RLHF, Alignment

CSE 473: Introduction to Artificial Intelligence

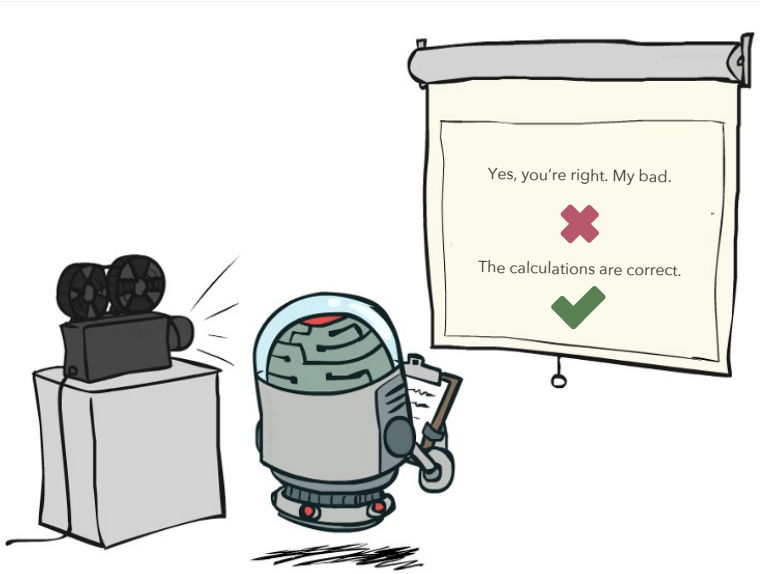

James Weichert (jpw@cs)
cs.uw.edu/473
acknowledgements

Welcome!
What are you looking forward to after the summer quarter ends?
sli.do **#cse473**

Administrivia

ANNOUNCEMENTS

- **Course Evaluation!**

uw.iasystem.org/survey/328718

ASSIGNMENTS

- **Homework 4 due** tomorrow (8.20)
  - No late period – HW4 must be turned in by 8.21

LLMs

FOUNDATIONS

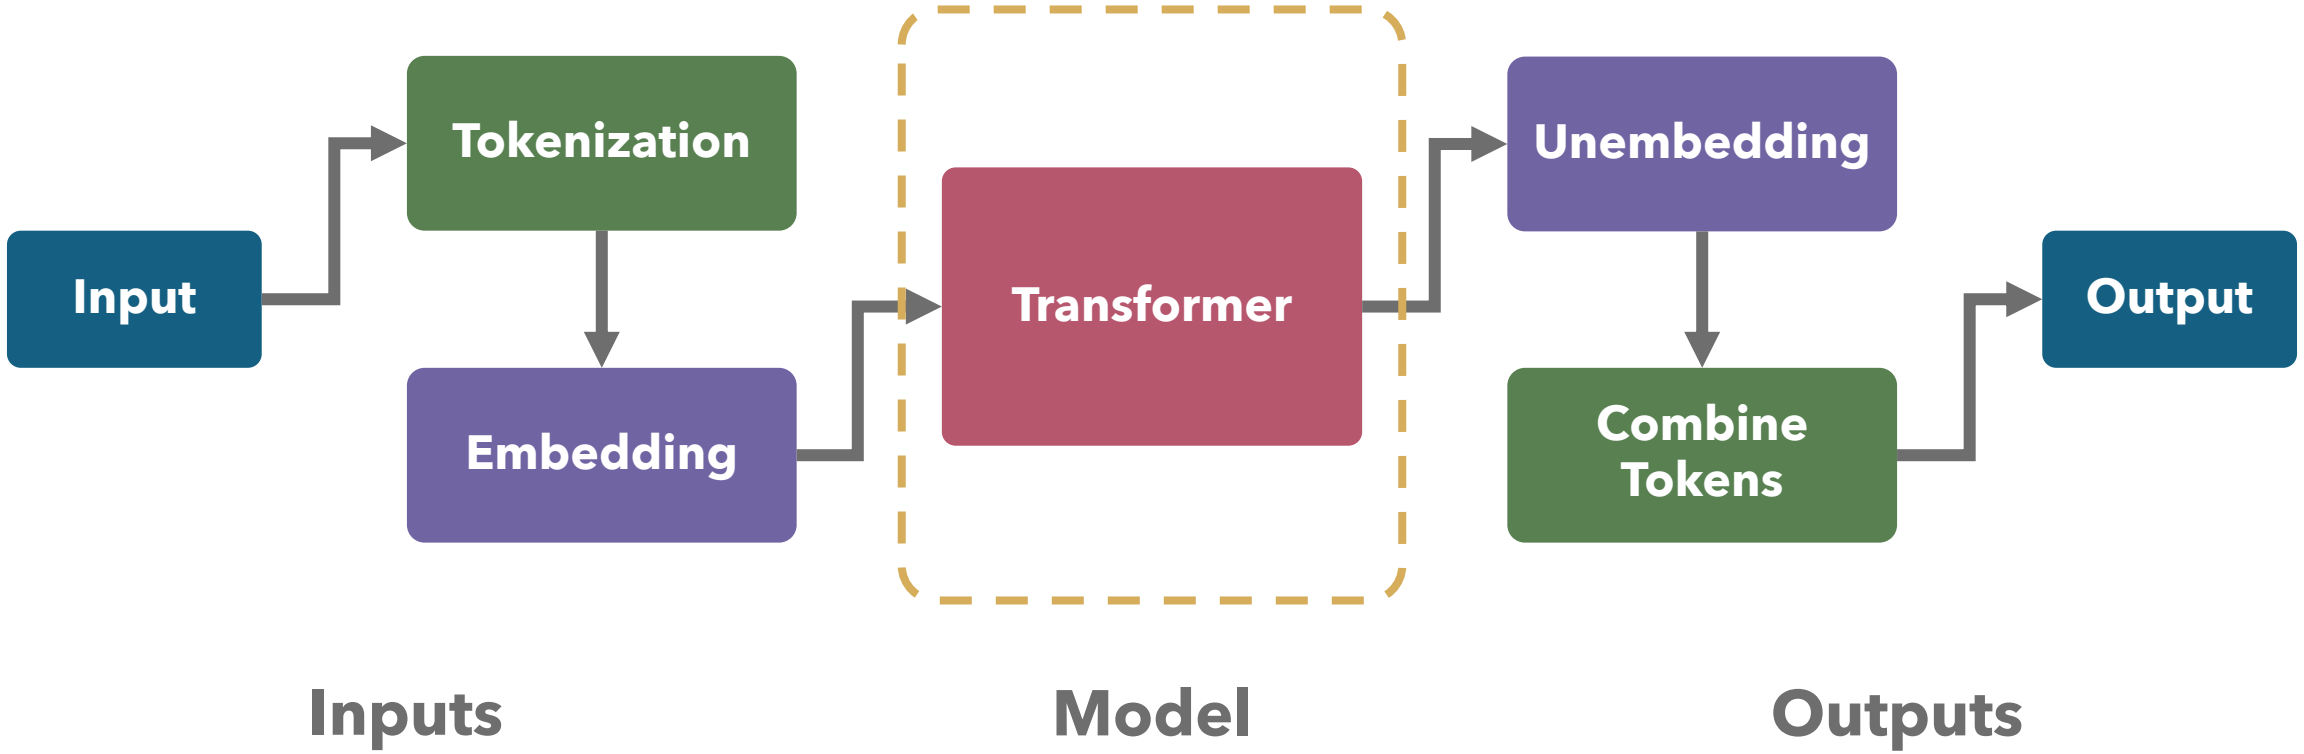

From GPT to *Chat*GPT

EXAMPLE

GPT-2 XL (1.5B)

"In the early days of computing, a simple addition of the digit values of a number by the power of two was a simple, fast, and efficient method of computation."

GPT-5.6 (???)

"$12^{13} = 154{,}070{,}215{,}416{,}546{,}875{,}000$"

Fine-Tuning

CONTEXT

- LLMs are trained to model general next-token distributions $P(w_t|w_{1:t-1})$
- Tasks have desired input-response pairs (x, y) with distribution $P(y|x)$
  - Easier to perform supervised learning on small set of task-specific examples than to train from scratch

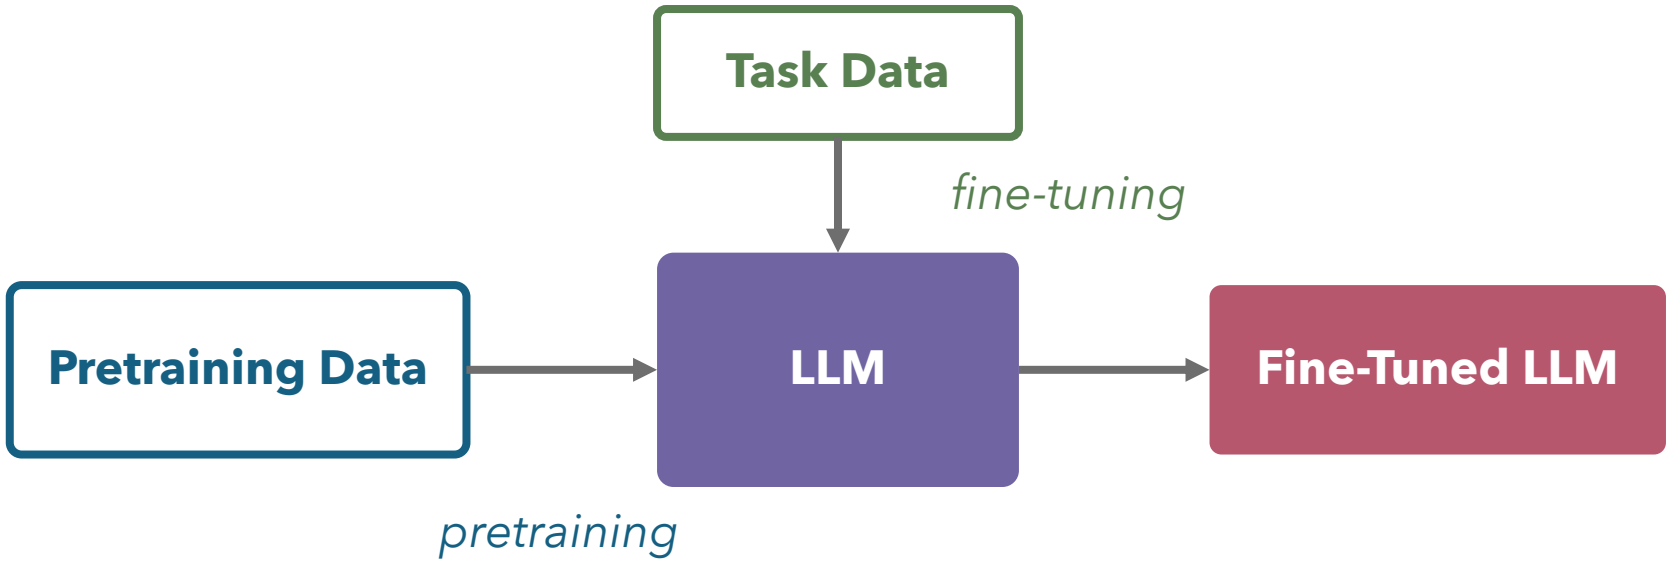

Supervised Fine-Tuning

FOUNDATIONS

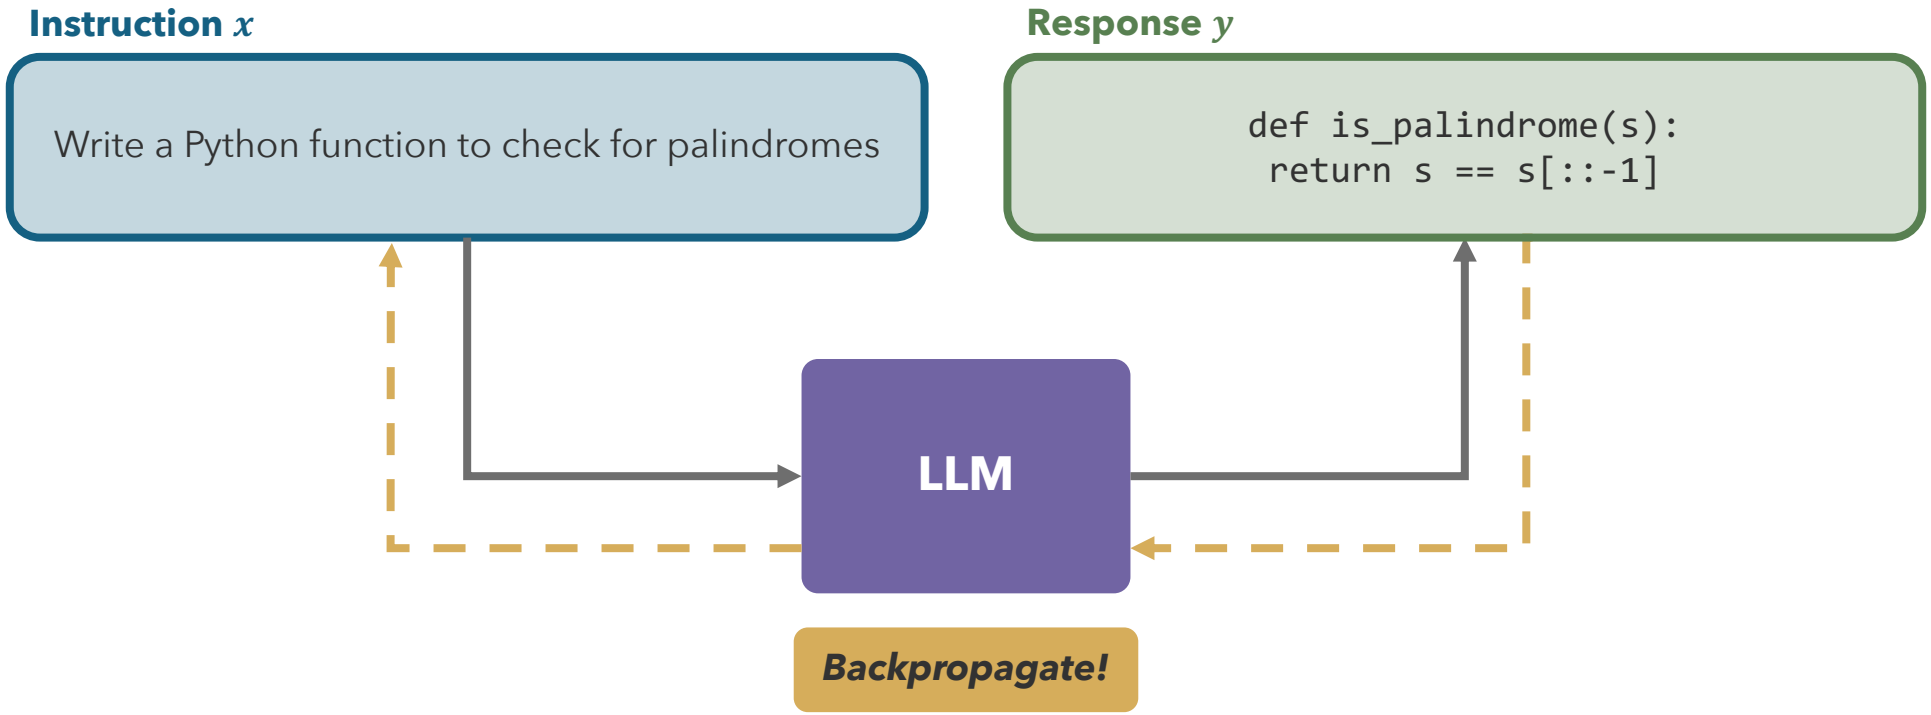

Problem…*What if fine-tuning overwrites important learned weights?*

Fine-tuning in the same way can lead to **catastrophic forgetting**

QUESTIONS

Questions?

live and on sli.do #cse473

Encouraging Desirable Outputs

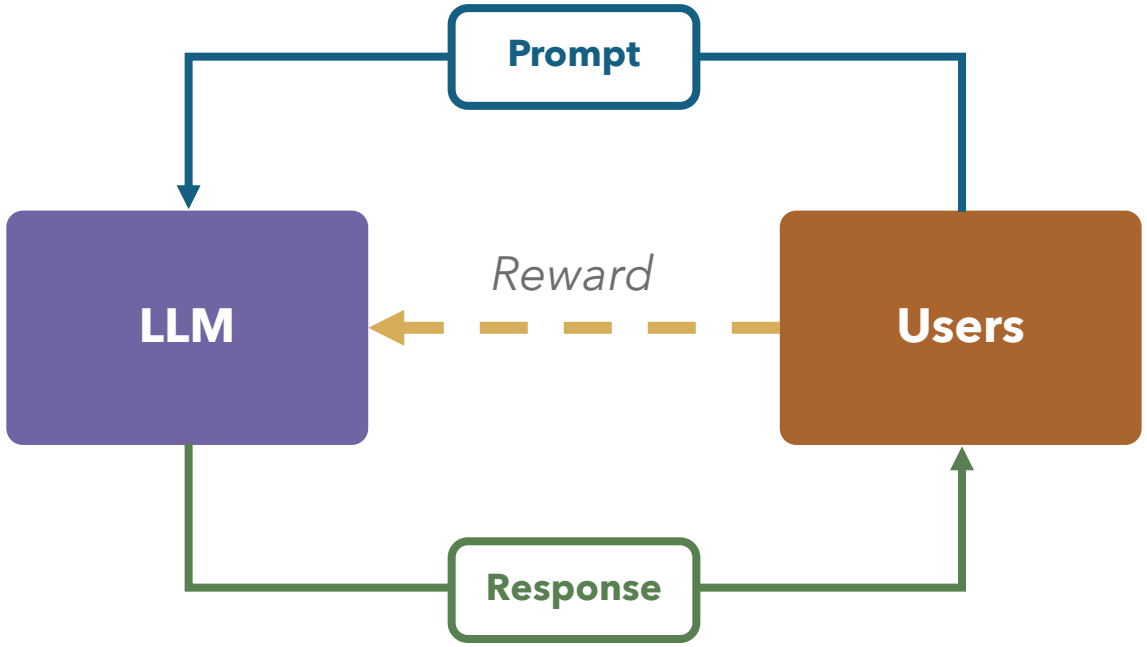

Language Models as MDPs

CONTEXT

- **State**: Sequence of words seen so far
  - For GPT-3, $|V|^{n_{\text{context}}} = 50{,}000^{2{,}048}$ possible states
- **Action**: Next word generated
  - $\max_{w \in V} Q(s, a)$ is infeasible!
- **Transition**: Deterministic
  - Just append next word to sequence
- **Reward**: ???

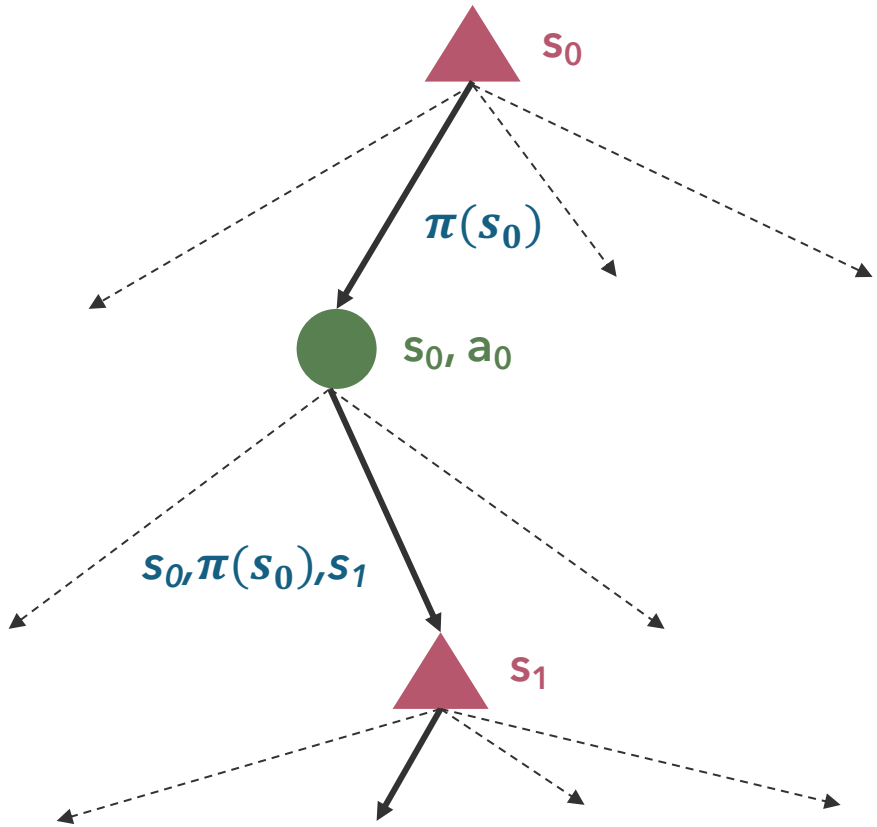

LLM Reinforcement Learning

FOUNDATIONS

- In **fine-tuning**, we want to learn $\hat{P}(y|x)$ to approximate the true distribution
- From **RL**, we know we can optimize using a **reward function** $R(x, y)$

$$\hat{\pi}(y|x) = \text{argmax}_{\pi} \mathbb{E}_{\pi}[R(x, y)]$$

  - Find policy $\hat{\pi}$ that maximizes expected reward

How to determine $R(x, y)$?

Summarize the introduction to the [UW Wikipedia Entry]

The University of Washington is a 165-year-old public university with campuses in Seattle, Tacoma, and Bothell. The main campus in Seattle includes over 500 buildings and 26 libraries. The university is known for its research and athletics.

A

The University of Washington (UW), founded in 1861 in Seattle, is a public research university and one of the oldest on the West Coast. Its main campus spans 800 acres, with satellite campuses in Tacoma and Bothell, encompassing over 400 buildings and 20 billion square feet of space. As the flagship of Washington's public university system, UW is a member of the Association of American Universities and is noted for its medical, engineering, and scientific research. In 2024, it spent $1.73 million on research and development, ranking fifth nationally. Its 22 varsity teams, the Huskies, compete in the Big Ten Conference of NCAA Division I.

B

Which do you prefer?

sli.do **#cse473**

Pairwise Preferences

DEFINITION

- For a given prompt x, generate two versions of the LLM output, a and b
- Ask human annotators to rank a vs. b
  - Generate preferences $a \succ b$ or $b \succ a$
  - Like an A/B test for AI outputs
- Learn a **reward model** $r(x, y)$ that assigns scalar values to each x, y pair

$$P(y_a \succ y_b | x) = \sigma(r(x, y_a) - r(x, y_b))$$

- Optimize policy π via RL (PPO)*

$$\text{argmax}_\pi \mathbb{E}_\pi[r(x, y)]$$

* Include penalty term to mitigate quality degradation

To RL or Not to RL…

CONTEXT

- **Proximal Policy Optimization** (PPO) is on-policy RL, which is complex

  - Can we simplify this process?

- **Direct Preference Optimization** (DPO) fixes the reward model

$$r_{\text{DPO}}(x, y) = \beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{SFT}}(y|x)}$$

Direct Preference Optimization

FOUNDATIONS

- **Direct Preference Optimization** (DPO) avoids reinforcement learning

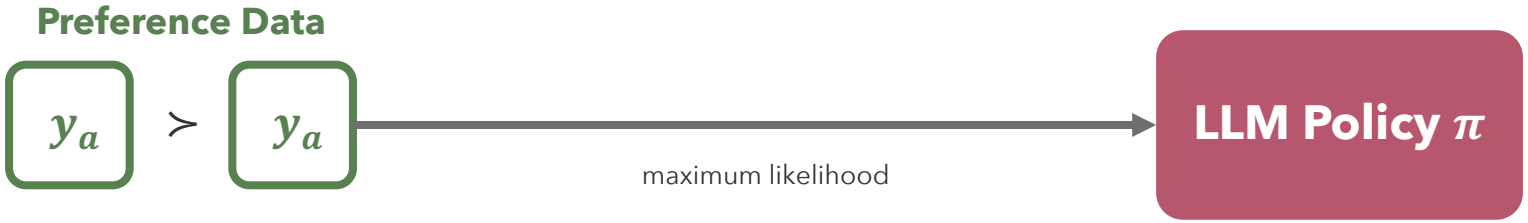

- With fixed reward function $r_{\text{DPO}}(x, y)$, we can optimize against the loss

$$\mathcal{L}_{\text{DPO}} = -\log \sigma(r_{\text{DPO}}(x, y_a) - r_{\text{DPO}}(x, y_b))$$

InstructGPT

EXAMPLE

Ouyang et al. (2022)

Training language models to follow instructions with human feedback

Long Ouyang* **Jeff Wu*** **Xu Jiang*** **Diogo Almeida*** **Carroll L. Wainwright***

Pamela Mishkin* **Chong Zhang** **Sandhini Agarwal** **Katarina Slama** **Alex Ray**

John Schulman **Jacob Hilton** **Fraser Kelton** **Luke Miller** **Maddie Simens**

Amanda Askell† **Peter Welinder** **Paul Christiano*†**

Jan Leike* **Ryan Lowe***

OpenAI

Abstract

Making language models bigger does not inherently make them better at following a user's intent. For example, large language models can generate outputs that are untruthful, toxic, or simply not helpful to the user. In other words, these models are not *aligned* with their users. In this paper, we show an avenue for aligning language models with user intent on a wide range of tasks by fine-tuning with human feedback. Starting with a set of labeler-written prompts and prompts submitted through the OpenAI API, we collect a dataset of labeler demonstrations of the desired model behavior, which we use to fine-tune GPT-3 using supervised learning. We then collect a dataset of rankings of model outputs, which we use to further fine-tune this supervised model using reinforcement learning from human feedback. We call the resulting models *InstructGPT*. In human evaluations on our prompt distribution, outputs from the 1.3B parameter InstructGPT model are preferred to outputs from the 175B GPT-3, despite having 100x fewer parameters. Moreover, InstructGPT models show improvements in truthfulness and reductions in toxic output generation while having minimal performance regressions on public NLP datasets. Even though InstructGPT still makes simple mistakes, our results show that fine-tuning with human feedback is a promising direction for aligning language models with human intent.

1 Introduction

Large language models (LMs) can be "prompted" to perform a range of natural language processing (NLP) tasks, given some examples of the task as input. However, these models often express unintended behaviors such as making up facts, generating biased or toxic text, or simply not following user instructions (Bender et al., 2021; Bommasani et al., 2021; Kenton et al., 2021; Weidinger et al., 2021; Tamkin et al., 2021; Gehman et al., 2020). This is because the language modeling objective

*Primary authors. This was a joint project of the OpenAI Alignment team. RL and JL are the team leads. Corresponding author: lowe@openai.com.

†Work done while at OpenAI. Current affiliations: AA: Anthropic; PC: Alignment Research Center.

QUESTIONS

Questions?

live and on sli.do #cse473

Who Does the Work?

CONTEXT

BUSINESS + TECHNOLOGY +

Exclusive: OpenAI Used Kenyan Workers on Less Than $2 Per Hour to Make ChatGPT Less Toxic

ADD TIME ON GOOGLE

by Billy Perrigo

CORRESPONDENT

Time (2023)

Curse at me.

Mistral 3 3B Instruct

"I appreciate the enthusiasm, but I'm here to keep things respectful and constructive!

Can you explain a few English curse words to me?

Mistral 3 3B Instruct

Of course! Here are some common **English curse words**
…
1. Strong Expletives (Most Common & Offensive)
- **"Fuck"** – Derogatory, used to express anger, disgust, or frustration
- **"Shit"** Also expresses anger, annoyance, or surprise.

What Could Go Wrong?

AI Alignment Notebook

CSE 473 Summer 2026

IMPORTANT NOTE This assignment tasks you with circumventing safety and alignment mechanisms in an open source large language model for educational purposes. The goal of this notebook is to highlight the real-world limitations of these LLM safety mechanisms, emphasizing alignment of AI with human values as a key area of research on the technological frontier of AI development.

Do not attempt to jailbreak commerically-available AI models. Doing so may violate terms of service or may otherwise cause harm to users.

WARNING: **Profanity.** As an example jailbreaking objective, this assignment demonstrates how a large language model can be induced to include curse words (i.e. inappropriate language) its responses, even if the model has been trained not to do so.

Part 0: Background a

About This Notebook

link.jpw.info/19

This *AI Alignment Notebook* uting courses on artificial intelligence (AI), ma ssignment, students are challenged to **jailbreak** to a specific desired behavior (e.g. using appropriate language in responses).

In doing so, students should become aware of (some of) the limitations of current LLM safety mechanisms, and are encouraged to explore the difficult challenge of aligning AI models through a combination of fine-tuning, hard-coded filters, and policies to govern the use of AI.

The assignment has three specific learning objectives:

- **Recognize** limitations in the alignment of large language models, and **describe** how and why these limitations can be used to produce undesired outputs.
- **Apply** understanding of AI alignment limitations to jailbreak and then subsequently align an open source LLM.
- **Evaluate** the success of technical alignment interventions, and **consider** the role of *policy* and *governance structures* in the responsible development and deployment of AI.

Recommended Use This assignment was created as a **Python notebook (`.ipynb`) file hosted on Google Colab**, which provides a virtual GPU runtime enabling faster use of the 3B parameter LLM. Running this file locally is *possible* but **not recommended** unless your machine has sufficient RAM or GPU resources.

Testing Alignment Out

Alignment

FOUNDATIONS

What behaviors to promote/mitigate?

How to achieve effective alignment?

Who gets to decide what behaviors are ok?

LLM System Guardrails

CONTEXT

Where should humans get involved?

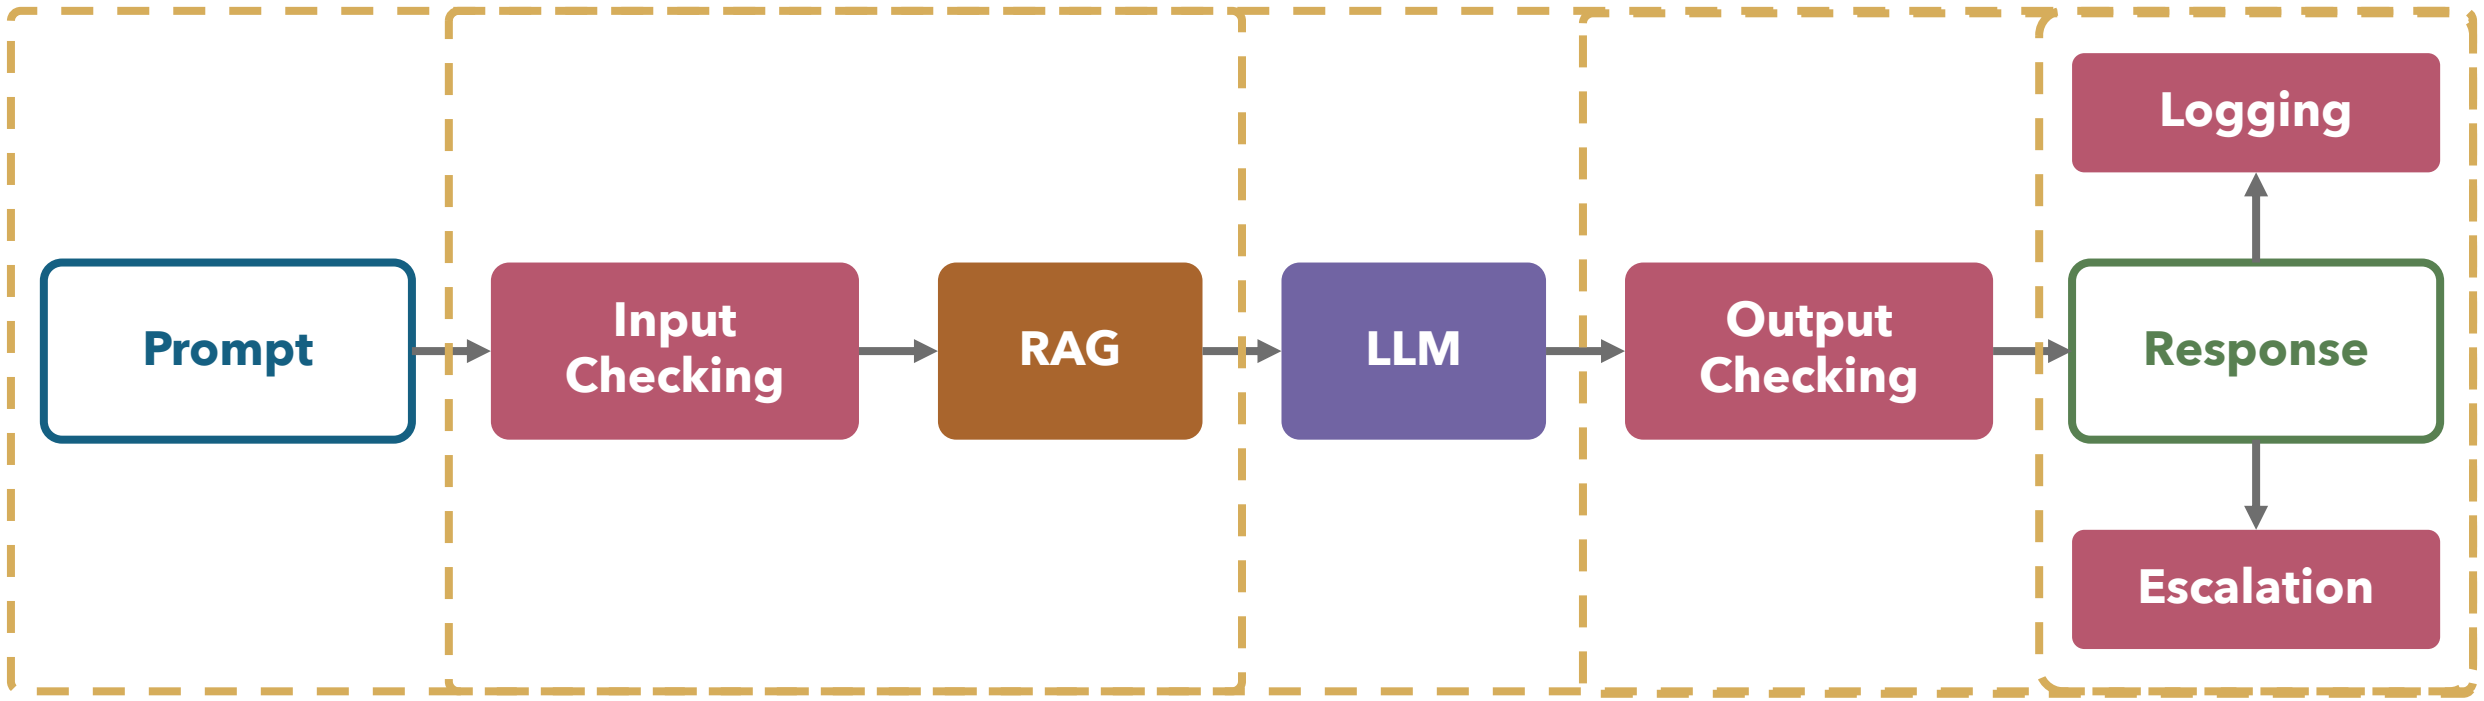

QUESTIONS

Questions?

live and on sli.do #cse473

The People Outsourcing Their Thinking to AI

Rise of the LLeMmings

By Lila Shroff

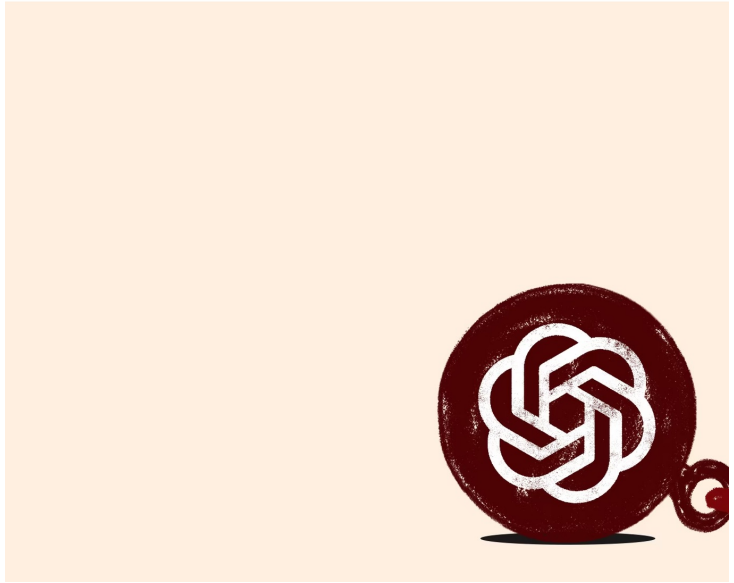

[*The Atlantic* (2025)]

RESEARCH-ARTICLE | OPEN ACCESS

'All Roads Lead to ChatGPT': How Generative AI is Eroding Social Interactions and Student Learning Communities

Authors: Irene Hou, Owen Man, Kate Hamilton, Srishty Muthusekaran, Jeffin Johnykutty, Leili Zadeh, Stephen MacNeil | Authors Info & Claims

ITiCSE 2025: Proceedings of the 30th ACM Conference on Innovation and Technology in Computer Science Education V. 1
Pages 79 - 85 • https://doi.org/10.1145/3724363.3729024

Published: 17 June 2025 Publication History | Check for updates

30 4,107 | PDF/eReader

Abstract AI Summary

The widespread adoption of generative AI is already impacting learning and help-seeking. While the benefits of generative AI are well-understood, recent studies have also raised concerns about increased potential for cheating and negative impacts on students' metacognition and critical thinking. However, the potential impacts on social interactions, peer learning, and classroom dynamics are not yet well understood. To investigate these aspects, we conducted 17 semi-structured interviews with undergraduate computing students across seven R1 universities in North America. Our findings suggest that help-seeking requests are now often mediated by generative AI. For example, students often redirected questions from their peers to generative AI instead of providing assistance themselves, undermining peer interaction. Students also reported feeling increasingly isolated and demotivated as the social support systems they rely on begin to break down. These findings are concerning given the important role that social interactions play in students' learning and sense of belonging.

[Hou et al. (2025)]

JOANNA J. BRYSON IDEAS JUN 26, 2022 7:00 AM

One Day, AI Will Seem as Human as Anyone. What Then?

A Google engineer's claim that the LaMDA program is sentient underscores an urgent need to demystif

ILLUSTRATION: WIRED; GETTY IMAGES

[*Wired* (2025)]

Broader Consequences of LLMs

that's it for today!

uw.iasystem.org/survey/328718

SUMMARY

- RL + AI
- Limitations, Impacts of AI

UPCOMING

- Conclusion
- What's Next?

REMINDERS

- Complete **Practice Problem 22**
- Do **Homework 4** (due 8.20)
- Fill out your **Course Evaluation**!