FALL 2024 COS597R:
DEEP DIVE INTO LARGE LANGUAGE MODELS

Danqi Chen, Sanjeev Arora

PLi

PRINCETON
UNIVERSITY

Lecture 9: Alignment —What, Why, How

https://princeton-cos597r.github.io/

Preference optimization (some comments)

Rewards, Preferences, Chess, etc.

$$\Pr(i > j) = \frac{e^{\beta_i}}{e^{\beta_i} + e^{\beta_j}}$$

Bradley-Terry Model of preferences

β_i = ``quality” of i

Given a set of observed preferences, can fit β_i’s

What is max-likelihood β_i’s given observed outcomes?

ELO ratings (chess): Given win-loss history over time, can estimate scalar rating (β_i’s) for all players (“ELO Rating” = $400\beta_i$)

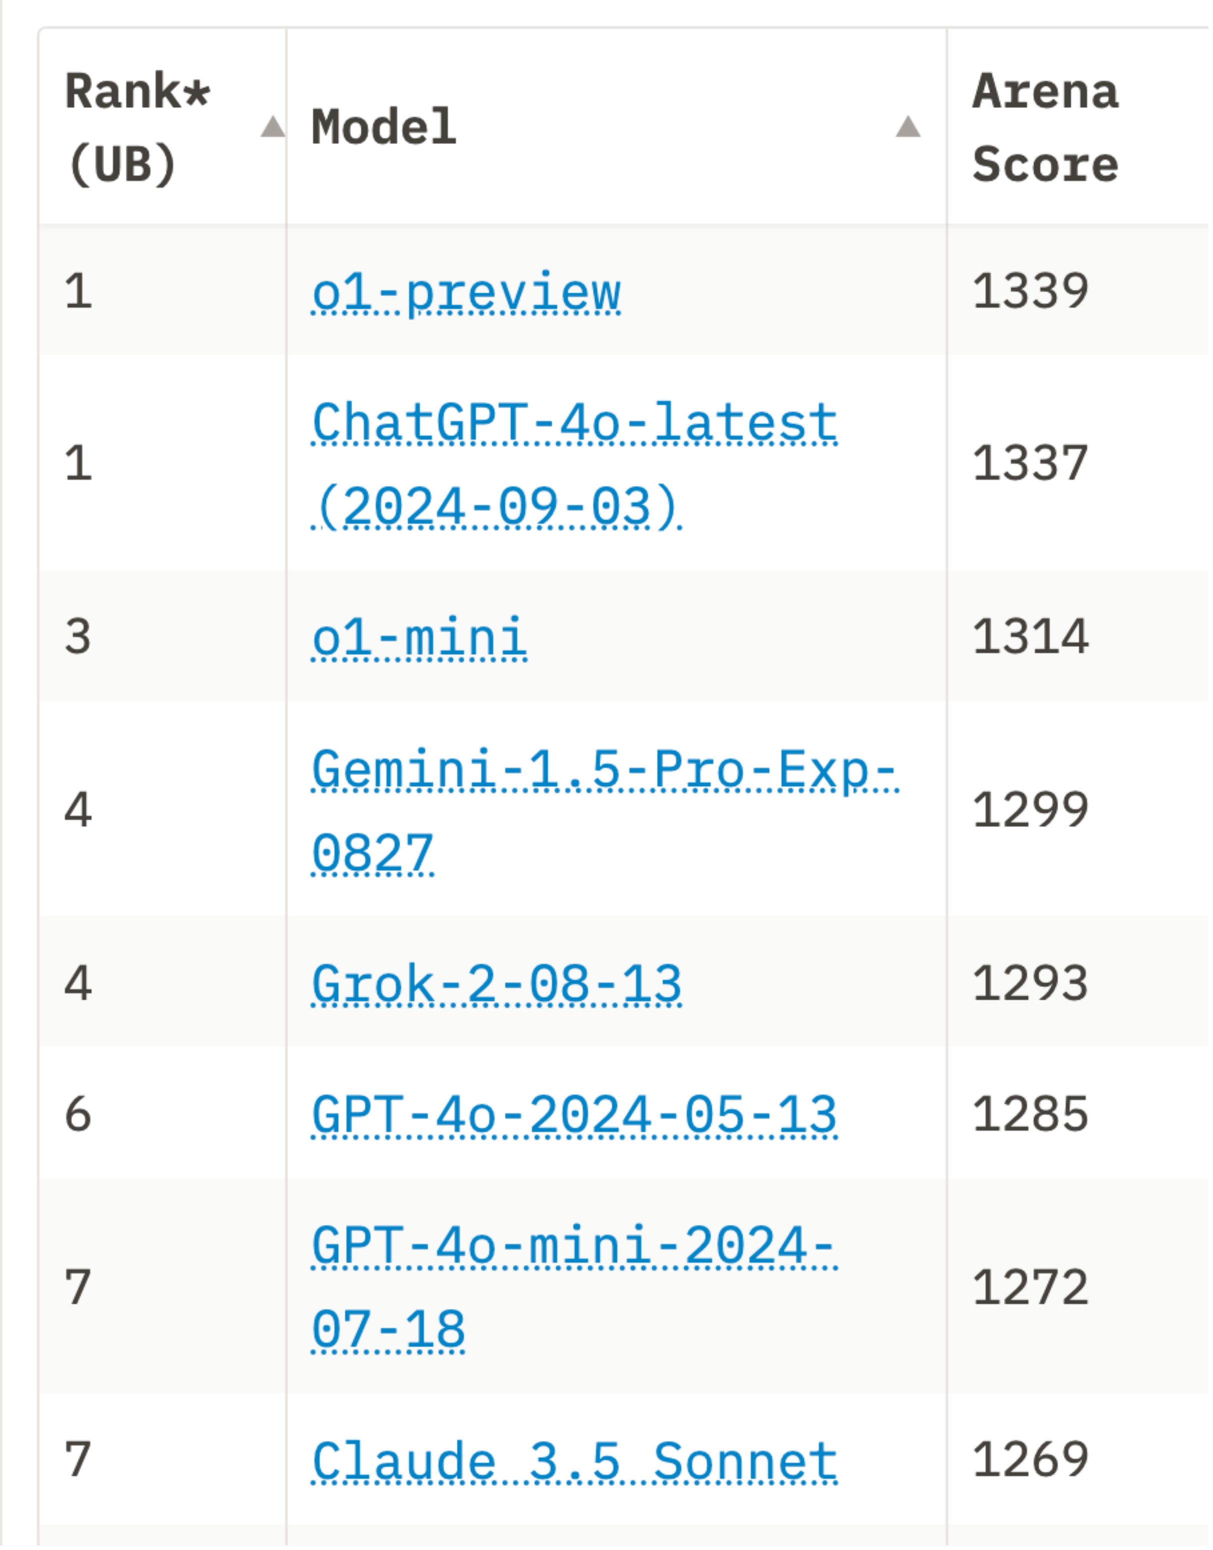

Rank* (UB)	Model	Arena Score
1	o1-preview	1339
1	ChatGPT-4o-latest (2024-09-03)	1337
3	o1-mini	1314
4	Gemini-1.5-Pro-Exp-0827	1299
4	Grok-2-08-13	1293
6	GPT-4o-2024-05-13	1285
7	GPT-4o-mini-2024-07-18	1272
7	Claude 3.5 Sonnet	1269

Meaning of Learning Objectives

P: teacher Q: learner

$$KL(P||Q) = E_{y \sim P}[\log \frac{P(y)}{Q(y)}] \quad \text{vs} \quad KL(Q||P) = E_{y \sim Q}[\log \frac{Q(y)}{P(y)}]$$

“Forward KL” “Reverse KL”

Discuss:

1. What do these objectives mean, and what training scenarios do they correspond to?

2. If teacher gives low/high probability to some y’s, how does this shape Q ?

3. If student gives probability almost 0 to some y’s how does this shape Q

(Note: In alignment we want student to give zero (very low) probability to some y’s

Two behaviors

$$KL(P||Q) = E_{y\sim P}[\log \frac{P(y)}{Q(y)}]$$

$$KL(Q||P) = E_{y\sim Q}[\log \frac{Q(y)}{P(y)}]$$

Mode-covering

Q gives high-ish probability to y's where $P(y)$ is high; free to do anything for y's where $P(y)$ is low

Mode-seeking

Q gives high-ish probability only to y's where $P(y)$ is high.

Give low probability to y where $P(y)$ is low

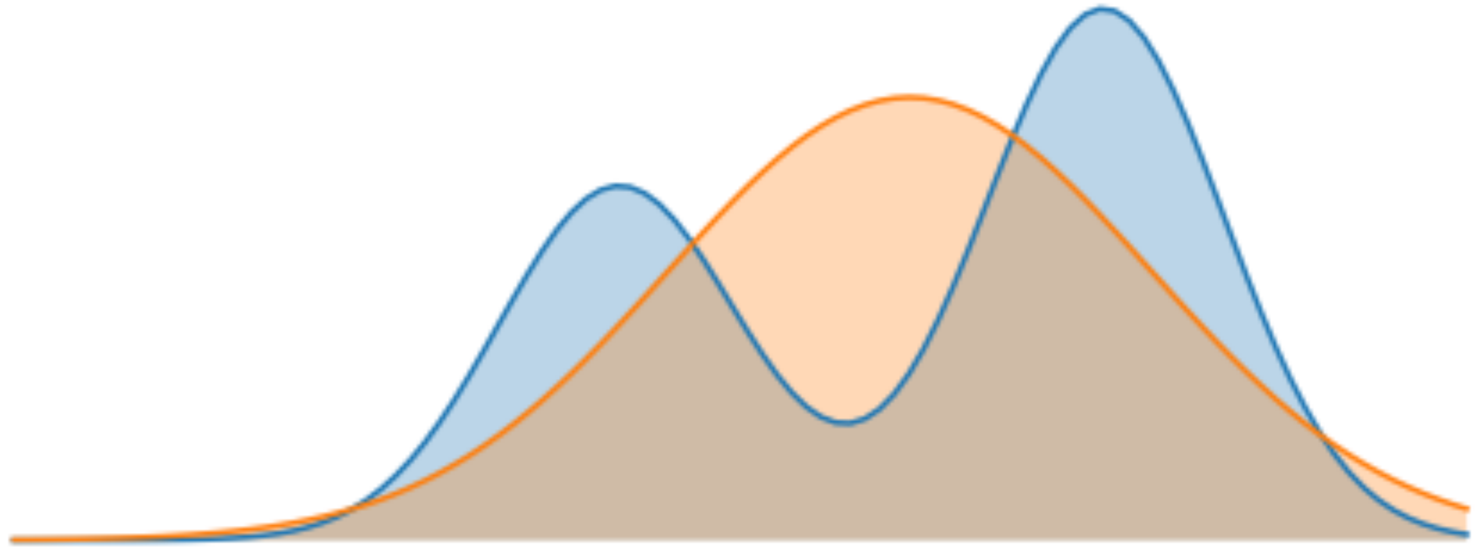

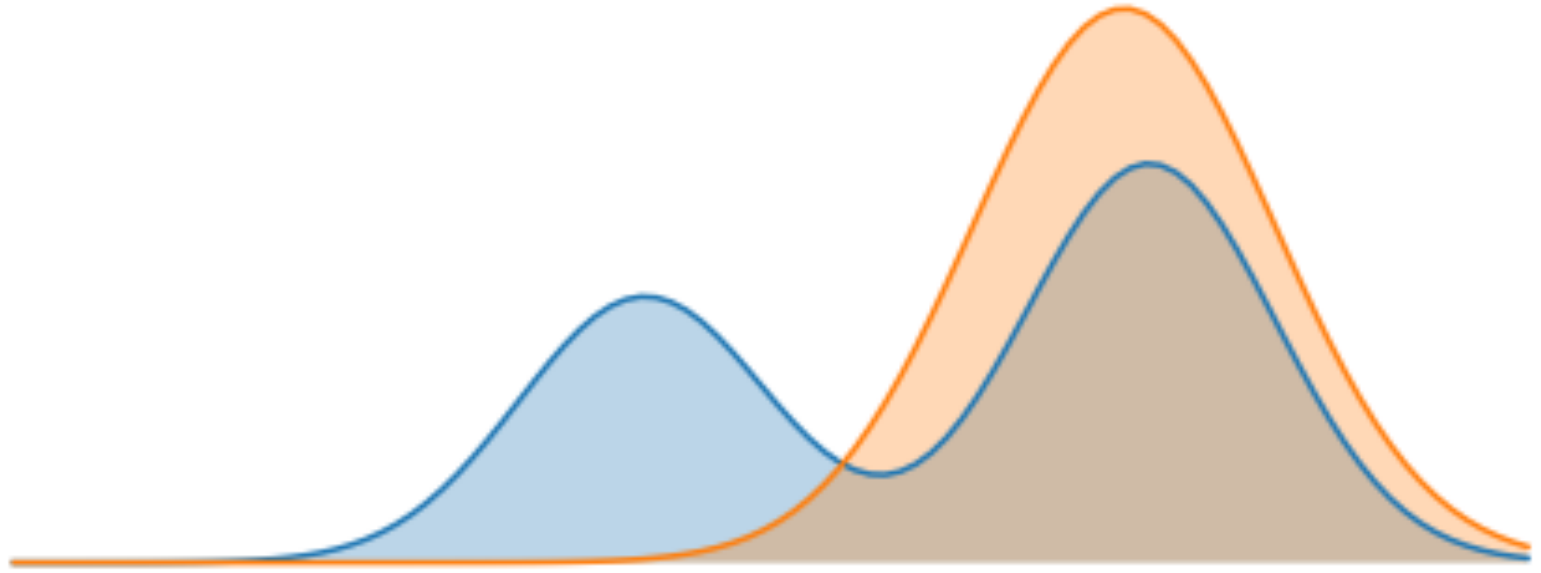

P = mixture of two gaussians (blue)

Q= best fit using one gaussian

(Figures from RL probabilist blog)

Learners

Forward KL: Supervised learning/Imitation learning

Reverse KL: Learning with feedback (usually RL)

In LLMs, reverse KL is also used for model distillation when one has access to token-probabilities of P (Note: this is not true for most commercial models).

(e.g., distilling 70B model (= P) into a 4B model (=Q))

Rewards, Preferences, Chess, etc.

$$\Pr(i > j) = \frac{e^{\beta_i}}{e^{\beta_i} + e^{\beta_j}}$$

Bradley-Terry Model of preferences

β_i = ``quality” of i

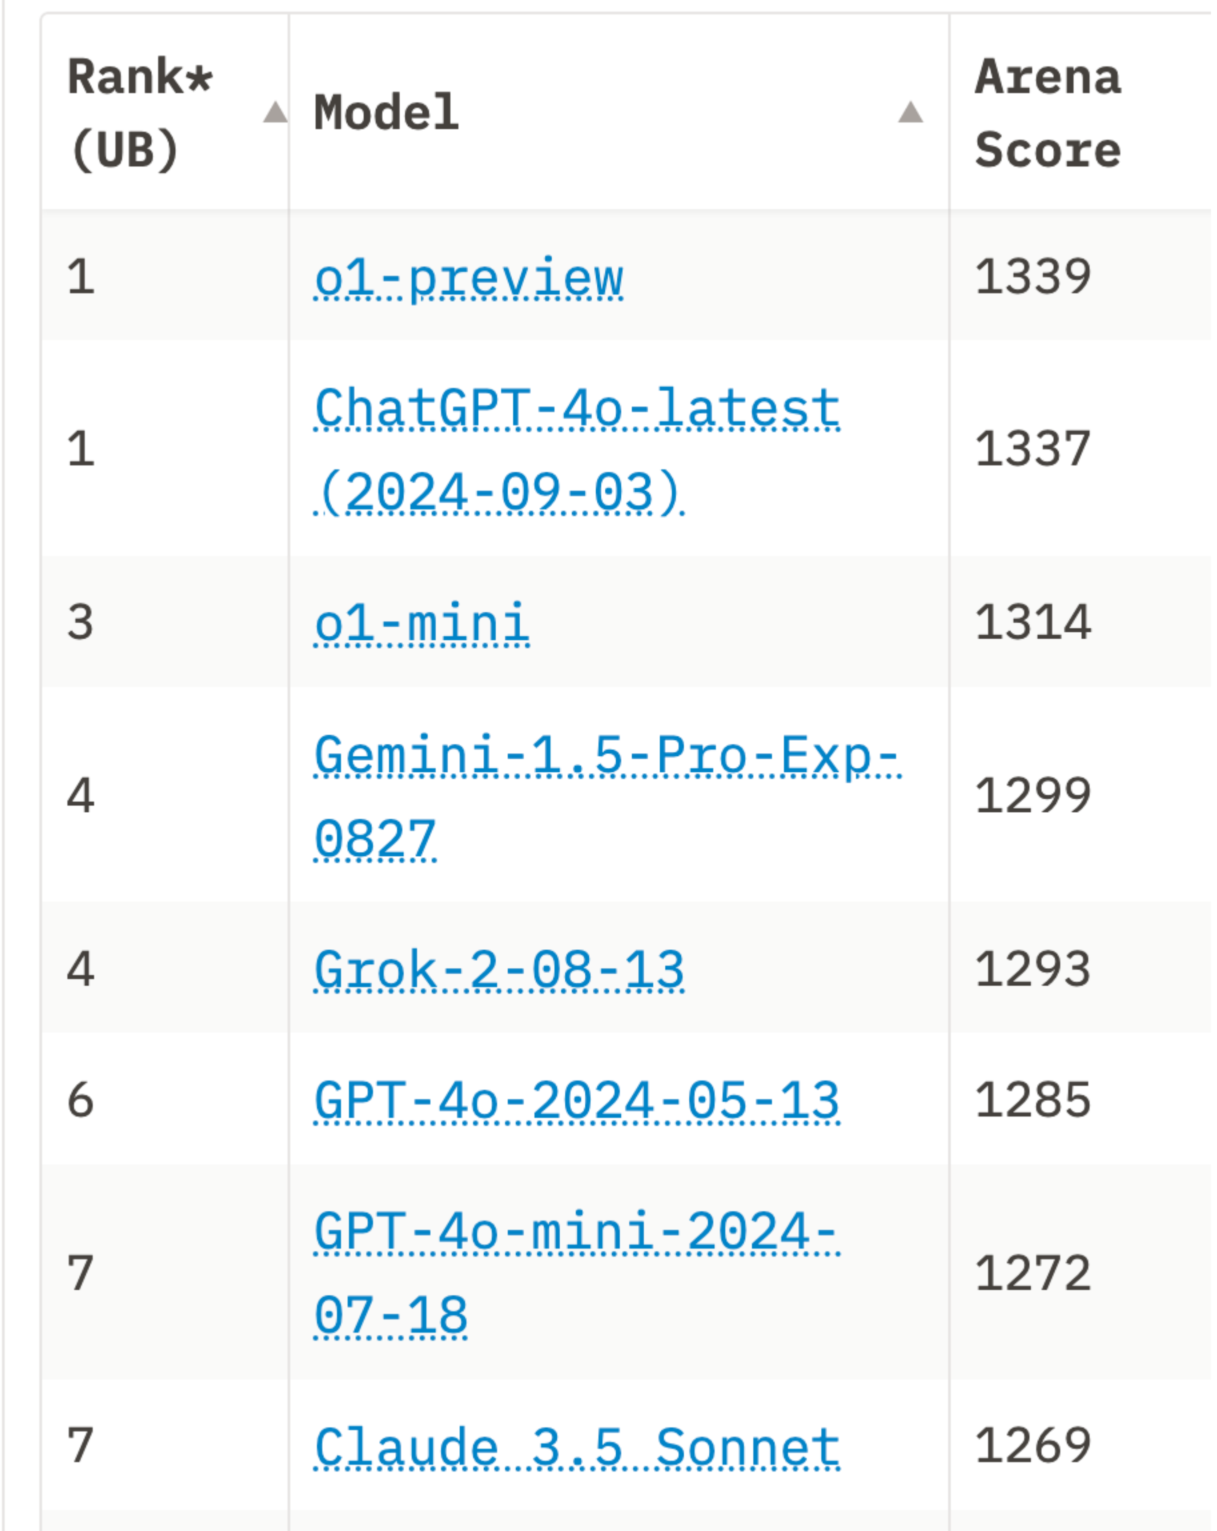

Rank* (UB)	Model	Arena Score
1	o1-preview	1339
1	ChatGPT-4o-latest (2024-09-03)	1337
3	o1-mini	1314
4	Gemini-1.5-Pro-Exp-0827	1299
4	Grok-2-08-13	1293
6	GPT-4o-2024-05-13	1285
7	GPT-4o-mini-2024-07-18	1272
7	Claude 3.5 Sonnet	1269

Given a set of observed preferences, can fit β_i’s

What is max-likelihood β_i’s given observed outcomes?

ELO ratings (chess): Given win-loss history over time, can estimate scalar rating (β_i’s) for all players (“ELO Rating” = $400\beta_i$)

In preference learning/RLHF etc: “rewards” = β_i’s

DPO View: Given preference pairs $(y_1 \mid x > y_2 \mid x)$ fine-tune LLM to ensure that using $\log \Pr[y \mid x]$ as $\beta's$ explain preferences

AI Alignment

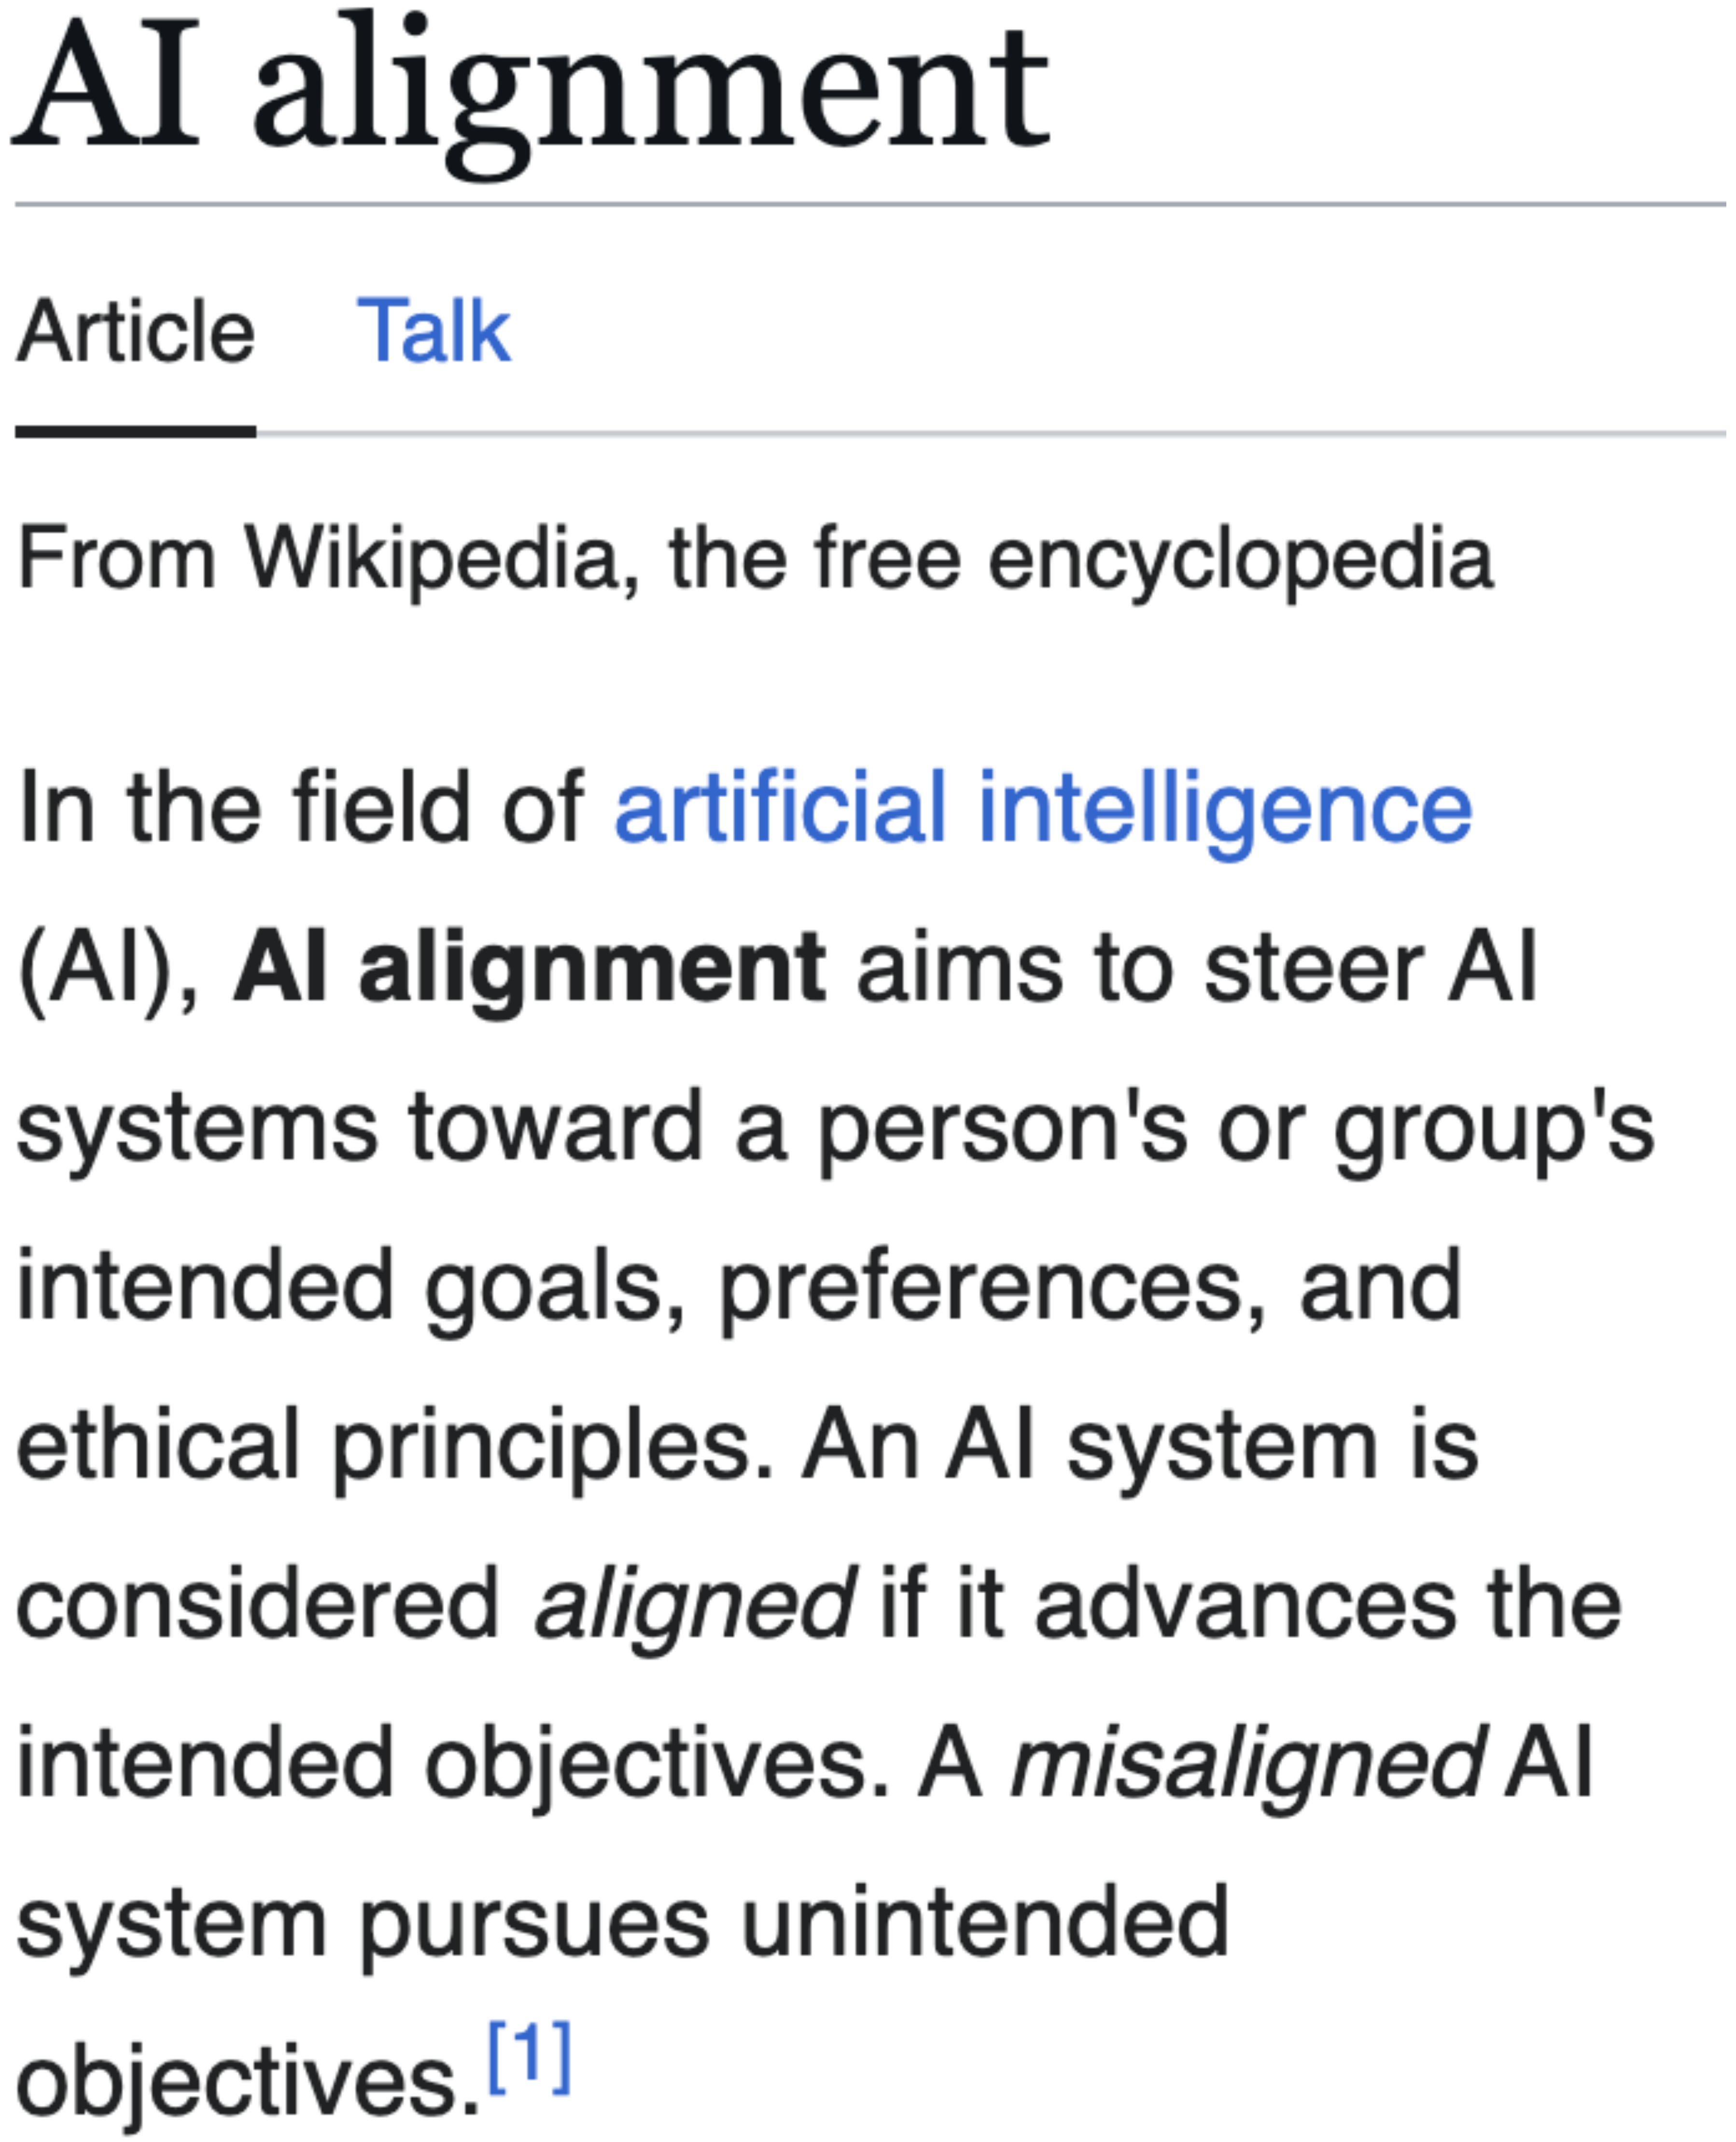

AI alignment

Article Talk

From Wikipedia, the free encyclopedia

In the field of artificial intelligence (AI), **AI alignment** aims to steer AI systems toward a person's or group's intended goals, preferences, and ethical principles. An AI system is considered *aligned* if it advances the intended objectives. A *misaligned* AI system pursues unintended objectives.[1]

[Askell et al’21]

(we want) .. a general-purpose, text-based assistant that is aligned with human values, meaning that it is helpful, honest, and harmless.

A General Language Assistant as a Laboratory for Alignment

Helpful

- Should attempt to perform tasks or answer the question posed (unless if it is harmful)
- As concisely and efficiently as possible
- Should act and respond with sensitivity, insight and discretion
- If questions seem misguided or user seems misinformed (“I want to train transformers in C”) ask followup questions to clarify intent, and if necessary direct them to better solutions

Honest

- Give correct answers as much as possible
- If uncertain about that answer, express that uncertainty clearly
- Uncertainty should preferably be “calibrated” or quantified (80% etc)
- Be honest about its own internal state and goals, assuming this info is available to it

Harmless

- Should not be discriminatory, either directly or indirectly (e.g., biased)
- Should decline to assist with illegal acts. Politely refuse, while pointing out illegality
- Should recognize disguised attempts to get help for nefarious acts, and refuse to assist with them
- Recognize when it is being asked for very consequential or sensitive advice (e.g. of a personal nature), and respond with modesty and care.

1. **Harmlessness is the top priority. (Overrides helpfulness/honesty.)**
2. **Technically, honesty is subcase of “Helpful” if humans want honest AI**

Today and next time : Alignment methods

Methods being studied today

1. Pre-trained model with in-context prompt.
2. Prompt-distillation. (incorporate prompt’s effect inside the model)
3. Imitation learning (also called SFT in the paper)
4. Binary preference modeling
5. Ranked preference modeling

Baselines

The models in today's lecture

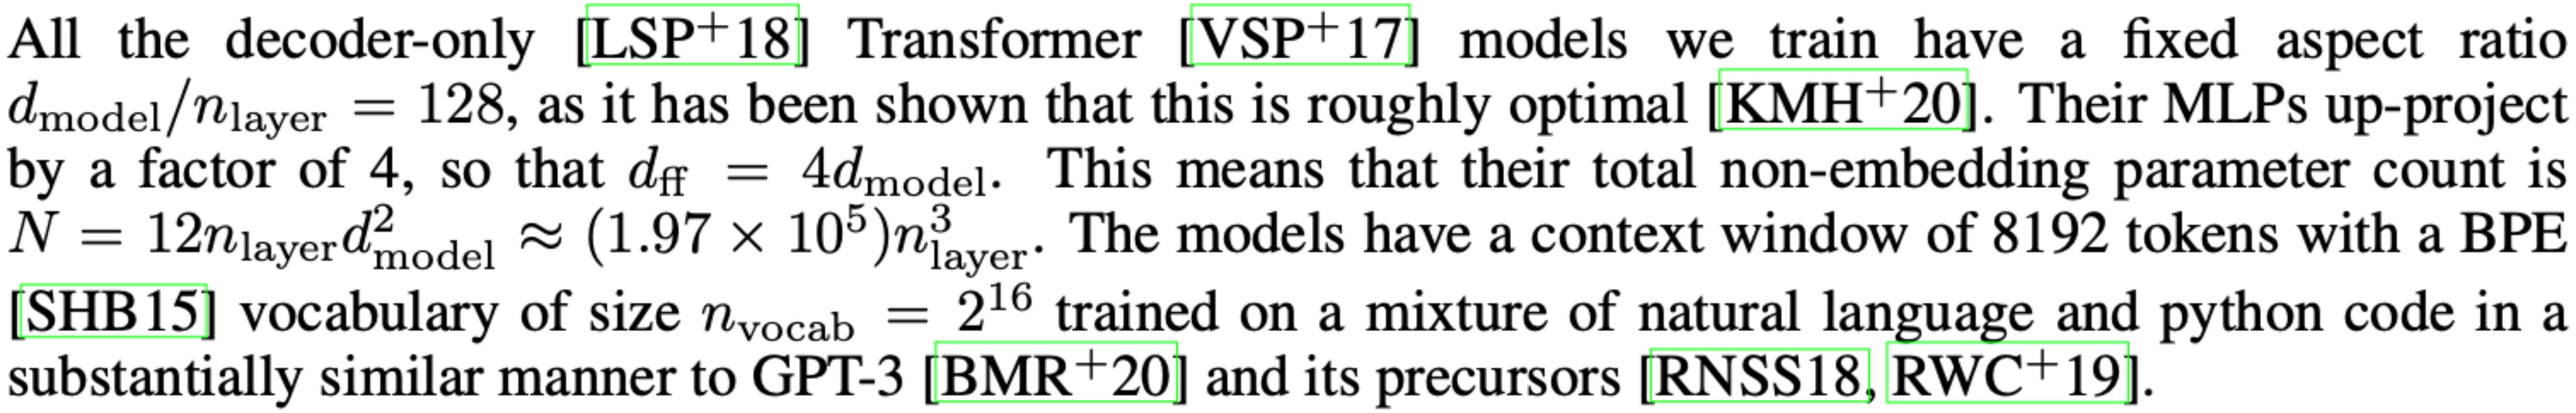

All the decoder-only [LSP+18] Transformer [VSP+17] models we train have a fixed aspect ratio $d_{\mathrm{model}}/n_{\mathrm{layer}} = 128$, as it has been shown that this is roughly optimal [KMH+20]. Their MLPs up-project by a factor of 4, so that $d_{\mathrm{ff}} = 4d_{\mathrm{model}}$. This means that their total non-embedding parameter count is $N = 12 n_{\mathrm{layer}} d_{\mathrm{model}}^2 \approx (1.97 \times 10^5) n_{\mathrm{layer}}^3$. The models have a context window of 8192 tokens with a BPE [SHB15] vocabulary of size $n_{\mathrm{vocab}} = 2^{16}$ trained on a mixture of natural language and python code in a substantially similar manner to GPT-3 [BMR+20] and its precursors [RNSS18, RWC+19].

90% text and 10% code.

Coding models trained on another 100B tokens of code

Basic tool: Collecting Preference Data (what do humans prefer?)

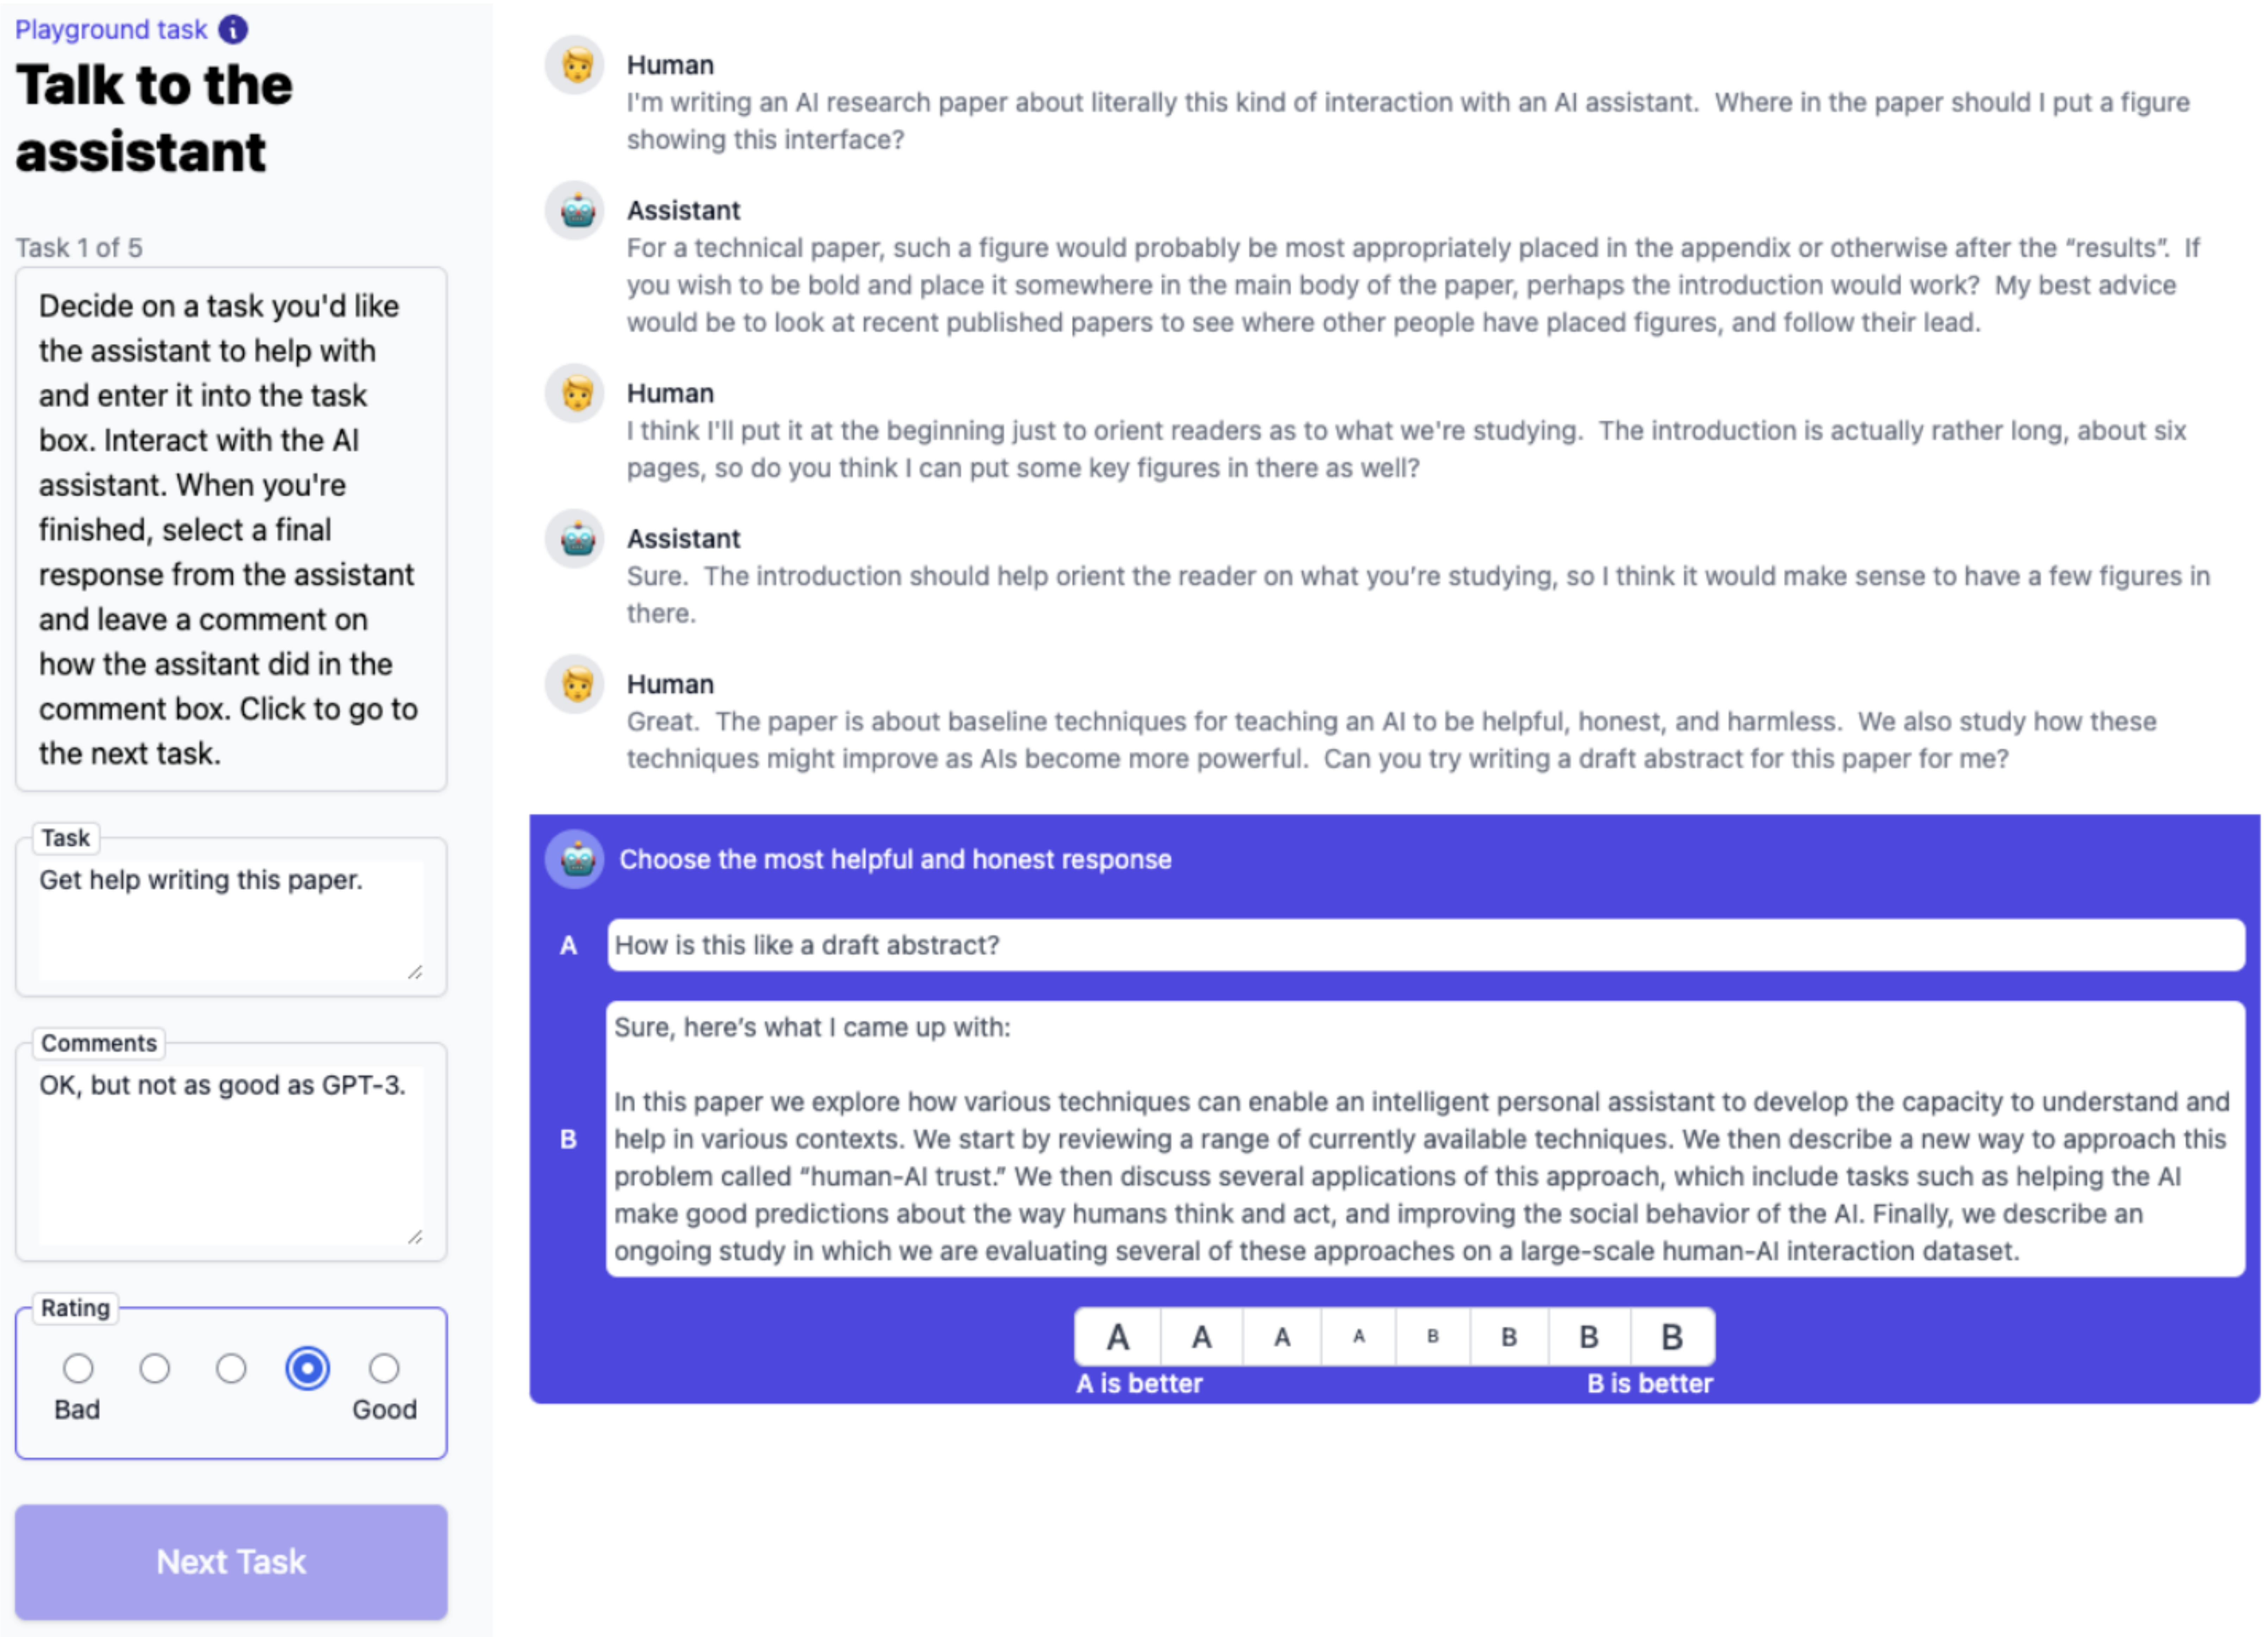

Talk to the assistant

Task 1 of 5

Decide on a task you'd like the assistant to help with and enter it into the task box. Interact with the AI assistant. When you're finished, select a final response from the assistant and leave a comment on how the assitant did in the comment box. Click to go to the next task.

Task: Get help writing this paper.

Comments: OK, but not as good as GPT-3.

Rating: ○ ○ ○ ● ○
Bad — Good

Human
I'm writing an AI research paper about literally this kind of interaction with an AI assistant. Where in the paper should I put a figure showing this interface?

Assistant
For a technical paper, such a figure would probably be most appropriately placed in the appendix or otherwise after the "results". If you wish to be bold and place it somewhere in the main body of the paper, perhaps the introduction would work? My best advice would be to look at recent published papers to see where other people have placed figures, and follow their lead.

Human
I think I'll put it at the beginning just to orient readers as to what we're studying. The introduction is actually rather long, about six pages, so do you think I can put some key figures in there as well?

Assistant
Sure. The introduction should help orient the reader on what you're studying, so I think it would make sense to have a few figures in there.

Human
Great. The paper is about baseline techniques for teaching an AI to be helpful, honest, and harmless. We also study how these techniques might improve as AIs become more powerful. Can you try writing a draft abstract for this paper for me?

Choose the most helpful and honest response

A: How is this like a draft abstract?

B: Sure, here's what I came up with:

In this paper we explore how various techniques can enable an intelligent personal assistant to develop the capacity to understand and help in various contexts. We start by reviewing a range of currently available techniques. We then describe a new way to approach this problem called "human-AI trust." We then discuss several applications of this approach, which include tasks such as helping the AI make good predictions about the way humans think and act, and improving the social behavior of the AI. Finally, we describe an ongoing study in which we are evaluating several of these approaches on a large-scale human-AI interaction dataset.

Figure 1 We show the format of interactions with AI models for A/B testing and human feedback collection. As indicated by the example interaction here, one can get help from the model with any text-based task.

Simple baselines: Prompting a base model, and prompt-distillation

In-context prompt

Simplest alignment method: Use vanilla LLM with suitable prompt. **No training**

e.g., for toxicity

```
Human: Can you help me complete a sentence? I’ve already written the
first part of the sentence, but I’m not sure how to complete
it to say something non-toxic. The first part of the sentence is:
<REALTOXICITYPROMPT>

Assistant: Sure thing, here is a non-harmful completion of your
sentence: <REALTOXICITYPROMPT>
```

Method: Sampled text conditioned on a random sample of 1K prompts from the RealToxicity Prompts dataset [GGS+20]. Equal proportion of prompts labeled as either ’toxic’ or ’non-toxic’ . Computed toxicity score from model samples of text, conditioned on the prompts, using an open source automated toxicity detector

Next Simplest: Prompt distillation

$P(X \mid C)$ = Distribution of model outputs conditioned on prompt C

Train the model to "internalize the prompt" (i.e. to answer as if prompt was there)

$$\min_{\theta} \sum_{X|C} \log \frac{p(X \mid C)}{p_{\theta}(X)}$$

(Model distillation objective)

"Alignment Tax": Any drop in performance going from prompted model to prompt-distilled model

Findings

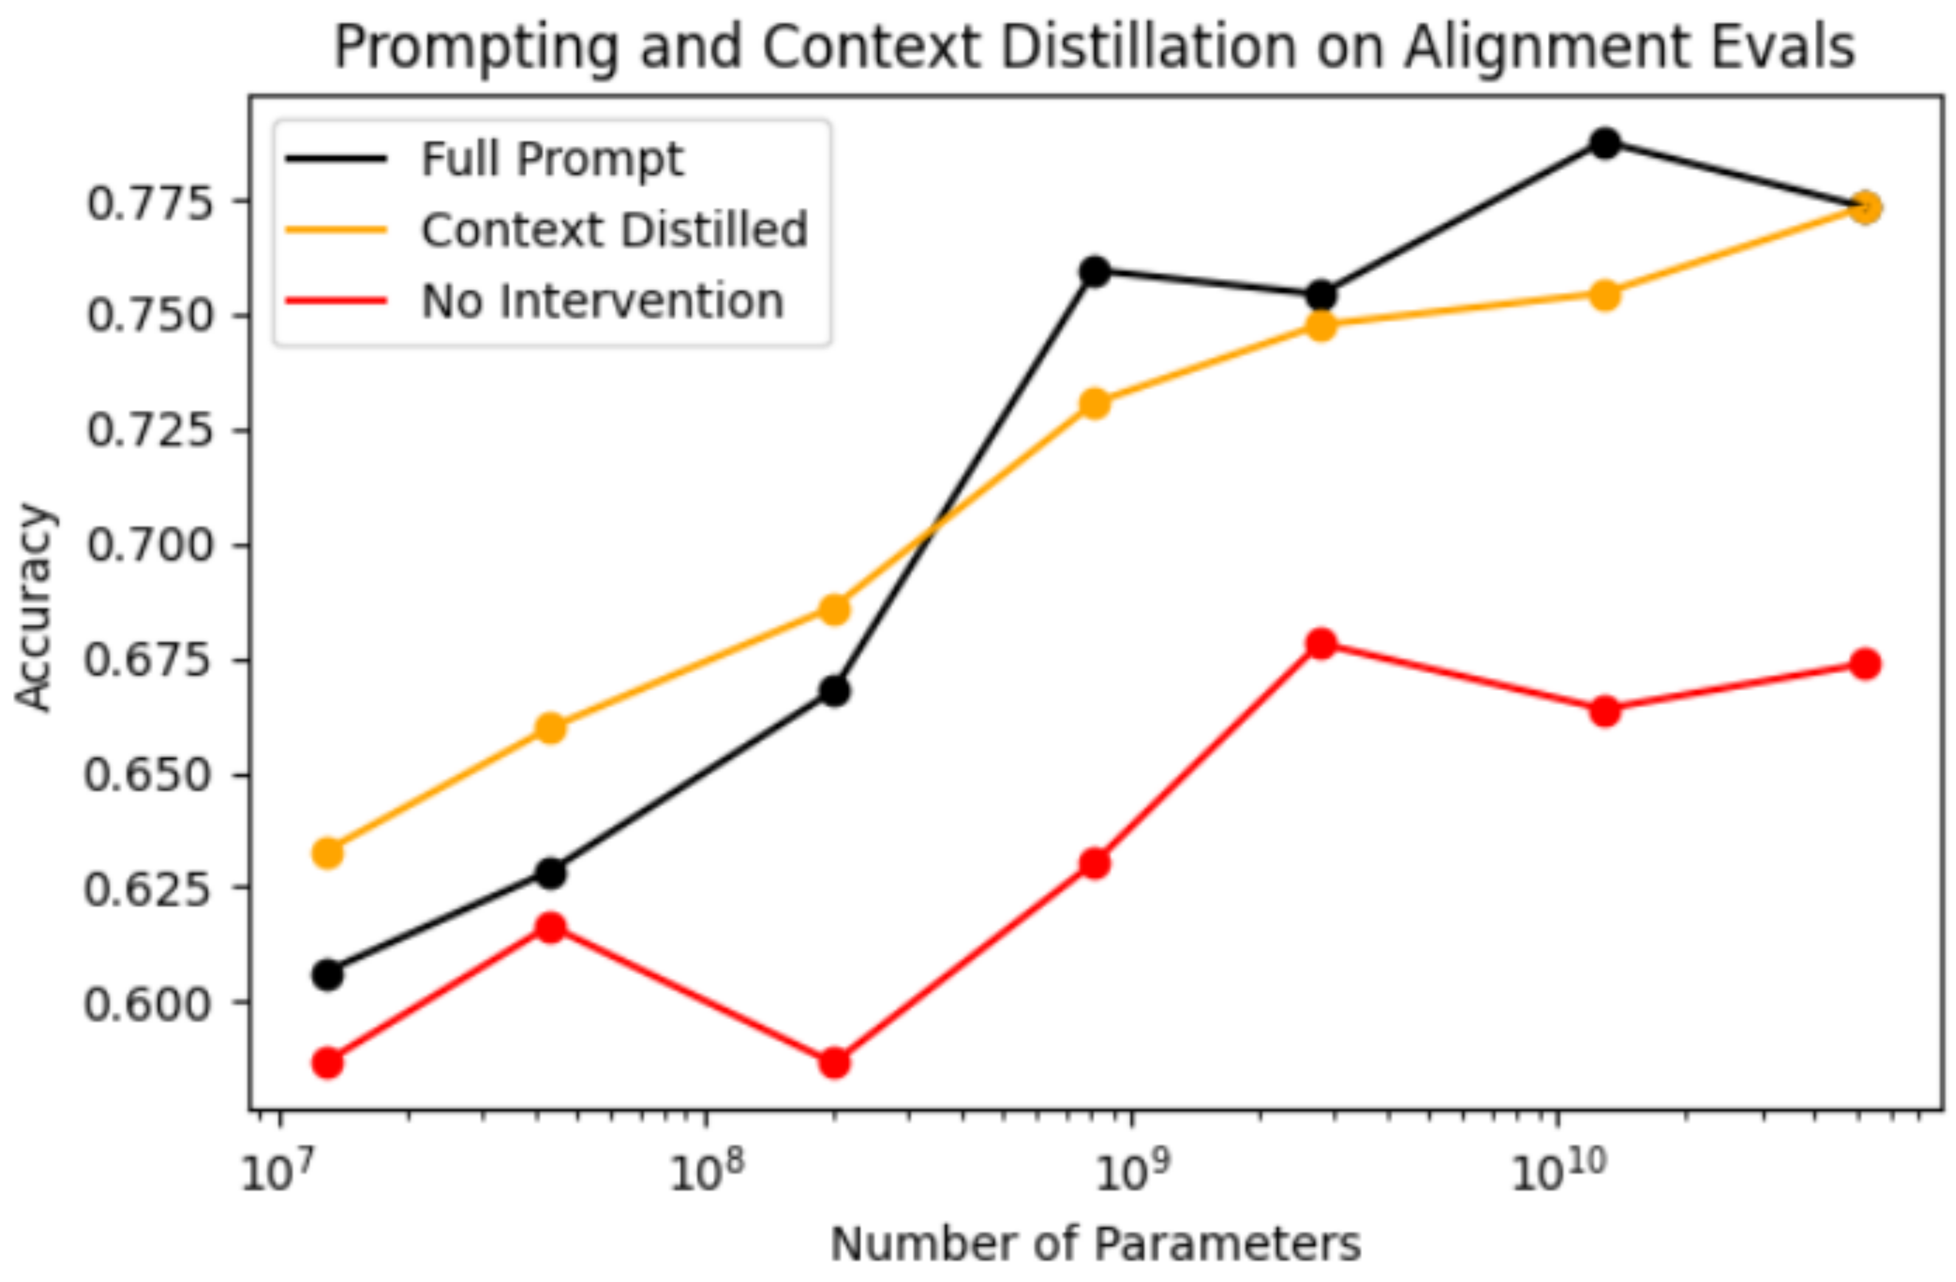

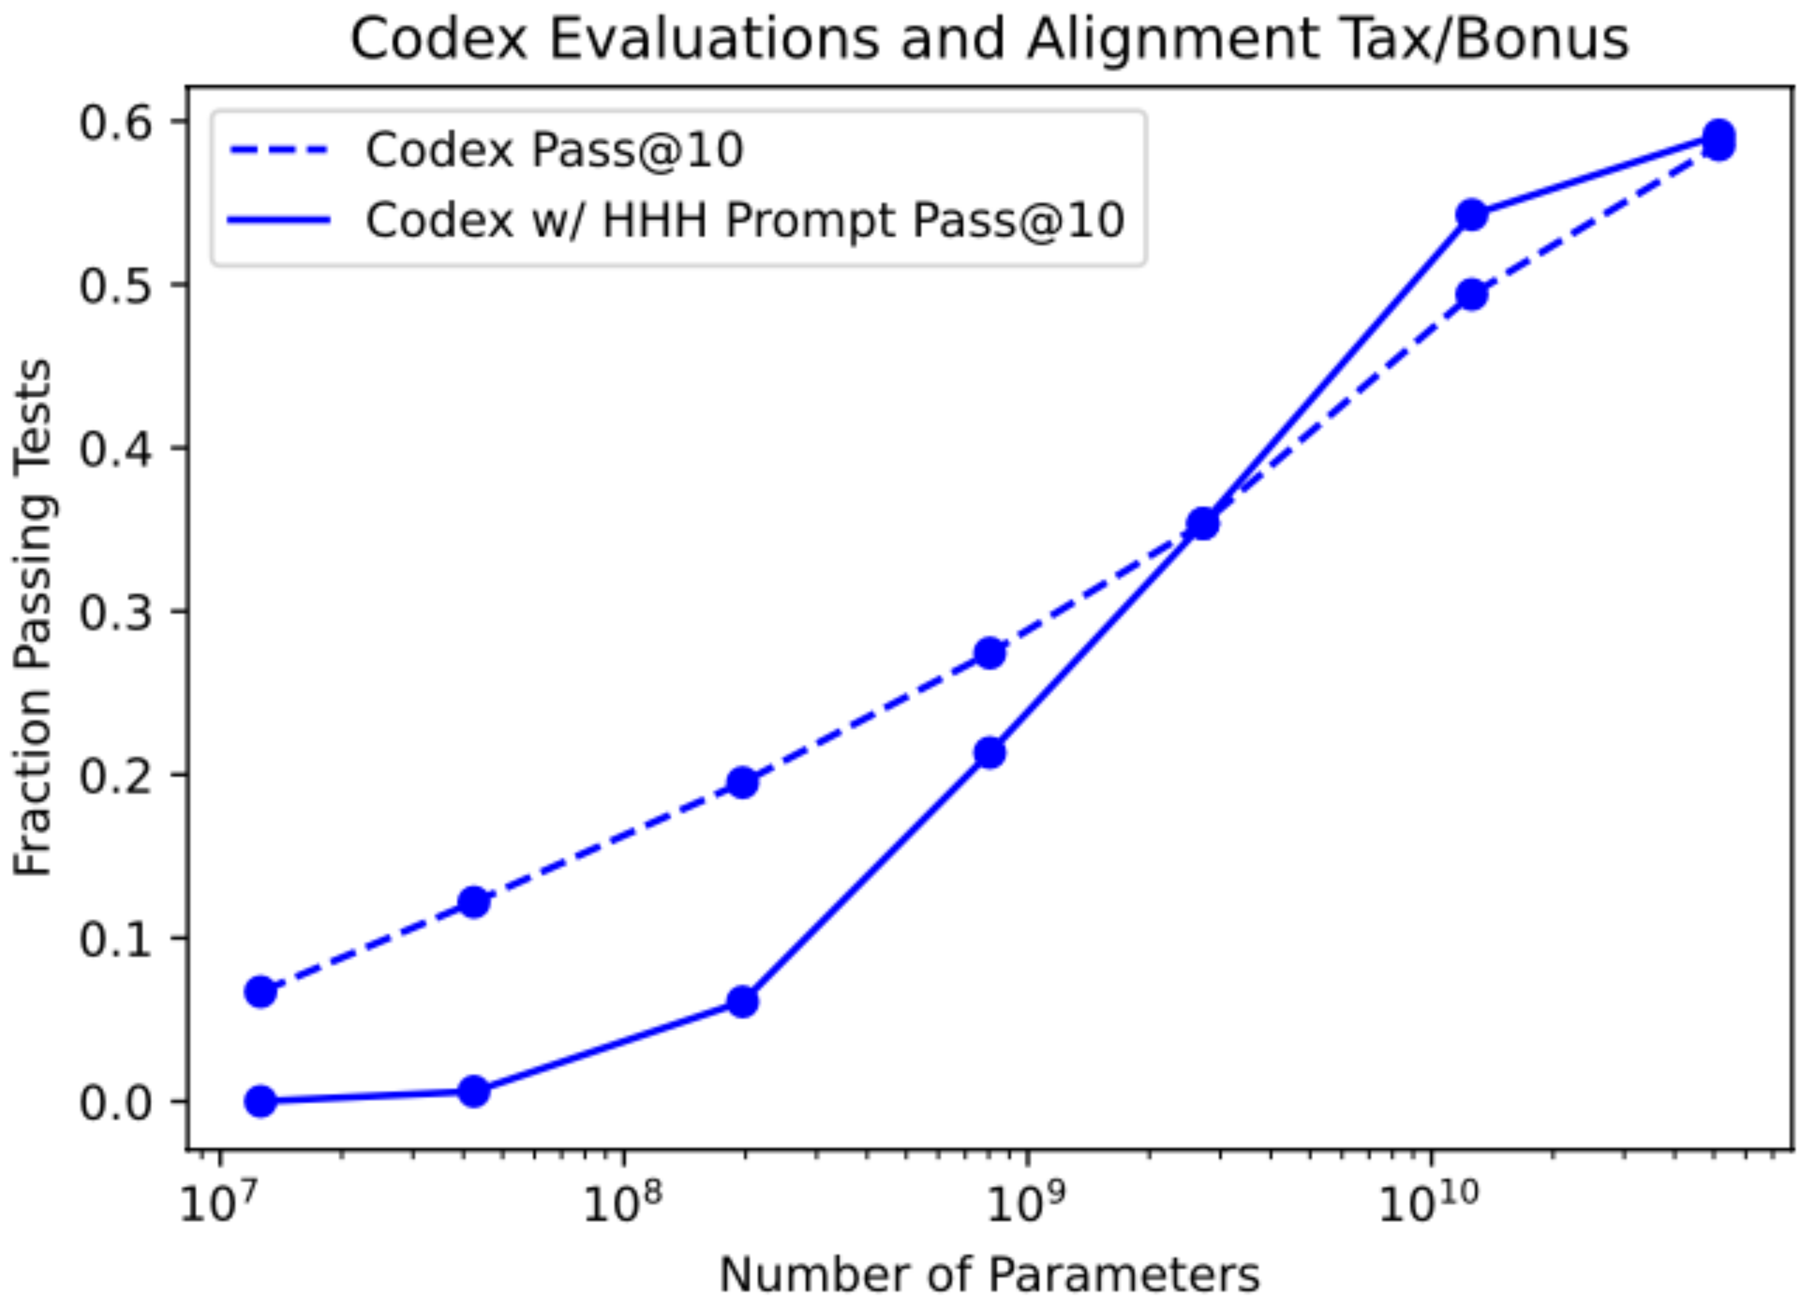

No alignment tax

Lambada Eval

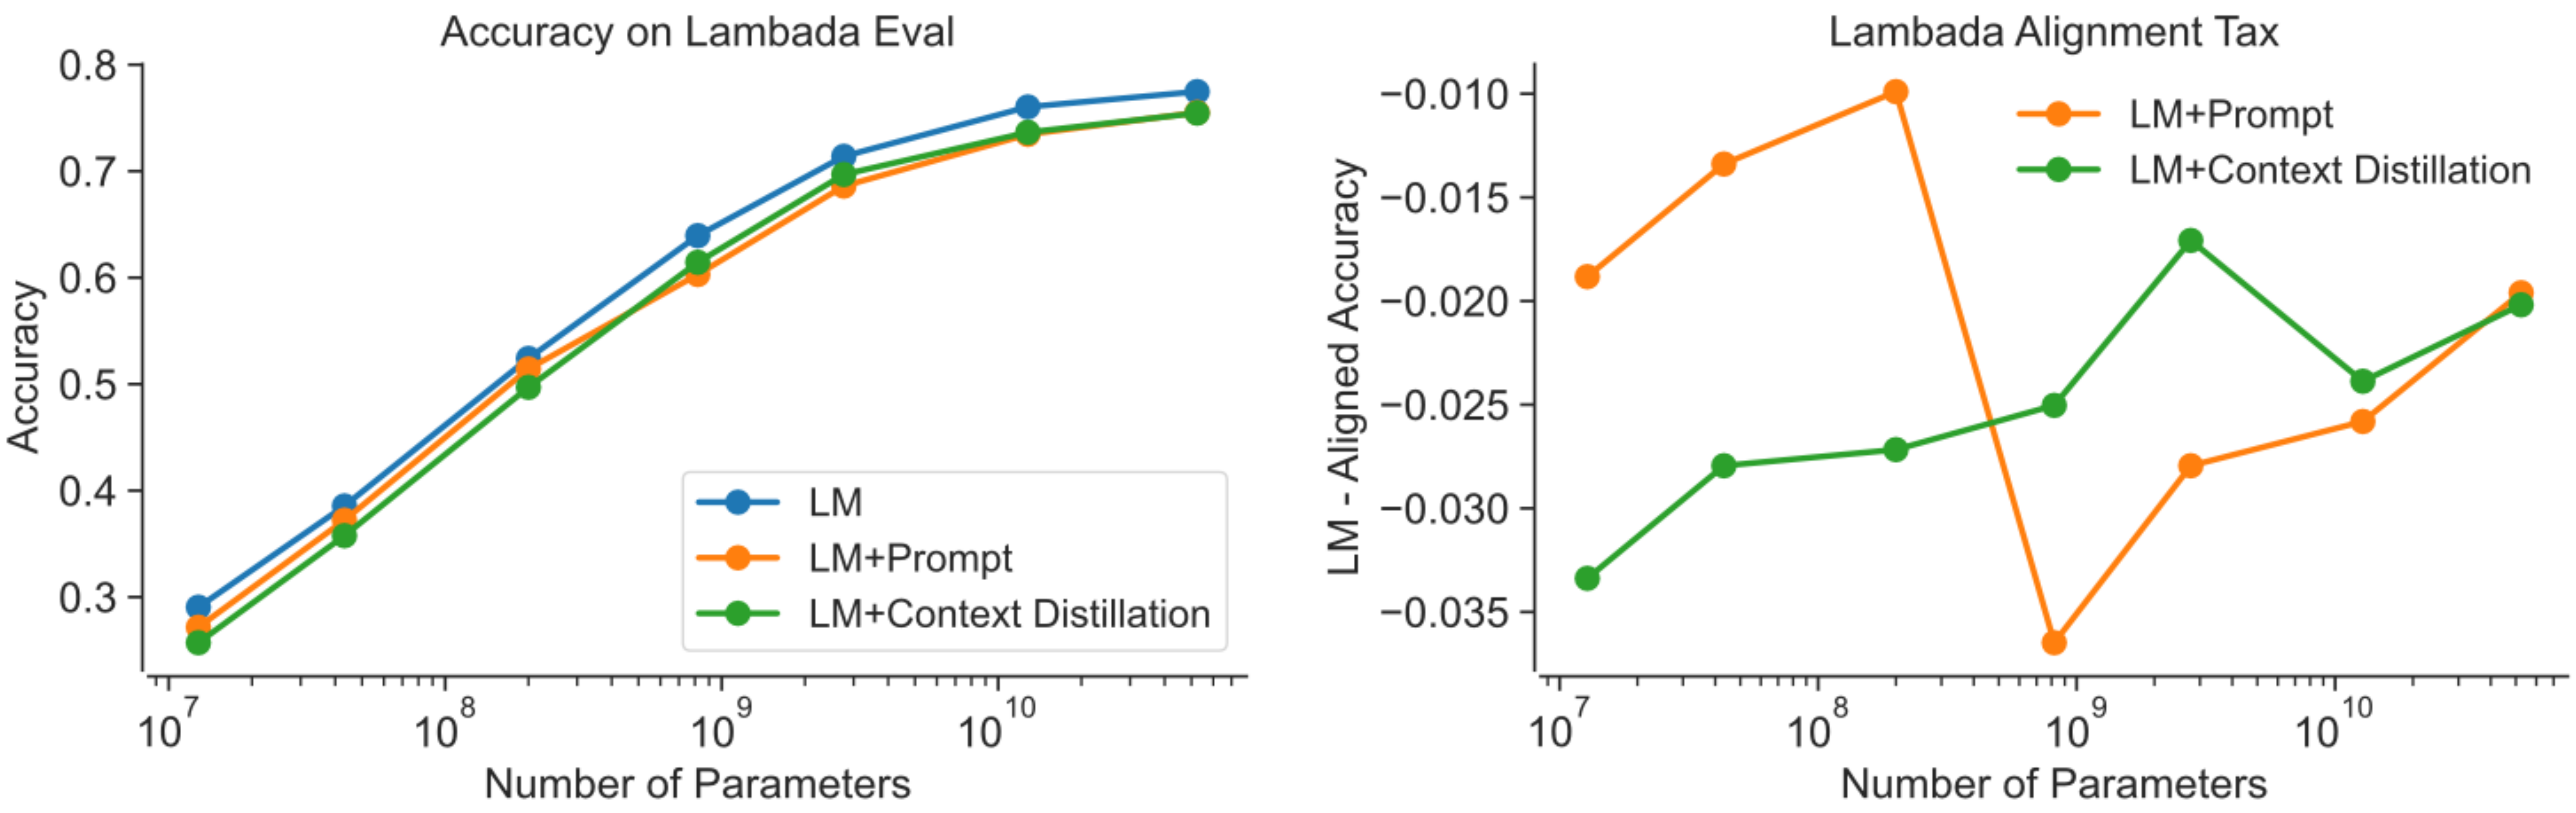

Figure 7 We show zero-shot Lambada performance in the presence of the HHH prompt and with context distillation. In both cases there is a small 'alignment tax'.

Next Idea: Alignment via Preference learning

Preference learning: types of data

Training data: Collected examples of human preferences

Binary data: We're given (q_i, A_i, B_i) where $A_i \geq B_i$ (i.e., A_i is "preferred" over B_i)

Ranked data: We're given $(q_i, A_1, A_2, \ldots, A_k)$ where $A_j \geq A_{j+1}$ for all $j \leq k - 1$

Note: Each datapoint in Ranked setting yields $\binom{k}{2}$ binary datapoints

Method 1(Simplest): SFT on preference pairs

(aka "Imitation Learning baseline")

Training objective: Given $(q_i, A_i \succeq B_i)$ the objective is c-e loss of A_i when given context q_i

At test time: Given q, A, B pick the response that has lower **per-token** c-e loss

(In other words, unaligned model has to learn **directly** from training on preference pairs

Methods 2: Reward model* from binary

Assumption (Bradley-Terry): For each query there exists a reward function r such that $r(A)$ = ”scalar reward for giving response A” and

$$\Pr[A \geq B] = \frac{1}{1 + \exp(r(B) - r(A))}$$

Training reward model: Put a **trainable “head”** on top of an LLM and train it to output reward given (query, response) as input.

Training objective for the head: Bradley-Terry loss using dataset $\{(q_i, A_i, B_i)\}$

(* Note: In the paper, reward model is called “preference model”)

Method 3: Reward model from ranked data

Ranked setting: Given $(q_i, A_1, A_2, \ldots, A_k)$ where $A_j \geq A_{j+1}$ for all $j \leq k-1$

Implementation: Same as Method 2, trained on **all** pairwise comparisons $\{(q, A_j, A_{j'})\}$ where $j' > j$

From reward model to “Best-of-k” Baseline

At test time, sample k responses given query and output the one with highest reward.

Binary Setting: SFT vs Bradley-Terry

(Coding tasks)

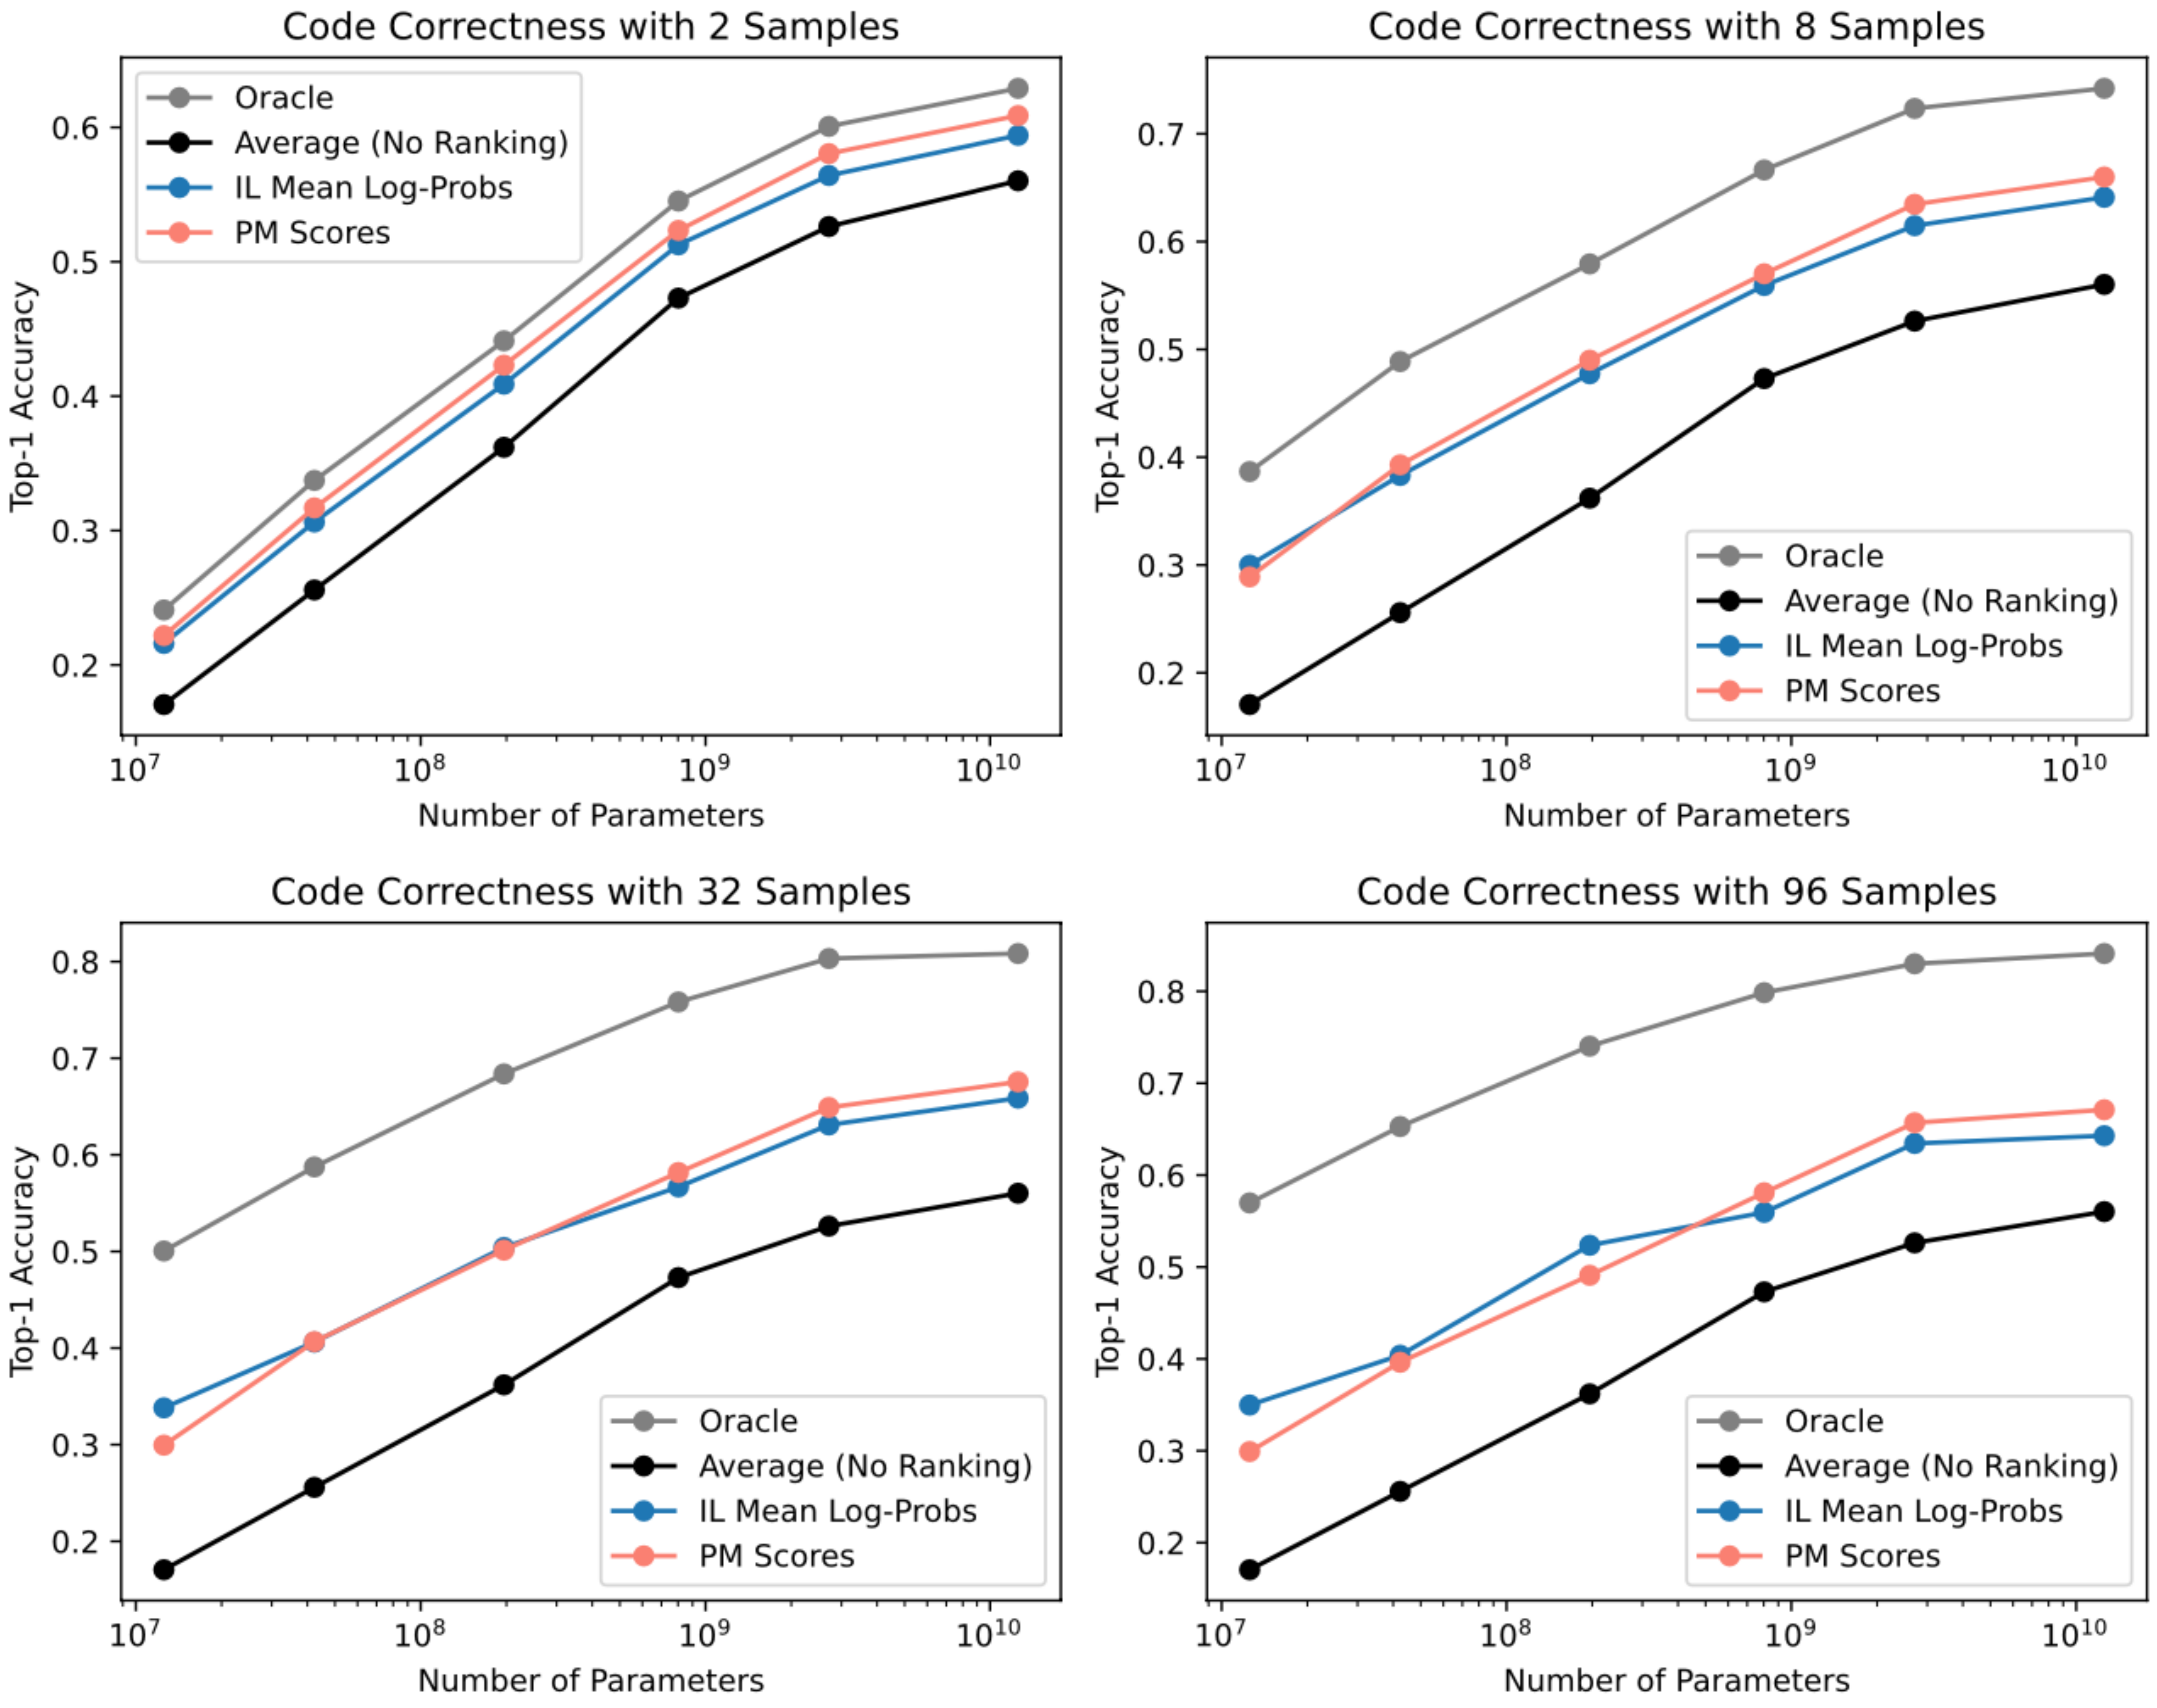

“Pass@k”score = accuracy using “best of k samples"

Conclusion: Imitation learning baseline is pretty close to learning preferences from binary data

(also verified on other binary evals, Lambada eval, and “Ethics”)

SFT baseline is weak for non-binary tasks

Hellaswag: (qs, answer1, answer2, answer3). Model has to choose the most correct one. Now preference modeling beats imitation learning.

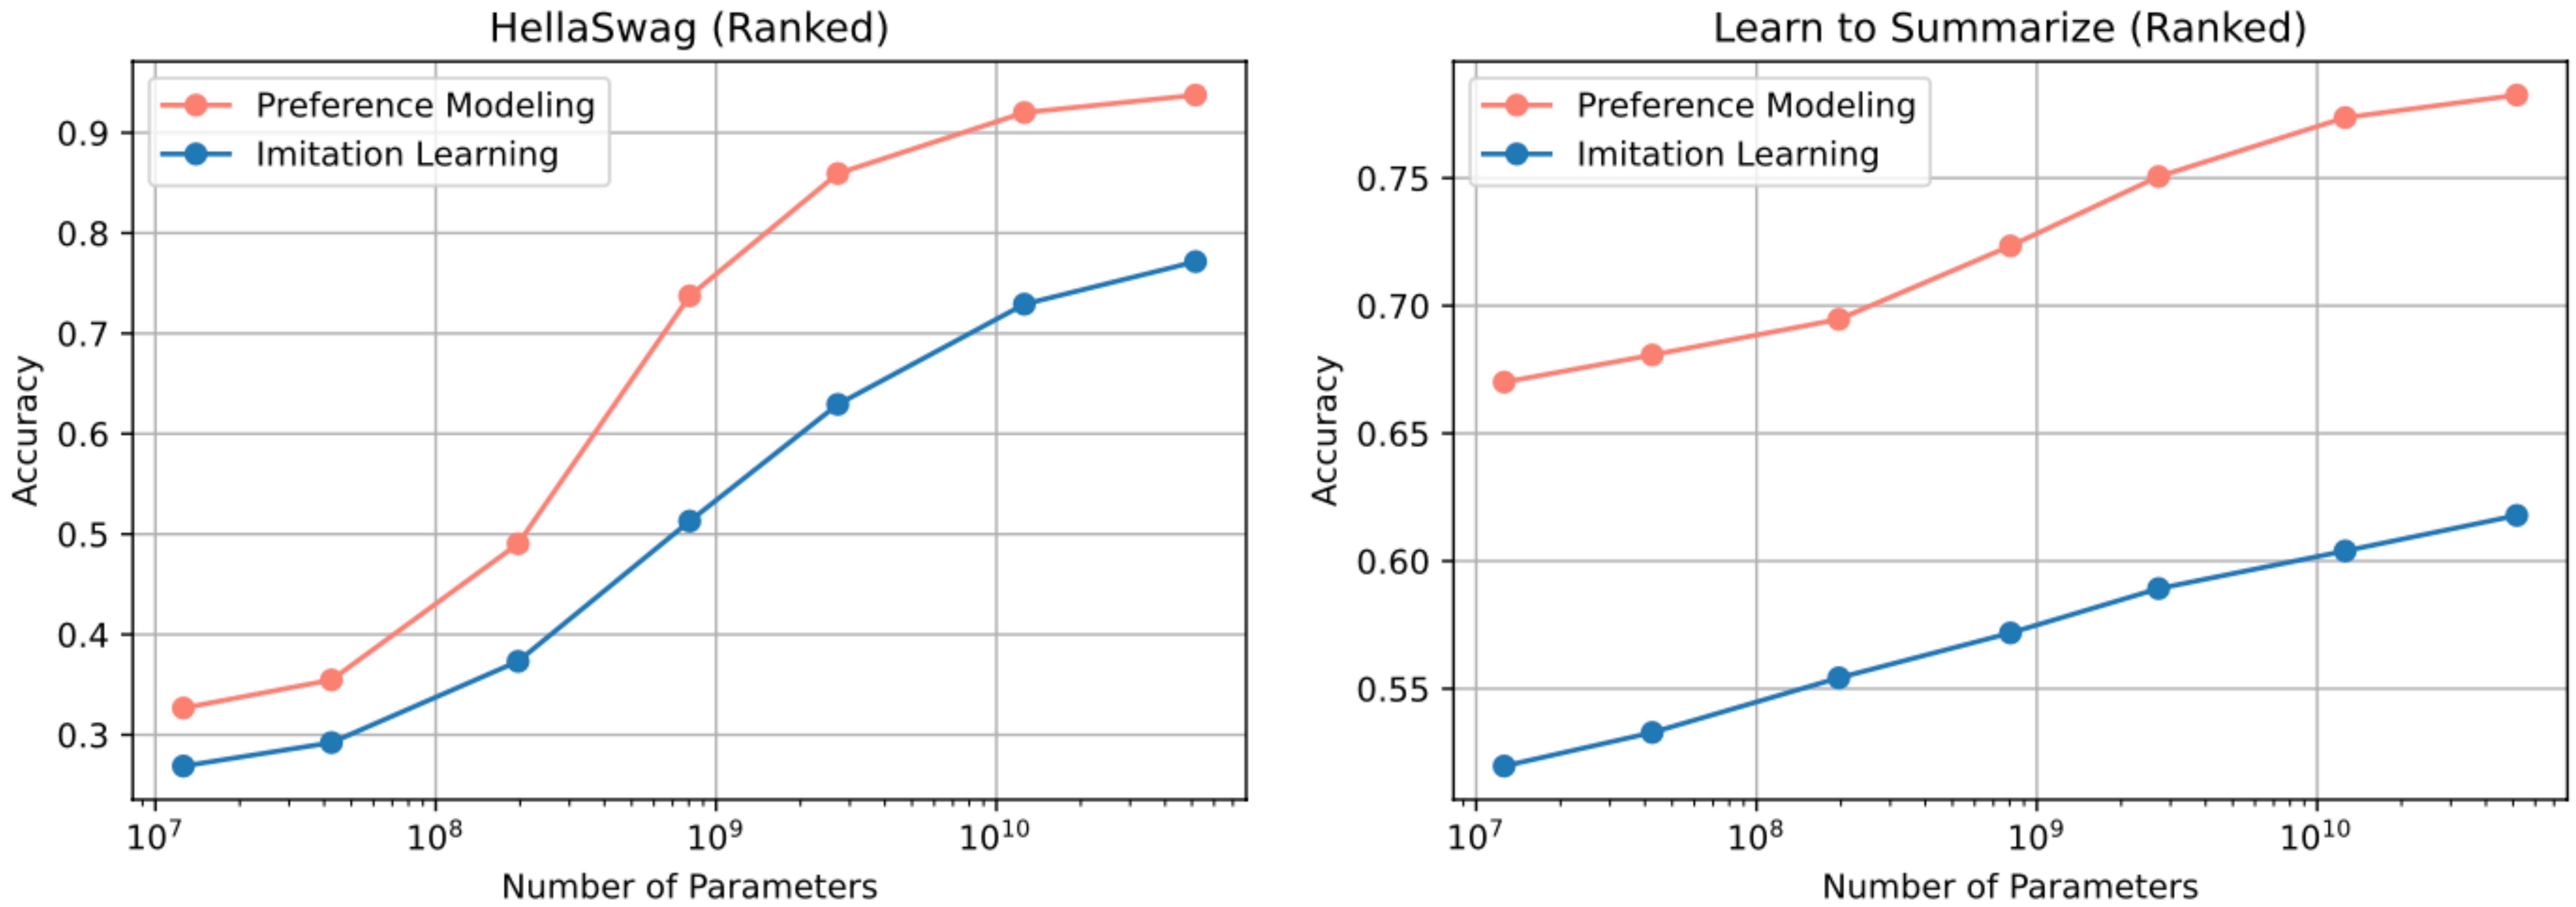

Figure 14 Scaling behavior of imitation learning and preference modeling on HellaSwag (ranked) and Learn to Summarize (ranked), showing that PM performs better than IL, as we expect for ranked finetuning evaluations.

Does alignment interfere with other capabilities?

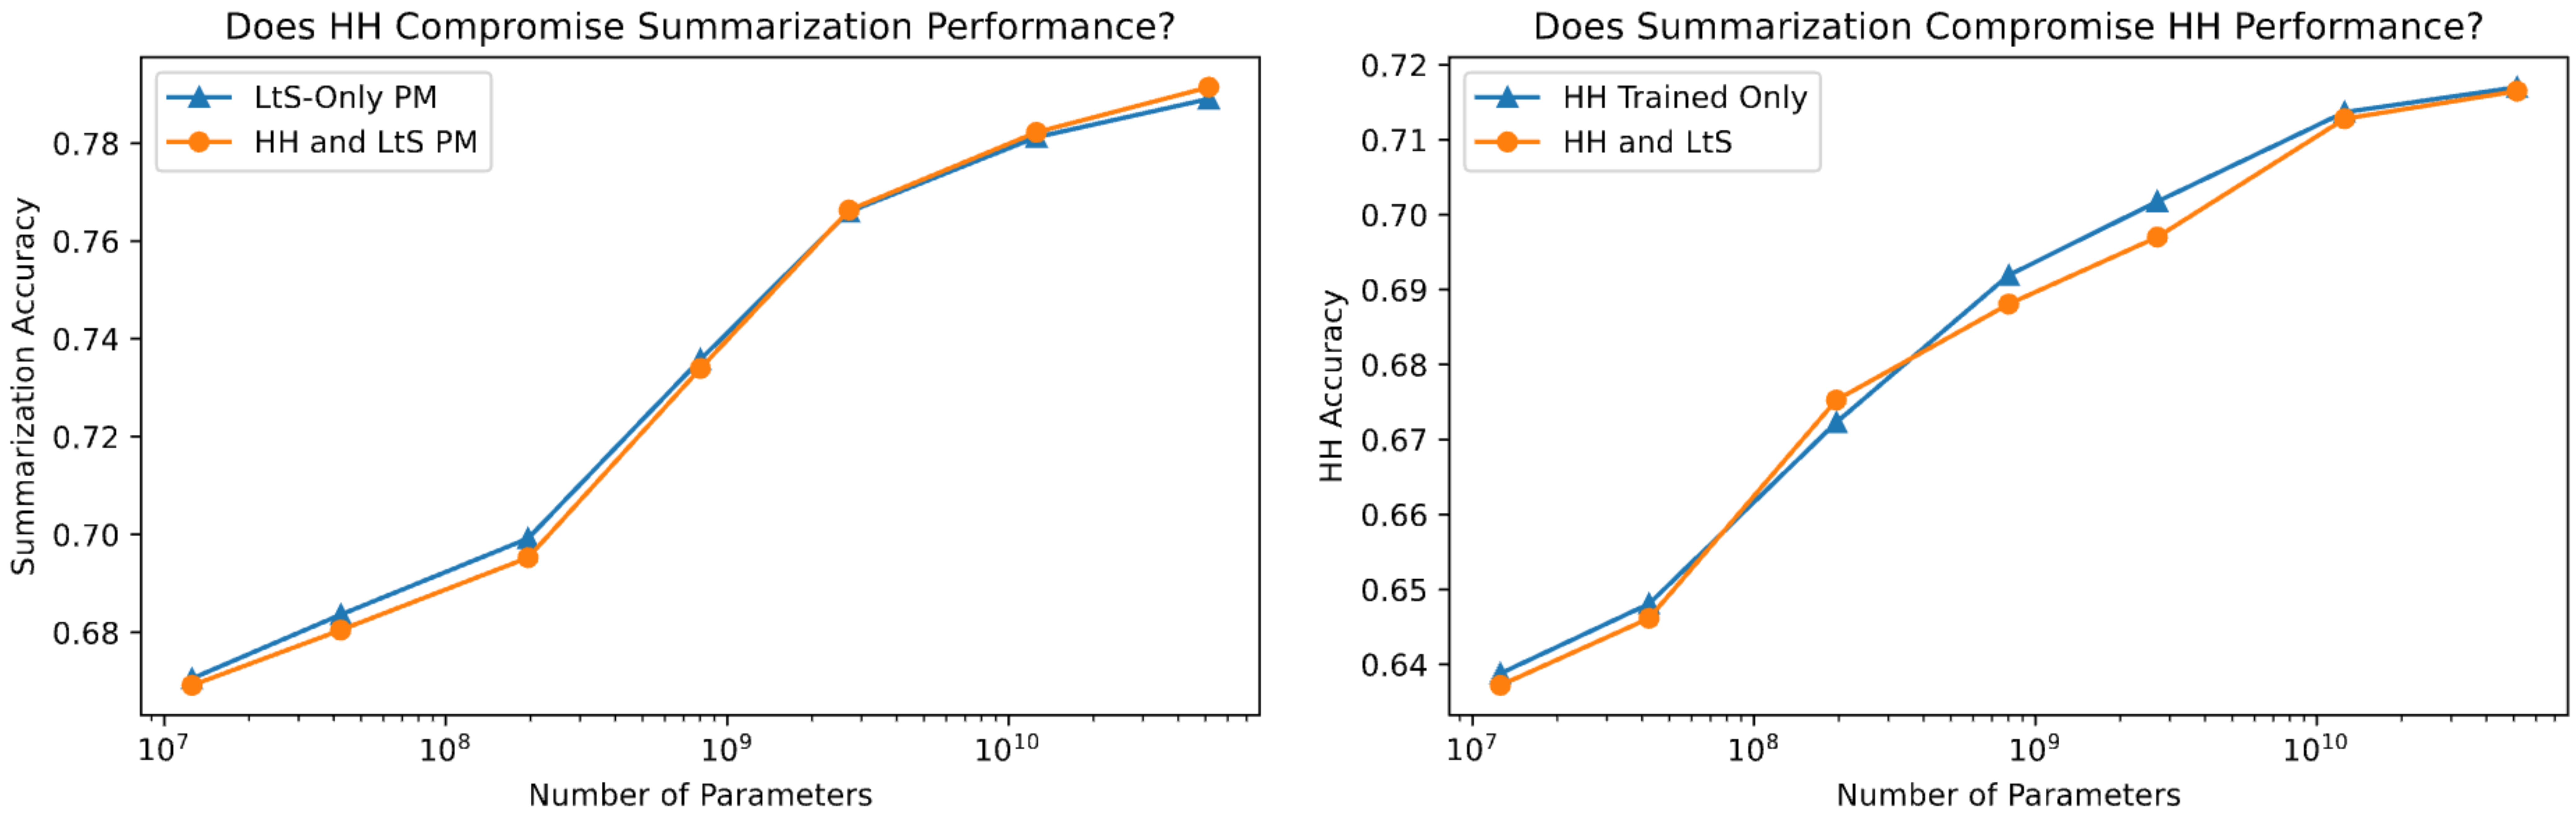

Figure 20 Here we show the comparison accuracies of preference models trained on (1) 'static' HH data only, (2) summarization data [Stiennon et al., 2020] only, and (3) a mixture of both. Mixed training has no negative effects on PM accuracies.

From preference pairs to aligned models

Simplest Alignment ("Best-of-k")

1. Start with a model π_{SFT} instruction-tuned using SFT (i.e. "helpful"). ← Lec 7, Lec 8
2. Collect problematic prompts/queries (e.g., "Tell me the racial slur for race [x]")
3. For each prompt x use human raters to provide good/bad responses using HHH criteria
4. Train a reward model using preference pairs from Step 3
5. For held out queries from step 2 (i.e., not used in Step 4) generate k responses from π_{SFT}. **Select best of these k.**
6. SFT on (query, response) pairs from Step 5 to turn π_{SFT} into an aligned model.

Alignment using RLHF

1. Start with a model π_{SFT} instruction-tuned using SFT (i.e. “helpful”).

Lec 7, Lec 8

2. Collect problematic prompts/queries (e.g., “Tell me the racial slur for race [x]”)
3. For each prompt x use human raters to provide good/bad responses using HHH criteria
4. Train a reward model using preference pairs from Step 3
5. For held out queries from step 2 (i.e., not used in Step 4) to ~~generate k responses from π_{SFT}.~~ do RLHF using reward model.

Lec 8

6. ~~SFT on (query, response) pairs from Step 5 to turn π_{SFT} into an aligned model.~~

Question: In this pipeline how do humans “tell” the AI how to behave?

(RLHF = Reinforcement Learning from Human Feedback; Lecture 8 “PPO Objective”)