LMs, the network $r_\phi(x, y)$ is often initialized from the SFT model $\pi^{\text{SFT}}(y \mid x)$ with the addition of a linear layer on top of the final transformer layer that produces a single scalar prediction for the reward value [51]. To ensure a reward function with lower variance, prior works normalize the rewards, such that $\mathbb{E}_{x,y \sim \mathcal{D}}\left[r_\phi(x, y)\right] = 0$ for all x.

Source: Rafael Rafailov et al, Direct Preference Optimization: Your Language model is secretly a Reward Model, 2023.

Dataset for Training Reward Models

From Human Feedback

From AI Feedback (e.g. LLM-as-the-judge, or Constitutional AI)

From other sources of ground truth, e.g. such as a simulator, a code interpreter, etc, for prompt related to Math-problem-solving or Programming tasks.

The overall RLHF Optimization

- With data and the reward model, we can now tune the language model !
- In principle, any “policy gradient algorithm would work.

BUT

- In practice, OpenAI used Proximal Policy Optimization (PPO), and others follow, (at least initially) and their approach has become synonymous with the flavor of RLHF for LLMs.
- Policy Gradient methods update model parameters to maximize expected reward.
    - PPO in particular clips objective in a range to ensure Stable Updates !

PPO at a Conceptual Level

A brief (and high level) intro to the various ideas in PPO..

Attempt 1: Policy gradients (variances are too high)

$$\nabla_\theta E_{p_\theta}[R(z)] = E_{p_\theta}[\mathrm{R(z)}\nabla_\theta \log p_\theta\ (z)]$$

Attempt 2: TRPO (Linearize the problem around the current policy)

$$\begin{aligned} \underset{\theta}{\text{maximize}} \quad & \hat{\mathbb{E}}_t\left[\frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}\hat{A}_t\right] \\ \text{subject to} \quad & \hat{\mathbb{E}}_t[\mathrm{KL}[\pi_{\theta_{\text{old}}}(\cdot \mid s_t), \pi_\theta(\cdot \mid s_t)]] \leq \delta. \end{aligned}$$

Attempt 3: PPO (Clip the ratios at some eps)

$$L(s, a, \theta_k, \theta) = \min\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}A^{\pi_{\theta_k}}(s,a),\ \ \mathrm{clip}\left(\frac{\pi_\theta(a|s)}{\pi_{\theta_k}(a|s)}, 1-\epsilon, 1+\epsilon\right)A^{\pi_{\theta_k}}(s,a)\right)$$

PPO at a Conceptual Level

The objective is basically all there is to it at a conceptual level

Algorithm 1 PPO-Clip

1: Input: initial policy parameters θ_0, initial value function parameters ϕ_0
2: **for** $k = 0, 1, 2, \ldots$ **do**
3: Collect set of trajectories $\mathcal{D}_k = \{\tau_i\}$ by running policy $\pi_k = \pi(\theta_k)$ in the environment.
4: Compute rewards-to-go $\hat{R}_t$.
5: Compute advantage estimates, $\hat{A}_t$ (using any method of advantage estimation) based on the current value function V_{ϕ_k}.
6: Update the policy by maximizing the PPO-Clip objective:

$$\theta_{k+1} = \arg\max_{\theta} \frac{1}{|\mathcal{D}_k|T} \sum_{\tau \in \mathcal{D}_k} \sum_{t=0}^{T} \min\left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_k}(a_t|s_t)} A^{\pi_{\theta_k}}(s_t, a_t), \quad g(\epsilon, A^{\pi_{\theta_k}}(s_t, a_t)) \right),$$

typically via stochastic gradient ascent with Adam.
7: Fit value function by regression on mean-squared error:

$$\phi_{k+1} = \arg\min_{\phi} \frac{1}{|\mathcal{D}_k|T} \sum_{\tau \in \mathcal{D}_k} \sum_{t=0}^{T} \left(V_\phi(s_t) - \hat{R}_t \right)^2,$$

typically via some gradient descent algorithm.
8: **end for**

[From - https://spinningup.openai.com/en/latest/algorithms/ppo.html]

Refer to

https://www.youtube.com/watch?v=xT4jxQUl0X8;

https://www.youtube.com/watch?v=8jtAzxUwDj0

for a brief but intuitive explanation on Policy Gradient optimization, PPO, GRPO by Julia Turc

PPO in Language Modeling

From InstructGPT

Reinforcement learning (RL). Once again following Stiennon et al. (2020), we fine-tuned the SFT model on our environment using PPO (Schulman et al., 2017). The environment is a bandit environment which presents a random customer prompt and expects a response to the prompt. Given the prompt and response, it produces a reward determined by the reward model and ends the episode. In addition, we add a per-token KL penalty from the SFT model at each token to mitigate over-optimization of the reward model. The value function is initialized from the RM. We call these models "PPO."

We also experiment with mixing the pretraining gradients into the PPO gradients, in order to fix the performance regressions on public NLP datasets. We call these models "PPO-ptx." We maximize the following combined objective function in RL training:

$$\begin{aligned}\text{objective}\,(\phi) =& E_{(x,y)\sim D_{\pi_\phi^{\mathrm{RL}}}}\left[r_\theta(x,y) - \beta \log\left(\pi_\phi^{\mathrm{RL}}(y \mid x)/\pi^{\mathrm{SFT}}(y \mid x)\right)\right] + \\ & \gamma E_{x\sim D_{\text{pretrain}}}\left[\log(\pi_\phi^{\mathrm{RL}}(x))\right]\end{aligned} \tag{2}$$

where π_ϕ^{RL} is the learned RL policy, π^{SFT} is the supervised trained model, and D_{pretrain} is the pretraining distribution. The KL reward coefficient, β, and the pretraining loss coefficient, γ, control the strength of the KL penalty and pretraining gradients respectively. For "PPO" models, γ is set to 0. Unless otherwise specified, in this paper InstructGPT refers to the PPO-ptx models.

More Details and Background from Stiennon

From "Learning to summarize from human feedback"

Reward models. To train our reward models, we start from a supervised baseline, as described above, then add a randomly initialized linear head that outputs a scalar value. We train this model to predict which summary $y \in \{y_0, y_1\}$ is better as judged by a human, given a post x. If the summary preferred by the human is y_i, we can write the RM loss as:

$$\text{loss}(r_\theta) = -E_{(x,y_0,y_1,i)\sim D}[\log(\sigma(r_\theta(x, y_i) - r_\theta(x, y_{1-i})))]$$

where $r_\theta(x, y)$ is the scalar output of the reward model for post x and summary y with parameters θ, and D is the dataset of human judgments. At the end of training, we normalize the reward model outputs such that the reference summaries from our dataset achieve a mean score of 0.

Human feedback policies. We want to use the reward model trained above to train a policy that generates higher-quality outputs as judged by humans. We primarily do this using reinforcement learning, by treating the output of the reward model as a reward for the entire summary that we maximize with the PPO algorithm [58], where each time step is a BPE token.[8] We initialize our policy to be the model fine-tuned on Reddit TL;DR. Importantly, we include a term in the reward that penalizes the KL divergence between the learned RL policy π_ϕ^{RL} with parameters ϕ and this original supervised model π^{SFT}, as previously done in [25]. The full reward R can be written as:

$$R(x, y) = r_\theta(x, y) - \beta \log[\pi_\phi^{\text{RL}}(y|x)/\pi^{\text{SFT}}(y|x)]$$

This KL term serves two purposes. First, it acts as an entropy bonus, encouraging the policy to explore and deterring it from collapsing to a single mode. Second, it ensures the policy doesn't learn to produce outputs that are too different from those that the reward model has seen during training.

For the PPO value function, we use a Transformer with completely separate parameters from the policy. This prevents updates to the value function from partially destroying the pretrained policy early in training (see ablation in Appendix G.1). We initialize the value function to the parameters of the reward model. In our experiments, the reward model, policy, and value function are the same size.

PPO in the Language Model Setting

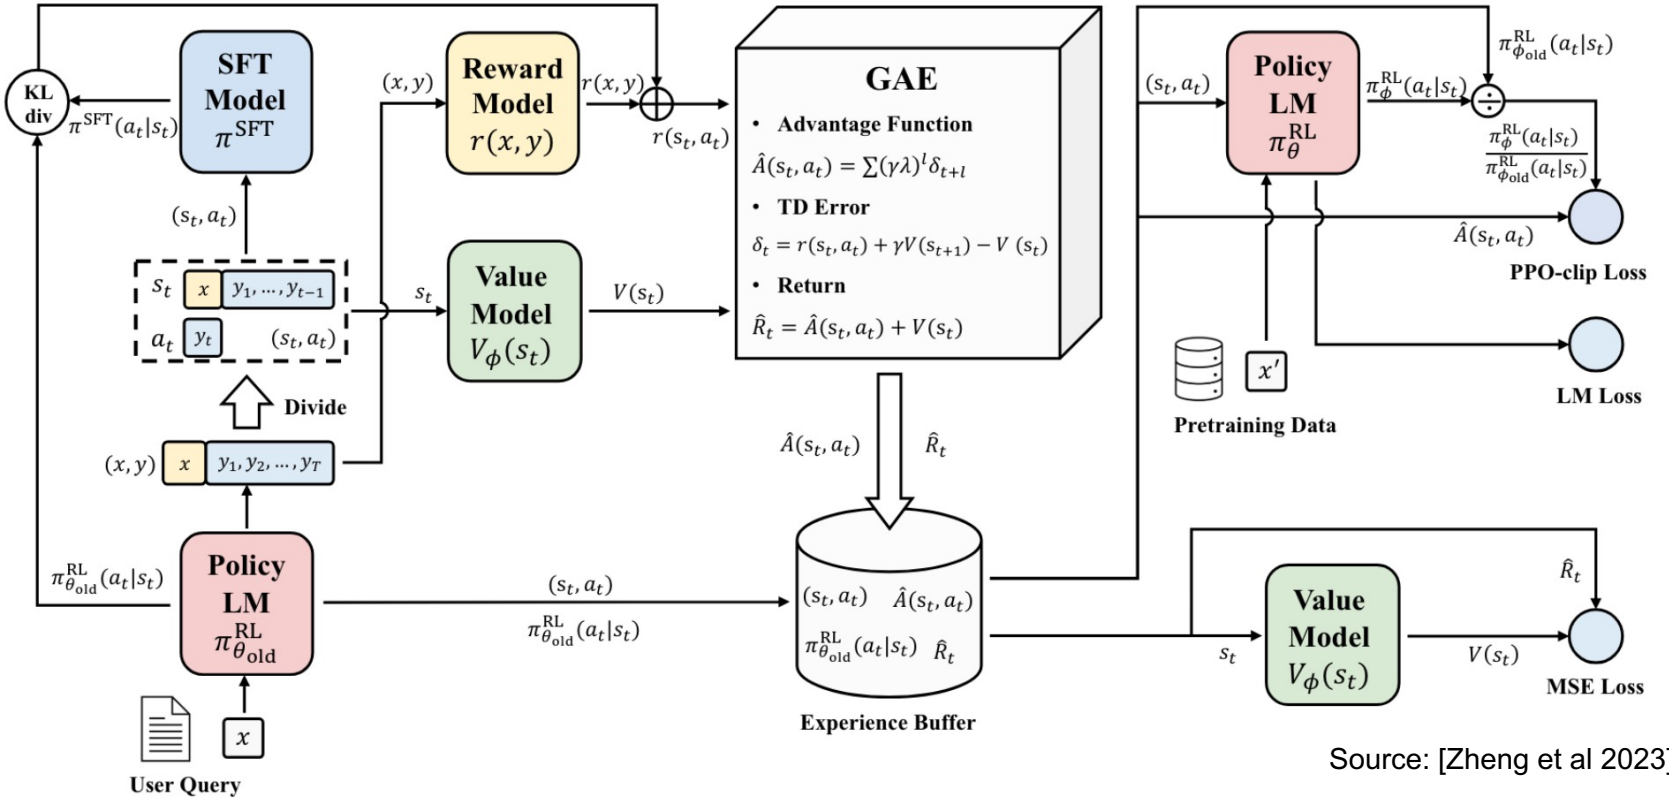

Source: [Zheng et al 2023]

- Similar to the RL formulation: Actions operate over tokens, large dense reward at the very end operating on full sequence.

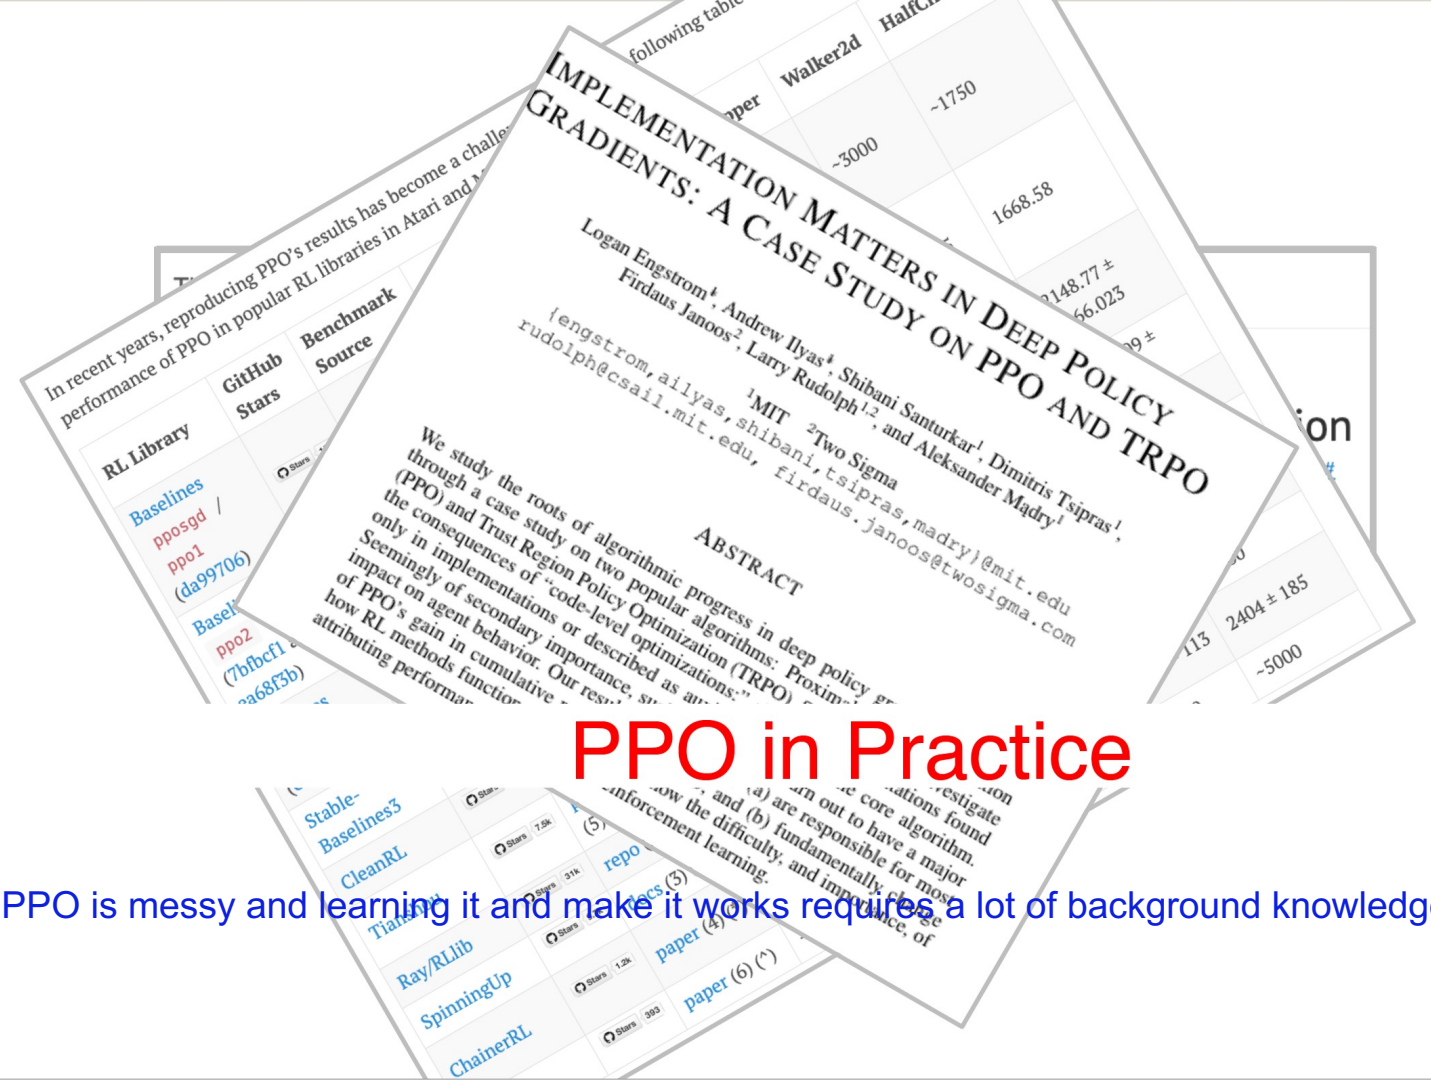

PPO in Practice

In Reality, PPO is messy and learning it and make it works requires a lot of background knowledge in RL !

Reality of RLHF via PPO

- ★ RLFH in 2024: PPO
  - 🤯 results are incredible
  - ❌ implementation is complicated
  - ❌ GPU memory (fit 4 models) 🤯

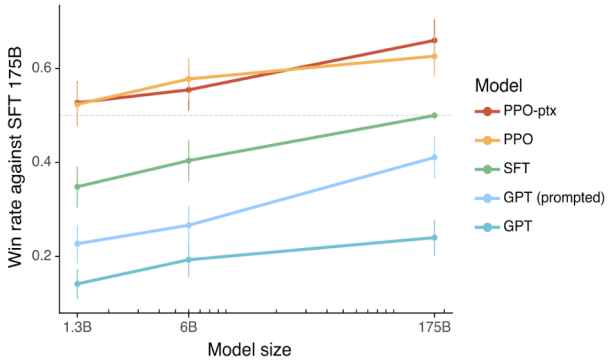

Figure 1: Human evaluations of various models on our API prompt distribution, evaluated by how often outputs from each model were preferred to those from the 175B SFT model. Our InstructGPT models (PPO-ptx) as well as its variant trained without pretraining mix (PPO) significantly outperform the GPT-3 baselines (GPT, GPT prompted); outputs from our 1.3B PPO-ptx model are preferred to those from the 175B GPT-3. Error bars throughout the paper are 95% confidence intervals.

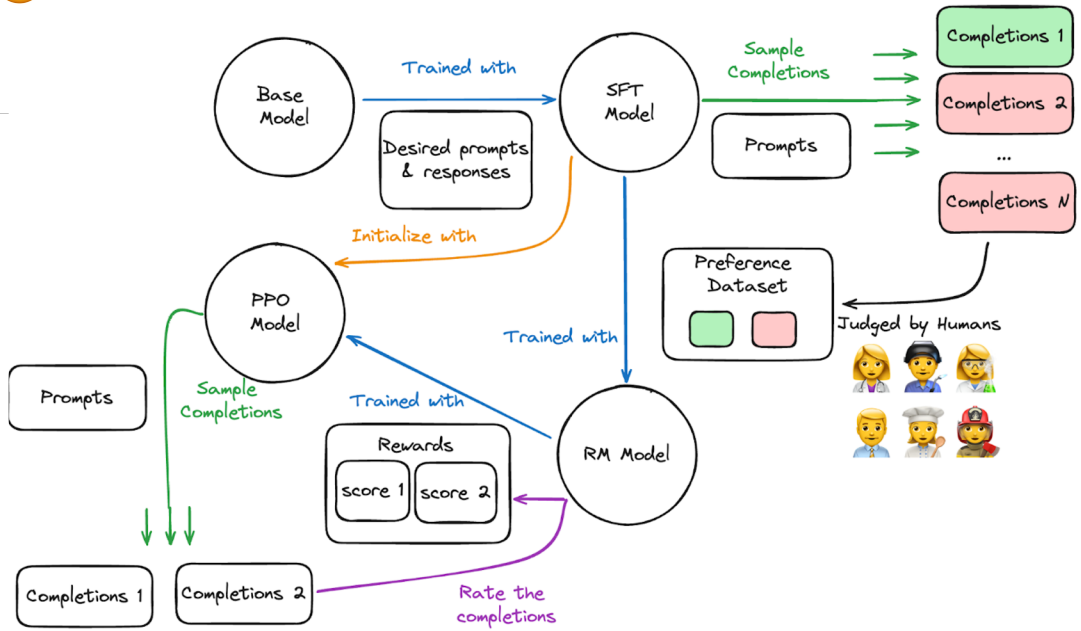

Things to watch out for in RLHF

Overoptimization / overfitting on the reward

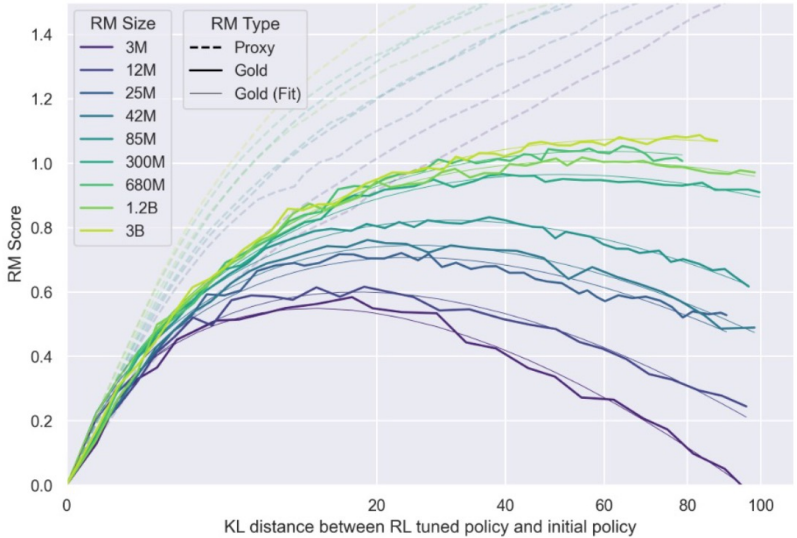

Mode collapse / entropy

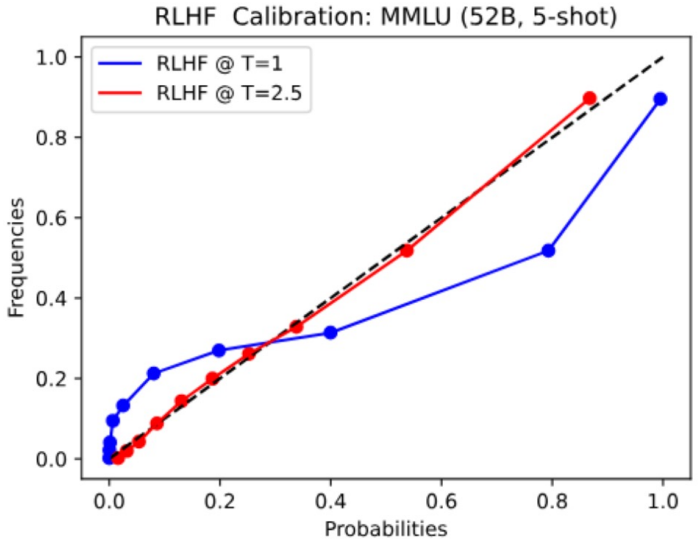

Things to watch out for in RLHF - Overoptimization

Across many different RLHF-style optimizers. Optimizing for reward overfits past a point

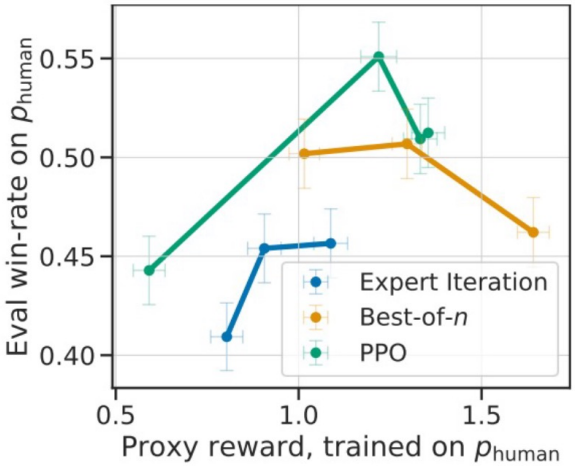

(a) Human preferences

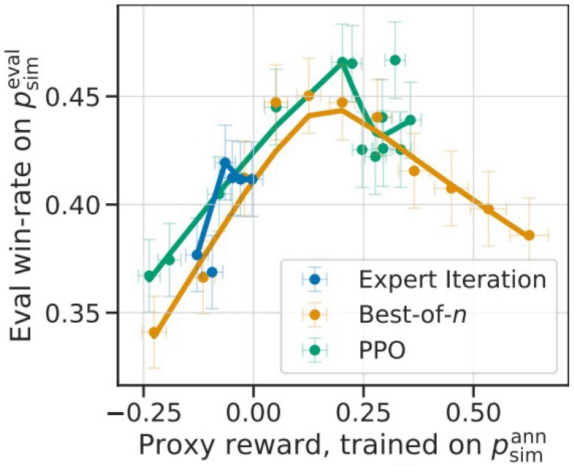

(b) AlpacaFarm

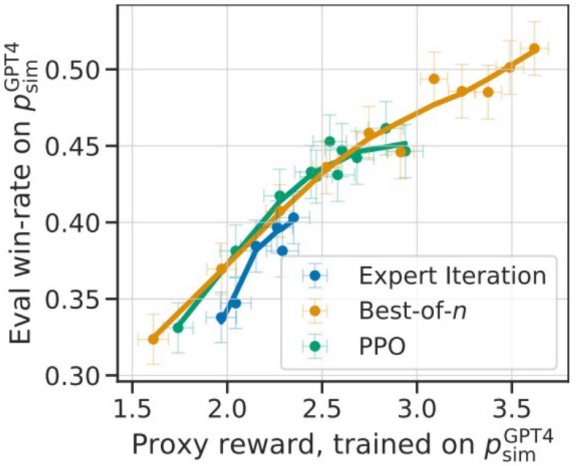

(c) Single-prompt GPT-4

Holds true for human pref (left), noisy LM pref (mid) but not noiseless LM pref (right)

Can we get rid of PPO ?

- The idea of RLHF was around since 2021, but did not really catch on (except in industry labs) until mid-2023
- Why ? RL, and PPO in particular, is notoriously difficult to get right.
- What if we could skip reward modeling and update the model with Preference data directly ?

Can we get rid of PPO ?

Question for many: can we avoid doing any ‘RL’ ? (i.e. on-policy RL algorithms)

Some reasonable stuff people thought about

- Train the model with a control token – (SFT on the pairs, prepend [GOOD] to chosen, [BAD] to not chosen)
- Train the model on only preferred output
- Train a reward model, get LM outputs, train on the preferred output
- Train a reward model, get 1024 LM outputs, take the best one.

Can we get rid of PPO ?

Most of these baselines turn out to just work worse than PPO on instruction-tuning

Method	Simulated win-rate (%)	Human win-rate (%)
GPT-4	79.0 ± 1.4	69.8 ± 1.6
ChatGPT	61.4 ± 1.7	52.9 ± 1.7
PPO	46.8 ± 1.8	55.1 ± 1.7
Best-of-n	45.0 ± 1.7	50.7 ± 1.8
Expert Iteration	41.9 ± 1.7	45.7 ± 1.7
SFT 52k (Alpaca 7B)	39.2 ± 1.7	40.7 ± 1.7
SFT 10k	36.7 ± 1.7	44.3 ± 1.7
Binary FeedME	36.6 ± 1.7	37.9 ± 1.7
Quark	35.6 ± 1.7	-
Binary Reward Conditioning	32.4 ± 1.6	-
Davinci001	24.4 ± 1.5	32.5 ± 1.6
LLaMA 7B	11.3 ± 1.1	6.5 ± 0.9

What if we just use gradient ascent on this equation?

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right] - \beta\mathbb{D}_{\mathrm{KL}}\left[\pi_\theta(y \mid x) \,||\, \pi_{\mathrm{ref}}(y \mid x)\right]$$

What if we just use gradient ascent on this equation?

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right] - \beta\mathbb{D}_{\mathrm{KL}}\left[\pi_\theta(y \mid x) \,||\, \pi_{\mathrm{ref}}(y \mid x)\right]$$

The answer, with some math, is:
Direct Preference Optimization (DPO)

Released on May 29th 2023

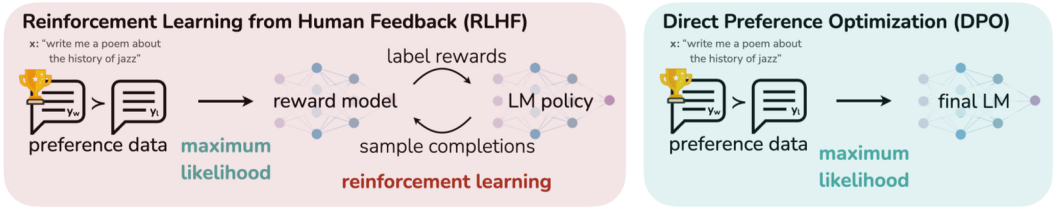

Figure 1: **DPO optimizes for human preferences while avoiding reinforcement learning.** Existing methods for fine-tuning language models with human feedback first fit a reward model to a dataset of prompts and human preferences over pairs of responses, and then use RL to find a policy that maximizes the learned reward. In contrast, DPO directly optimizes for the policy best satisfying the preferences with a simple classification objective, fitting an *implicit* reward model whose corresponding optimal policy can be extracted in closed form.

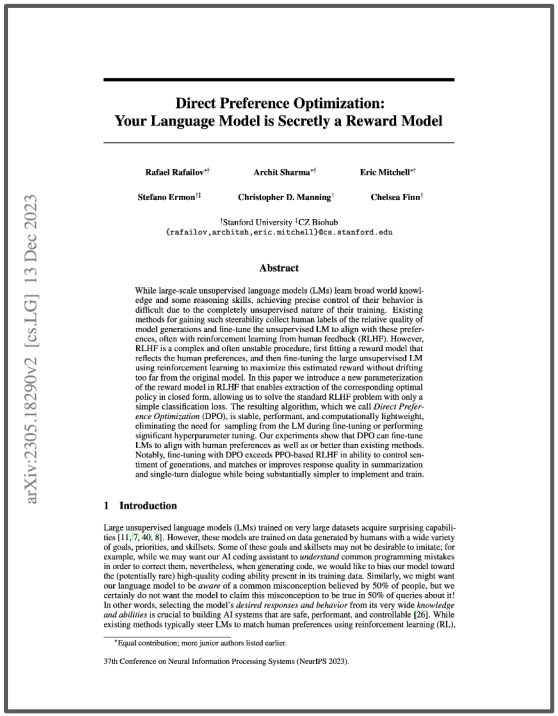

arXiv:2305.18290v2 [cs.LG] 13 Dec 2023

Direct Preference Optimization:
Your Language Model is Secretly a Reward Model

Rafael Rafailov*† Archit Sharma*† Eric Mitchell*†
Stefano Ermon†‡ Christopher D. Manning† Chelsea Finn†
†Stanford University ‡CZ Biohub
{rafailov,architsh,eric.mitchell}@cs.stanford.edu

Abstract

While large-scale unsupervised language models (LMs) learn broad world knowledge and some reasoning skills, achieving precise control of their behavior is difficult due to the completely unsupervised nature of their training. Existing methods for gaining such steerability collect human labels of the relative quality of model generations and fine-tune the unsupervised LM to align with these preferences, often with reinforcement learning from human feedback (RLHF). However, RLHF is a complex and often unstable procedure, first fitting a reward model that reflects the human preferences, and then fine-tuning the large unsupervised LM using reinforcement learning to maximize this estimated reward without drifting too far from the original model. In this paper we introduce a new parameterization of the reward model in RLHF that enables extraction of the corresponding optimal policy in closed form, allowing us to solve the standard RLHF problem with only a simple classification loss. The resulting algorithm, which we call *Direct Preference Optimization* (DPO), is stable, performant, and computationally lightweight, eliminating the need for sampling from the LM during fine-tuning or performing significant hyperparameter tuning. Our experiments show that DPO can fine-tune LMs to align with human preferences as well as or better than existing methods. Notably, fine-tuning with DPO exceeds PPO-based RLHF in ability to control sentiment of generations, and matches or improves response quality in summarization and single-turn dialogue while being substantially simpler to implement and train.

1 Introduction

Large unsupervised language models (LMs) trained on very large datasets acquire surprising capabilities [11, 7, 40, 8]. However, these models are trained on data generated by humans with a wide variety of goals, priorities, and skillsets. Some of these goals and skillsets may not be desirable to imitate; for example, while we may want our AI coding assistant to *understand* common programming mistakes in order to correct them, nevertheless, when generating code, we would like to bias our model toward the (potentially rare) high-quality coding ability present in its training data. Similarly, we might want our language model to be *aware* of a common misconception believed by 50% of people, but we certainly do not want the model to claim this misconception to be true in 50% of queries about it! In other words, selecting the model's *desired responses and behavior* from its very wide *knowledge and abilities* is crucial to building AI systems that are safe, performant, and controllable [26]. While existing methods typically steer LMs to match human preferences using reinforcement learning (RL),

*Equal contribution; more junior authors listed earlier.

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

Rafailov, Sharma, Mitchell et al. 2023

Outstanding Paper Runner-up in NeurIPS 2023

DPO – RLHF without Tears ?

Try to simplify PPO by..

- Getting rid of the reward model
- Getting rid of any on-policy stuff (rollouts, outer loops etc)

Instead

- Take gradient steps on log-loss of good stuff
- Take negative gradient steps on bad stuff (appropriately weighted).

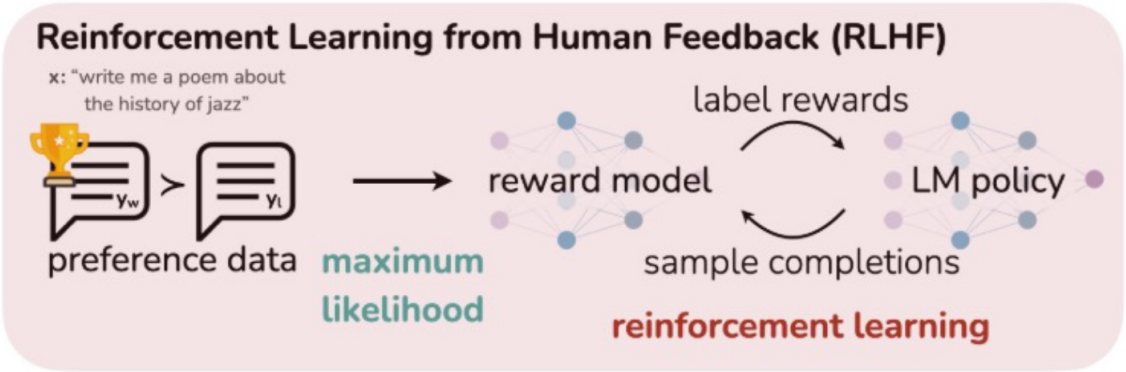

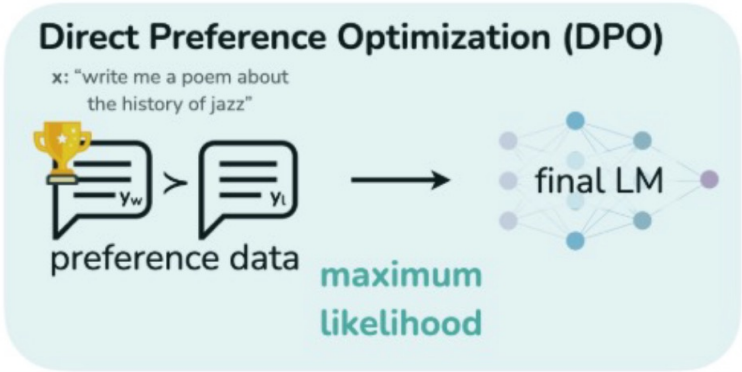

DPO – Your language model is secretly a reward model

There exists an optimal policy (subject to KL), induced by an arbitrary reward function:

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right] - \beta\mathbb{D}_{\mathrm{KL}}\left[\pi_\theta(y \mid x) \,||\, \pi_{\mathrm{ref}}(y \mid x)\right]$$

** Closed-form optimal policy

$$\pi_r(y \mid x) = \frac{1}{Z(x)}\pi_{\mathrm{ref}}(y \mid x)\exp\left(\frac{1}{\beta}r(x,y)\right)$$

$Z(x)$ crossed out because it's just a normalizing term to make π a proper distribution.

Details on deriving ** in the DPO objective (1/2)

In this appendix, we will derive Eq. 4. Analogously to Eq. 3, we optimize the following objective:

$$\max_{\pi} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi}\left[r(x,y)\right] - \beta\mathbb{D}_{\mathrm{KL}}\left[\pi(y|x)||\pi_{\mathrm{ref}}(y|x)\right] \tag{11}$$

under any reward function $r(x, y)$, reference model π_{ref} and a general non-parametric policy class. We now have:

$$\begin{aligned}
&\max_{\pi} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi}\left[r(x,y)\right] - \beta\mathbb{D}_{\mathrm{KL}}\left[\pi(y|x) \,||\, \pi_{\mathrm{ref}}(y|x)\right] \\
&\qquad = \max_{\pi} \mathbb{E}_{x\sim\mathcal{D}}\mathbb{E}_{y\sim\pi(y|x)}\left[r(x,y) - \beta\log\frac{\pi(y|x)}{\pi_{\mathrm{ref}}(y|x)}\right] \\
&\qquad = \min_{\pi} \mathbb{E}_{x\sim\mathcal{D}}\mathbb{E}_{y\sim\pi(y|x)}\left[\log\frac{\pi(y|x)}{\pi_{\mathrm{ref}}(y|x)} - \frac{1}{\beta}r(x,y)\right] \\
&\qquad = \min_{\pi} \mathbb{E}_{x\sim\mathcal{D}}\mathbb{E}_{y\sim\pi(y|x)}\left[\log\frac{\pi(y|x)}{\frac{1}{Z(x)}\pi_{\mathrm{ref}}(y|x)\exp\left(\frac{1}{\beta}r(x,y)\right)} - \log Z(x)\right]
\end{aligned} \tag{12}$$

where we have partition function:

$$Z(x) = \sum_{y} \pi_{\mathrm{ref}}(y|x)\exp\left(\frac{1}{\beta}r(x,y)\right).$$

Source: Rafael Rafailov et al, Direct Preference Optimization: Your Language model is secretly a Reward Model, 2023.

Details on deriving ** the DPO objective (2/2)

Note that the partition function is a function of only x and the reference policy π_{ref}, but does not depend on the policy π. We can now define

$$\pi^*(y|x) = \frac{1}{Z(x)}\pi_{\text{ref}}(y|x)\exp\left(\frac{1}{\beta}r(x,y)\right),$$

which is a valid probability distribution as $\pi^*(y|x) \geq 0$ for all y and $\sum_y \pi^*(y|x) = 1$. Since $Z(x)$ is not a function of y, we can then re-organize the final objective in Eq 12 as:

$$\min_\pi \mathbb{E}_{x\sim\mathcal{D}}\left[\mathbb{E}_{y\sim\pi(y|x)}\left[\log\frac{\pi(y|x)}{\pi^*(y|x)}\right] - \log Z(x)\right] = \tag{13}$$

$$\min_\pi \mathbb{E}_{x\sim\mathcal{D}}\left[\mathbb{D}_{\text{KL}}(\pi(y|x)\ ||\ \pi^*(y|x)) - \log Z(x)\right] \tag{14}$$

Now, since $Z(x)$ does not depend on π, the minimum is achieved by the policy that minimizes the first KL term. Gibbs' inequality tells us that the KL-divergence is minimized at 0 if and only if the two distributions are identical. Hence we have the optimal solution:

$$\pi(y|x) = \pi^*(y|x) = \frac{1}{Z(x)}\pi_{\text{ref}}(y|x)\exp\left(\frac{1}{\beta}r(x,y)\right) \tag{15}$$

for all $x \in \mathcal{D}$. This completes the derivation.

Source: Rafael Rafailov et al, Direct Preference Optimization: Your Language model is secretly a Reward Model, 2023.

DPO – Your language model is secretly a reward model

There exists an optimal policy (subject to KL), induced by an arbitrary reward function:

$$\max_{\pi_\theta} \mathbb{E}_{x\sim\mathcal{D},y\sim\pi_\theta(y|x)}\left[r_\phi(x,y)\right] - \beta\mathbb{D}_{\text{KL}}\left[\pi_\theta(y \mid x) \,||\, \pi_{\text{ref}}(y \mid x)\right]$$

Closed-form optimal policy

$$\pi_r(y \mid x) = \frac{1}{Z(x)}\pi_{\text{ref}}(y \mid x)\exp\left(\frac{1}{\beta}r(x,y)\right)$$

Rearrange terms

$$r(x,y) = \beta\log\frac{\pi_r(y \mid x)}{\pi_{\text{ref}}(y \mid x)} + \beta\log Z(x)$$

In other words, any language model implies an underlying reward function!

DPO – Train a LLM like a Preference Classifier

Recall Bradley-Terry model:

$$p^*(y_1 \succ y_2 \mid x) = \frac{\exp\left(r^*(x, y_1)\right)}{\exp\left(r^*(x, y_1)\right) + \exp\left(r^*(x, y_2)\right)}$$

$$r(x, y) = \beta \log \frac{\pi_r(y \mid x)}{\pi_{\text{ref}}(y \mid x)} + \beta \log Z(x)$$

Plug into Bradley-Terry

$$p^*(y_1 \succ y_2 \mid x) = \frac{1}{1 + \exp\left(\beta \log \frac{\pi^*(y_2|x)}{\pi_{\text{ref}}(y_2|x)} - \beta \log \frac{\pi^*(y_1|x)}{\pi_{\text{ref}}(y_1|x)}\right)}$$

Literally train like a binary classifier

$$\mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(y_w \mid x)}{\pi_{\text{ref}}(y_w \mid x)} - \beta \log \frac{\pi_\theta(y_l \mid x)}{\pi_{\text{ref}}(y_l \mid x)}\right)\right]$$

What does the DPO update do ?

- Reduce to “+ve gradient on Good ; -ve gradient on Bad”

What does the DPO update do? For a mechanistic understanding of DPO, it is useful to analyze the gradient of the loss function $\mathcal{L}_{\text{DPO}}$. The gradient with respect to the parameters θ can be written as:

$$\nabla_\theta \mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) =$$

$$-\beta \mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}}\Big[\underbrace{\sigma(\hat{r}_\theta(x, y_l) - \hat{r}_\theta(x, y_w))}_{\text{higher weight when reward estimate is wrong}}\Big[\underbrace{\nabla_\theta \log \pi(y_w \mid x)}_{\text{increase likelihood of } y_w} - \underbrace{\nabla_\theta \log \pi(y_l \mid x)}_{\text{decrease likelihood of } y_l}\Big]\Big],$$

- Scaled by “prediction error” of the Implied Reward model.

Implementation of DPO

- Simple implementation in 10 lines of code:

```
pi_logratios = policy_chosen_logps - policy_rejected_logps
if reference_free:
    ref_logratios = 0
else:
    ref_logratios = reference_chosen_logps - reference_rejected_logps

pi_logratios = pi_logratios.to(self.accelerator.device)
ref_logratios = ref_logratios.to(self.accelerator.device)
logits = pi_logratios - ref_logratios

# The beta is a temperature parameter for the DPO loss, typically something in the range of 0.1 to 0.5.
# We ignore the reference model as beta -> 0. The label_smoothing parameter encodes our uncertainty about the labels and
# calculates a conservative DPO loss.
if self.loss_type == "sigmoid":
    losses = (
        -F.logsigmoid(self.beta * logits) * (1 - self.label_smoothing)
        - F.logsigmoid(-self.beta * logits) * self.label_smoothing
    )
```

Source: https://github.com/huggingface/trl/blob/ef441ea02818a08b5cd26536446f19c3678dd58b/trl/trainer/dpo_trainer.py

DPO core facts

1. Extremely **simple** to implement.
2. **Scales nicely** with existing distributed training libraries.
3. Trains an implicit reward function.

DPO is easier to iterate on, but slightly underperforms online RL methods on absolute potential.

```python
import torch.nn.functional as F

def dpo_loss(pi_logps, ref_logps, yw_idxs, yl_idxs, beta):
    """
    pi_logps: policy logprobs, shape (B,)
    ref_logps: reference model logprobs, shape (B,)
    yw_idxs: preferred completion indices in [0, B-1], shape (T,)
    yl_idxs: dispreferred completion indices in [0, B-1], shape (T,)
    beta: temperature controlling strength of KL penalty
    Each pair of (yw_idxs[i], yl_idxs[i]) represents the
    indices of a single preference pair.
    """

    pi_yw_logps, pi_yl_logps = pi_logps[yw_idxs], pi_logps[yl_idxs]
    ref_yw_logps, ref_yl_logps = ref_logps[yw_idxs], ref_logps[yl_idxs]

    pi_logratios = pi_yw_logps - pi_yl_logps
    ref_logratios = ref_yw_logps - ref_yl_logps

    losses = -F.logsigmoid(beta * (pi_logratios - ref_logratios))
    rewards = beta * (pi_logps - ref_logps).detach()

    return losses, rewards
```

Example code.
Rafailov, Sharma, Mitchell et al. 2023

DPO model proliferation

DPO gave us:

1. A **faster, easier to implement baseline** for iterating on post-training.
2. A **massive adoption/ proliferation of preference finetuning** models and methods in the open.

Continue through today.

DPO model proliferation

HuggingFace H4's **Zephyr Beta**

- First model to make a splash with DPO!
- Fine-tune of Mistral 7B with UltraFeedback dataset.
- Discovered weird low learning rates that are now standard (~5e-7).

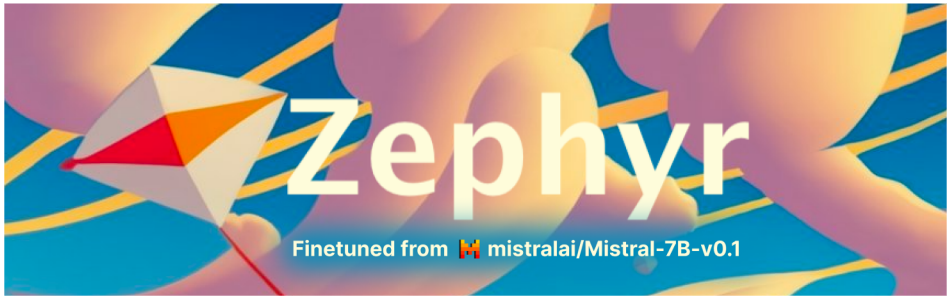

UltraFeedback: https://arxiv.org/abs/2310.01377
Model: https://huggingface.co/HuggingFaceH4/zephyr-7b-beta

DPO model proliferation

Allen AI's **Tülu 2 70B**

- First to scale DPO to 70B parameters.
- State-of-the-art open model on external benchmarks.
- Open models began to match and surpass GPT-4.

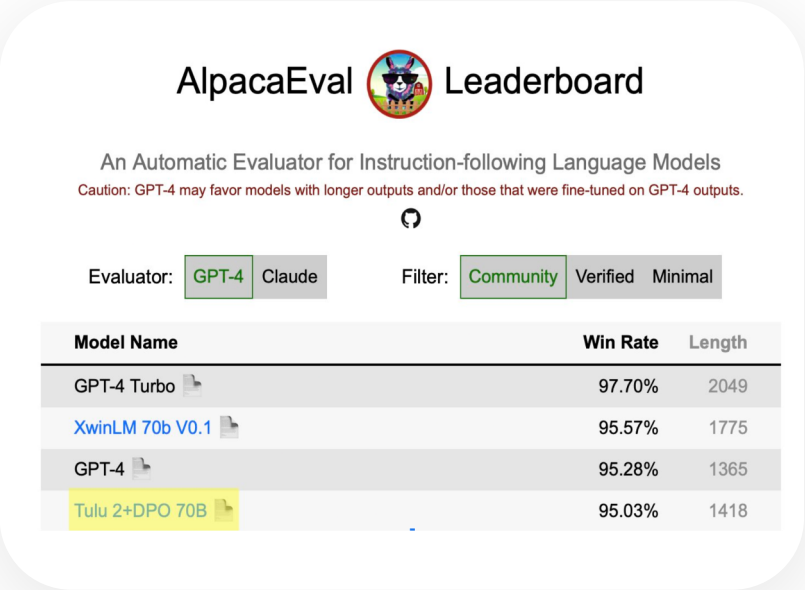

Ivison et al. 2023. Tulu 2.

DPO models proliferation

T	Model	Average	ARC	HellaSwag	MMLU	TruthfulQA	Winogrande	GSM8K
🟦	udkai/Turdus	74.66	73.38	88.56	64.52	67.11	86.66	67.7
🟦	fblgit/UNA-TheBeagle-7b-v1	73.87	73.04	88	63.48	69.85	82.16	66.72
🟦	argilla/distilabeled-Marcoro14-7B-slerp	73.63	70.73	87.47	65.22	65.1	82.08	71.19
🟦	mlabonne/NeuralMarcoro14-7B	73.57	71.42	87.59	64.84	65.64	81.22	70.74
🔶	abideen/NexoNimbus-7B	73.5	70.82	87.86	64.69	62.43	84.85	70.36
🟦	Neuronovo/neuronovo-7B-v0.2	73.44	73.04	88.32	65.15	71.02	80.66	62.47
🟦	argilla/distilabeled-Marcoro14-7B-slerp-full	73.4	70.65	87.55	65.33	64.21	82	70.66
🟦	CultriX/MistralTrix-v1	73.39	72.27	88.33	65.24	70.73	80.98	62.77
🟦	ryandt/MusingCaterpillar	73.33	72.53	88.34	65.26	70.93	80.66	62.24
🟦	Neuronovo/neuronovo-7B-v0.3	73.29	72.7	88.26	65.1	71.35	80.9	61.41
🟦	CultriX/MistralTrixTest	73.17	72.53	88.4	65.22	70.77	81.37	60.73
🔶	samir-fama/SamirGPT-v1	73.11	69.54	87.04	65.3	63.37	81.69	71.72
🔶	SanjiWatsuki/Lelantos-DPO-7B	73.09	71.08	87.22	64	67.77	80.03	68.46

DPO
DPO (& UNA)
DPO
DPO
Merge (of DPO models)
DPO
DPO
DPO
DPO
No info but prob DPO, given
Merge (incl. DPO)
DPO

From Chris Manning – Most ‘top open-source’ RLHF models are DPO’d

Summary on Alignment via DPO

★ RLFH in 2024: DPO

✅ results are very good

✅ Implementation is easy

✅ Fits in GPU memory well (2 models)

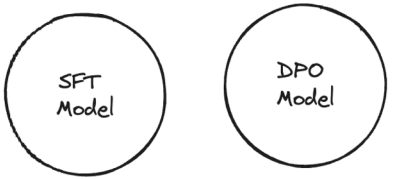

$$\nabla_\theta \mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) =$$
$$-\beta \mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}}\Big[\underbrace{\sigma(\hat{r}_\theta(x,y_l)-\hat{r}_\theta(x,y_w))}_{\text{higher weight when reward estimate is wrong}}\Big[\underbrace{\nabla_\theta \log \pi(y_w \mid x)}_{\text{increase likelihood of } y_w} - \underbrace{\nabla_\theta \log \pi(y_l \mid x)}_{\text{decrease likelihood of } y_l}\Big]\Big],$$

where $\hat{r}_\theta(x,y) = \beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)}$ is the reward implicitly defined by the language model π_θ and refer-

ZEPHYR: DIRECT DISTILLATION OF LM ALIGNMENT

Lewis Tunstall,* Edward Beeching,* Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf
The H4 (Helpful, Honest, Harmless, Huggy) Team
https://huggingface.co/HuggingFaceH4
lewis@huggingface.co

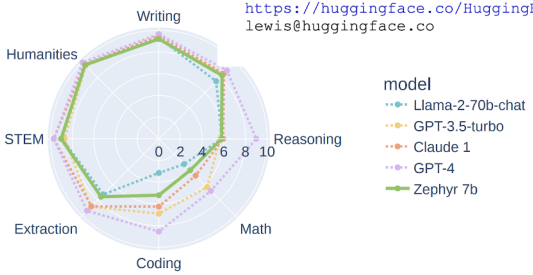

Direct Preference Optimization: Your Language Model is Secretly a Reward Model

Rafael Rafailov*† **Archit Sharma***† **Eric Mitchell***†

Stefano Ermon†‡ **Christopher D. Manning**† **Chelsea Finn**†

†Stanford University ‡CZ Biohub
{rafailov,architsh,eric.mitchell}@cs.stanford.edu

Revisiting Reinforcement Style Optimization for Alignment

★ RLFH in 2024: Is PPO/on-policy RL out yet?

Reka Chat Models

Reka Flash and Reka Edge base models are instruction tuned and then RLHFed with PPO using Reka Flash as the reward model. We conduct a series of human evaluations to evaluate our chat models.

erences from human raters and trained a reward function under the Bradley-Terry model (Bradley and Terry, 1952), similarly to Gemini. The policy was trained to optimize this reward function using a variant of REINFORCE (Williams, 1992) with a Kullback–Leibler regularization term towards the initially tuned model. Similar to the

Back to Basics: Revisiting REINFORCE Style Optimization for Learning from Human Feedback in LLMs

Arash Ahmadian, *Cohere For AI*
Chris Cremer, *Cohere*
Matthias Gallé, *Cohere*
Marzieh Fadaee, *Cohere For AI*
Julia Kreutzer, *Cohere For AI*
Olivier Pietquin, *Cohere*
Ahmet Üstün, *Cohere For AI*
Sara Hooker, *Cohere For AI*

Method	TL;DR	HH (Pythia)	HH (Llama)
RLOO (k=4)	**77.9**	**43.7**	**64.1**
RAFT (k=4)	73.2	42.1	63.3
RLOO (k=2)	**74.2**	**47.6**	**62.2**
RAFT (k=2)	72.1	37.7	58.4
REINFORCE. w/ BASELINE	**70.7**	**37.9**	**55.3**
VANILLA PG	70.4	36.4	52.3
PPO	67.6	29.2	32.0
DPO	66.6	39.0	61.9

Summary: DPO vs RL (PPO, REINFORCE, …)

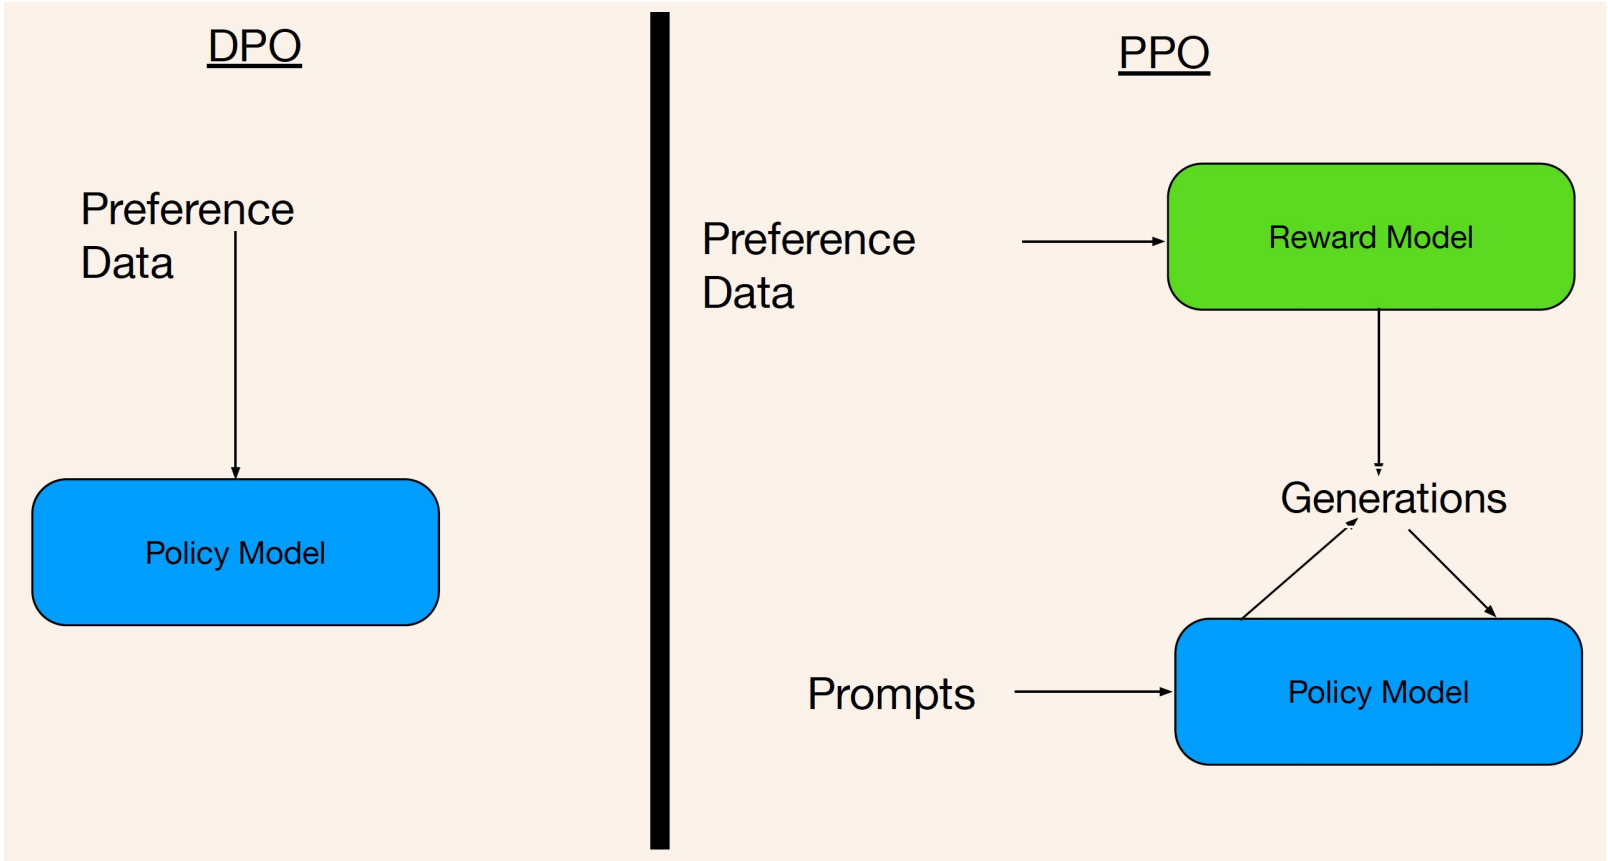

Summary: DPO vs RL (PPO, REINFORCE, …)

PPO consistently outperforms DPO, but at the cost of:

- Implementation complexity
- Memory usage, and
- Throughput

Unpacking DPO and PPO: Disentangling Best Practices for Learning from Preference Feedback

Hamish Ivison[♣♠] Yizhong Wang[♣♠] Jiacheng Liu[♣♠]
Zeqiu Wu[♠] Valentina Pyatkin[♣♠] Nathan Lambert[♣]
Noah A. Smith[♣♠] Yejin Choi[♣♠] Hannaneh Hajishirzi[♣♠]

[♣]Allen Institute for AI [♠]University of Washington
hamishiv@cs.washington.edu

Summary: DPO vs RL (PPO, REINFORCE, …)

PPO consistently outperforms DPO, but at the cost of:

- Implementation complexity
- Memory usage, and
- Throughput

Normally can get ~1% improvement from switching from DPO to PPO

Data / Model	Alg.	Factuality	Reasoning	Coding	Truthfulness	Safety	Inst. Foll.	Average
Llama 2 base	-	52.0	37.0	30.7	32.7	32.7	-	-
TÜLU 2 (SFT)	-	55.4	47.8	45.1	56.6	91.8	44.2	56.8
StackExchange	DPO	**55.3**	**47.8**	42.4	**56.2**	92.0	46.7	56.7
	PPO	55.1	**47.8**	**46.4**	54.2	**92.6**	**47.4**	**57.3**
ChatArena (2023)	DPO	**55.4**	**50.2**	45.9	**58.5**	67.3	**50.8**	54.7
	PPO	55.2	49.2	**46.4**	55.8	**79.4**	49.7	**55.9**
HH-RLHF	DPO	**55.2**	47.6	44.2	**60.0**	**93.4**	46.6	57.8
	PPO	54.9	**48.6**	**45.9**	58.0	92.8	**47.0**	**57.9**
Nectar	DPO	**55.6**	45.8	39.0	**68.1**	**93.3**	**48.4**	58.4
	PPO	55.2	**51.2**	**45.6**	60.1	92.6	47.4	**58.7**
UltraFeedback (FG)	DPO	55.3	50.9	45.9	69.3	**91.9**	52.8	61.0
	PPO	**56.0**	**52.0**	**47.7**	**71.5**	91.8	**54.4**	**62.2**
Avg. Δ b/w PPO & DPO		-0.1	+1.3	+2.9	-2.5	+2.3	+0.1	+0.7

Table 2: **DPO vs PPO:** Average performance of 13B models trained using DPO and PPO across different datasets, along with the performance difference between DPO and PPO (Δ). Blue indicates improvements over the SFT baseline, orange degradations. All datasets are downsampled to 60,908 examples (except ChatArena, which is made up of 20,465 responses). PPO outperforms DPO by an average of 0.7 points, where most improvements are in reasoning, coding, and chat capabilities.

Summary: DPO vs RL (PPO, REINFORCE, …)

- DPO and PPO are very different optimizers.
- It is learning directly from preferences vs. using RL update rules.
- It is also not really online vs offline RL, but that is more muddled.
- Efforts to develop better Direct Preference Optimization approaches are still ongoing, e.g.,
  - Y. Meng et al, “SimPO: Simple Preference Optimization with a Reference-free Reward, NeurIPS 2024, https://arxiv.org/pdf/2405.14734
  - A. Chen et al, “Preference Learning Algorithms Do Not Learn Preference Rankings”,NeurIPS 2024, https://arxiv.org/pdf/2405.19534

More discussion:
https://twitter.com/srush_nlp/status/1729896568956895370,
https://www.interconnects.ai/p/the-dpo-debate,
https://www.youtube.com/watch?v=YJMCSVLRUNs

What components matter for LLMs ?

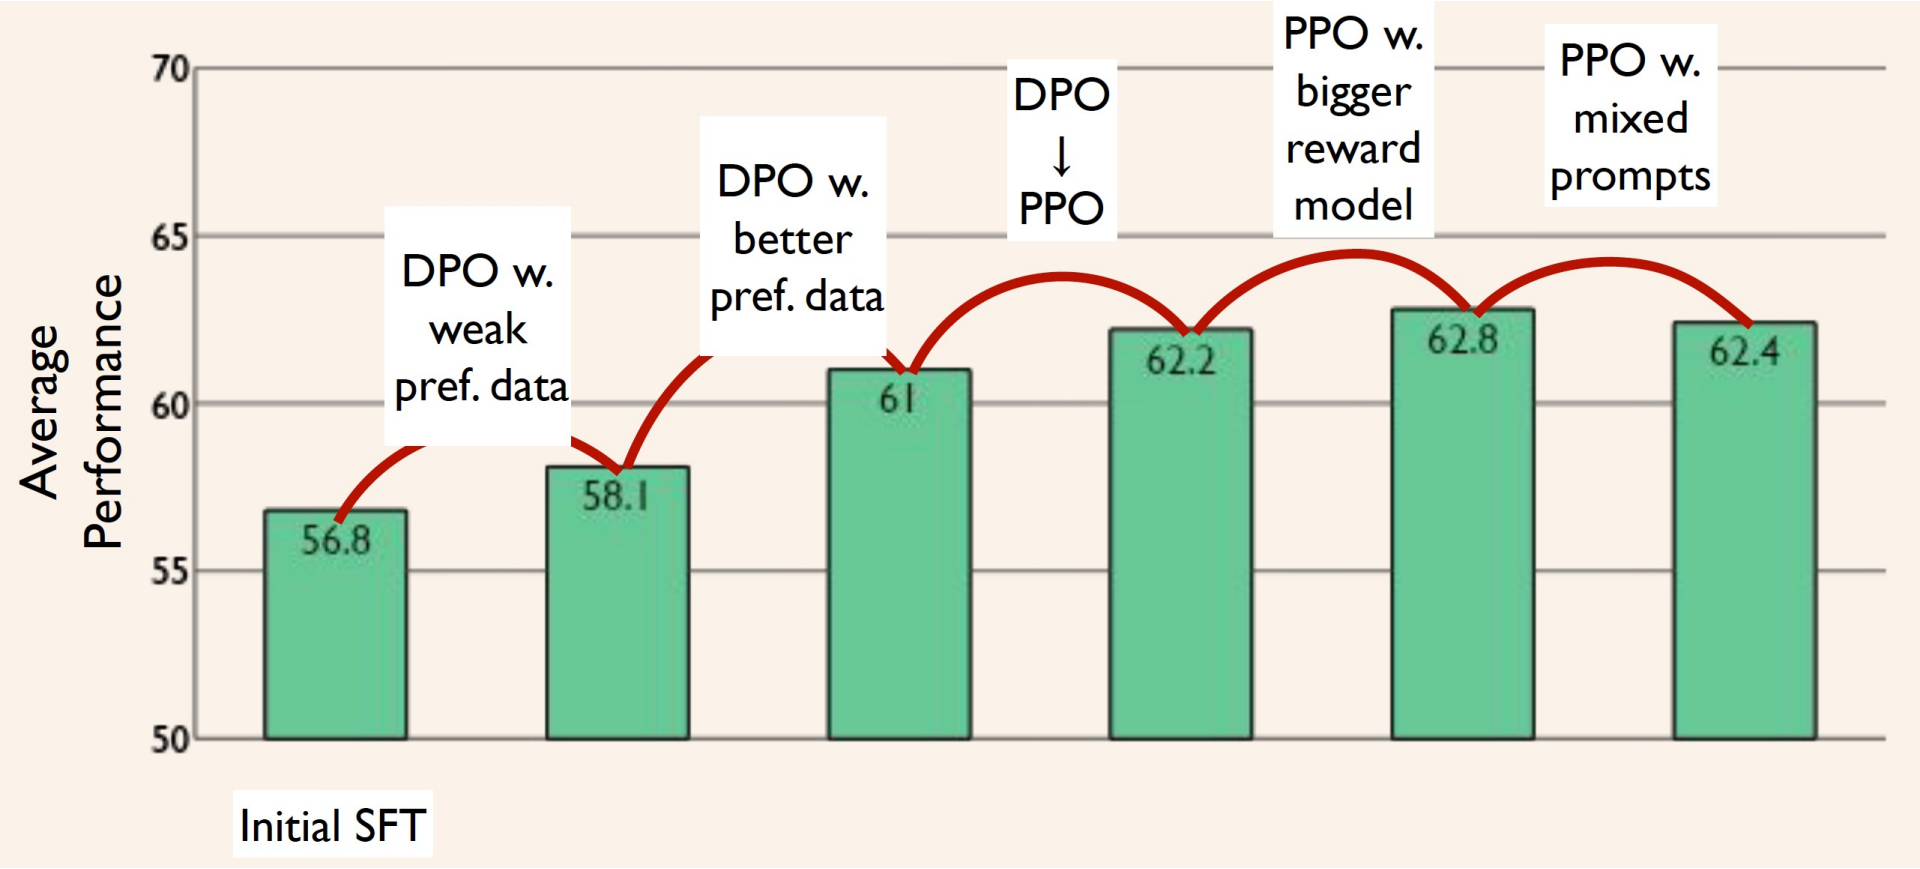

What components matter for LLMs ?

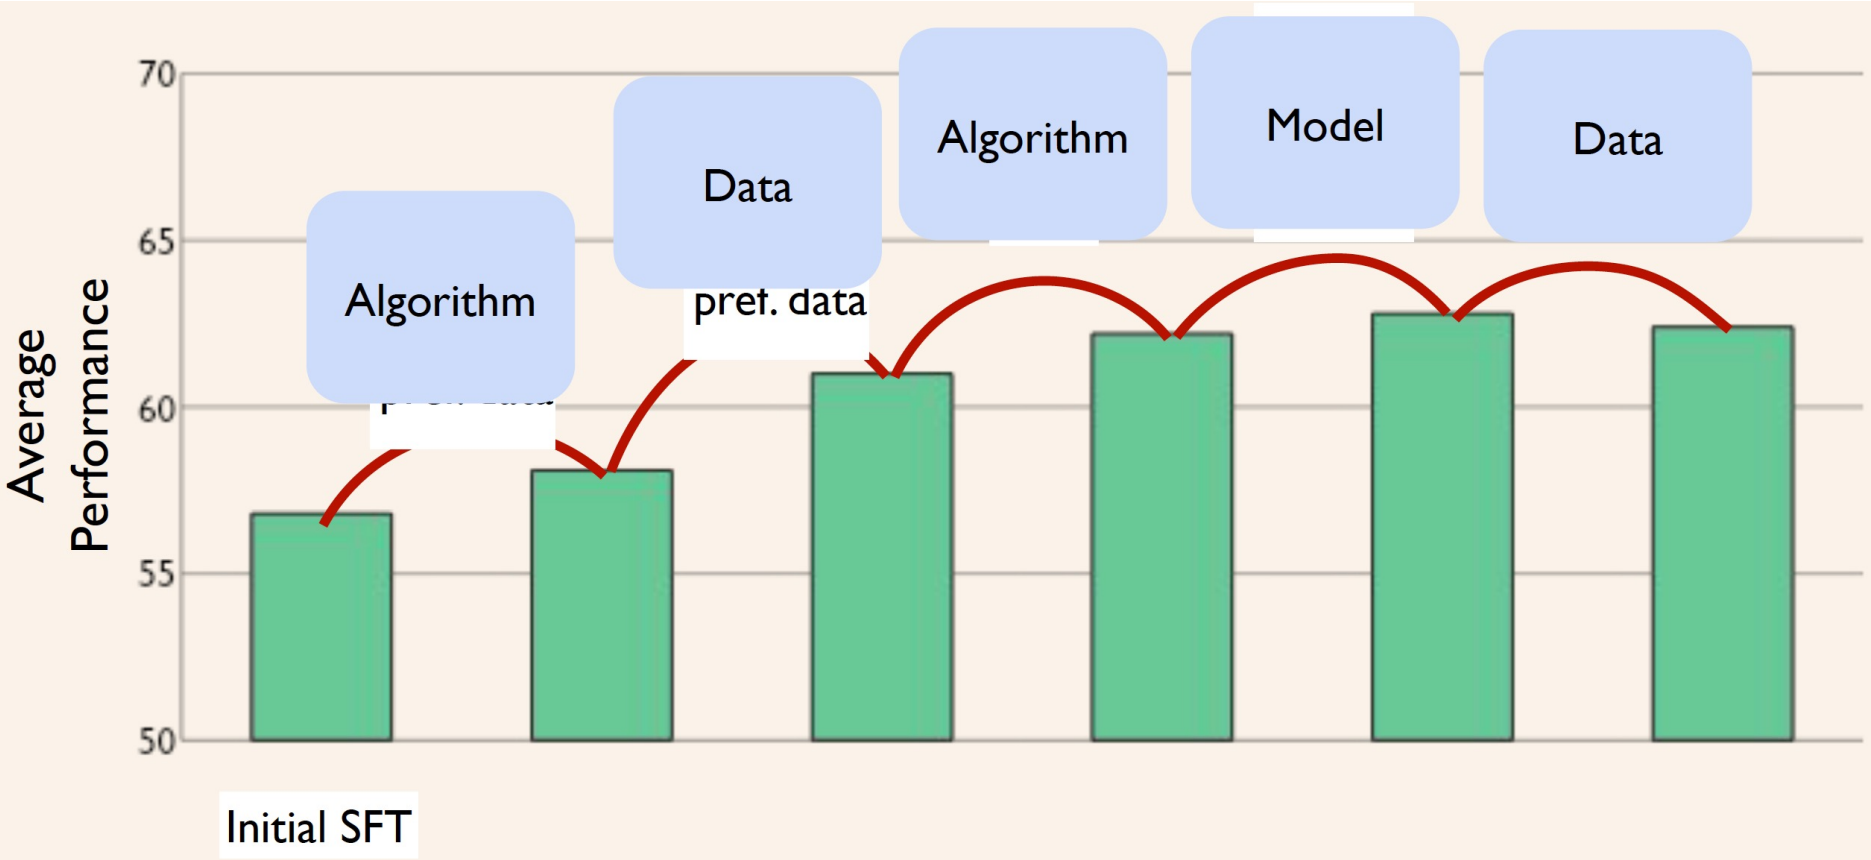

Ivison*, Wang* et al. 2023; Ivison, Wang et al. 2024

Increasingly, labeling is a human-machine collaboration…



- **LLMs can reference and follow the labeling instructions just as humans can.**
- => LLMs can create drafts, for humans to slice together into a final label.
- => LLMs can review and critique labels based on the instructions.
- …

From RLHF to RLAIF (AI Feedback)

- Do we really need human feedback? What if we ask an aligned language model for feedback?
- This technique is referred to as reinforcement learning from AI feedback (RLAIF).

ZEPHYR: DIRECT DISTILLATION OF LM ALIGNMENT

Lewis Tunstall,* Edward Beeching,* Nathan Lambert, Nazneen Rajani,
Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra,
Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero,
Alexander M. Rush, and Thomas Wolf
The H4 (Helpful, Honest, Harmless, Huggy) Team
https://huggingface.co/HuggingFaceH4
lewis@huggingface.co

The RLAIF (AI Feedback) Approach

1. Given a set of prompts, feed into multiple language models (e.g. Llama, Falcon, Vicuna, Claude…) to generate n responses.

2. Ask a Teacher model (e.g. GPT4o) to rate all of the n responses.

3. The response with the highest score is chosen as the winning response and the losing response is randomly chosen.

4. Run DPO on this dataset !

LM-generated RLHF Data & Evaluate with Alpaca Farm

Using LM-generated pairwise preferences is a standard, low-cost way to study RLHF

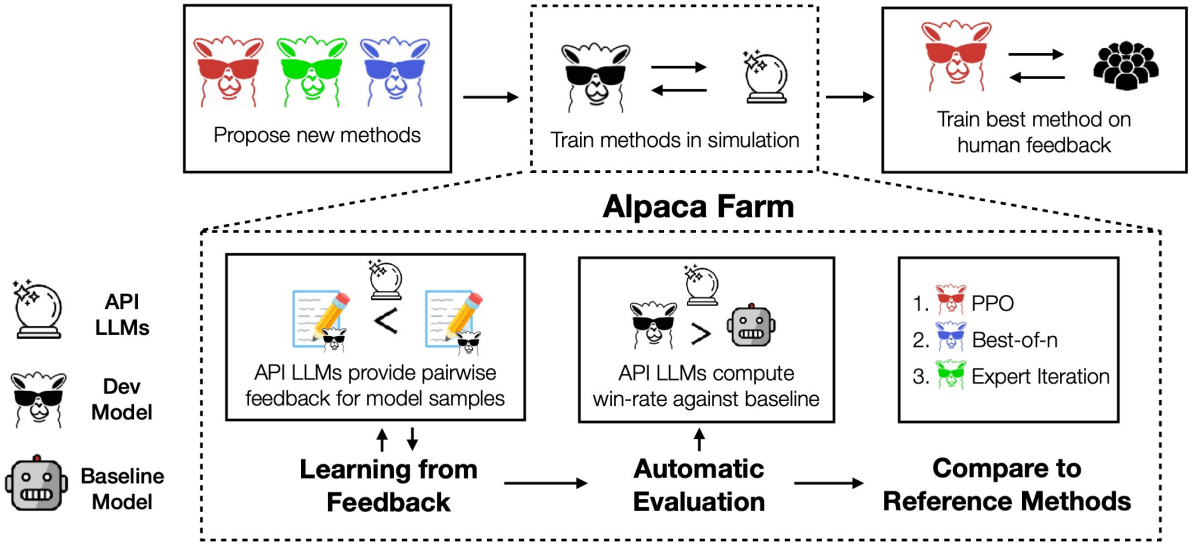

	Feedback Source	Experiment Cost	Iteration time	Reproducible
Alpaca Farm		$70	Hours	✅
Human Feedback		$3,150	Days	❌

Source: [Dubois et al, 2023]

This has also become a very common fine-tuning setup for non-frontier models

Near-perfect Rank Correlation at System Level

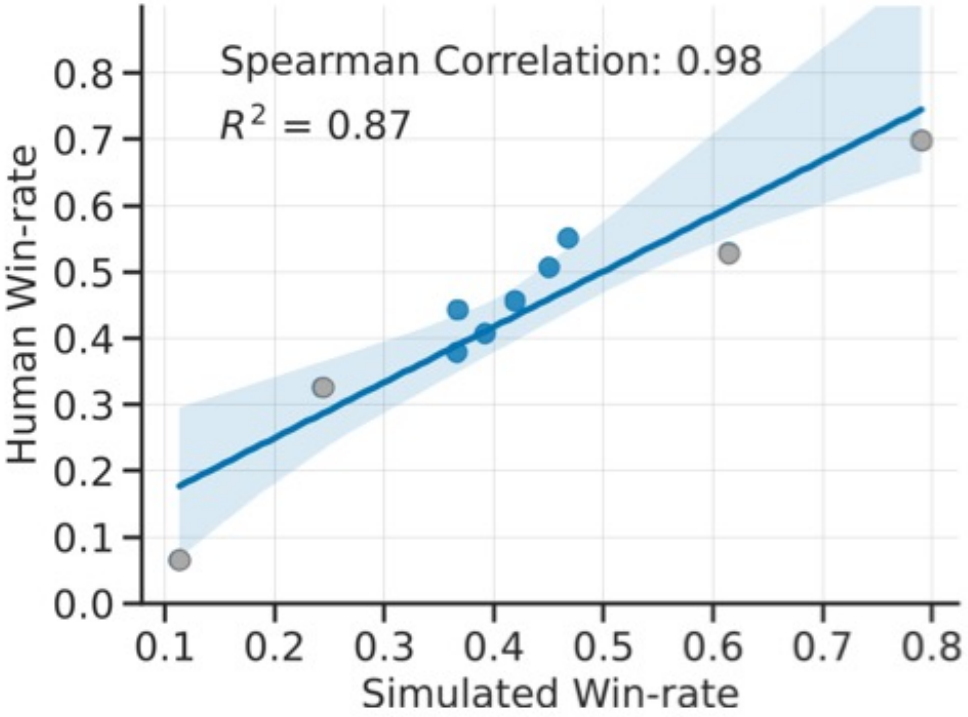

Figure 3: The ranking of methods trained and evaluated in AlpacaFarm matches that of methods trained and evaluated in the human-based pipeline. Each point represents one method M (e.g. PPO). The x-axis shows the simulated evaluation (win-rates measured by $p_{\text{sim}}^{\text{eval}}$) on methods trained in simulation M_{sim}. The y-axis shows human evaluation (win-rates measured by p_{human}) on methods trained with human feedback M_{human}. Gray points show models that we did not train, so their x and y values only differ in the evaluation (simulated vs human). Without those points, we have $R^2 = 0.83$ and a Spearman Correlation of 0.94.

GPT4 is a very good pairwise feedback system !

Source: [Dubois et al, 2023]

LM-generated RLHF and Data

- At the lower-end of the Cost-vs-Quality trade-offs, using GPT4 or other Frontier LLMs has become the Standard !

Ultrafeedback

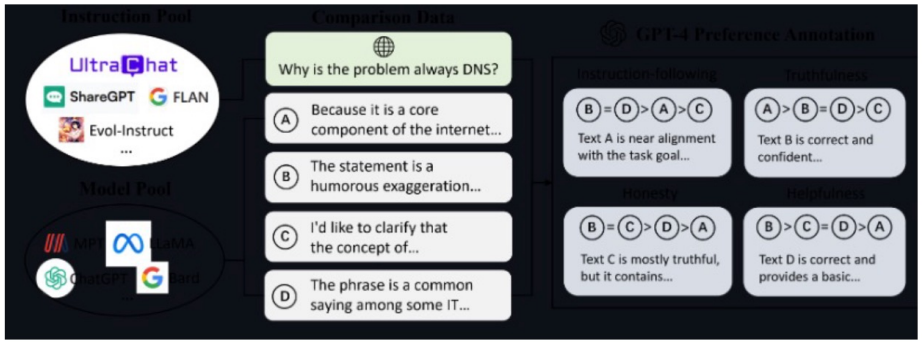

Zephyr 7b

> ***One of the primary contributions of ZEPHYR was the incorporation of AI feedback via teacher models for the alignment tasks. Why did you choose this approach over more established human feedback mechanisms?***

Earlier in the year, we had actually experimented with collecting human feedback from a data vendor, but found the process was both time consuming and costly to oversee. Based on this experience, we felt AI feedback was a more accessible route for both our small team and as a means to popularize a method that the community could also adopt.

Used in e.g. Olmo, Zephyr, etc

Source: https://thesequence.substack.com/p/the-sequence-chat-hugging-faces-lewis

Leading synthetic preference data method: UltraFeedback

Key aspects:

- Diverse model pool for completions.
- Diverse prompt pool.
- Extra performance is use on-policy generations from SFT checkpoint(s)

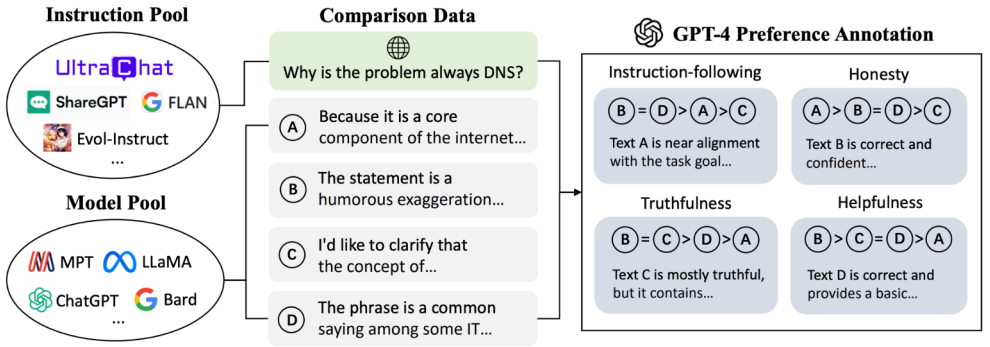

Figure 1. ULTRAFEEDBACK construction process. We sample instructions and models from large pools to guarantee diversity, then query GPT-4 with detailed illustrations for fine-grained and high-quality annotations in both textual and numerical formats.

Cui, Ganqu, et al. 2024. UltraFeedback.

Human preferences vs LLM-as-a-judge

Both sources of preference data are used extensively.

In frontier labs:

- Human data used extensively as foundation of PreFT.
- Synthetic data used to enhance behaviors (e.g. Constitutional AI).

In open research:

- Synthetic data dominates due to price.
  - One LLM-as-a-judge label costs <1 cent.
  - One human datapoint costs $5-20.

Human preferences vs LLM-as-a-judge

Human data: High noise, low bias.

Synthetic data: Low noise, high bias.

Constitutional AI (another form of RL w/ AI Feedback)

- Replaces ~10,000s of human labels with ~10 human-written principles (aka the “constitution”)!
- Perform RL-finetuning using AI feedback derived from a “constitution” describing desired behavior. Humans don’t need to be in the loop, except to write the constitution!

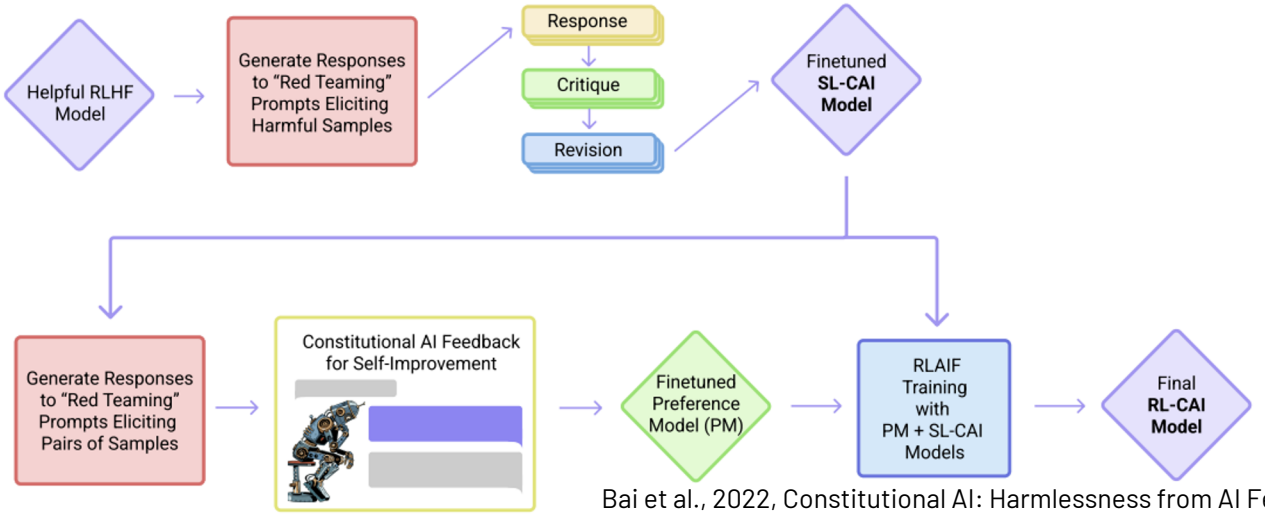

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback (Used in Anthropic’s Claude)

An Example Constitution with 3 Principles

#1
```
CritiqueRequest: Identify specific ways in which the assistant's last response
    is harmful, unethical, racist, sexist, toxic, dangerous, or illegal.
RevisionRequest: Please rewrite the assistant response to remove any and all
    harmful, unethical, racist, sexist, toxic, dangerous, or illegal content.
```

#2
```
CritiqueRequest: Take a look at the assistant's last response and note if it was
    in any way misogynistic or otherwise gender biased. Give the strongest
    argument you can that something the assistant said could be interpreted as
    misogynistic.
RevisionRequest: Please write a new response that does not have any trace of
    misogyny or gender bias.
```

#3
```
CritiqueRequest: Did the assistant's preceding response include anything that
    may be inappropriate for young children? If so, discuss how it can be
    revised to be more appropriate.
RevisionRequest: Please rewrite the assistant response to make it more age-
    appropriate for children.
```

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Constitutional AI: Self-Improving AI using AI Feedback

In "Constitutional AI", the LLM follows a Constitution (set of principles written by a human) to generate feedback for self-improvement

1. **Supervised learning (SL):** Finetune LLM on data generated by self-critique and revisions

   Example:

   Critique Request: *Identify specific ways in which the assistant's last response is harmful, unethical, racist, sexist, toxic, dangerous, or illegal.*

   Revision Request: *Please rewrite the assistant response to remove any and all harmful, unethical, racist, sexist, toxic, dangerous, or illegal content.*

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Constitutional AI Supervised Learning Results

Harmlessness improves and helpfulness declines as a function of the number of revisions. Overall HH (Helpfulness + Harmlessness) improves monotonically.

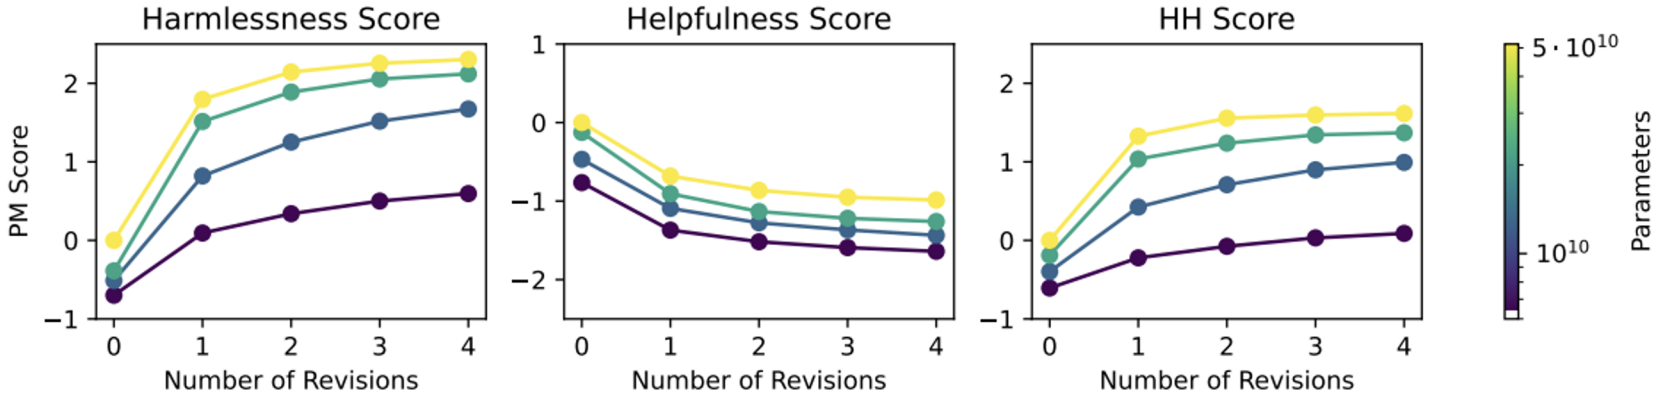

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Constitutional AI: Self-Improving AI Using AI Feedback

In "Constitutional AI", the LLM follows a Constitution (set of principles written by a human) to generate feedback for self-improvement

1. **Supervised learning (SL):** Finetune LLM on data generated by self-critique and revisions

1. **Reinforcement learning from AI feedback (RLAIF):**
   a. Train a preference model based on LLM (from Step 1) responses and the constitution
   b. Finetune the LLM to maximize the preference model

Example: *Choose the response that answers the human in the most thoughtful, respectful and cordial manner.*

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Constitutional AI: A Step Towards Self-improving AI

- Higher quality supervision:
  - AI has already surpassed humans in a growing number of applications, including games (Go and Dota), chip floorplanning, diagnostics imaging (retinal scans, skin screening)
- More transparency:
  - Using AI and the constitution can enable a more objective approach for data creation and ranking
- Scalable:
  - AI models automatically and in-parallel perform data collection and model finetuning
- Cost-friendly:
  - Using AI eliminates the high cost of collecting high-quality labeled traces and human evaluations

Constitutional AI and RLAIF Further Improve Finetuning

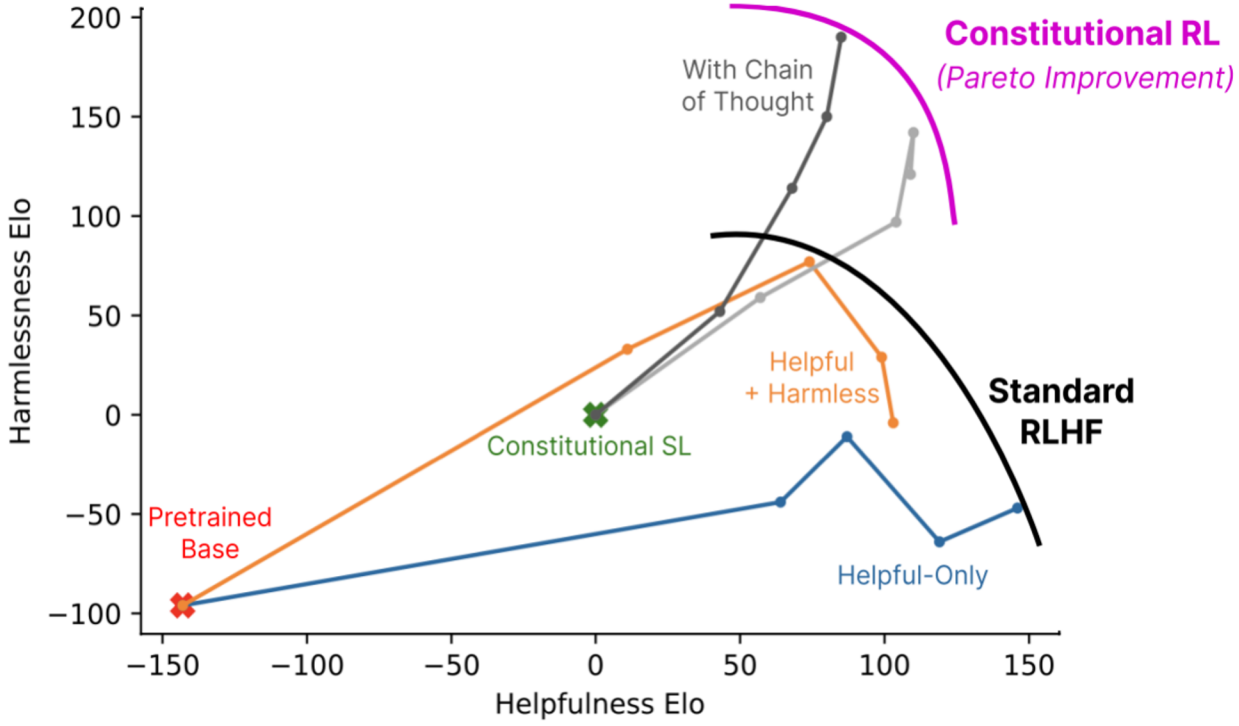

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Constitutional AI Improves Scaling

Similar to RLHF, except that we replace human preferences with AI feedback

The AI model itself evaluates responses according to a set of constitutional principles

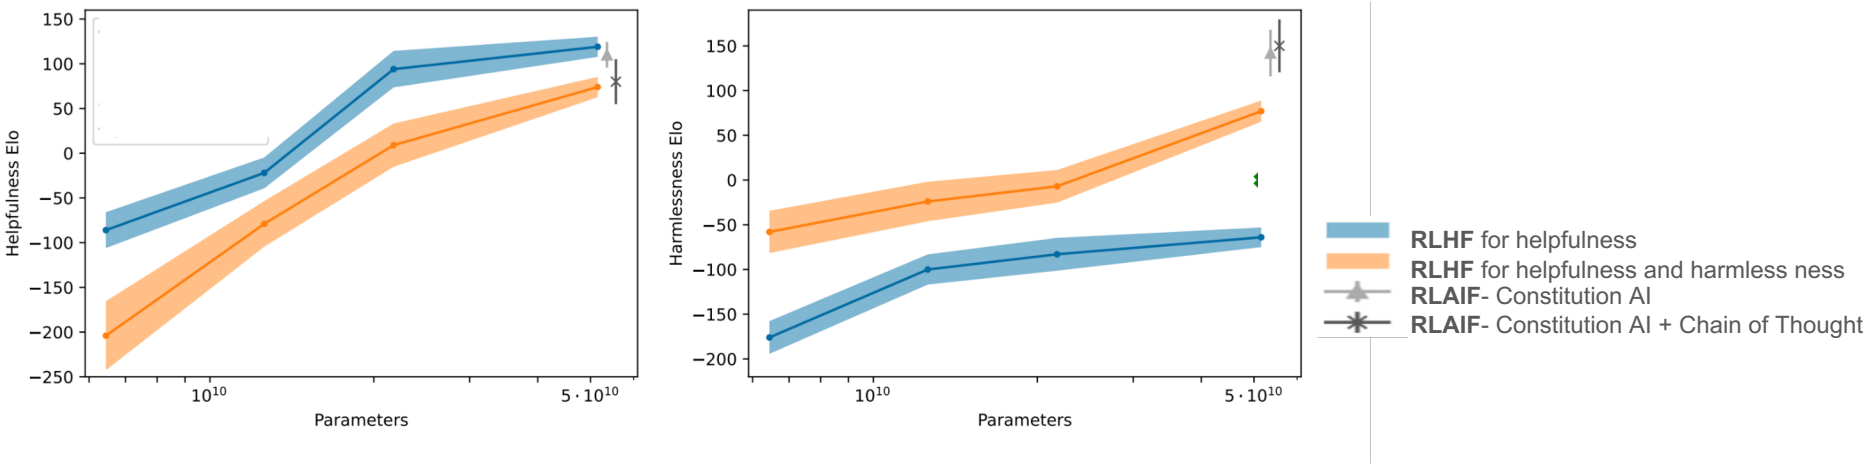

Bai et al., 2022, Constitutional AI: Harmlessness from AI Feedback

Other Ways to Improve Finetuning

Quality of finetuing improves with:

- Increasing the number of tasks the model is finetuned on
- Scaling model size
- Using chain of thought instructions in the finetuning mix

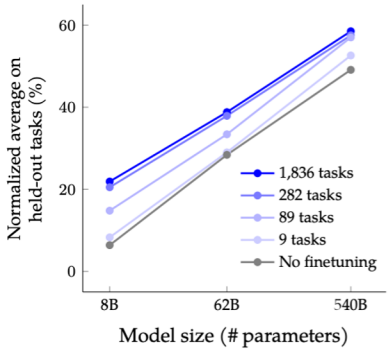

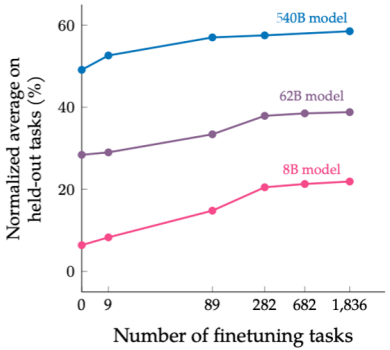

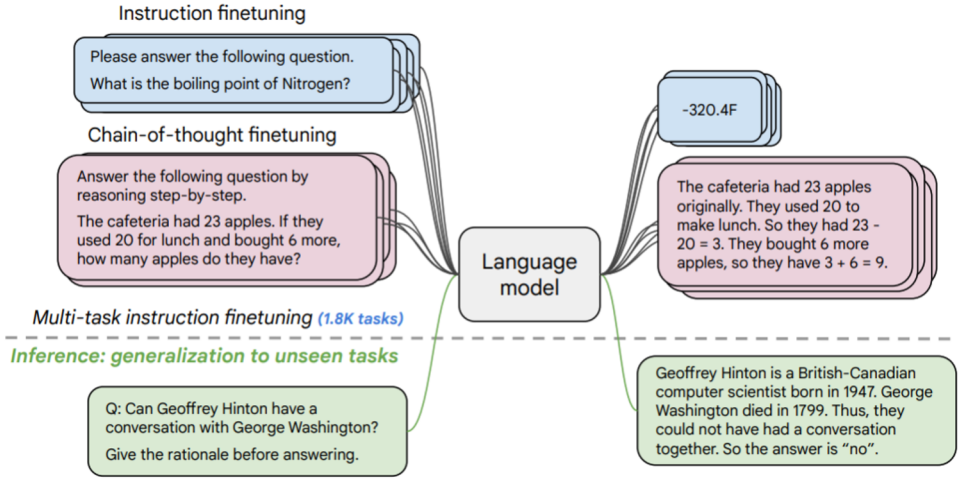

Chung, Hou, Longpre et. al, Scaling Instruction-Finetuned Language Models, 2022

Latest: RL finetuning

Reinforcement learning as a training objective other than just for Human Preferences:

- OpenAI's o1 and related models trained with "large-scale RL" for reasoning
- Finetuning based on verifiable outputs:
  - Tülu 3's Reinforcement Learning with Verifiable Rewards (RLVR) or
  - OpenAI's Reinforcement Finetuning API
  - Extensive research in specific domains: Code verification, VinePPO for math, Quiet STaR, DeepSeek R1, etc.

How to Implement
RL with Verifiable Reward (one form of RL finetuning) ?

Start with:
Standard RLHF

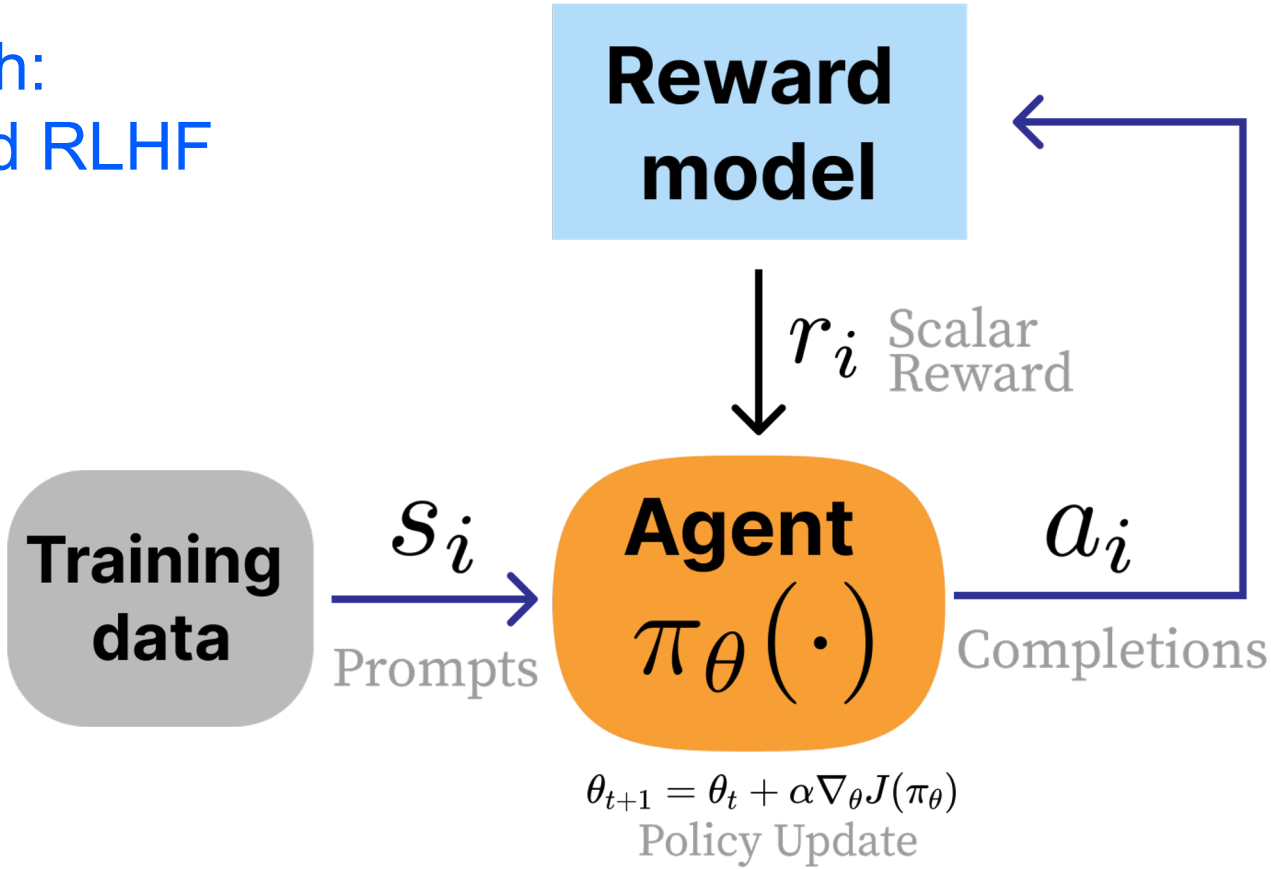

Lambert, Nathan et al. 2024. Tülu 3.

Standard RLHF (with RL details)

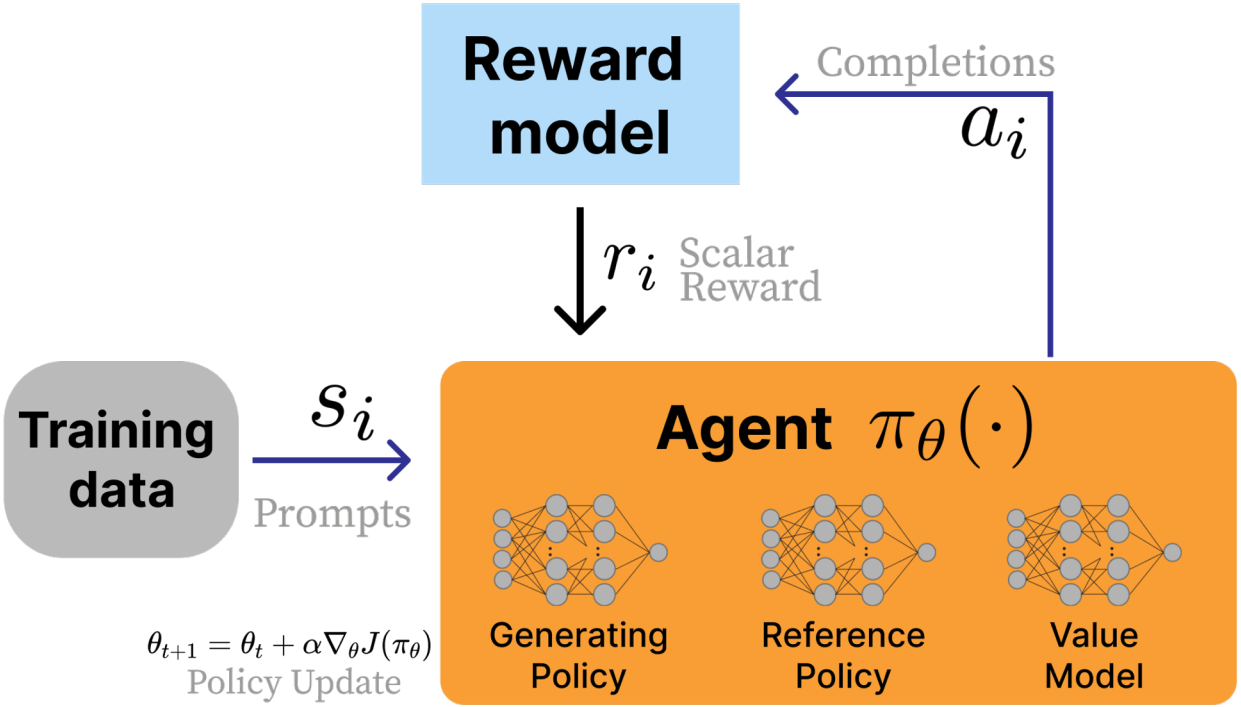

Lambert, Nathan et al. 2024. Tülu 3.

Adding Verifiable Rewards

Can we reward the policy for being correct?

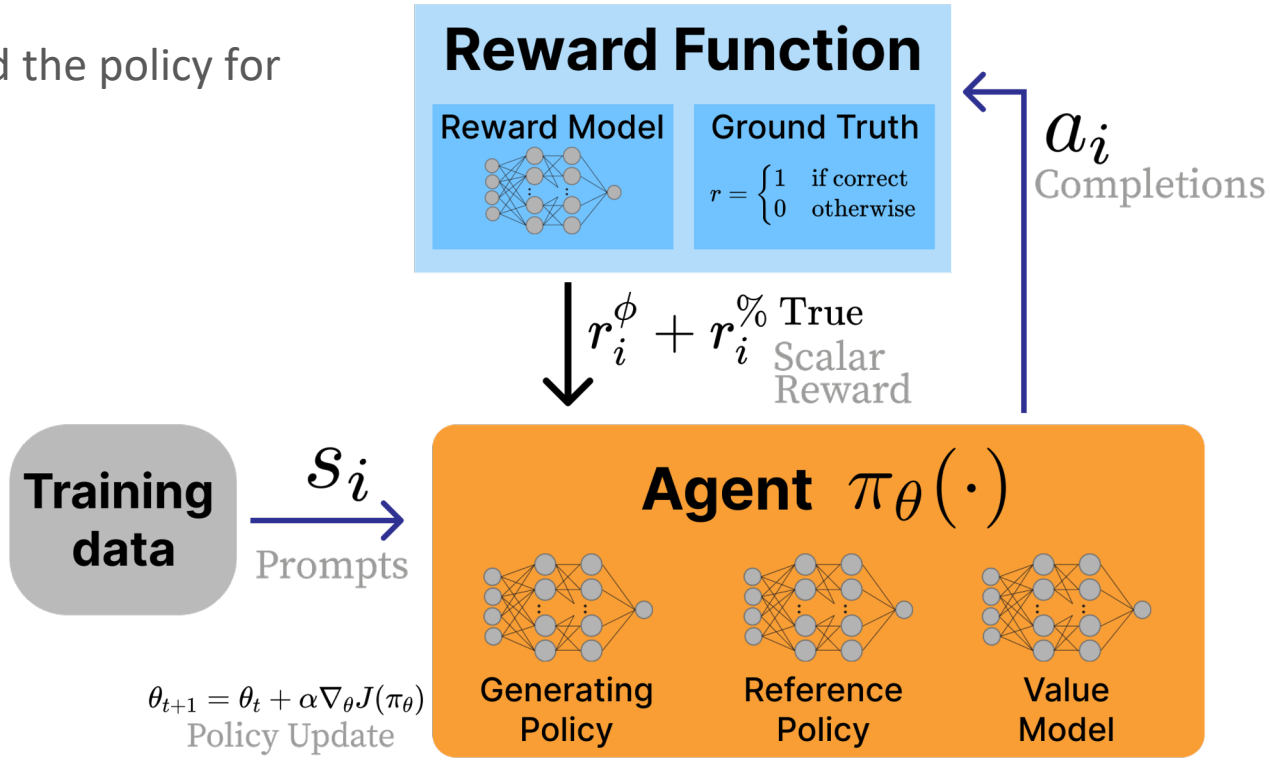

Lambert, Nathan et al. 2024. Tülu 3.

RL with Verifiable Rewards details

1. We do not use a reward model! Just "environment" reward.

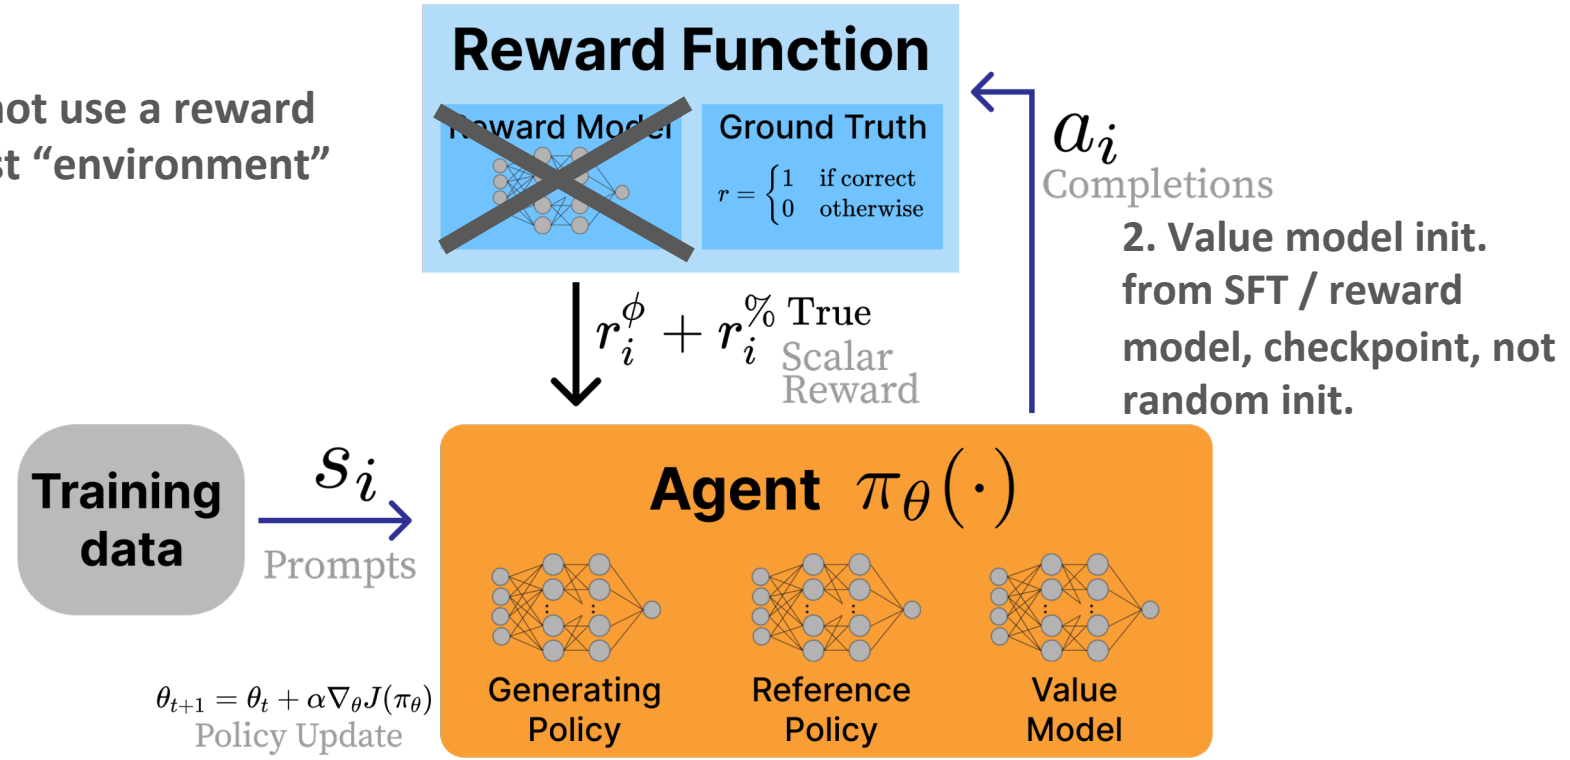

Lambert, Nathan et al. 2024. Tülu 3.

Ground Truth RL

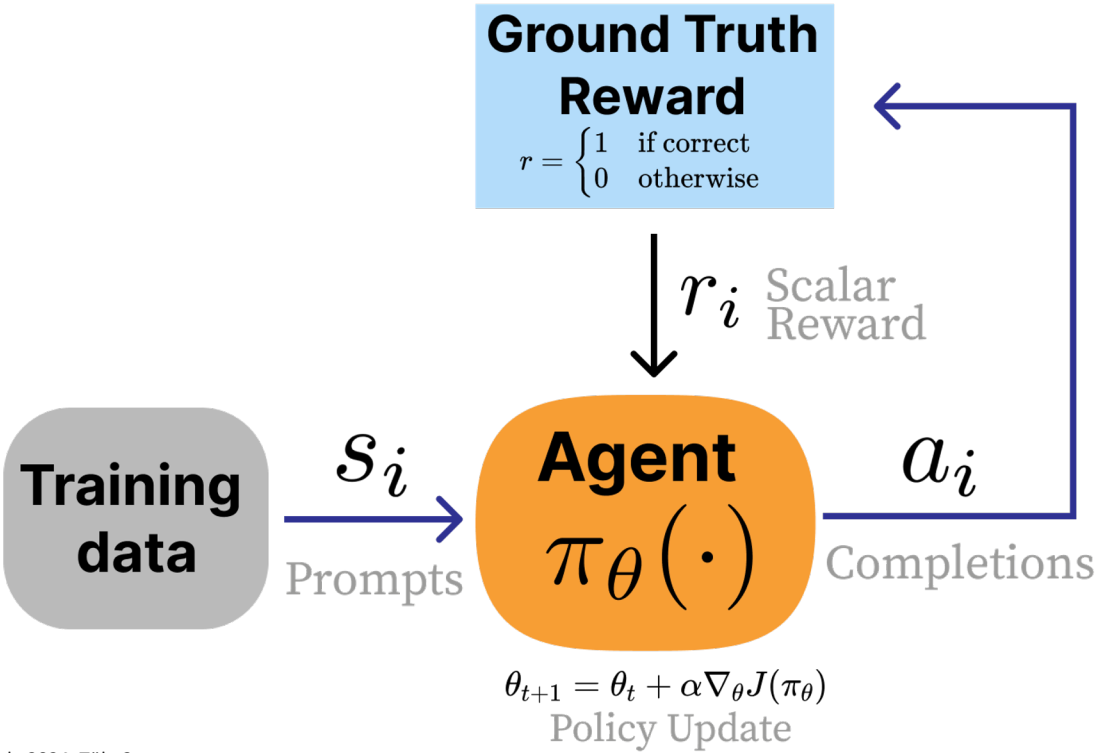

Lambert, Nathan et al. 2024. Tülu 3.

RL Fine-tuning (e.g. RLVR) vs. with classical RL

- RL with verifiable rewards and all forms of RL finetuning are extremely close to the classical definition of RL problems with reward given to correct behavior.
- Main innovation is in **stability of implementation without degradation across out-of-domain capabilities.**

Summary on RLHF / RLAIF

- RLHF data collection is HARD ! Many Confounding Factors
- RLHF algorithms, especially PPO, are more complex than SFT
- DPO is an inexpensive, much more simple alternative to PPO, at the expense of slight degradation of performance (~1%)
- Be careful of the effect of (over) optimizing for rewards, aka “reward hacking”
- Hybrid use of Human Feedback + AI Feedback becomes reality !

Summary: Taxonomy of Post-Training approaches for

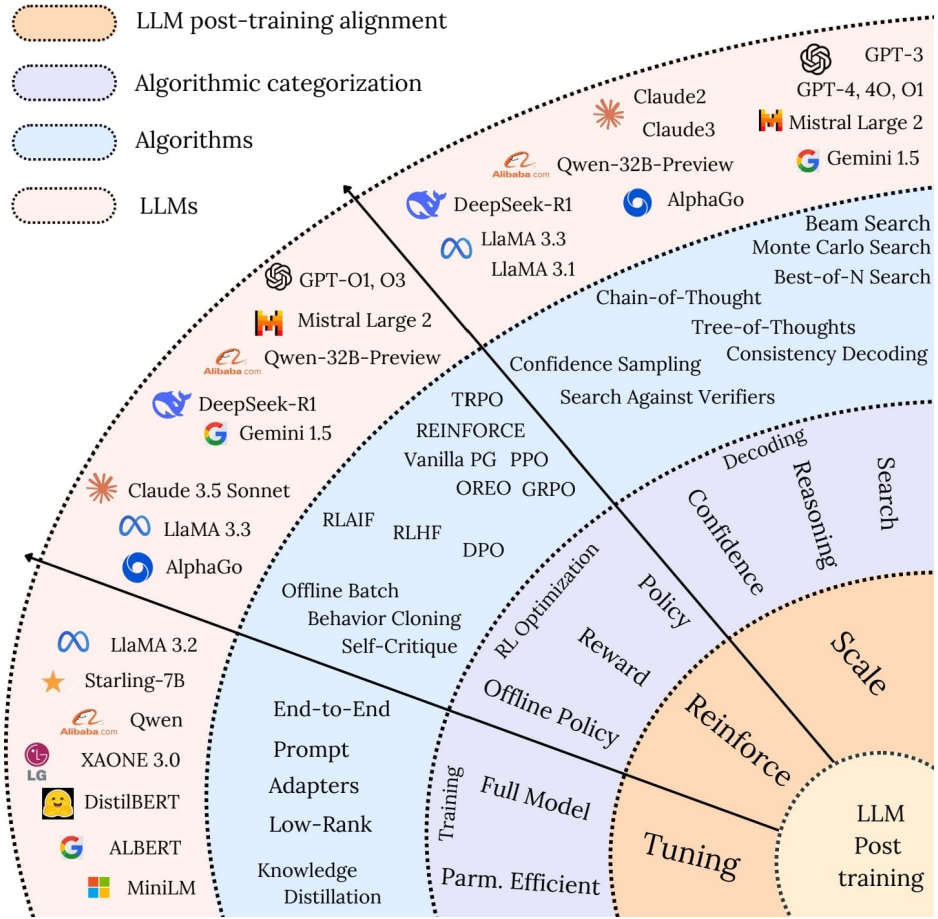

Source: https://github.com/mbzuai-oryx/Awesome-LLM-Post-training/blob/main/Images/teasor.jpg

The Big Picture of Pre-training vs. Post-training

- In many ways is very aligned with Yann's cake metaphor.
  - Foundation / pretraining is unsupervised learning.
  - SFT / PreFT is supervised learning.
  - RL is cherry on top to finalize behavior.
- The boundary between Pre-training and Post-training (also the recent introduction of Mid-training), as well as their relationship to other LLM behaviors, are still undergoing rapid changes !