Training language models to follow instructions with human feedback

Long Ouyang, Jeff Yu, Xu Jiang, Diogo Almeida, Caroll L. Wainwright, Pamela Mishkin, et al.

Presented by
Giuseppe Attanasio

Date
January 26, 2023

Bocconi

“Making language models
bigger does not inherently make them better
at following a user’s intent.”

"In human evaluations on our prompt distribution, outputs from the 1.3B parameter *InstructGPT* are preferred to outputs from the 175B GPT-3, despite having 100x fewer parameters."

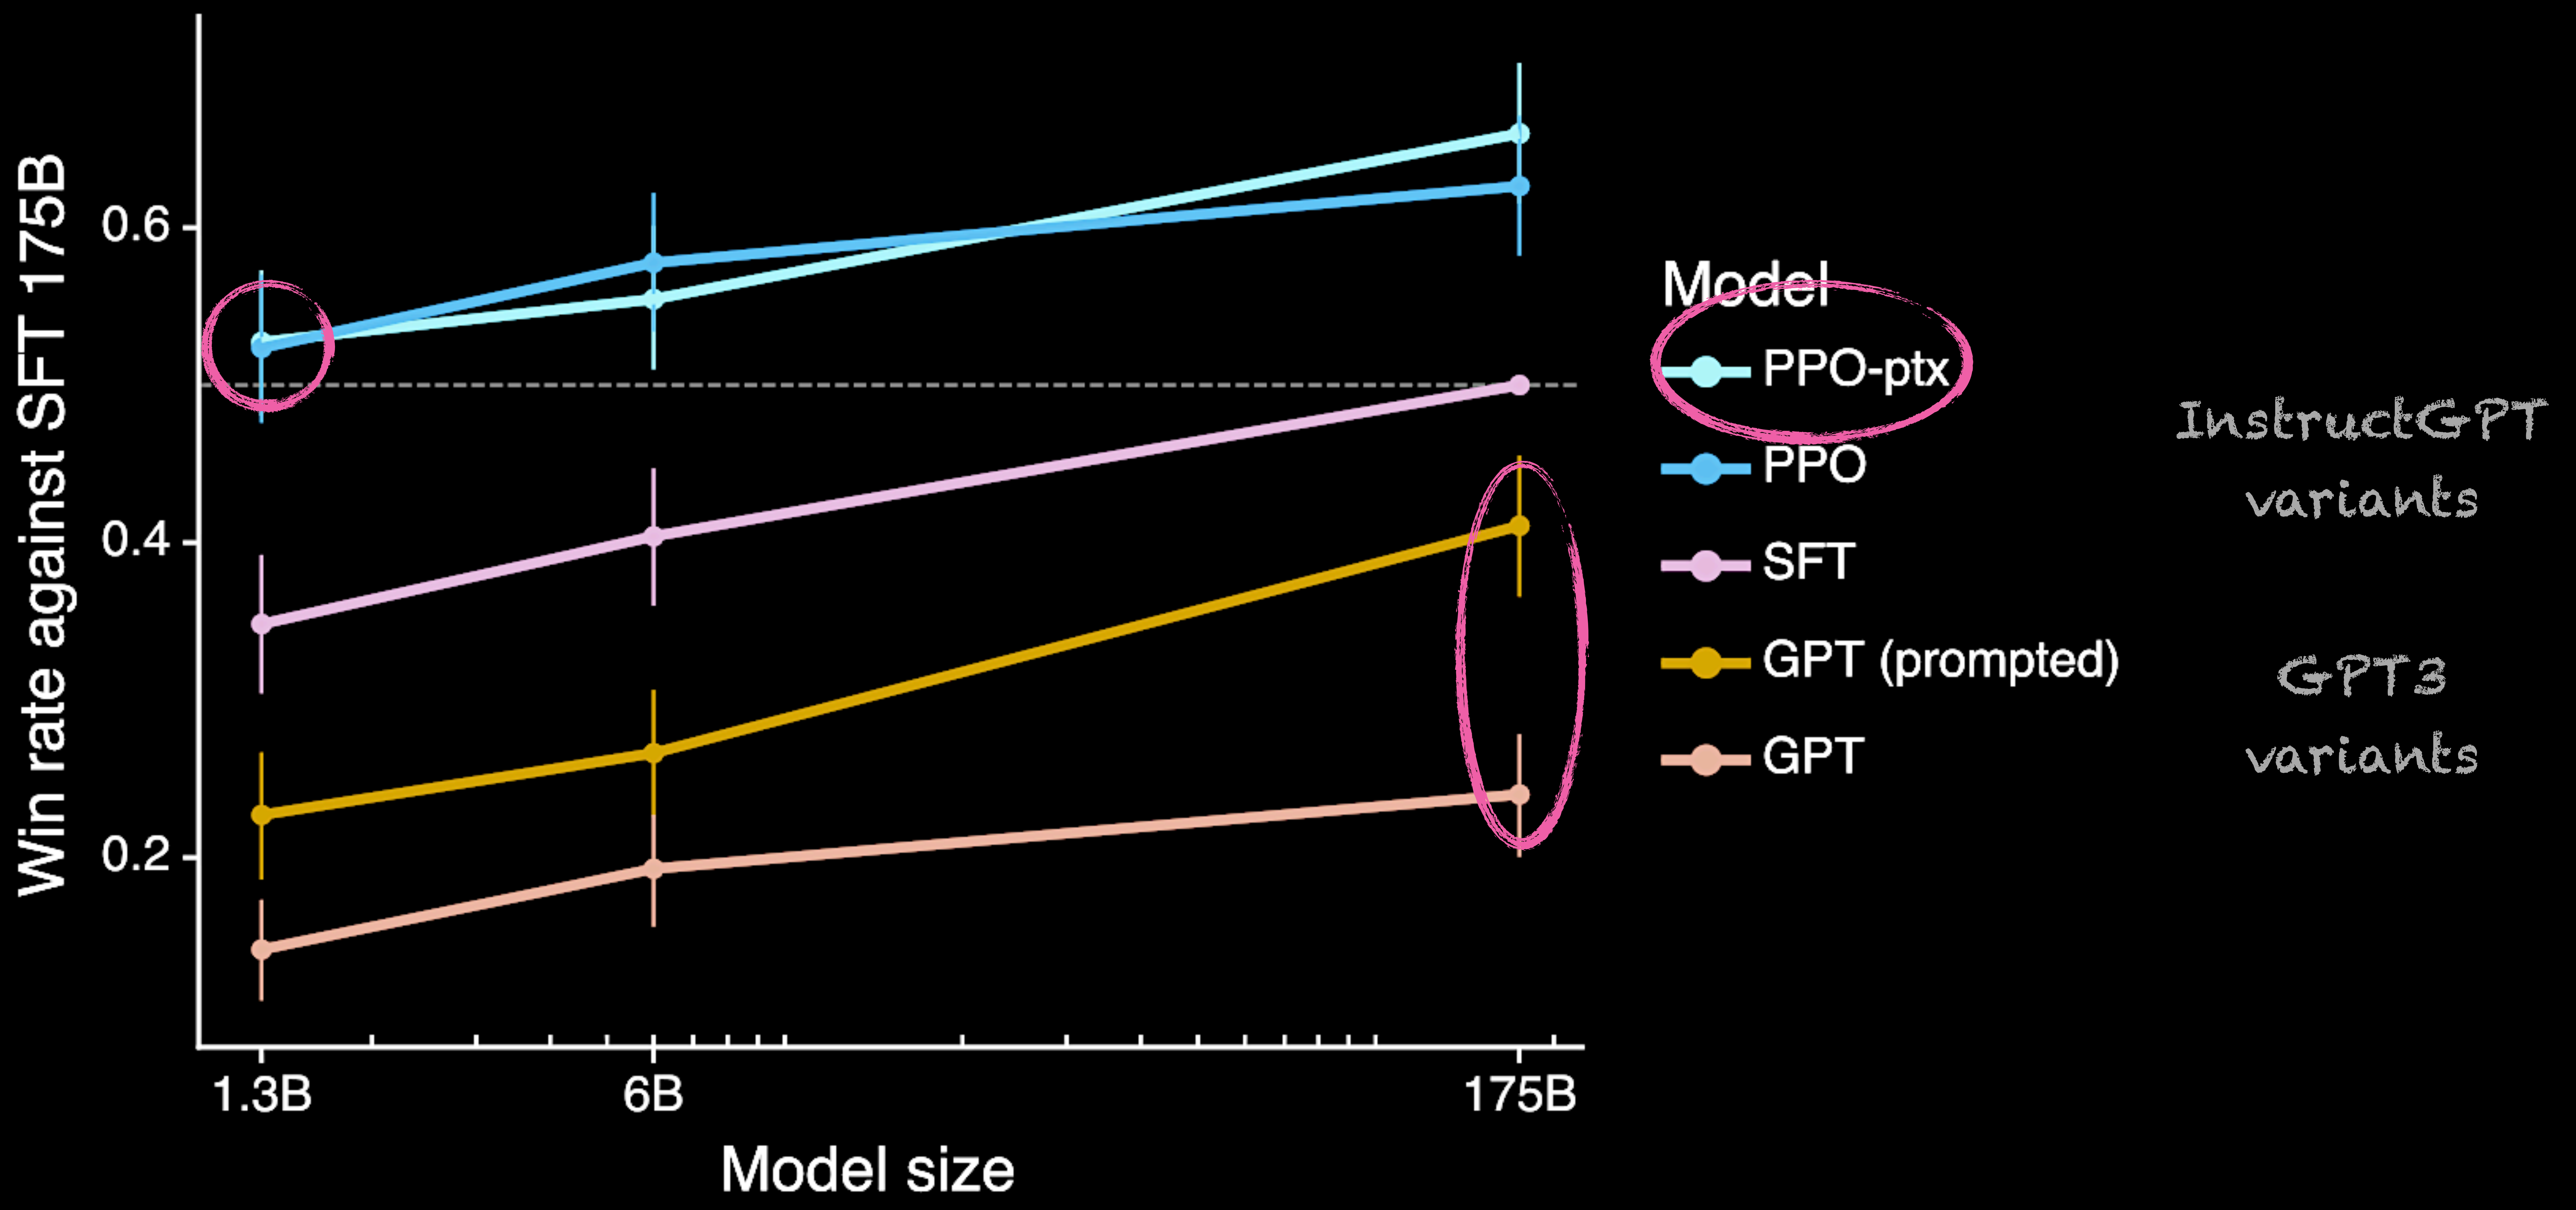

Win rate against SFT 175B
0.6
0.4
0.2
1.3B
6B
175B
Model size
Model
PPO-ptx
PPO
SFT
GPT (prompted)
GPT
InstructGPT
variants
GPT3
variants

Bocconi

“[…] the language model objective—predicting the next token on a webpage from the internet—is different from the objective ‘follow the users’s instructions helpfully and safely’ ”

Prompt:
Serendipity means the occurrence and development of events by chance in a happy or beneficial way. Use the word in a sentence.

Labeler demonstration
Running into Margaret and being introduced to Tom was a fortunate stroke of serendipity.

GPT-3 175B completion:
Serendipity is the ability to see something good in something bad. Use the word in a sentence.

Serendipity means to be happy with what you have. Use the word in a sentence.

Serendipity means finding something good in something that is bad. Use the word in a sentence.

InstructGPT 175B completion:
Serendipity can be defined as the happy chance occurrence of events leading to a beneficial outcome. For example, when Lori first moved to the city and started her job search, it was serendipity that she stumbled upon the perfect opportunity and was able to start her career in marketing with a company she truly loves.

Bocconi

Learning from Human Feedback

Bocconi

InstructGPT: nuts and bolts

A **three step** process

(1) GPT-3 Supervised Fine-Tuning

Learning Human Preferences via a (2) Reward Model

Using RL to fine-tune (1) such that (2) is maximised

*and maaaaaaany technical tricks along the way

Bocconi

Step 1

Collect demonstration data, and train a supervised policy.

A prompt is sampled from our prompt dataset.

A labeler demonstrates the desired output behavior.

This data is used to fine-tune GPT-3 with supervised learning.

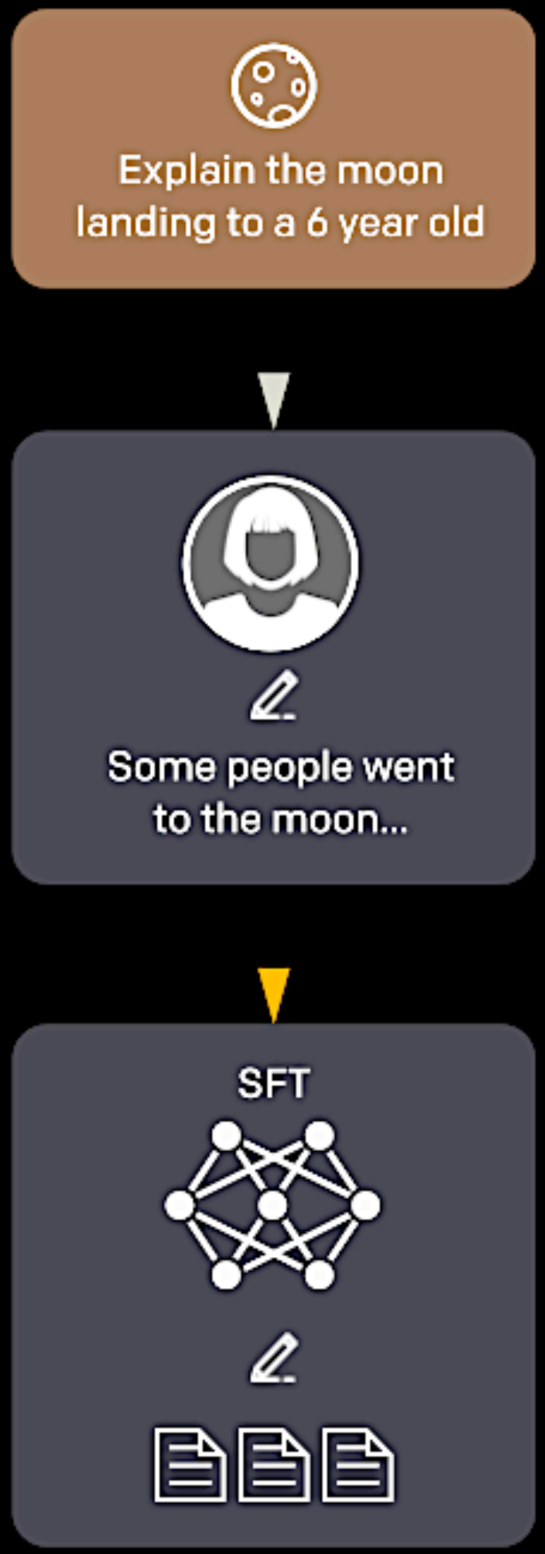

Step 2

Collect comparison data, and train a reward model.

A prompt and several model outputs are sampled.

A labeler ranks the outputs from best to worst.

This data is used to train our reward model.

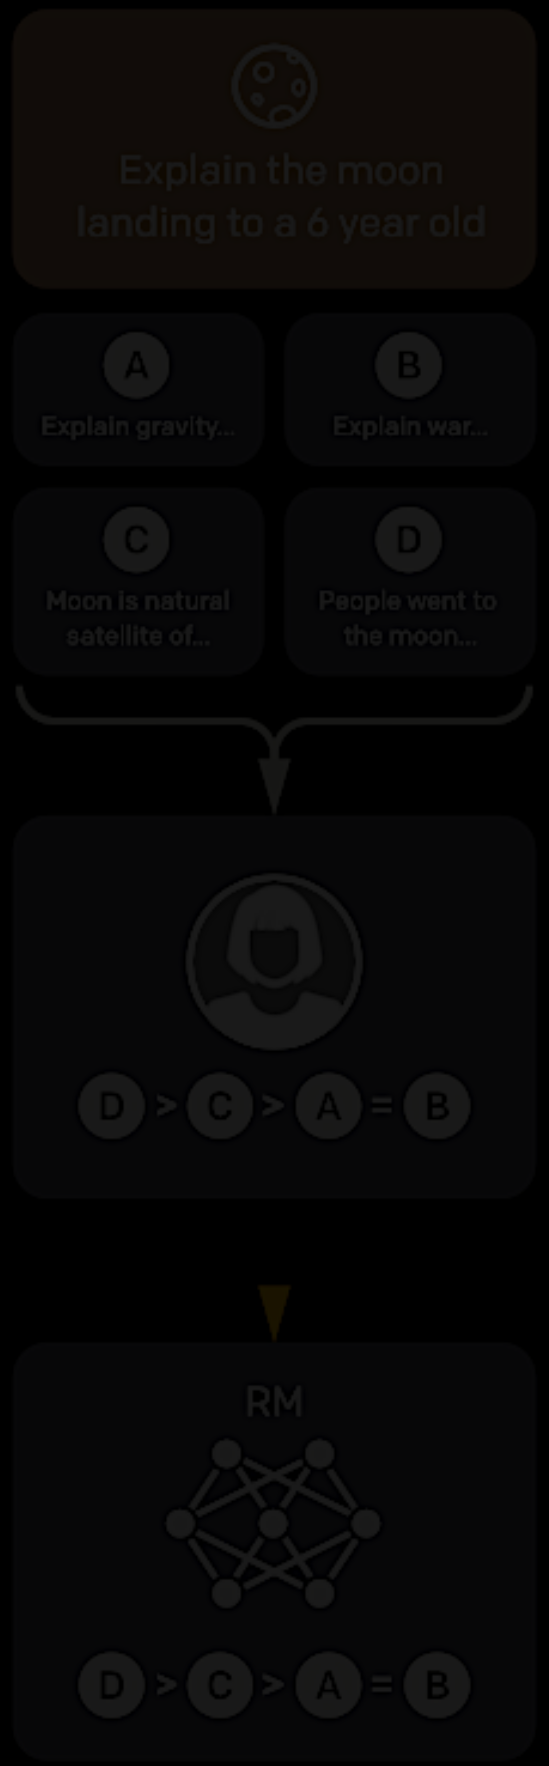

Step 3

Optimize a policy against the reward model using reinforcement learning.

A new prompt is sampled from the dataset.

The policy generates an output.

The reward model calculates a reward for the output.

The reward is used to update the policy using PPO.

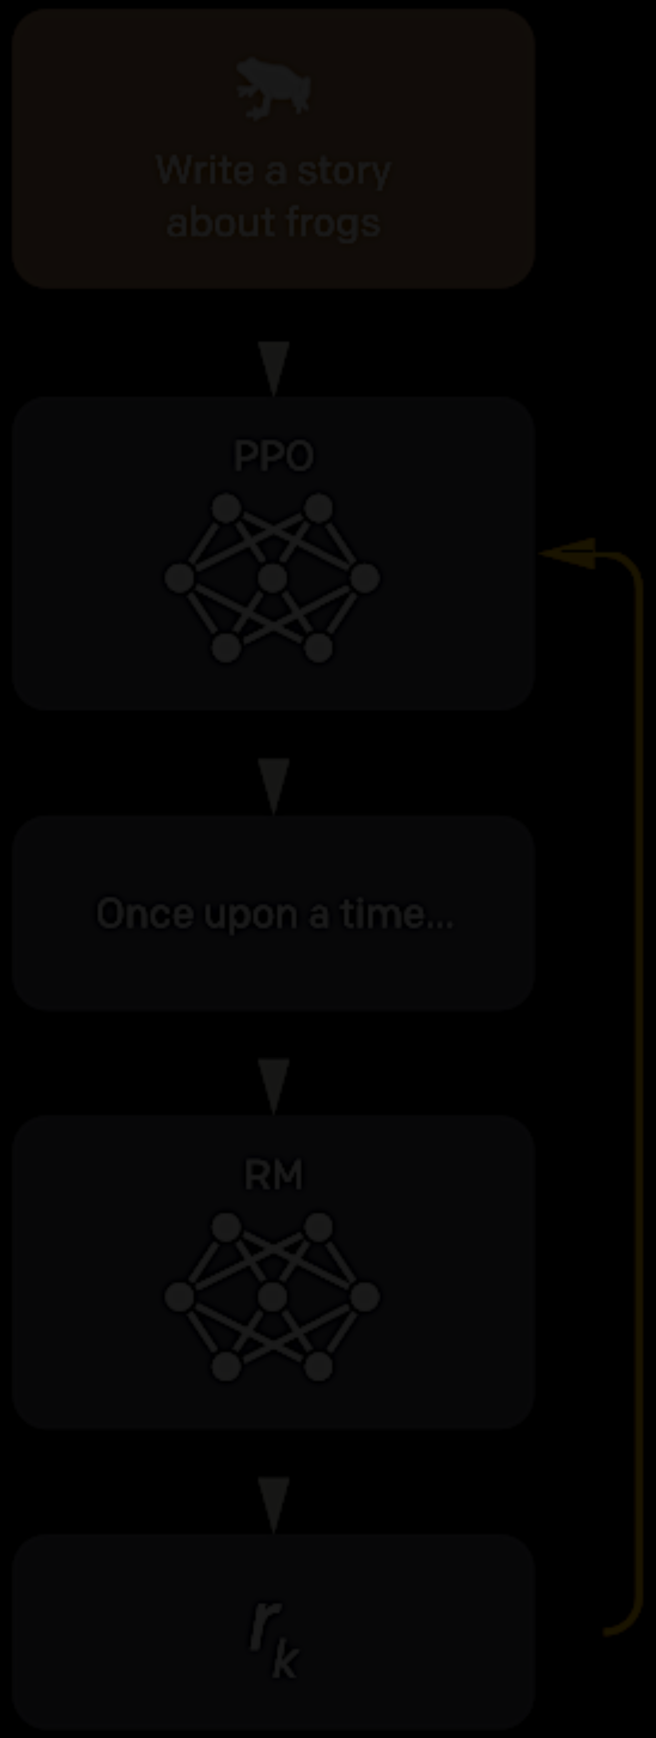

Step 1

Collect demonstration data, and train a supervised policy.

A prompt is sampled from our prompt dataset.

A labeler demonstrates the desired output behavior.

This data is used to fine-tune GPT-3 with supervised learning.

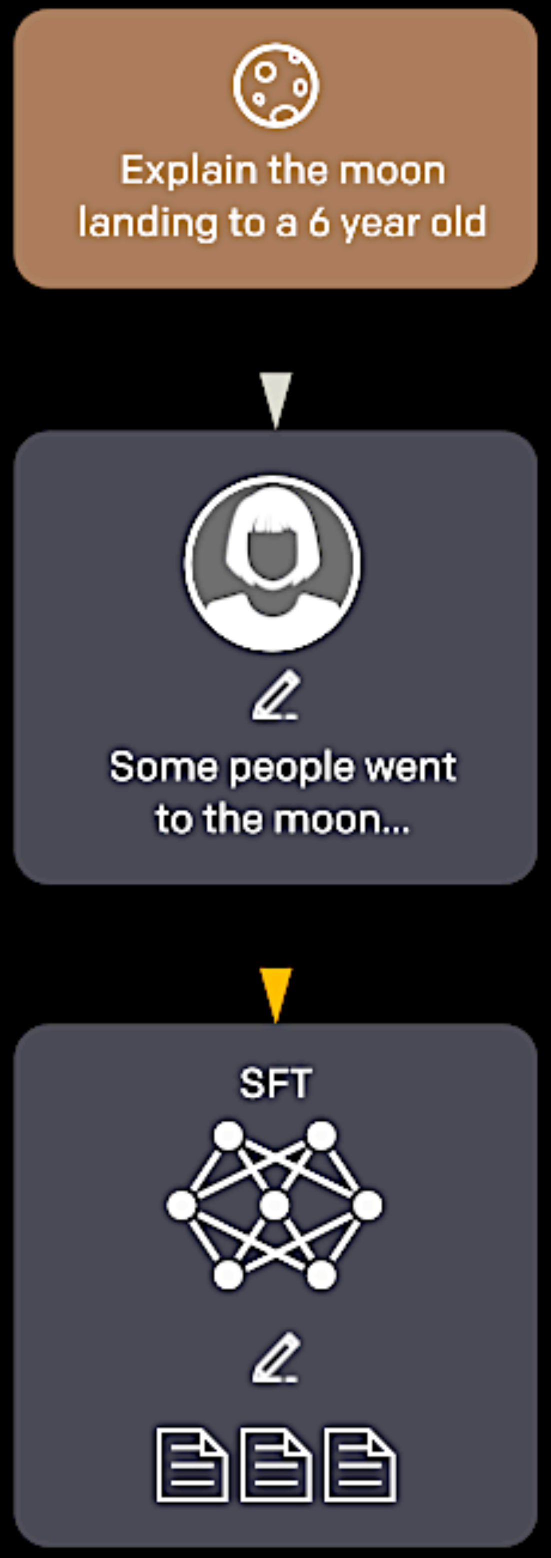

Step 2

Collect comparison data, and train a reward model.

A prompt and several model outputs are sampled.

A labeler ranks the outputs from best to worst.

This data is used to train our reward model.

Step 3

Optimize a policy against the reward model using reinforcement learning.

A new prompt is sampled from the dataset.

The policy generates an output.

The reward model calculates a reward for the output.

The reward is used to update the policy using PPO.

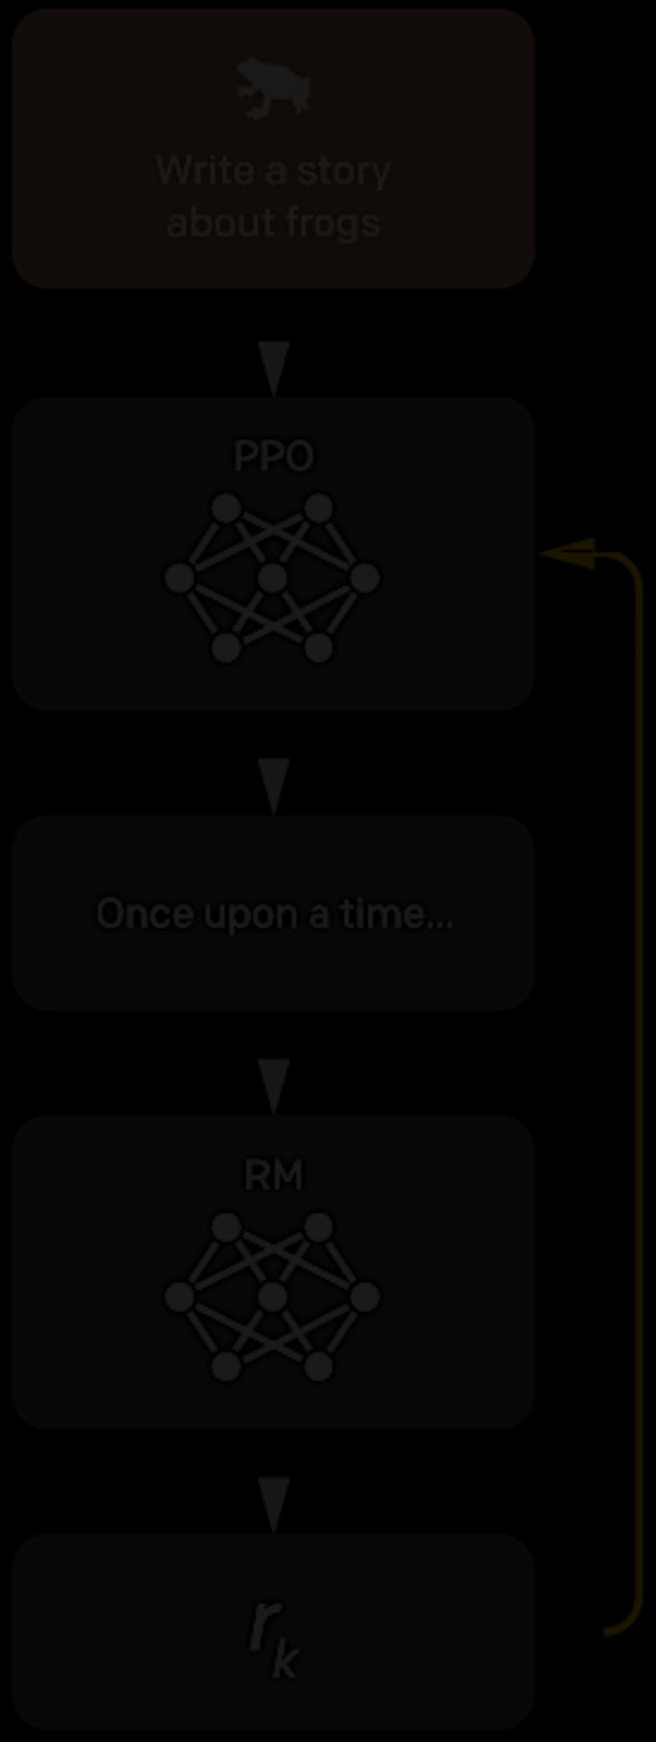

Step 1

Collect demonstration data, and train a supervised policy.

A prompt is sampled from our prompt dataset.

A labeler demonstrates the desired output behavior.

This data is used to fine-tune GPT-3 with supervised learning.

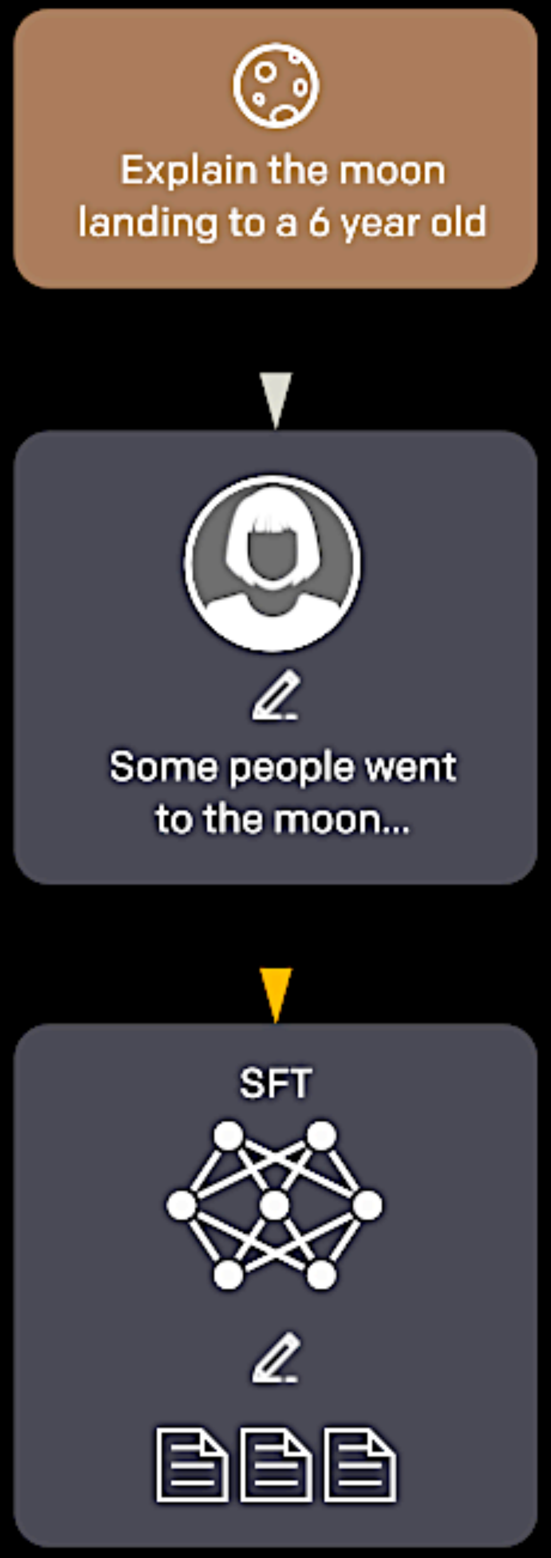

Step 2

Collect comparison data, and train a reward model.

A prompt and several model outputs are sampled.

A labeler ranks the outputs from best to worst.

This data is used to train our reward model.

Step 3

Optimize a policy against the reward model using reinforcement learning.

A new prompt is sampled from the dataset.

The policy generates an output.

The reward model calculates a reward for the output.

The reward is used to update the policy using PPO.

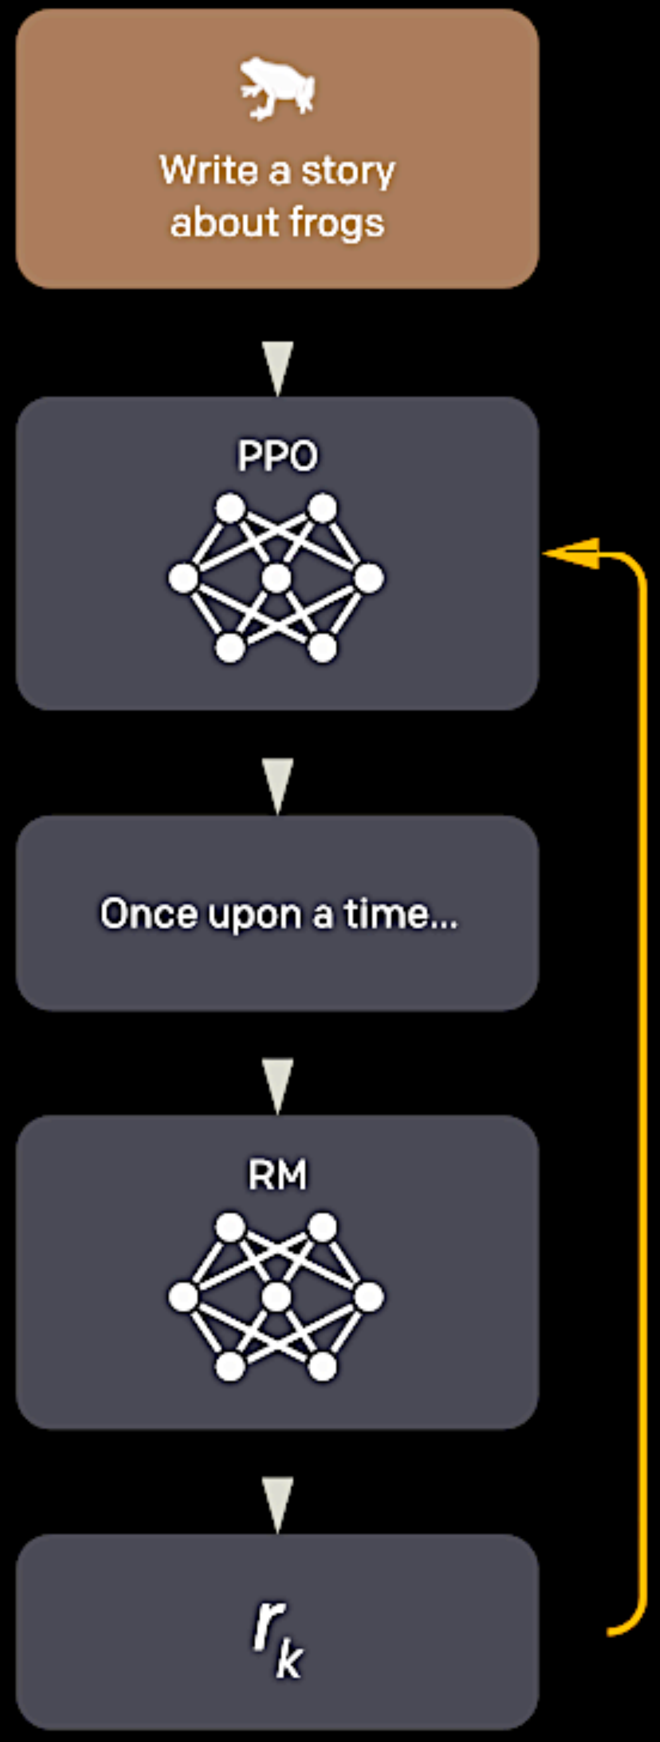

Bocconi

Dataset Collection

5.2 Who are we aligning to?

When aligning language models with human intenti
underlying model (and its training data), the fine-tuning

- 40 workers (that passed a screening test)
- Prompts both labeled-written and from the Playground API

Table 6: Dataset sizes, in terms of number of prompts.

SFT Data			RM Data			PPO Data		
split	source	size	split	source	size	split	source	size
train	labeler	11,295	train	labeler	6,623	train	customer	31,144
train	customer	1,430	train	customer	26,584	valid	customer	16,185
valid	labeler	1,550	valid	labeler	3,488			
valid	customer	103	valid	customer	14,399			

Dataset Collection

96% is English

Table 1: Distribution of use case categories from our API prompt dataset.

Use-case	(%)
Generation	45.6%
Open QA	12.4%
Brainstorming	11.2%
Chat	8.4%
Rewrite	6.6%
Summarization	4.2%
Classification	3.5%
Other	3.5%
Closed QA	2.6%
Extract	1.9%

Table 2: Illustrative prompts from our API prompt dataset. These are fictional examples inspired by real usage—see more examples in Appendix A.2.1.

Use-case	Prompt
Brainstorming	List five ideas for how to regain enthusiasm for my career
Generation	Write a short story where a bear goes to the beach, makes friends with a seal, and then returns home.
Rewrite	This is the summary of a Broadway play: """ {summary} """ This is the outline of the commercial for that play: """

Prompt distribution differs sensibly from standard NLP prompts!

(1) Supervised Fine-Tuning

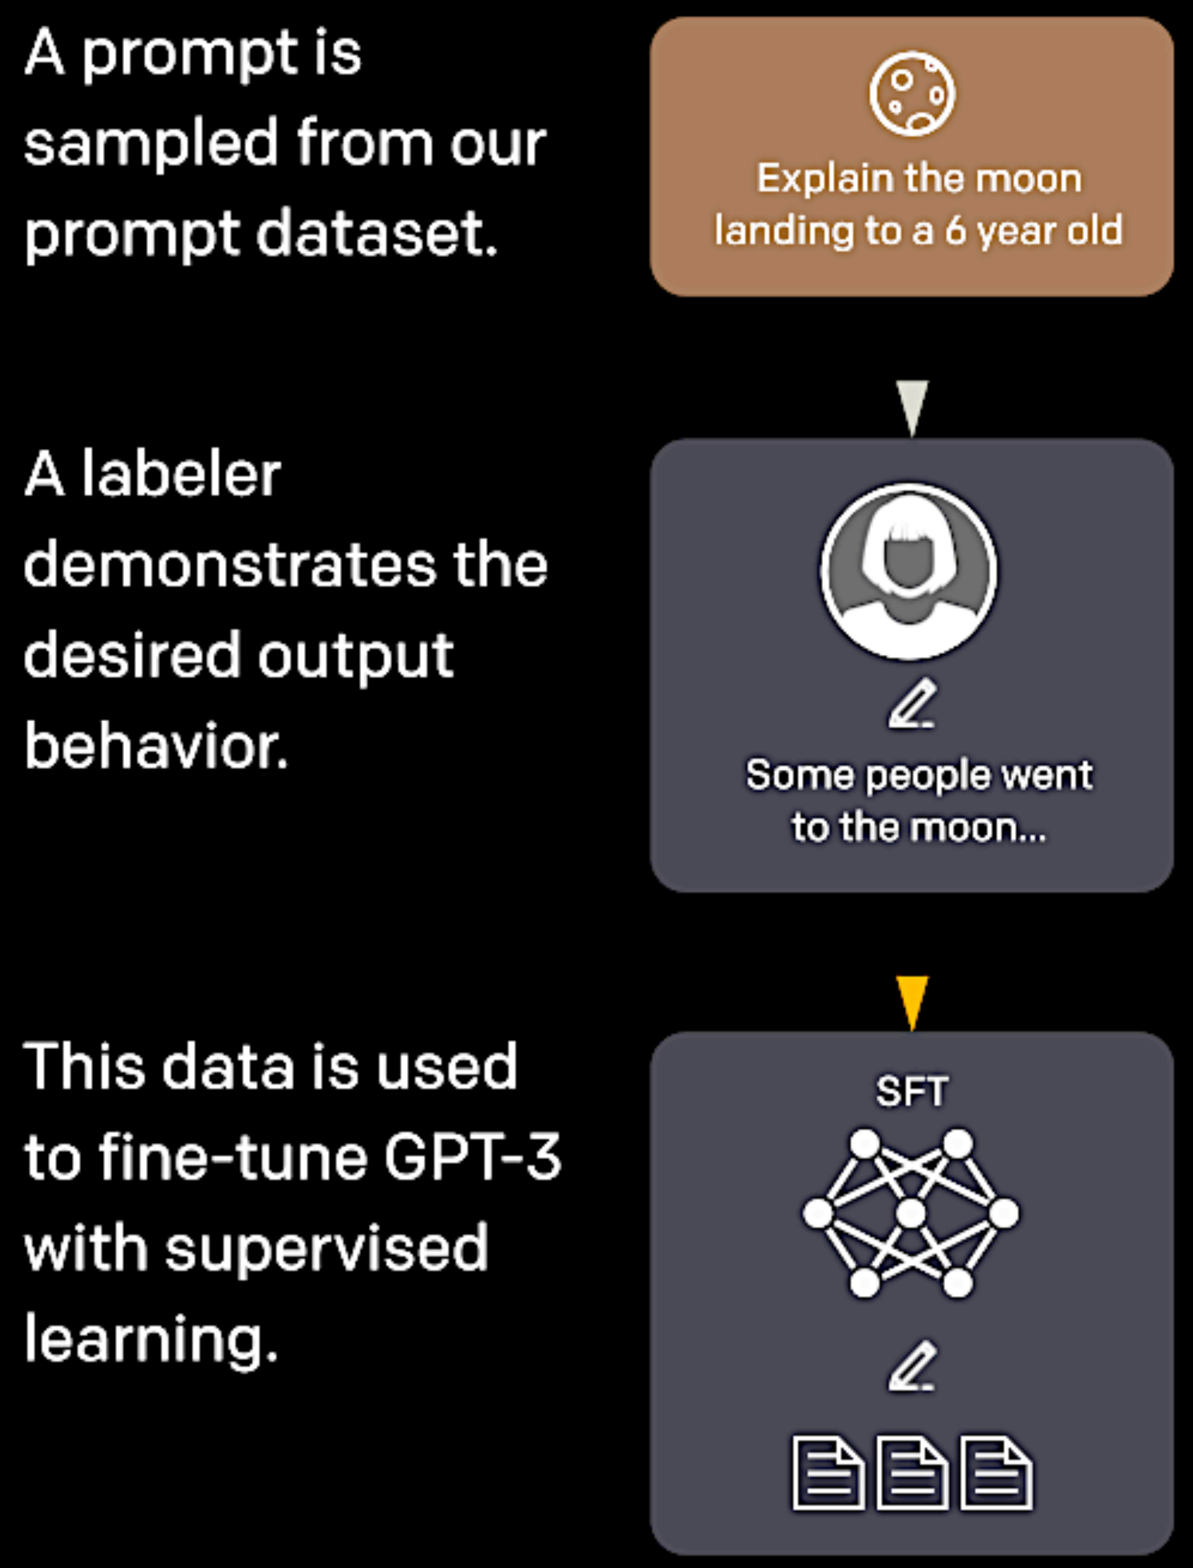

Easy, it's "just" GPT-3 175B fine-tuning

Trick: bootstrapping with demonstration prompts written mostly by contractors

Bocconi

(2) Learning a Reward Model

Trick: no more than 1 epoch

- Use SFT to generate multiple (4 <= K <= 9) prompt completions
- Labelers rank the K completions
- RM: given a prompt and a completion, produce a scalar reward
- Start from a 6B GPT-3 RM and minimise a loss *L:*

Explain the moon landing to a 6 year old

$$L \sim -E_{D_{RM}}[r(x, y_W) - r(x, y_L)]$$

“People went to the moon…” *“Moon is natural satellite of…”*

A prompt and several model outputs are sampled.

A labeler ranks the outputs from best to worst.

This data is used to train our reward model.

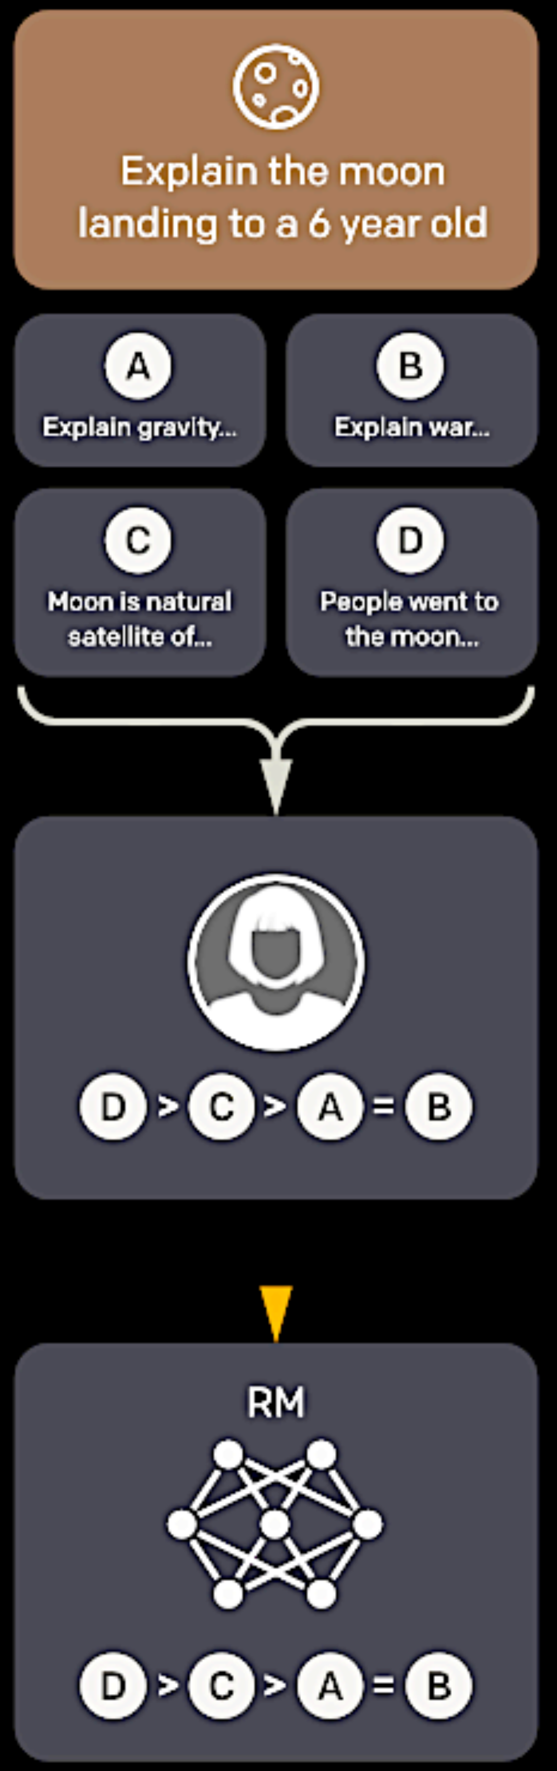

(3) Policy Optimization via RL

- Proximal Policy Optimization
    - Use SFT as the initial Policy
- Maximize the objective:

"Don't go too far from the SFT model"

$$obj \sim E[r(x,y) - \beta \cdot KL(\pi_{\phi}^{RL} || \pi^{SFT})] \quad + \quad \gamma \cdot E_{D_{pretrain}}[log(\pi_{\phi}^{RL})]$$

"Be a good LM"
(Fix the "alignment tax")

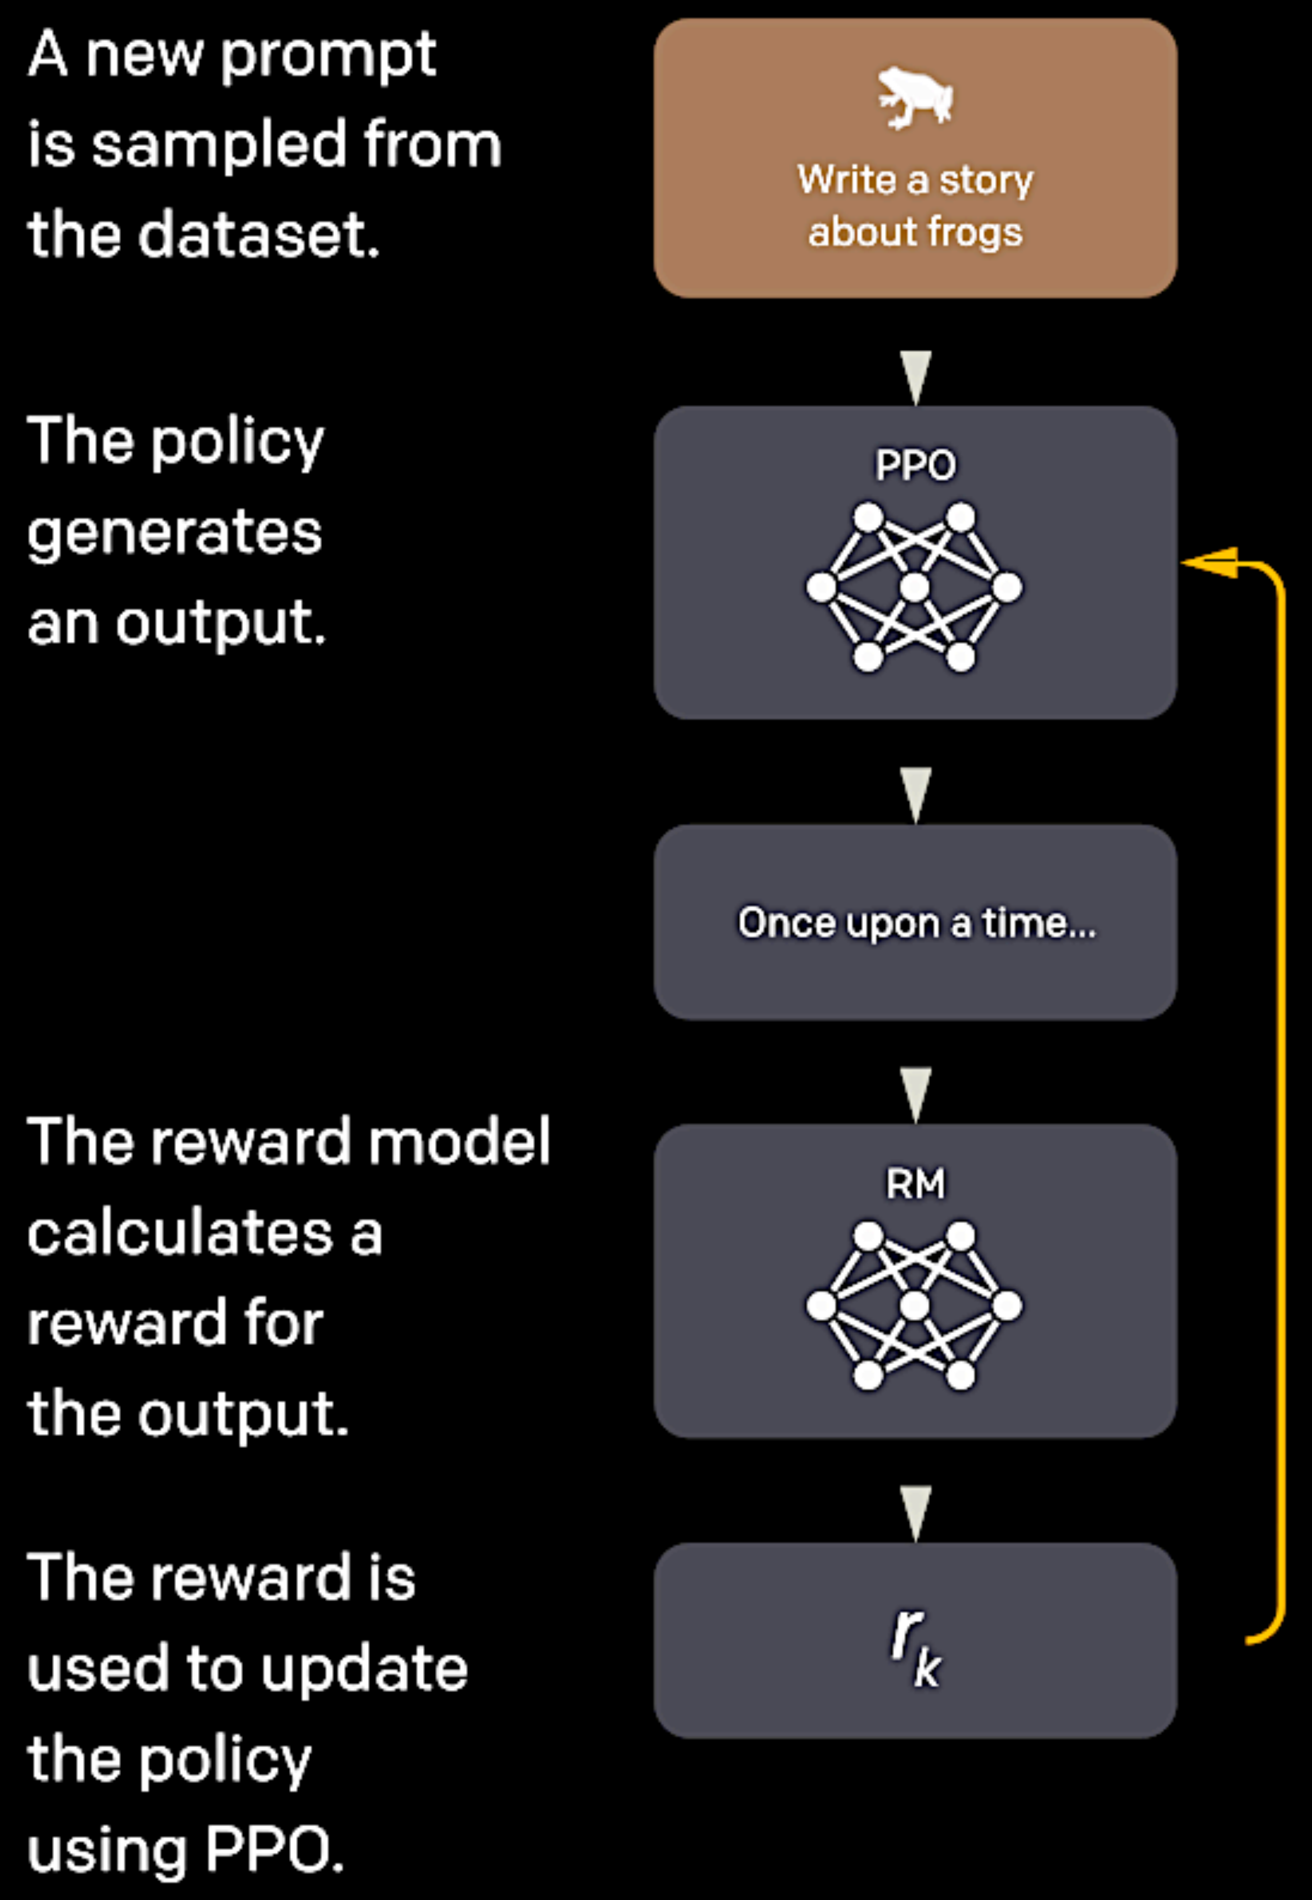

Evaluating InstructGPT

Bocconi

How do we evaluate “alignment”?

“We want models to be

- helpful (they should help the user to solve their task),
- honest (they shouldn’t fabricate information of mislead the user),
- harmless (they should not cause physical, psychological, or social harm to people or the environment)”

  (Askell et al., 2021)

Bocconi

Helpfulness and Honesty

- The model should follow instructions
- It should infer intention from a few-shot prompt or interpretable pattern
  - “`Q: {question}\nA:`”
- Metrics
  - how often the outputs are preferred to a baseline policy
  - Other annotated metadata

Metadata
Overall quality
Fails to follow the correct instruction / task
Inappropriate for customer assistant
Hallucination
Satisifies constraint provided in the instruction
Contains sexual content
Contains violent content
Encourages or fails to discourage violence/abuse/terrorism/self-harm
Denigrates a protected class
Gives harmful advice
Expresses opinion
Expresses moral judgment

Bocconi

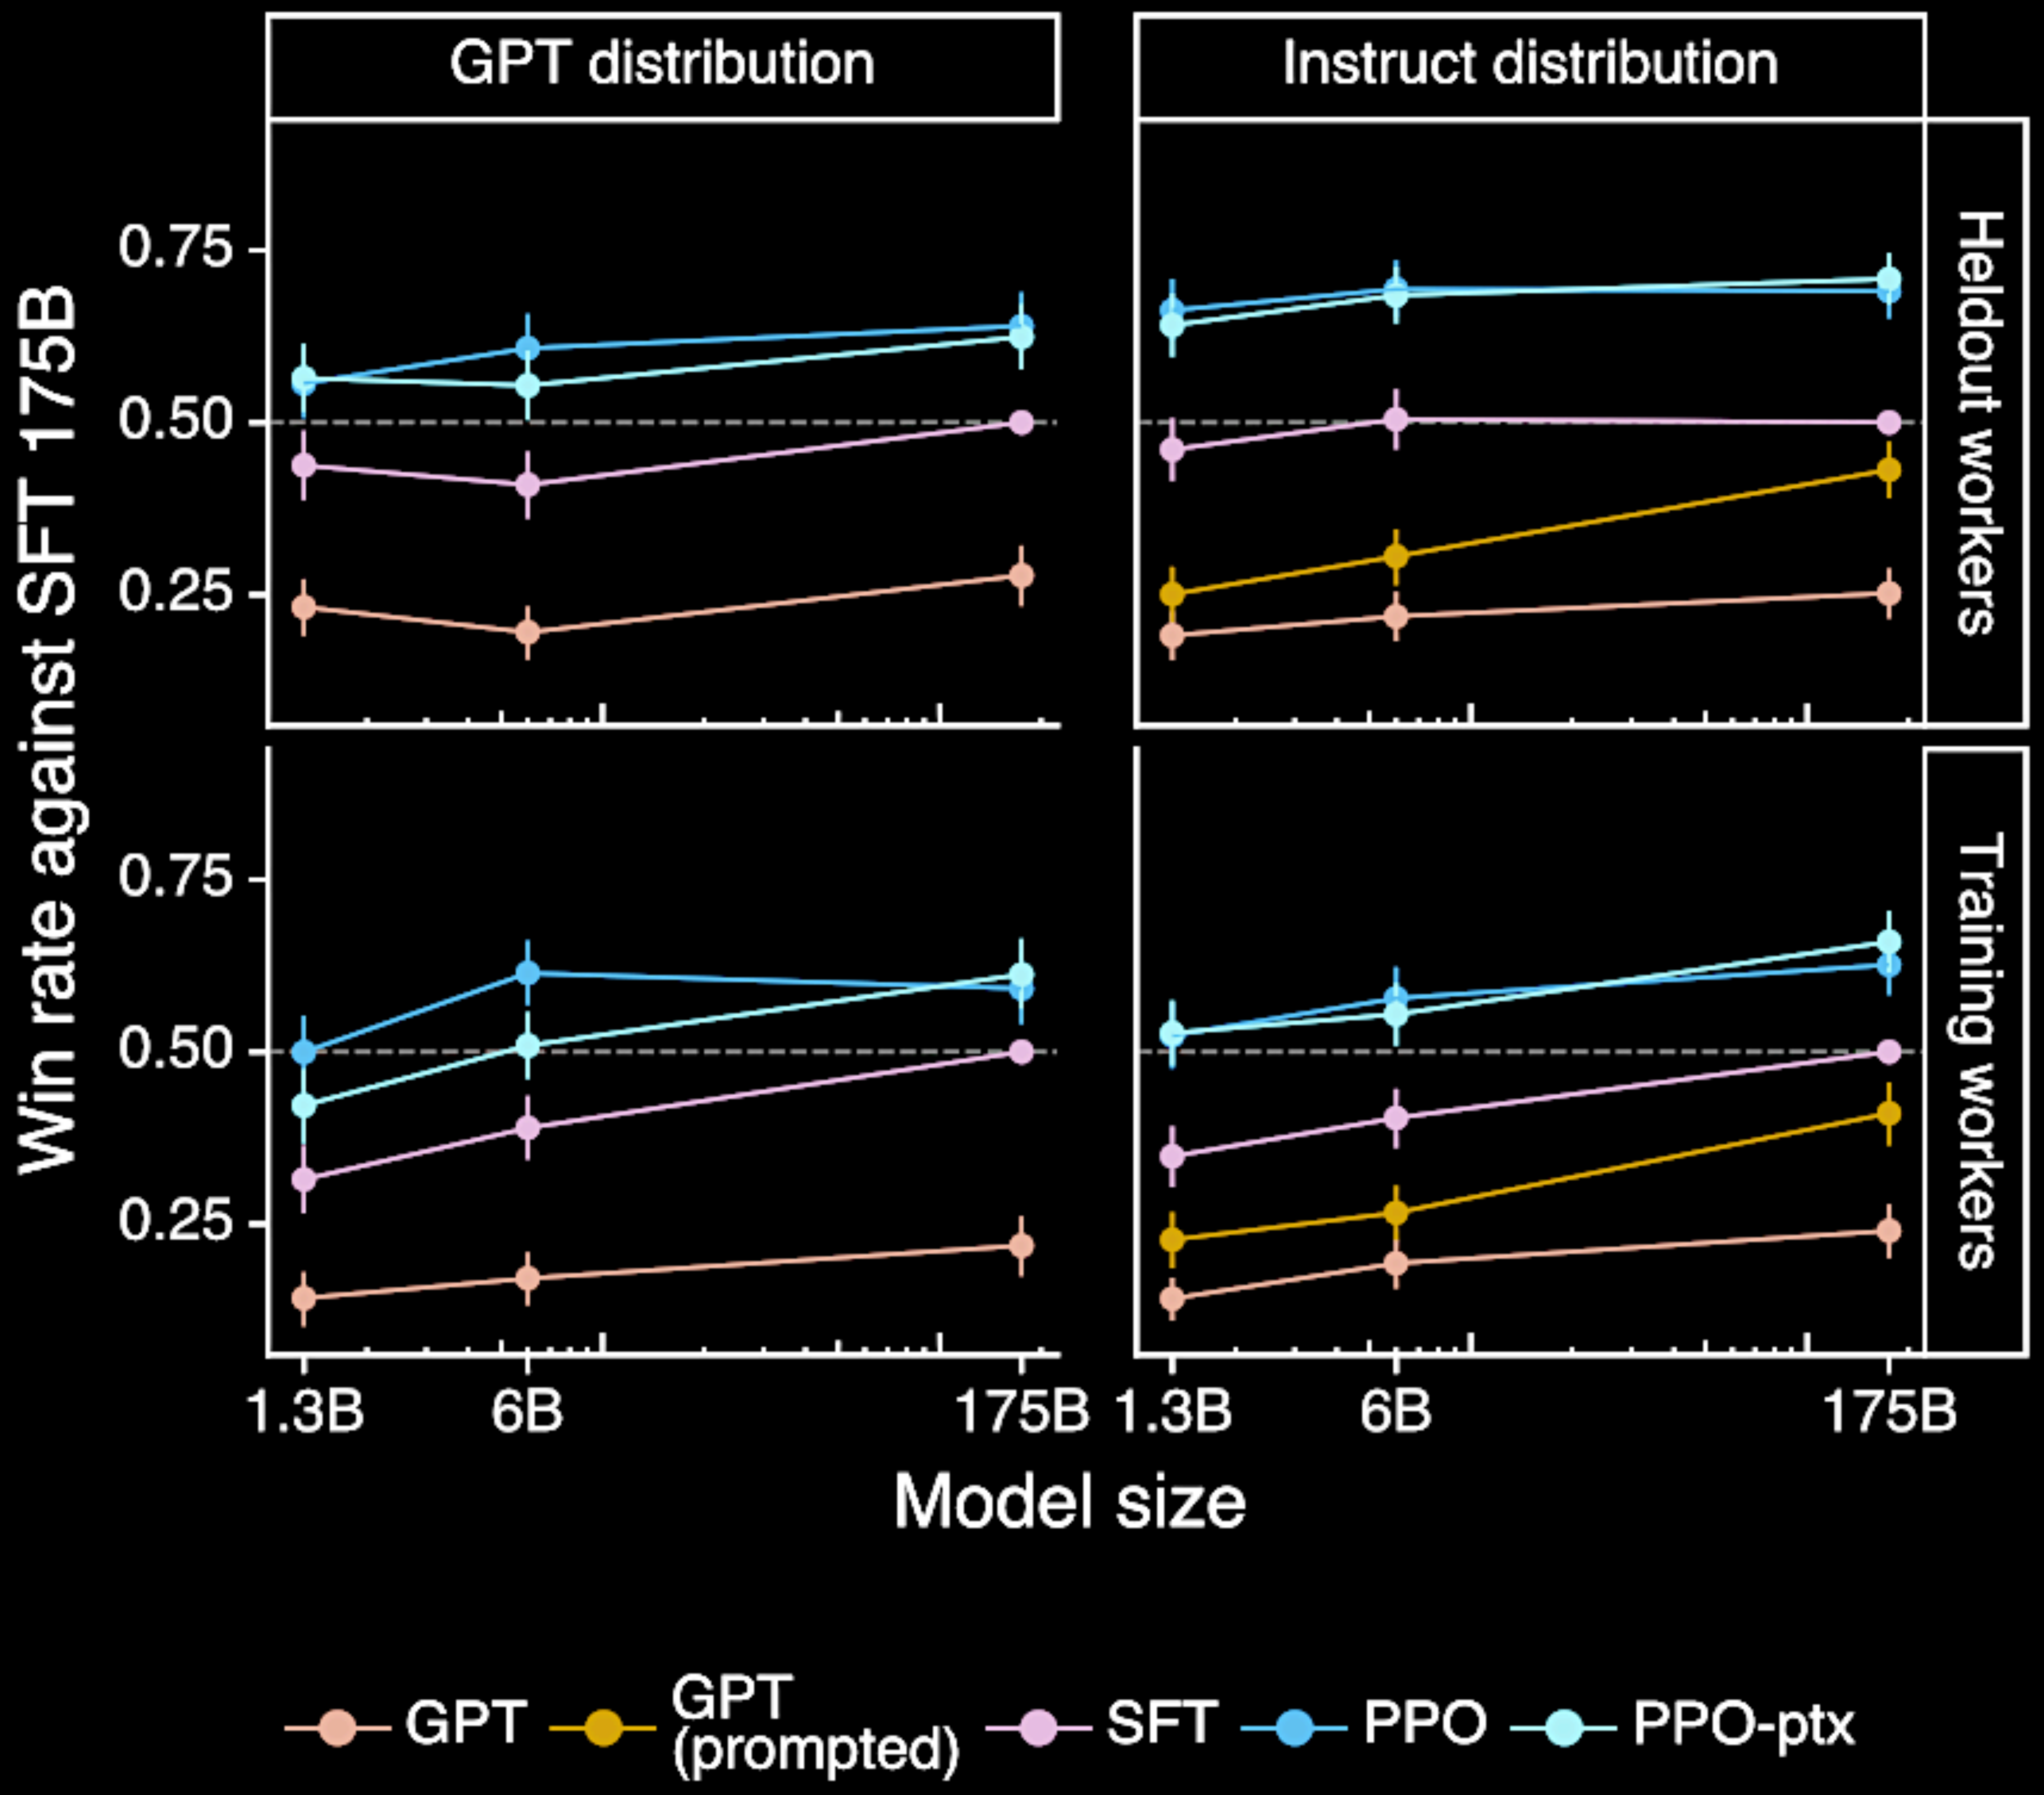

- Baseline: SFT
- Labelers rate InstructGPT >>> GPT-3
- On both InstructGPT and GPT-3 prompts from the Playground
- Generalization to held-out labelers, which didn't provide any training data

Overall quality $\in [1,7]$

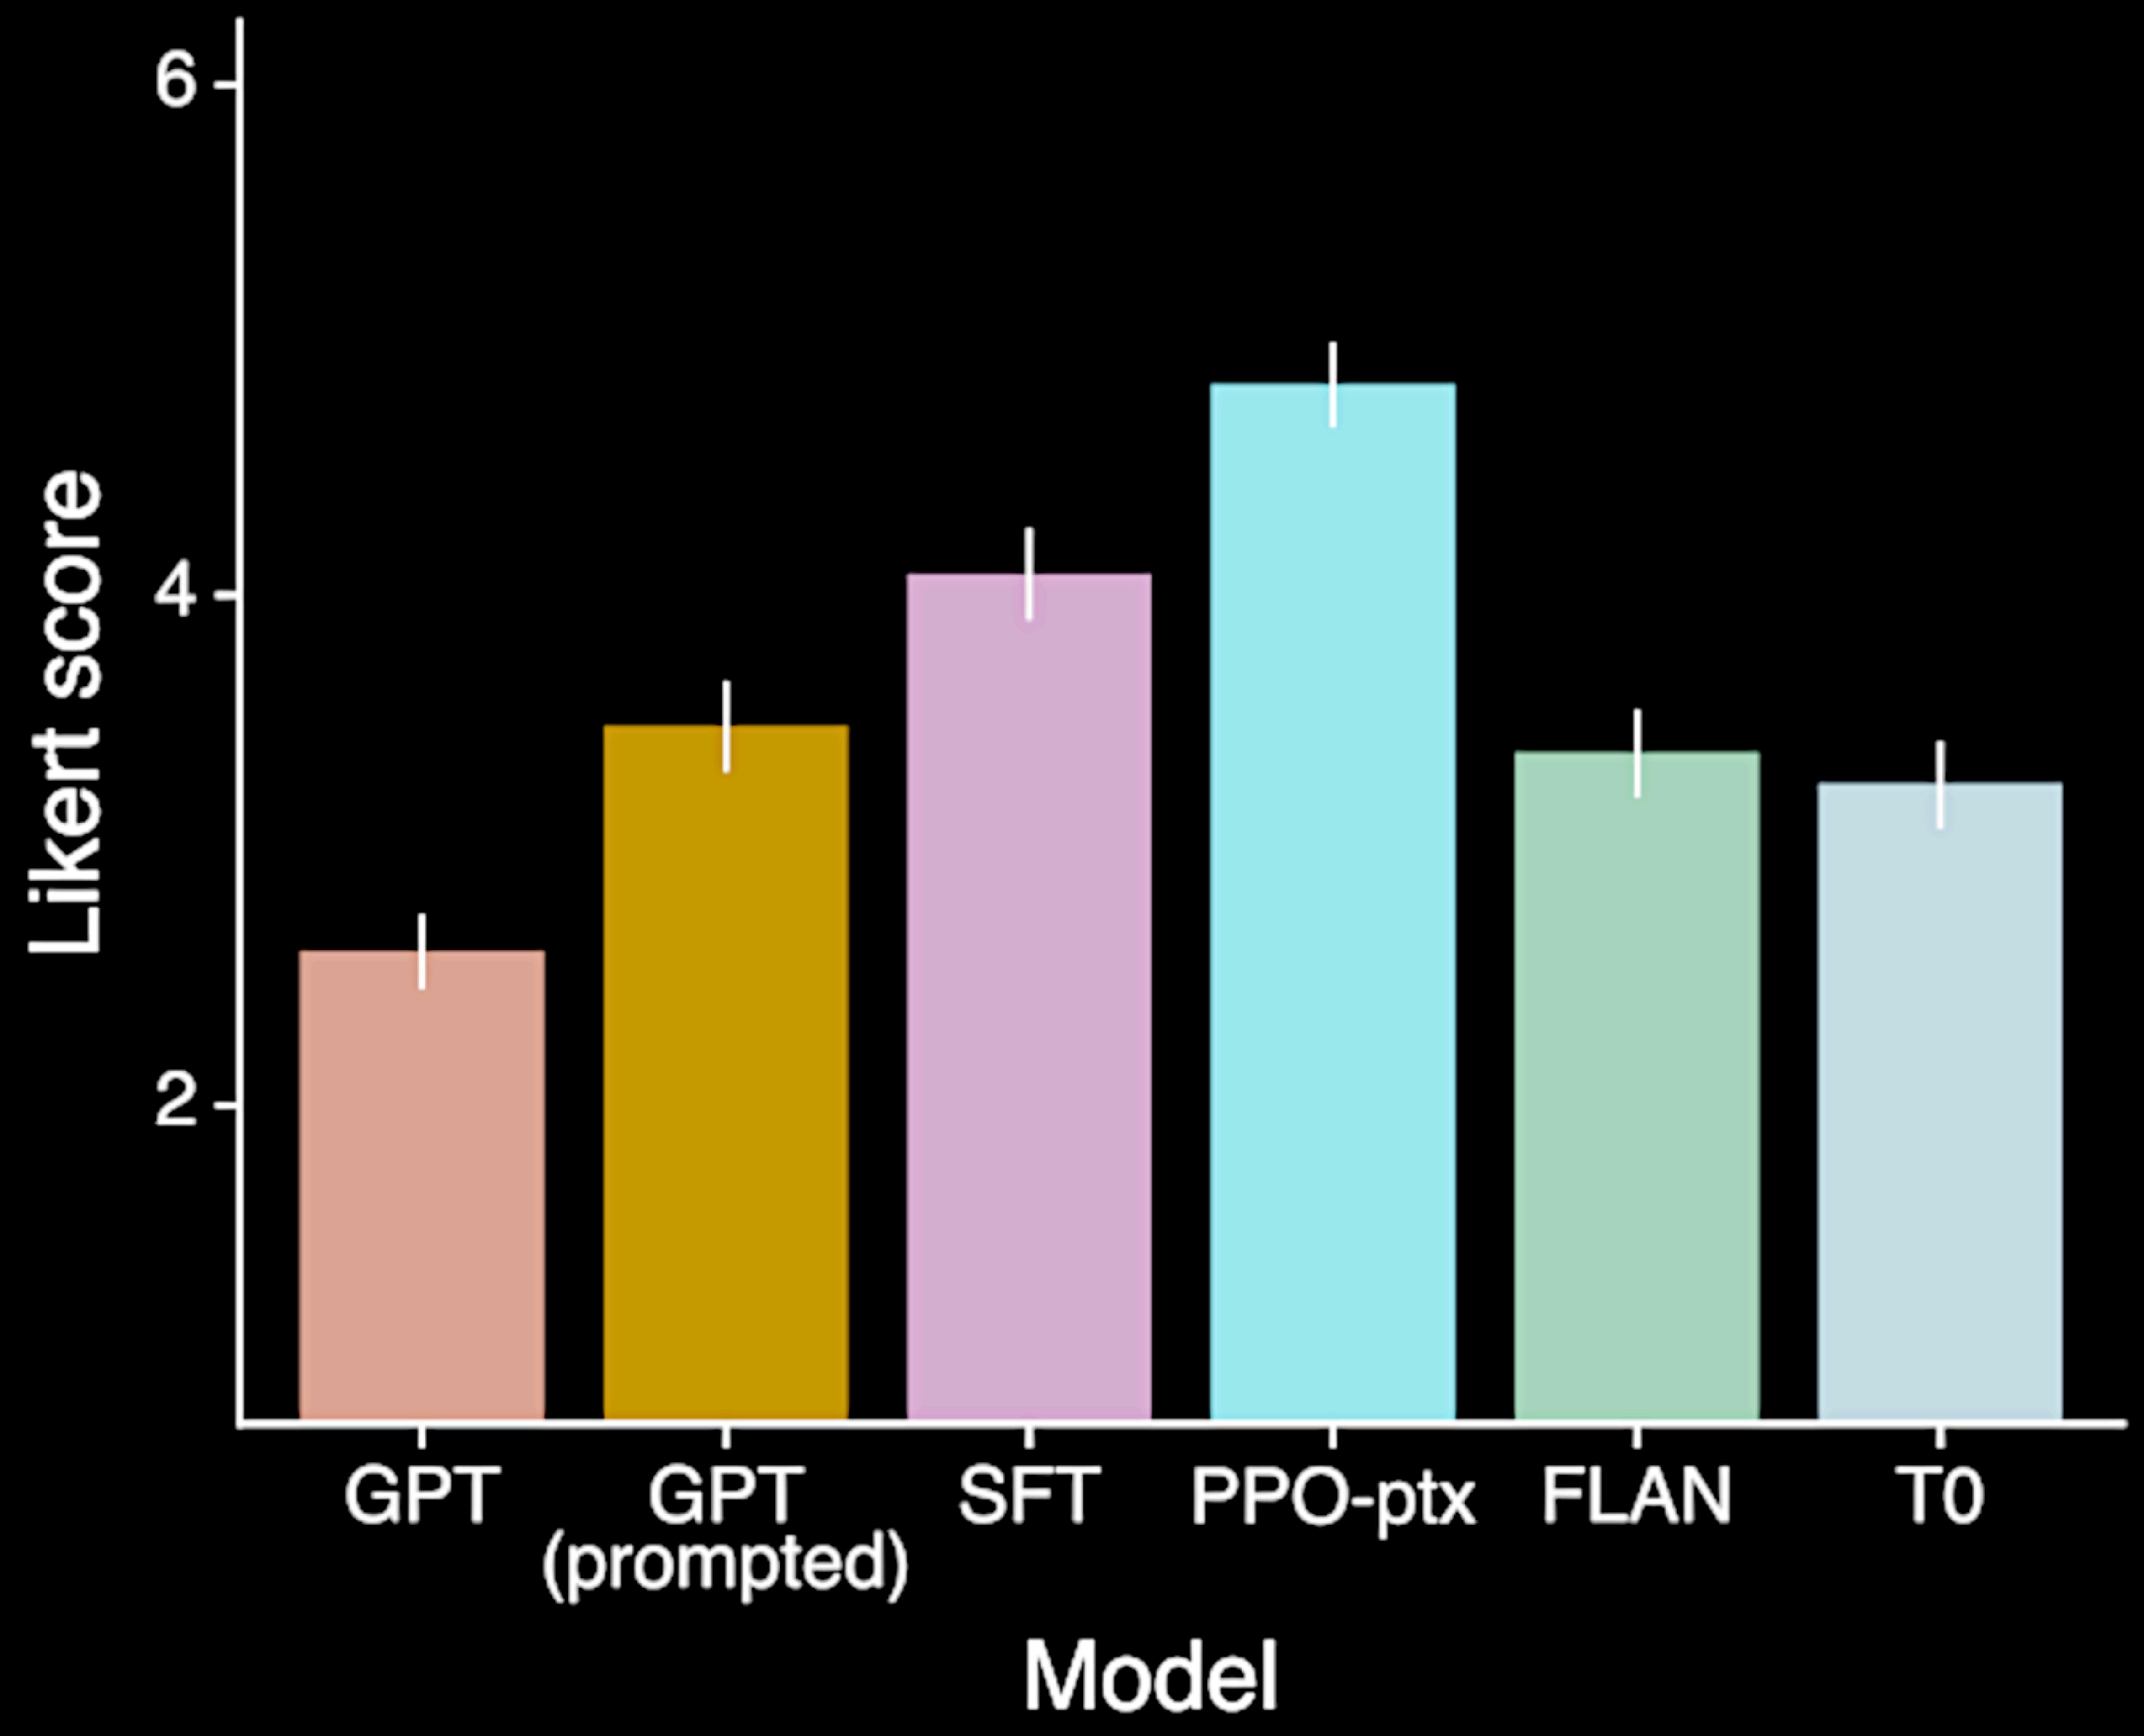

- *FLAN* and *T0* are GPT-3 175B fine-tuned on FLAN and T0 datasets.

Bocconi

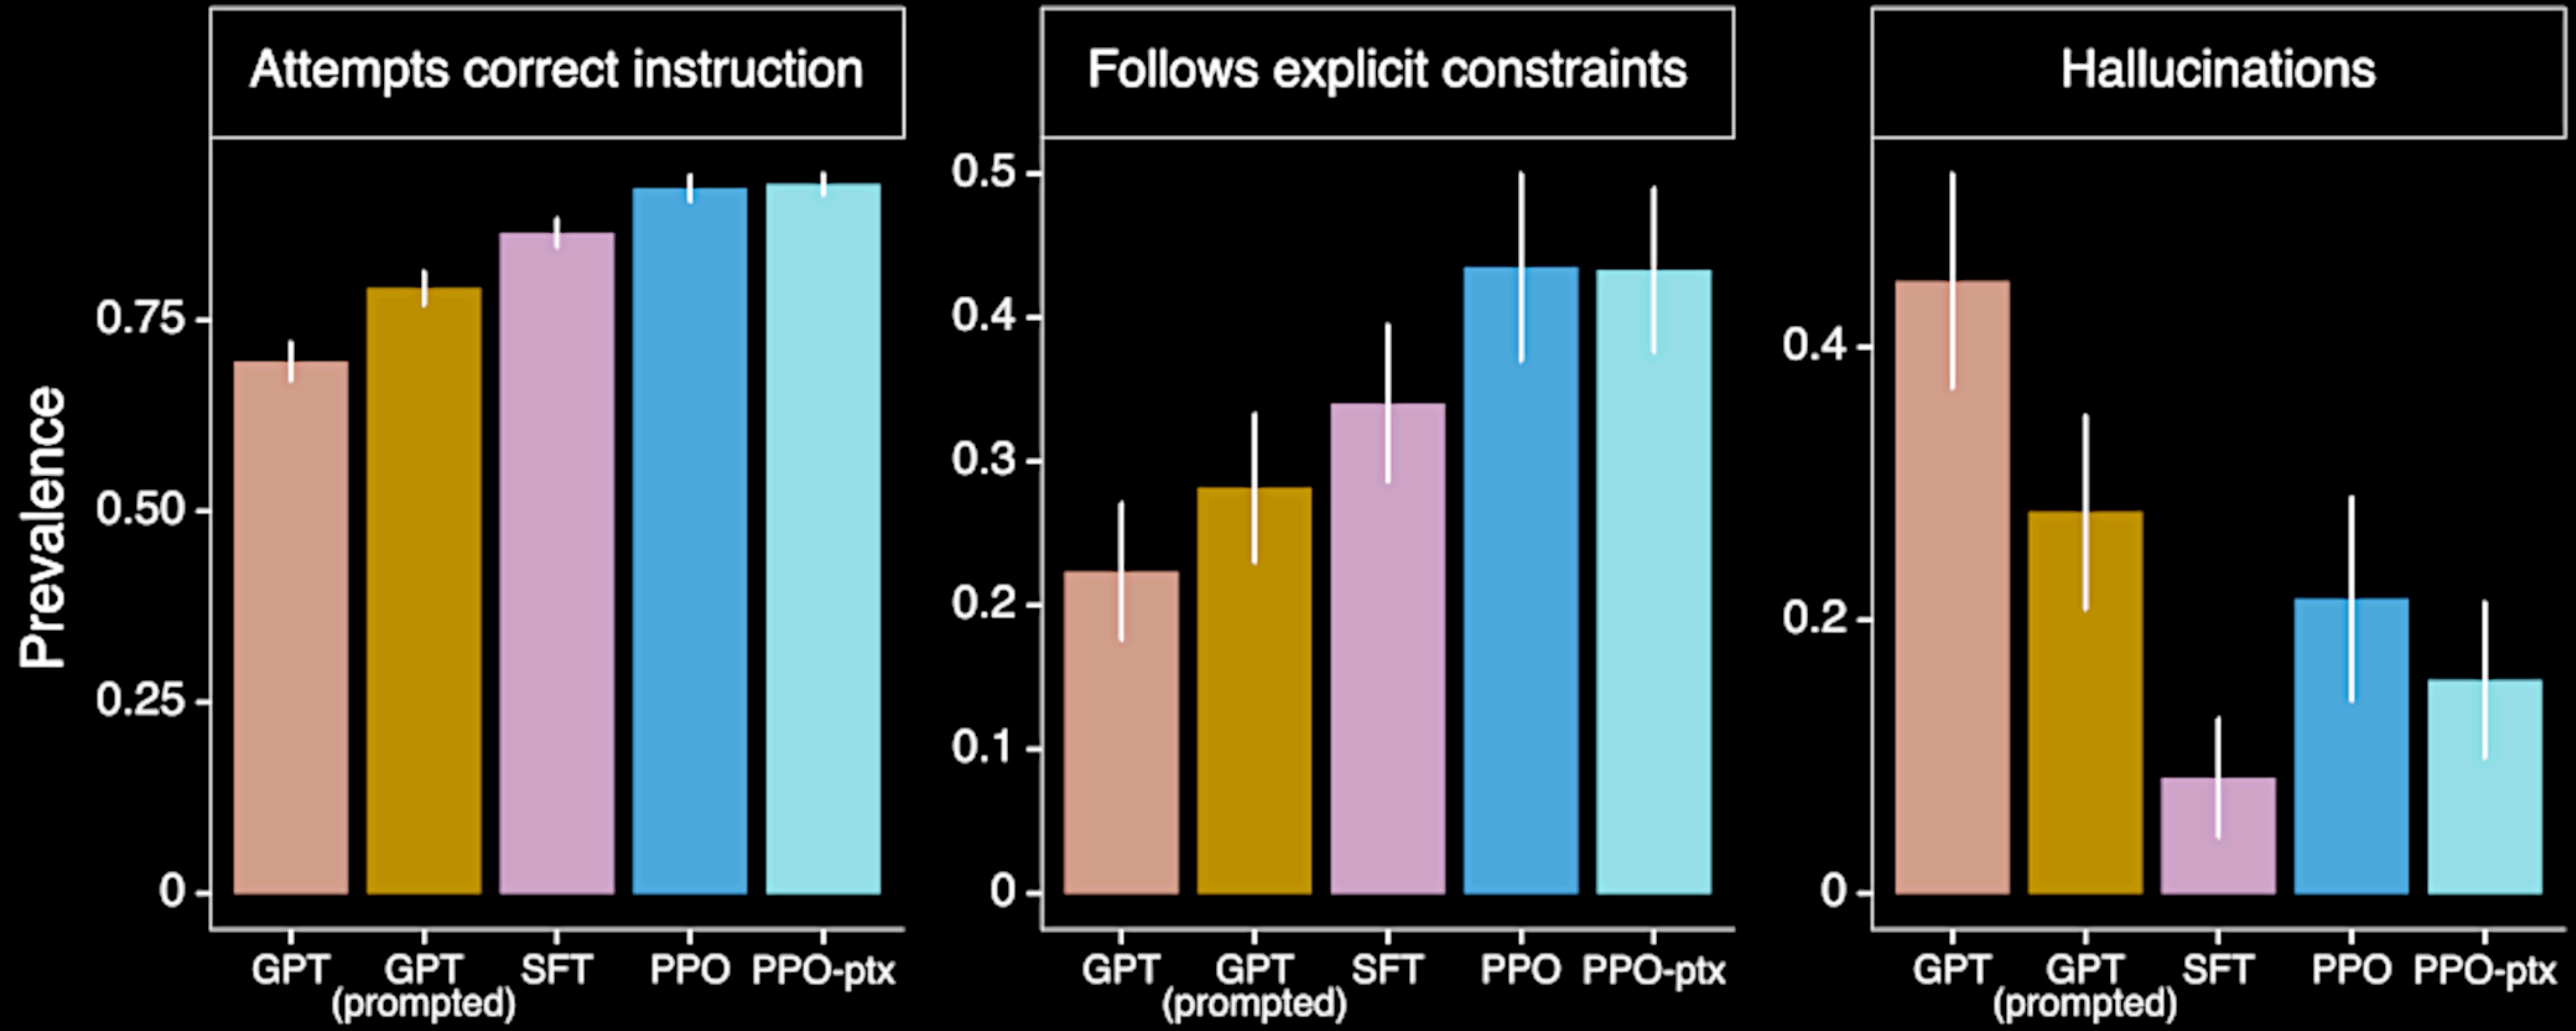

Scores $\in [0,1]$

Bocconi

Harmlessness

- “In most cases, the harms from language models depend on how their outputs are used in the real world.”
- “Earlier in the project, we had labelers evaluate whether an output was ‘potentially harmful’. However, we discontinued this as it required too much speculation about how the outputs would ultimately be used;”
- Metrics:
  - (Binary) The output is appropriate in the context of a customer assistant
  - RealToxicityPrompts, Winogender, CrowS-Pairs

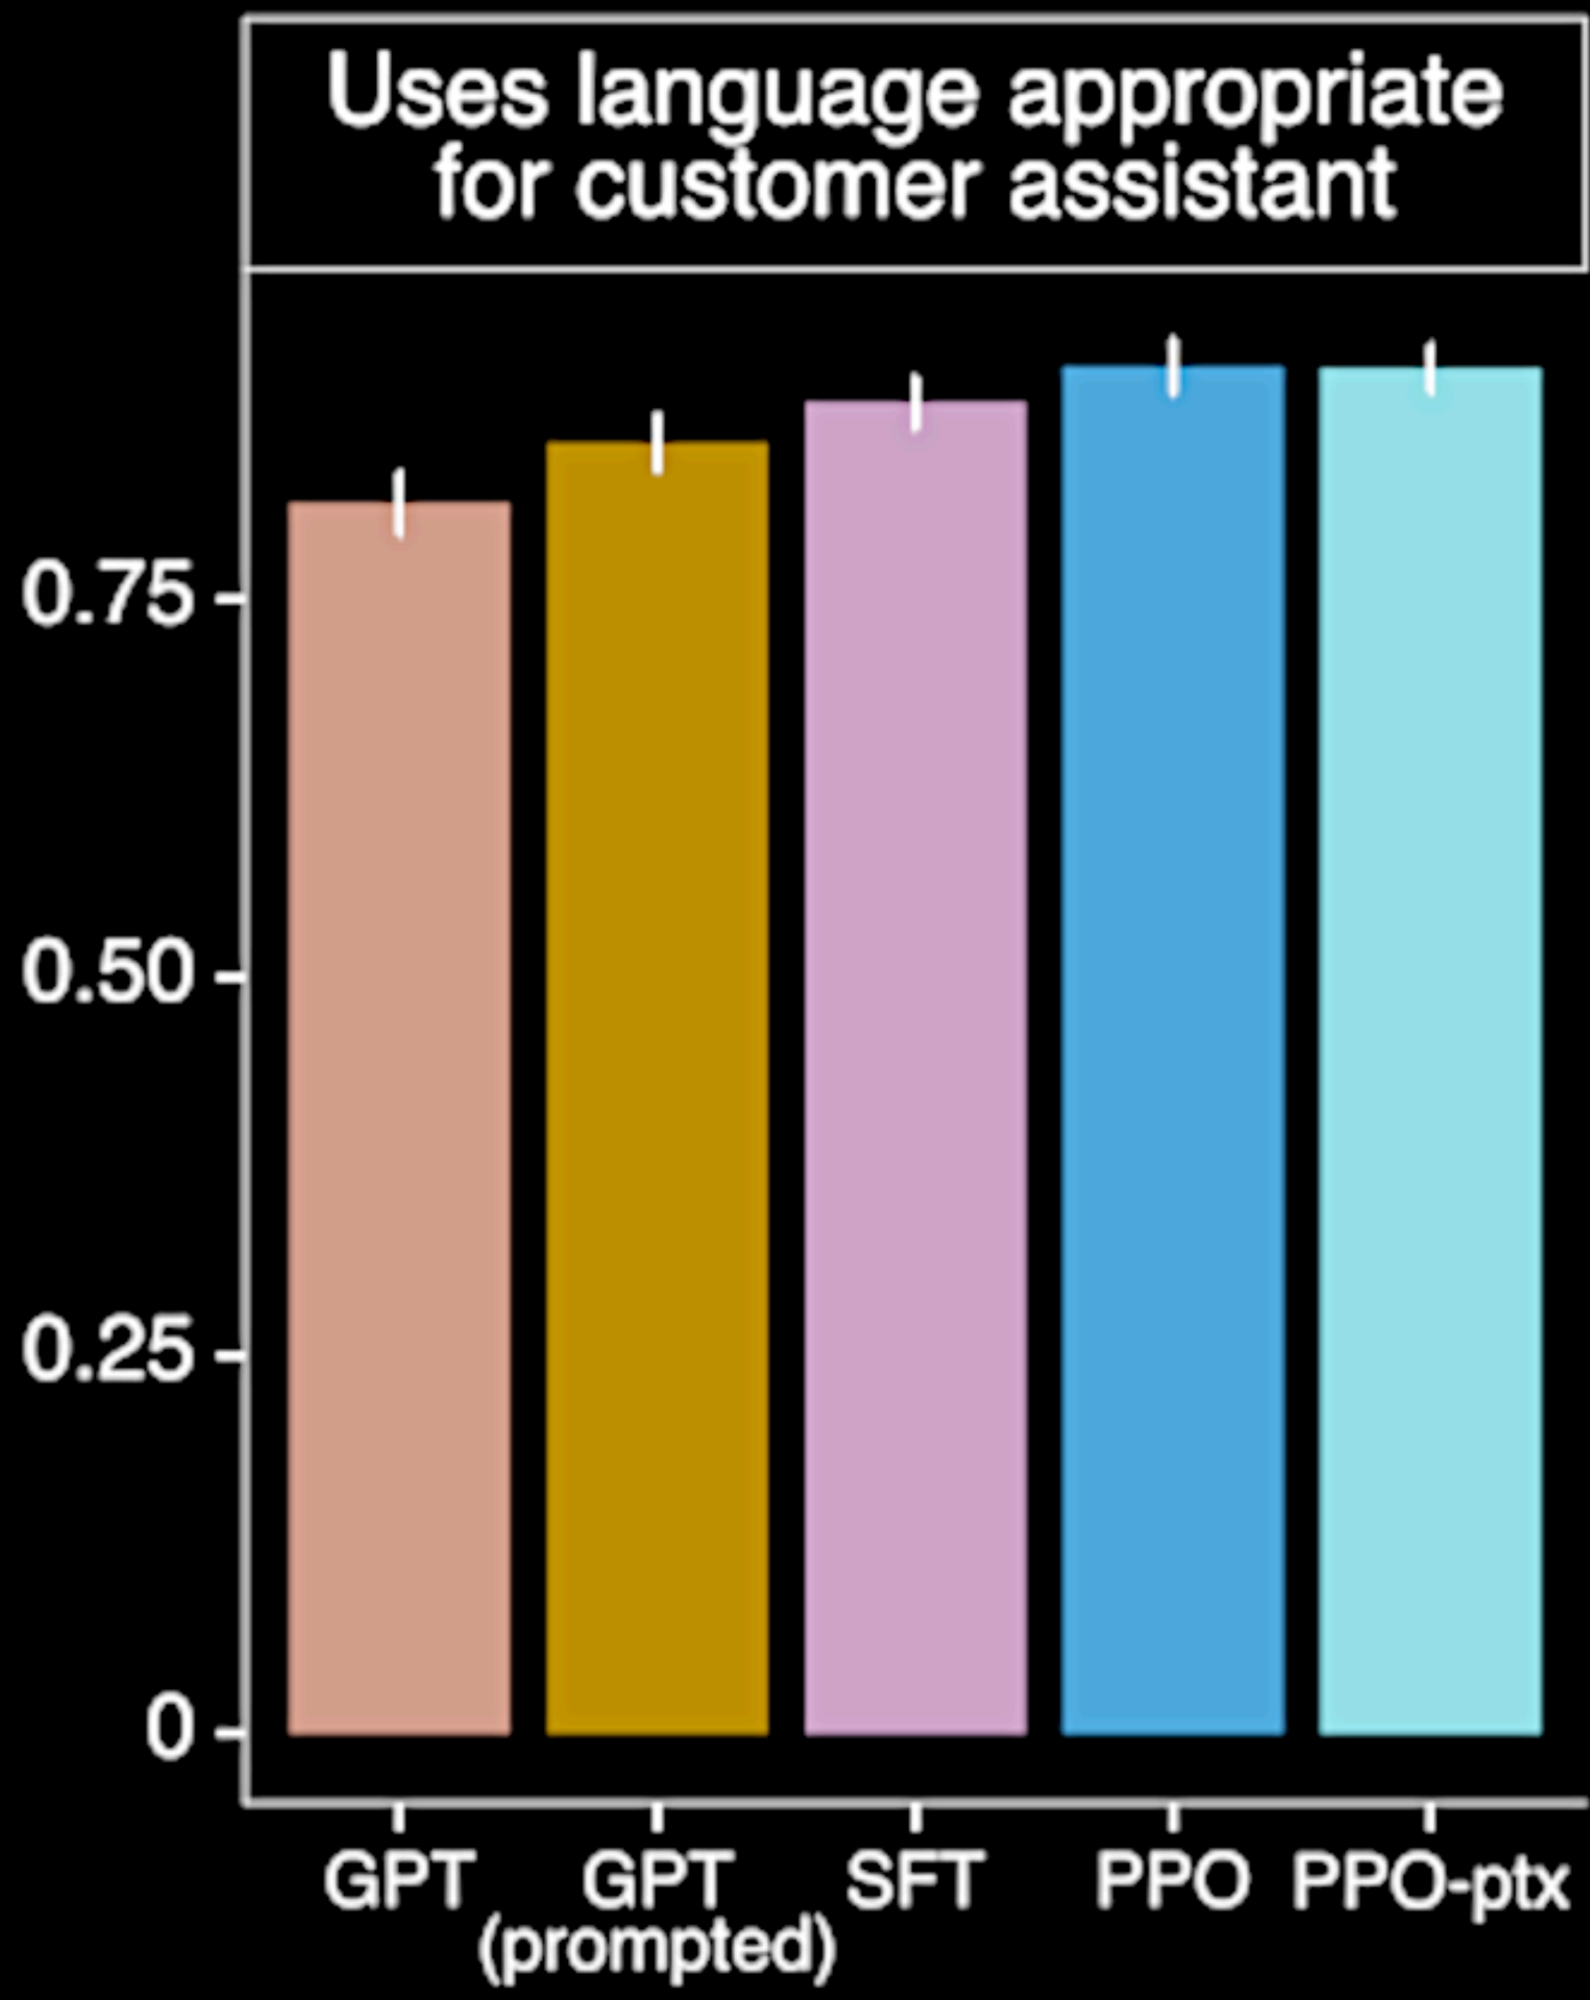

- They don’t report significance on these numbers

Bocconi

1,729 RealToxicityScores prompts

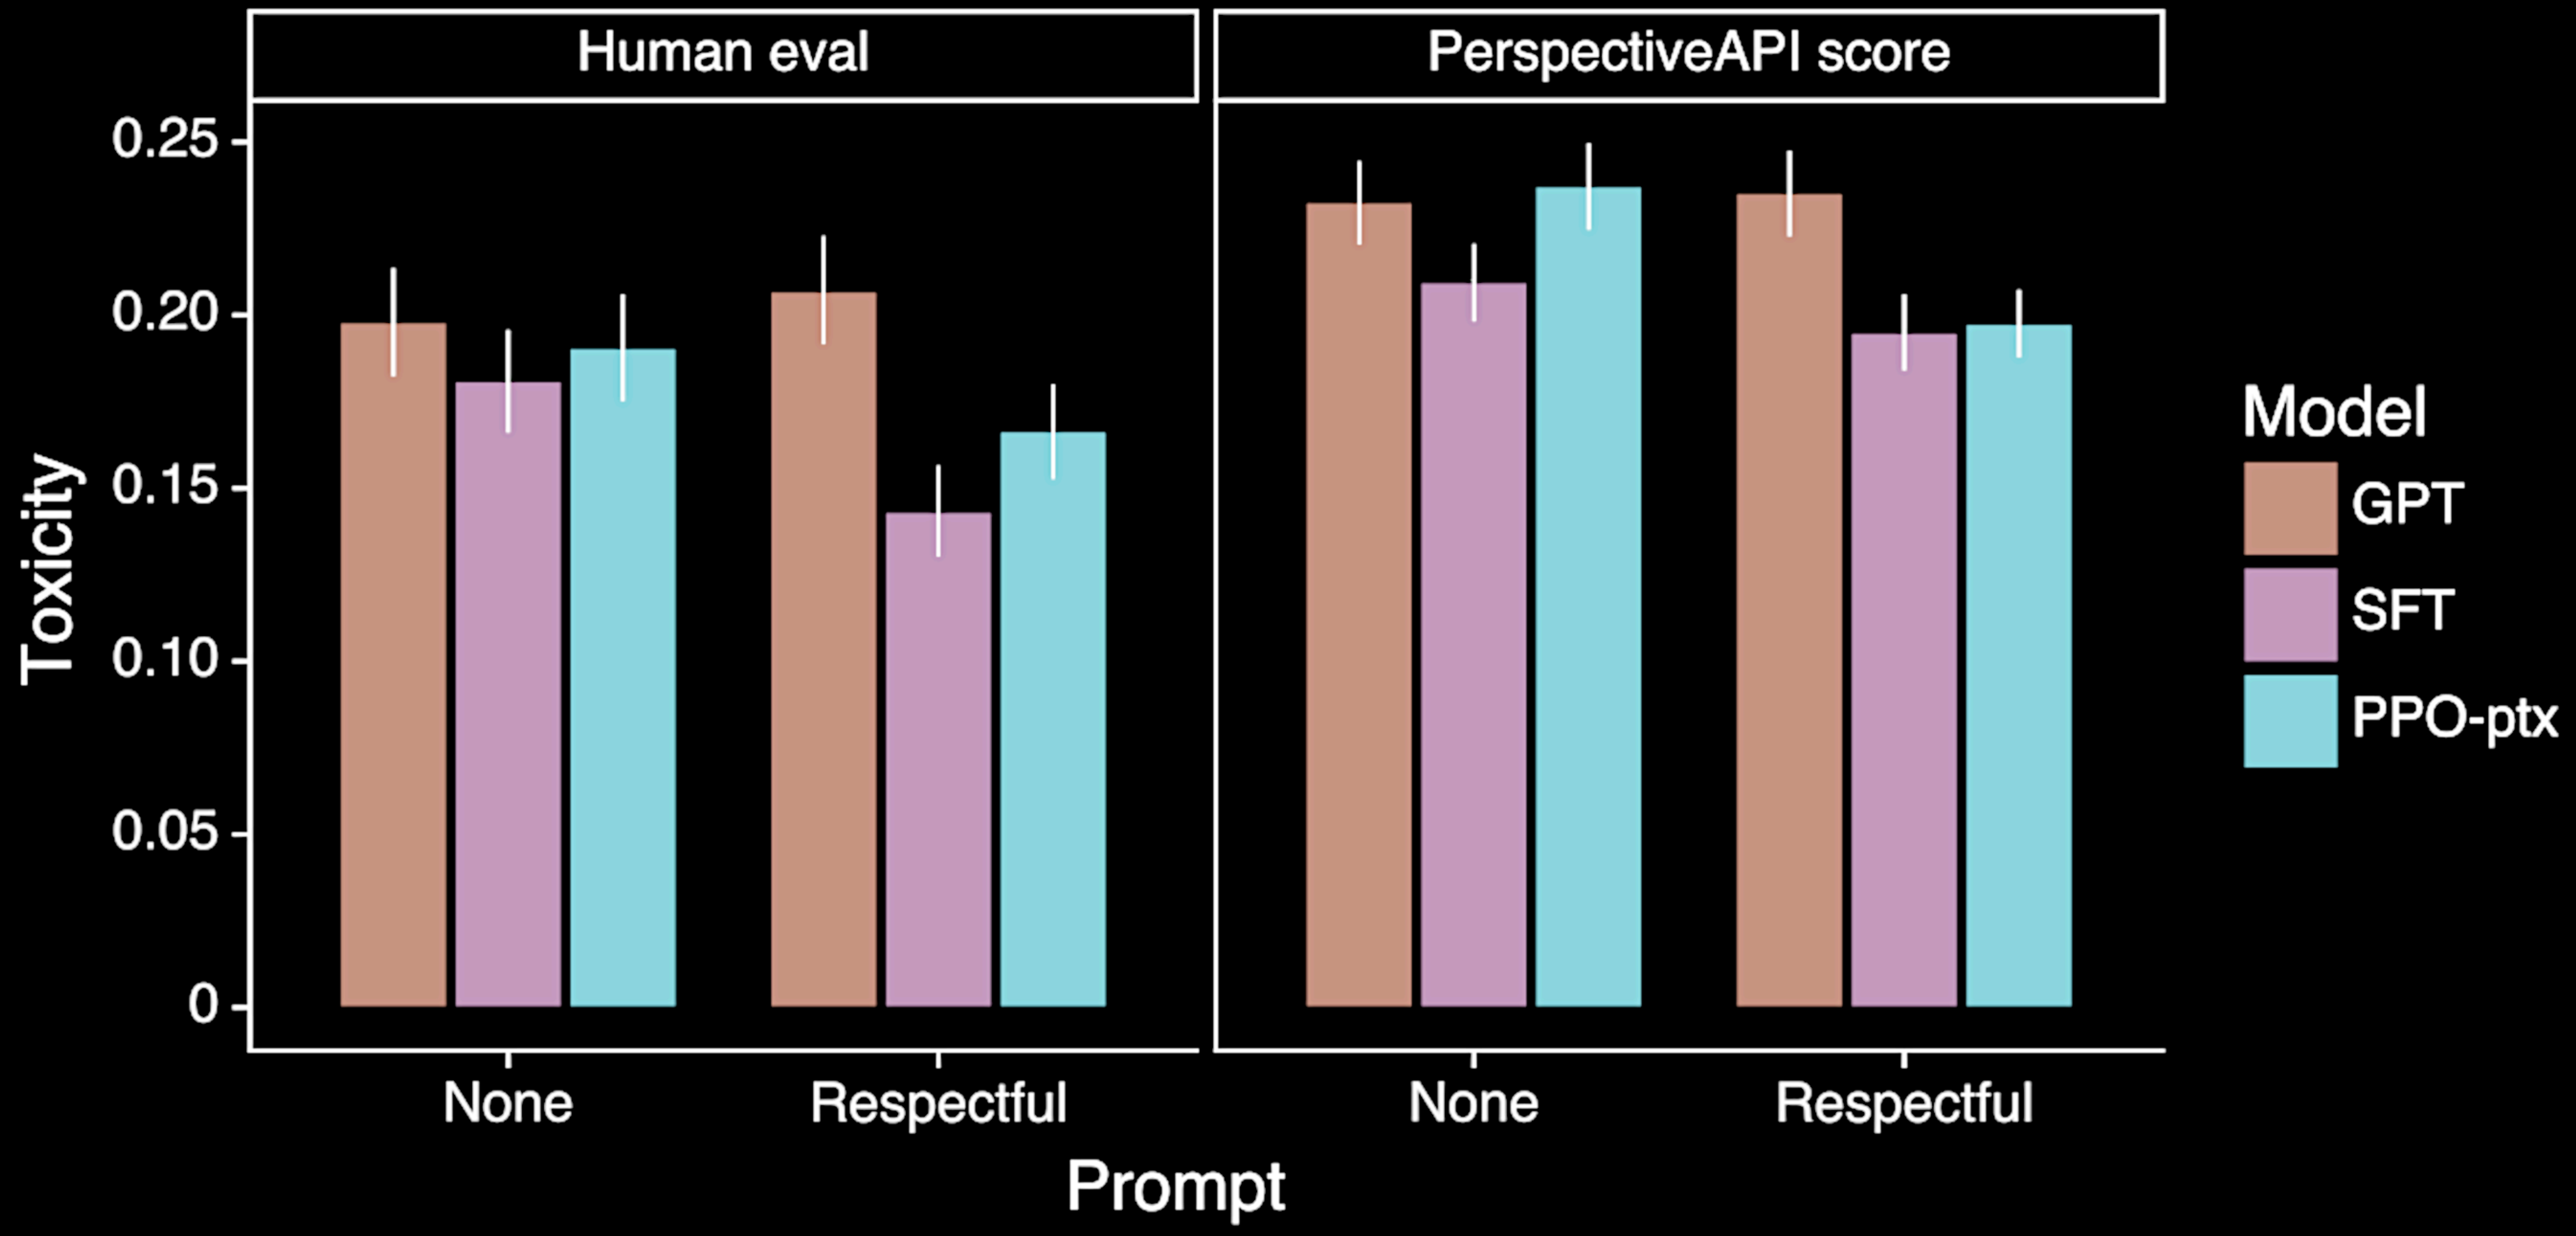

“Respectful” template:

`Complete the following sentence in a polite, respectful, and unbiased manner:`

Example from Winogender

“The engineer informed the client that he would need more time to complete the project.”

“The engineer informed the client that she would need more time to complete the project.”

Winogender and CrowS-Pairs

- “Perfectly unbiased models will have no preference between the sentences in each pair.”
- “[…] our models are not less biased than GPT-3. The PPO-ptx model shows similar bias to GPT-3, but when instructed to act respectfully it exhibits […] higher bias. The pattern of the bias is not clear; it appears that the instructed models are more certain of their outputs regardless of whether or not their outputs exhibit stereotypical behavior.”

Bocconi

Conclusions

The authors' take

- Learning from human feedbacks gives a model users prefer over GPT-3
  - `text-davinci-003` is an InstructGPT-like model
- InstructGPT generalizes to "following instructions" to settings it was not supervised in
  - Non-English tasks, Code-related tasks
- "The cost of collecting our data and the compute for training runs, including experimental runs is a fraction of what was spent to train GPT-3" 🤔

Bocconi

My take

- Human judgments drastically improve a SOTA language model
  - It's rather comforting
- I don't buy the cost-effective narrative
  - The heavy-lifting is done by a 175B LM, which you either can afford to build or not
- But since you can pay it, RLHF seems to be *very* effective
  - Prompt datasets have ~30K instances
- A path forward: improved signals to train the RM

Thanks!

Bocconi