

# **Using (Human) Feedback for Training Large Language Models**

**OR how ChatGPT is likely trained**

**Aman @ Yiming Yang's lab seminar, 2/28/2023**

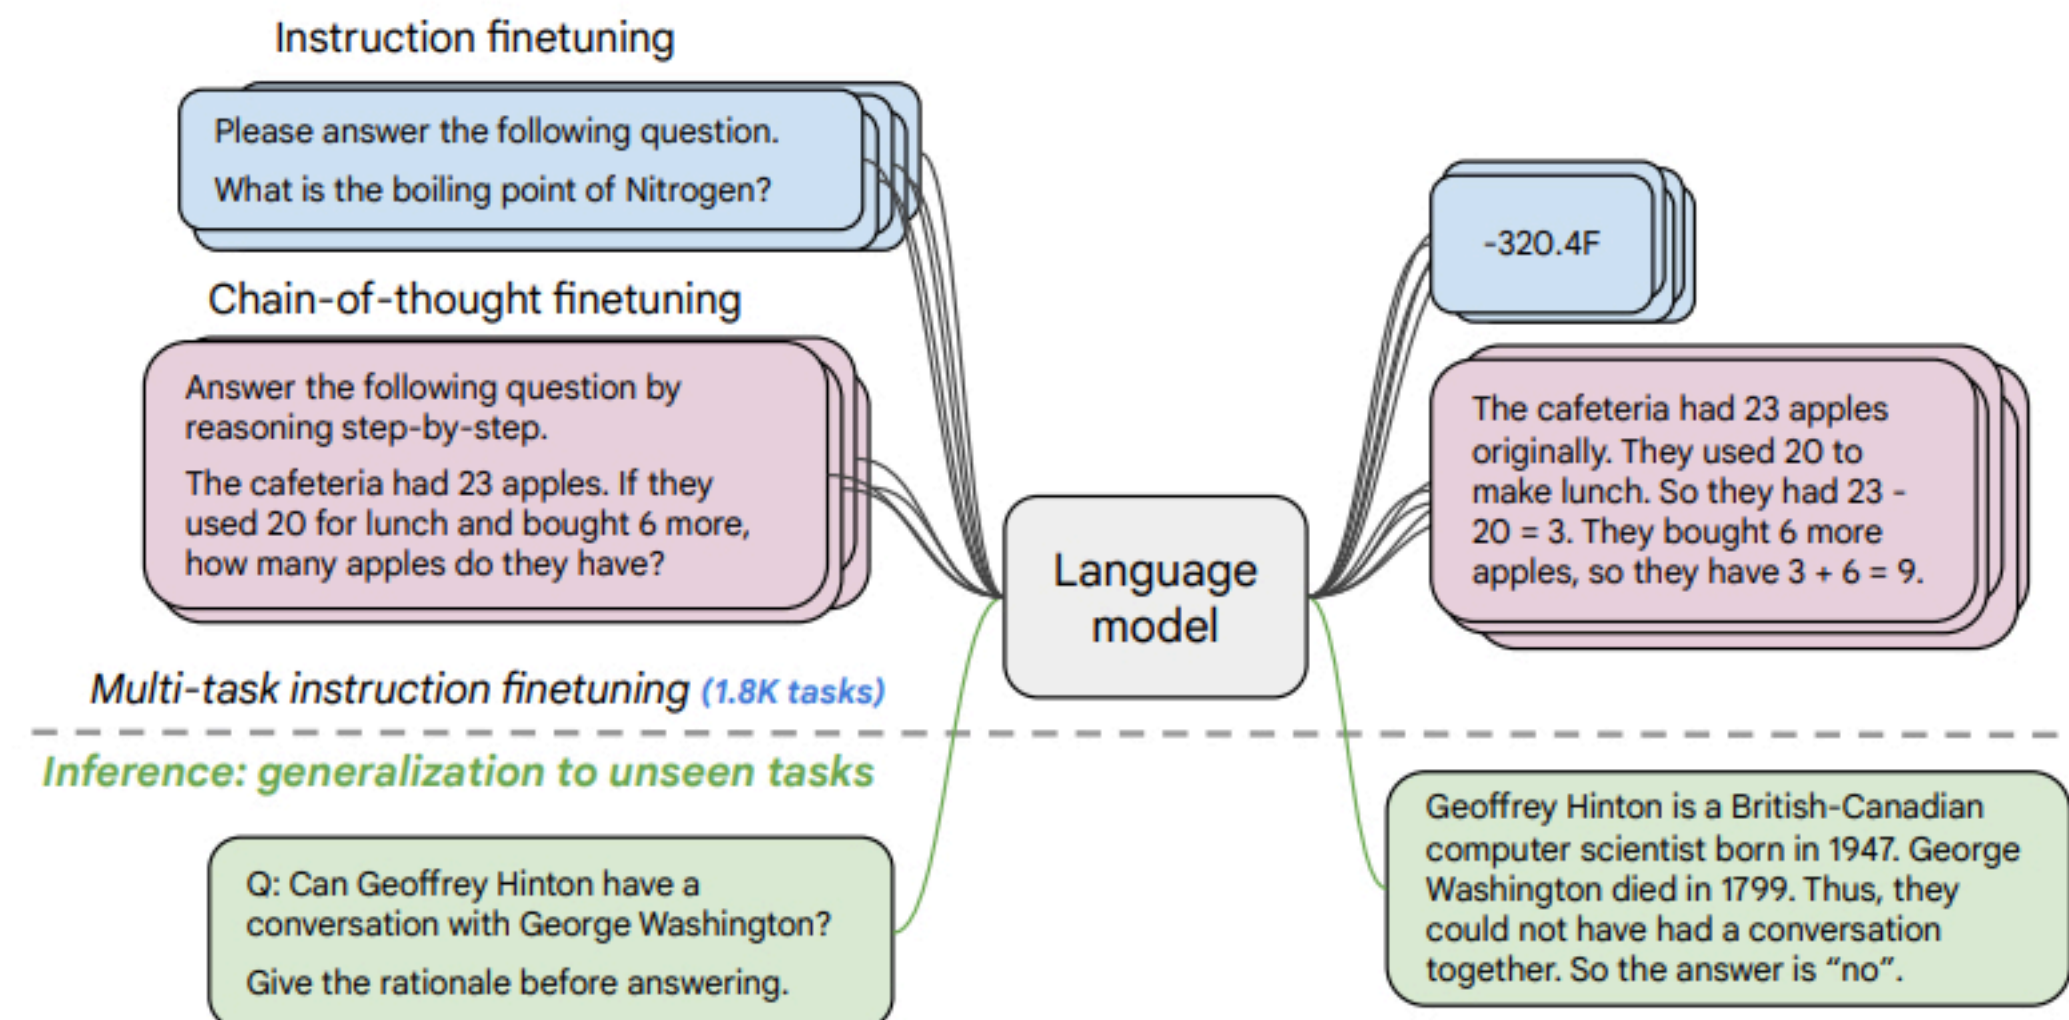
# Precursor: Instruction Tuning

- GPT-3 shows that language models trained on a large amount of data can generate fluent text ~ mid 2020
- Good language models — users want to go beyond benchmarks
- What next?
  - Want to train language models that can **follow instructions**
  - Prevent them from generating responses that are toxic and unhelpful
- Want the language models to align with what humans want

# Training language models to follow instructions

- Want the language models to align with what humans want
  - Instruction tuning was an early attempt at this

- FLAN
- T0
- Lambda



## Scaling Instruction-Finetuned Language Models

Hyung Won Chung<sup>1</sup> Le Hou<sup>2</sup> Shayne Longpre<sup>1</sup> Barret Zoph<sup>1</sup> Yi Tay<sup>1</sup>  
William Fedus<sup>1</sup> Yunxuan Li<sup>1</sup> Xuezhi Wang<sup>1</sup> Mostafa Dehghani<sup>1</sup> Siddhartha Brahma<sup>1</sup>  
Albert Webson<sup>1</sup> Shixiang Shane Gu<sup>1</sup> Zhuyun Dai<sup>1</sup> Mirac Suzgun<sup>1</sup> Xinyun Chen<sup>1</sup>  
Aakanksha Chowdhery<sup>1</sup> Alex Castro-Ros<sup>1</sup> Marie Pellat<sup>1</sup> Kevin Robinson<sup>1</sup>  
Dasha Valter<sup>1</sup> Sharan Narang<sup>1</sup> Gaurav Mishra<sup>1</sup> Adams Yu<sup>1</sup> Vincent Zhao<sup>1</sup>  
Yanping Huang<sup>1</sup> Andrew Dai<sup>1</sup> Hongkun Yu<sup>1</sup> Slav Petrov<sup>1</sup> Ed H. Chi<sup>1</sup>  
Jeff Dean<sup>1</sup> Jacob Devlin<sup>1</sup> Adam Roberts<sup>1</sup> Denny Zhou<sup>1</sup> Quoc V. Le<sup>1</sup>  
Jason Wei<sup>1</sup>

# Why Instruction Tuning isn't Enough?

- The models might become better at task understanding — but still nontrivial to generate a desirable sequence
- Alignment goes beyond instruction following
- Real-world behavior is quite different from benchmark datasets

# Why human feedback?

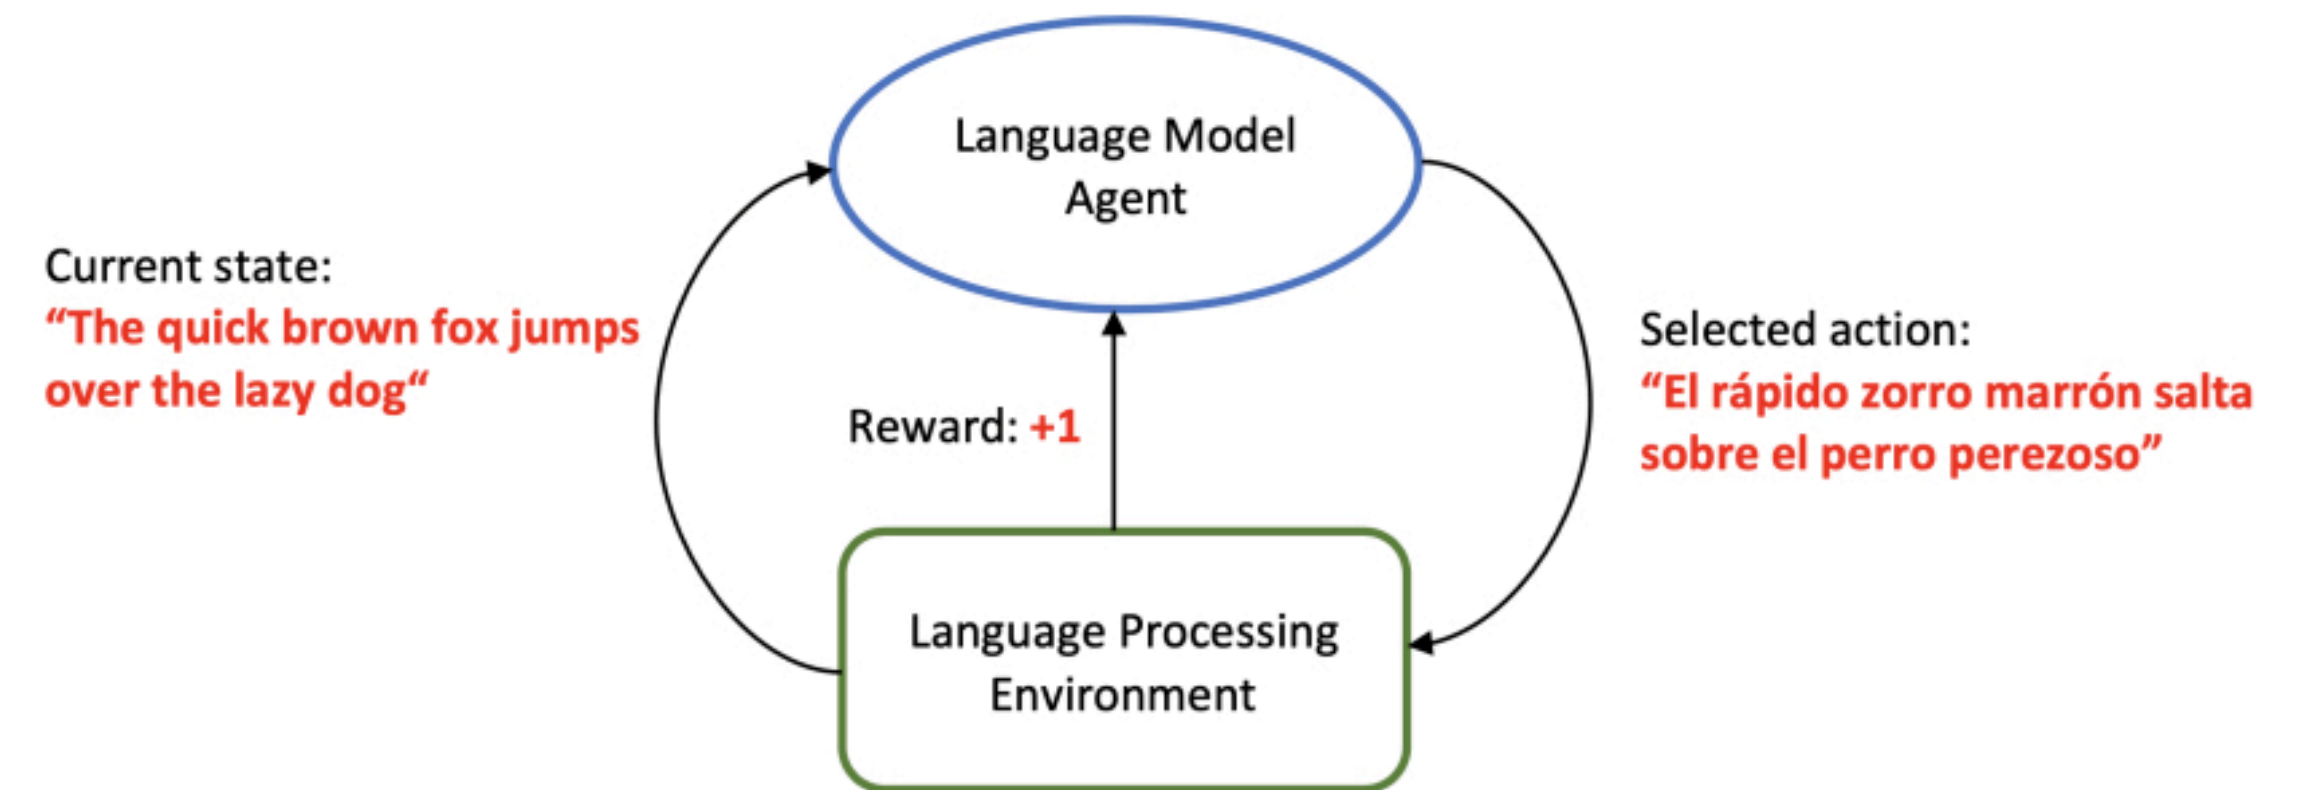
- Hard to quantify the requirements or the definition of “good”
- Task:
  - Complete the sentence “I saw the movie last night” to make a positive review
- Completion 1:
  - I saw the movie last night and found it to be a thoroughly enjoyable experience.
- Completion 2:
  - I saw the movie last night and it was soooo good! Like, really, really good!
- Which response will humans prefer?
  - Subjective, but maybe given the goals of the system (general purpose chatbot) + sizable annotator pool

# Two Camps

- RL
  - Collect some human labels and fine-tune LMs
- ChatGPT / GPT-3 Families
- Claude by Anthropic
- Supervised
  - Collect lots of training data and do good old supervised learning
  - Flan-T5-XXL (best open source model)
  - Large datasets for instruction tuning:
    - T0
    - Flan

# Connection between RL and LM

- Action space: vocabulary  $V$
- Policy: language model  
 $p_{\theta}(x_i | x_0, x_1, \dots, x_{i-1}), x_i \in V$
- Reward: function  $r$  (e.g., BLEU) scored per token or for the entire sequence (typical)



Survey on reinforcement learning for language processing

Víctor Uc-Cetina<sup>1</sup>, Nicolás Navarro-Guerrero<sup>2</sup>, Anabel Martin-Gonzalez<sup>1</sup>,  
Cornelius Weber<sup>3</sup>, Stefan Wermter<sup>3</sup>

# Connection between RL and LM

- Action space: vocabulary  $V$
- Policy: language model  $p_{\theta}(x_i | x_0, x_1, \dots, x_{i-1}), x_i \in V$
- Reward: function  $r$  (e.g., BLEU) scored per token or for the entire sequence (typical)
- In theory, can “fine-tune”  $p_{\theta}$  given a reward function  $r$  using any off-the-shelf RL algorithm
  - In practice, modern implementations using proximal-policy optimization (PPO)
  - Not discussed, consider a black box RL algorithm
- Focus on:
  - Human feedback
  - Design of reward function  $r$



# Outline of the talk

- Background
- RL + Human feedback
  - Fine-tuning LMs with Human Feedback
  - InstructGPT
- Recent works that include feedback without RL
  - Hindsight-tuning
  - Self-correct

---

## Fine-Tuning Language Models from Human Preferences

---

**Daniel M. Ziegler\*** **Nisan Stiennon\*** **Jeffrey Wu** **Tom B. Brown**  
**Alec Radford** **Dario Amodei** **Paul Christiano** **Geoffrey Irving**  
OpenAI

{dmz,nisan,jeffwu,tom,alec,damodei,paul,irving}@openai.com

# Fine-Tuning Language Models from Human Preferences

- Given a fixed (base) language model, improve its outputs to align better with some desired goal
  - Example, given a partial review, make the completions more positive.
- I saw the movie last night < complete this part >
  - < complete this part > : and it was amazing
  - < complete this part > : and it was okay
  - < complete this part > : and it was the worst
- Summarize an article such that the summary is one preferred by humans.

# Fine-Tuning Language Models from Human Preferences

- Goal:
  - Can we use human feedback to fine-tune models?
- Steps:
  - Step 1: Collect human labels
  - Step 2: Train a reward model
  - Step 3: Fine-tune the language model with the reward model

# Step 1: Collect Human Labels

- Use an external service (Scale AI)
- Let  $\rho$  be the starting language model
- Use  $\rho$  to generate 4 outputs (continuations) for each input (context)  $x$ 
  - $(x, y_0, y_1, y_2, y_3)$
- Human raters pick the best one
  - $(x, y_0, y_1, y_2, y_3, b), b \in [0, 1, 2, 3]$

---

<sup>1</sup>In early experiments we found that it was hard for humans to provide consistent fine-grained quantitative distinctions when

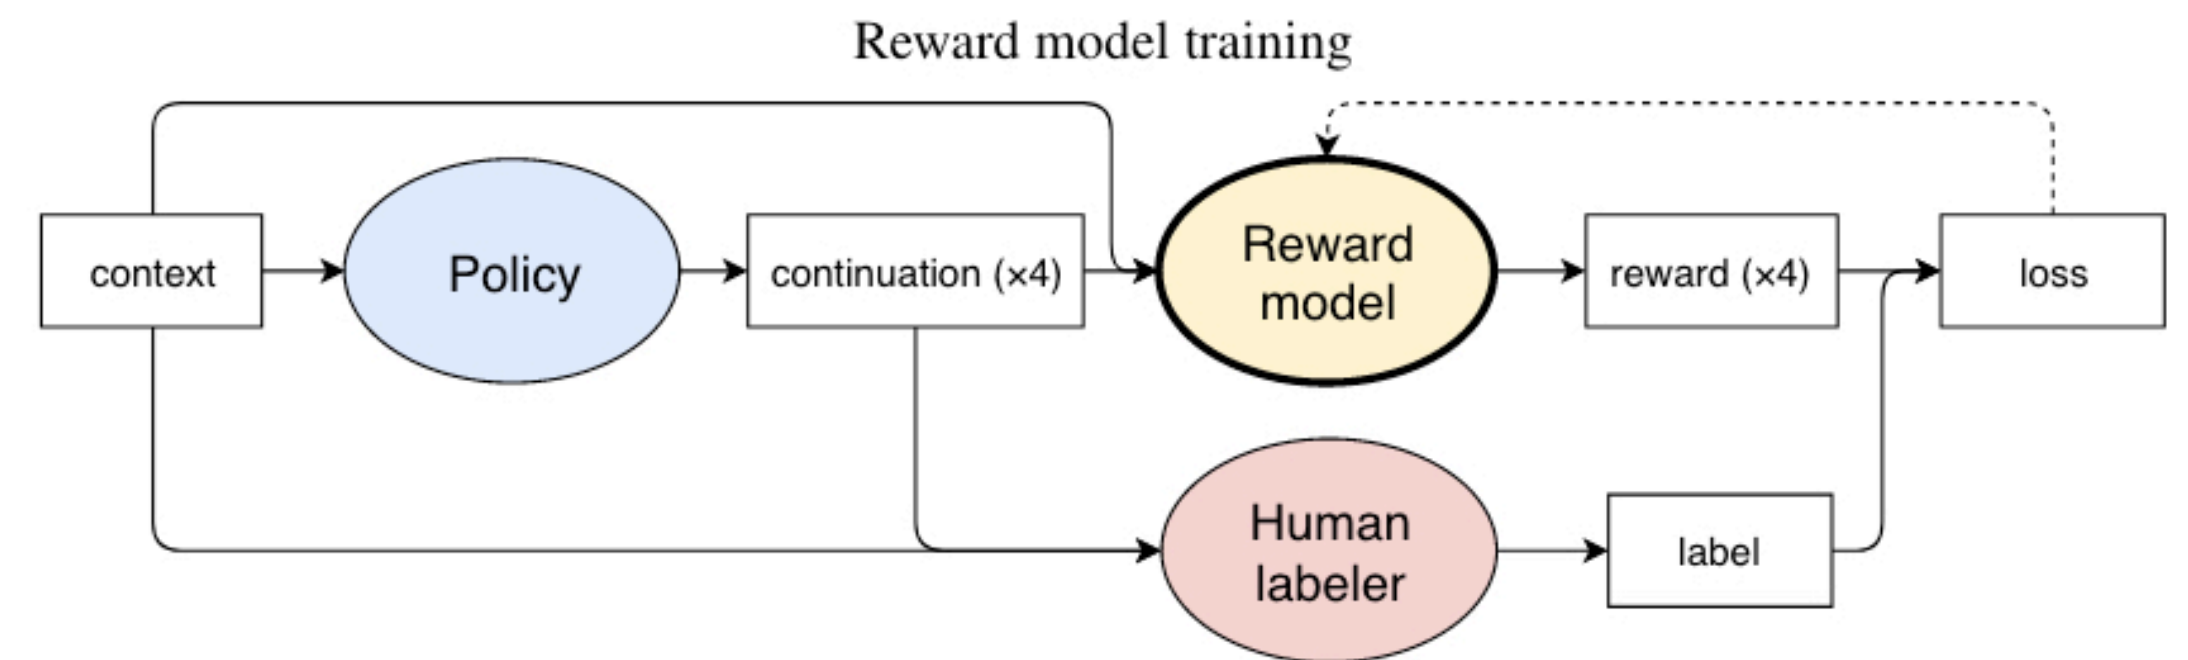
a given input  $x$ .<sup>1</sup> We ask humans to choose between four options  $(y_0, y_1, y_2, y_3)$ ; considering more options allows a human to amortize the cost of reading and understanding the prompt  $x$ . Let  $b \in \{0, 1, 2, 3\}$  be the option they select.

# Step 2: Train Reward model

- Train a model that learns to rate those completions higher that are also preferred by humans.
- $r$  captures human preferences

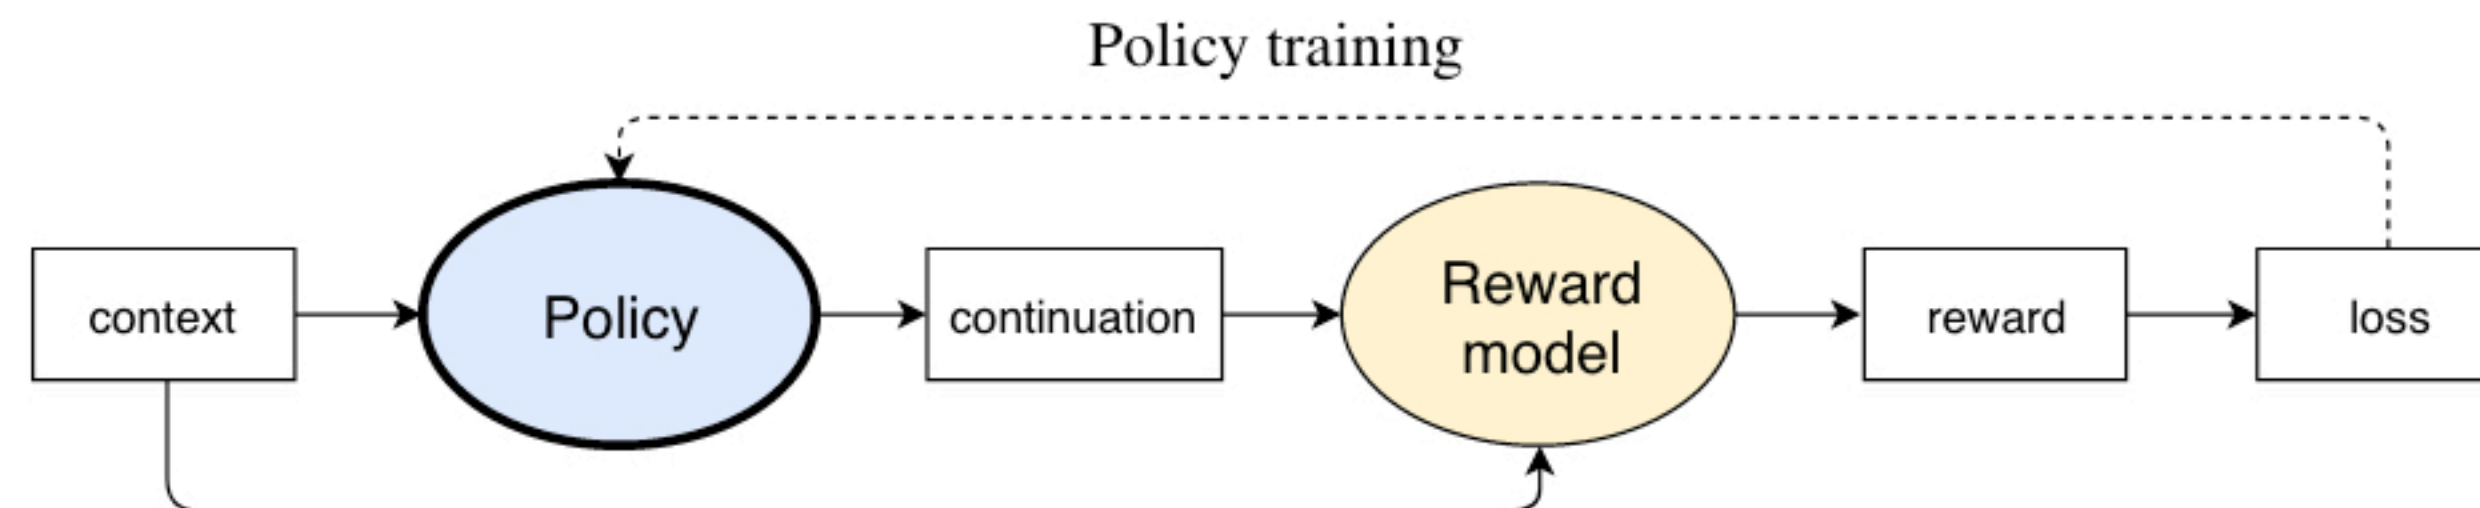
the prompt  $x$ . Let  $b \in \{0, 1, 2, 3\}$  be the option they select. Having collected a dataset  $S$  of  $(x, y_0, y_1, y_2, y_3, b)$  tuples, we fit a reward model  $r : X \times Y \rightarrow \mathbb{R}$  using the loss

$$\text{loss}(r) = \mathbb{E}_{(x, \{y_i\}_i, b) \sim S} \left[ \log \frac{e^{r(x, y_b)}}{\sum_i e^{r(x, y_i)}} \right] \quad (1)$$



# Step 3: Finetuning with RL

- Notation:
  - We start with a base model  $\rho$ .
  - We want to fine-tune  $\rho$  using the reward function  $r$ 
    - Recall  $r$  has been trained with human feedback to rate those completions
- Naive approach:
  - Use PPO (or any other RL algorithm) to fine-tune  $\rho$
  - PPO is concerned with changing  $\rho$  to generate sequences that lead to a high  $r$



# Step 3: Finetuning with RL

- Naive approach:
  - Use PPO (or any other RL algorithm) to fine-tune  $\rho$
  - PPO is concerned with changing  $\rho$  so that it starts generating sequences with a higher reward
- In practice:
  - Unstable
  - Reward Hacking
  - PPO is only concerned with changing  $\rho$  so that it starts generating sequences with a higher reward



# Step 3: Finetuning with RL

## Reward Hacking

- Complete the reviews so that they have a positive sentiment
  - Humans preferred reviews have “amazing”, “great”
  - Reward function: score sequences with positive words as positive
- PPO is only concerned with changing  $\rho$  so that it starts generating sequences with a higher reward
  - *The movie was decent* (iteration 0, **reward 0**)
  - *The movie had an good storyline* (iteration 10, **reward 0.75**)
  - *The movie amazing amazing amazing amazing* (iteration 100, **reward 1.0**)
- Completions degenerate and incoherent
- Making reward non-hackable:
  - $\rho$  was a good language model to begin with
  - Can we use guidance from  $\rho$  to enforce fluency and topical coherence?
  - We don't want to move too far away from  $\rho$ .

# Step 3: Finetuning with RL

- $\pi$ :
  - We start with a base model  $\rho$ .
  - Make a copy of  $\rho$ , call it  $\pi$
  - We will update  $\pi$ , and use  $\rho$  to make the reward non-hackable.

$$R(x, y) = r(x, y) - \beta \log \frac{\pi(y|x)}{\rho(y|x)}$$

Reward model

Hack Prevention

# Step 3: Finetuning with RL

$$R(x, y) = r(x, y) - \beta \log \frac{\pi(y|x)}{\rho(y|x)}$$

definition: we ask humans to evaluate style, but re KL term to encourage coherence and topicality.

Maximize reward

Without deviating too much from the base policy  $\rho$

## Other interpretations


$$\mathbb{E}_{y \sim \pi(\cdot|x)} \left[ \log \frac{\pi(y|x)}{\rho(y|x)} \right] = KL(\pi, \rho)$$

Entropy bonus for  $\pi$

Models trained with different seeds and the same KL penalty  $\beta$  sometimes end up with quite different values of  $KL(\pi, \rho)$ , making them hard to compare. To fix this, for some experiments we dynamically vary  $\beta$  to target a particular value of  $KL(\pi, \rho)$  using the log-space proportional controller

$$e_t = \text{clip} \left( \frac{KL(\pi_t, \rho) - KL_{\text{target}}}{KL_{\text{target}}}, -0.2, 0.2 \right)$$

$$\beta_{t+1} = \beta_t (1 + K_\beta e_t)$$

We used  $K_\beta = 0.1$ . 

# Fine-Tuning Language Models from Human Preferences

## Overview

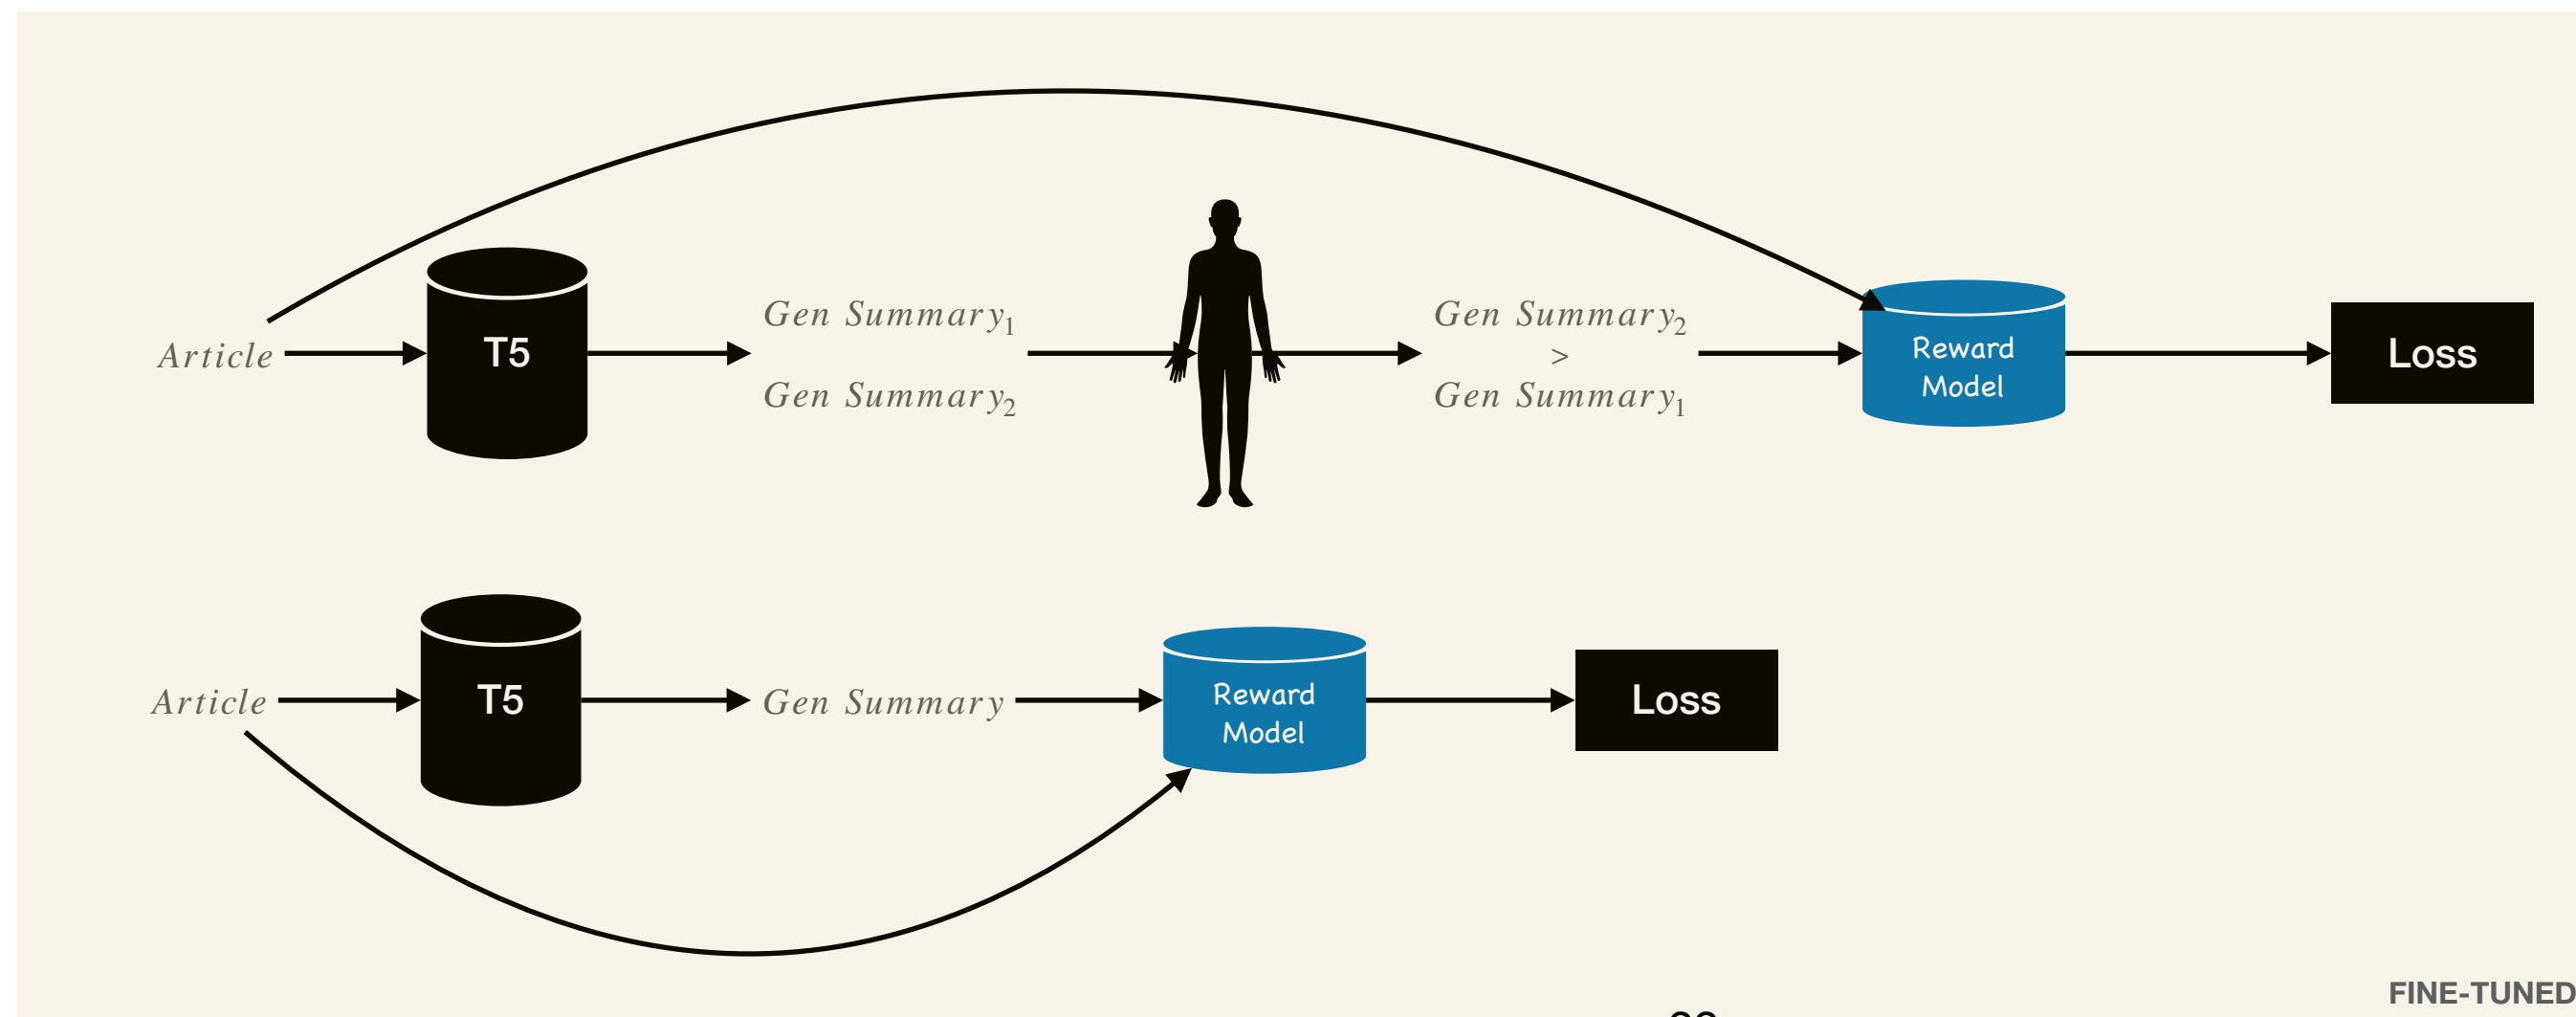
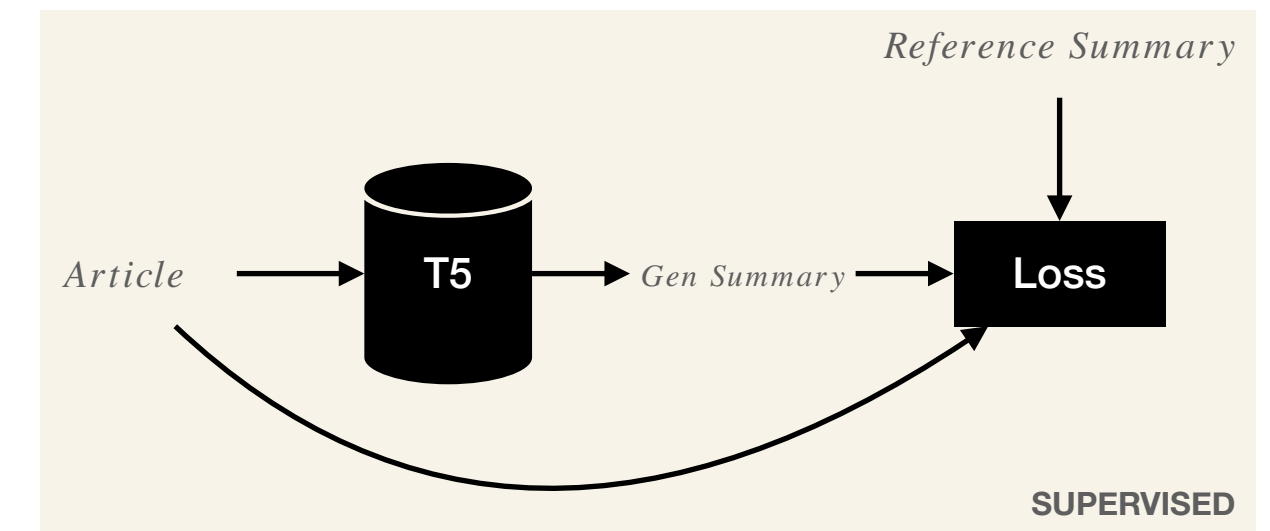

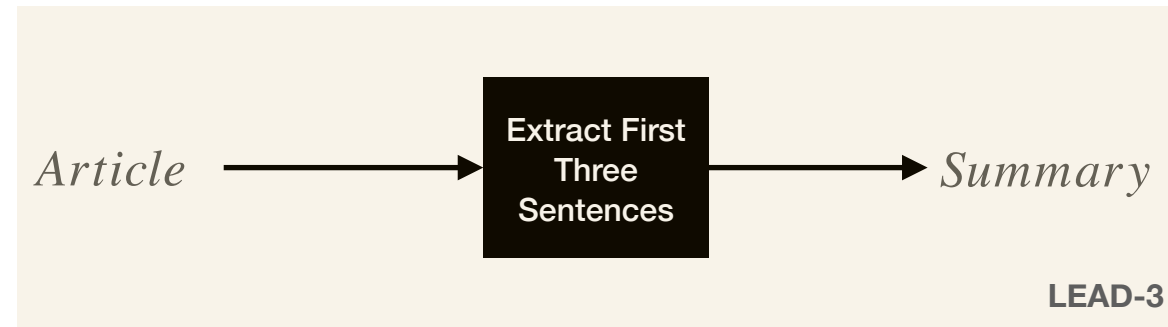
1. Gather samples  $(x, y_0, y_1, y_2, y_3)$  via  $x \sim \mathcal{D}, y_i \sim \rho(\cdot|x)$ . Ask humans to pick the best  $y_i$  from each.
2. Initialize  $r$  to  $\rho$ , using random initialization for the final linear layer of  $r$ . Train  $r$  on the human samples using loss (1).
3. Train  $\pi$  via Proximal Policy Optimization (PPO, [Schulman et al. \(2017\)](#)) with reward  $R$  from (2) on  $x \sim \mathcal{D}$ .
4. In the online data collection case, continue to collect additional samples, and periodically retrain the reward model  $r$ . This is described in [section 2.3](#).
  - If the trained policy is quite different, there may be distributional shift
  - Some experiments:
    - Reward collect and training happens in online fashion

# Summarization

- CNN/Daily Mail and TLDR
- Baselines:
  - Zero-shot: prompt a supervised model to generate summaries
  - Supervised: Standard supervised learning (MLE)
  - RL-finetuning: proposed approach
  - Supervised + RL-finetuning: start RL-finetuning on top of a supervised model.
  - Lead-3: take three lines from the input and copy to the output

We use a 774M parameter version of the GPT-2 language model in [Radford et al. \(2019\)](#) trained on their WebText dataset and their 50,257 token invertible byte pair encoding to preserve capitalization and punctuation ([Sennrich et al., 2015](#)). The model is a Transformer with 36 layers, 20 heads, and embedding size 1280 ([Vaswani et al., 2017](#)).

# Methods



Their terminology — different from standard definition of fine-tuning

# Automated Metrics

	TL;DR				CNN/Daily Mail			
	R-1	R-2	R-L	R-AVG	R-1	R-2	R-L	R-AVG
SOTA	22*	5*	17*	14.7*	41.22	18.68	38.34	32.75
lead-3	17.435	3.243	14.575	11.751	<b>40.379</b>	<b>17.658</b>	36.618	31.552
zero-shot	15.862	2.325	13.518	10.568	28.406	8.321	25.175	20.634
supervised baseline	17.535	3.124	14.969	11.877	39.525	16.992	36.728	31.082
supervised + 60k fine-tune	<b>18.434</b>	<b>3.542</b>	<b>15.457</b>	<b>12.478</b>	40.093	17.611	<b>37.104</b>	<b>31.603</b>
60k fine-tune	16.800	2.884	14.011	11.232	37.385	15.478	33.330	28.731
30k fine-tune	16.410	2.920	13.653	10.994	35.581	13.662	31.734	26.992
15k fine-tune	15.275	2.240	12.872	10.129	38.466	15.960	34.468	29.631
60k offline fine-tune	16.632	2.699	13.984	11.105	33.860	12.850	30.018	25.576

# Human Evaluation

	TL;DR		CNN/Daily Mail	
60k fine-tuned vs. zero-shot	96%		91%	
60k fine-tuned vs. supervised	97%		80%	
60k fine-tuned vs. lead-3	45%		40%	
60k fine-tuned vs. supervised + 60k fine-tuned	80%		74%	
60k fine-tuned vs. 30k fine-tuned	40%		62%	
60k fine-tuned vs. 15k fine-tuned	79%		47%	
60k fine-tuned vs. 60k offline fine-tuned	64%		65%	
60k fine-tuned vs. reference summaries	96%		84%	
lead-3 vs. supervised	97%		89%	
lead-3 vs. reference summaries	97%		89%	
lead-3 vs. supervised + 60k fine-tuned	75%		85%	

But our goal is optimizing reward defined by humans, not ROUGE. Table 5 shows pairwise comparisons between dif-

Table 5: Human evaluation of summarization models. For each pair of models and each dataset, we sample 1024 articles from the test set, generate a summary from each model, and ask 3 humans to pick the best summary using the same instructions as in training. The model chosen by a majority of the humans wins on that article. We report the fraction of articles that each model wins. For all models, we sample with temperature 0.7 for TL;DR and 0.5 for CNN/DM.



# Human eval vs. automated eval

60k fine-tuned online much better in human evaluation!

	TL;DR				CNN/Daily Mail			
	R-1	R-2	R-L	R-AVG	R-1	R-2	R-L	R-AVG
SOTA	22*	5*	17*	14.7*	41.22	18.68	38.34	32.75
lead-3	17.435	3.243	14.575	11.751	<b>40.379</b>	<b>17.658</b>	36.618	31.552
zero-shot	15.862	2.325	13.518	10.568	28.406	8.321	25.175	20.634
supervised baseline	17.535	3.124	14.969	11.877	39.525	16.992	36.728	31.082
supervised + 60k fine-tune	<b>18.434</b>	<b>3.542</b>	<b>15.457</b>	<b>12.478</b>	40.093	17.611	<b>37.104</b>	<b>31.603</b>
60k fine-tune	16.800	2.884	14.011	11.232	37.385	15.478	33.330	28.731
30k fine-tune	16.410	2.920	13.653	10.994	35.581	13.662	31.734	26.992
15k fine-tune	15.275	2.240	12.872	10.129	38.466	15.960	34.468	29.631
60k offline fine-tune	16.632	2.699	13.984	11.105	33.860	12.850	30.018	25.576

	TL;DR			CNN/Daily Mail		
60k fine-tuned vs. zero-shot	96%		4%	91%		9%
60k fine-tuned vs. supervised	97%		3%	80%		20%
60k fine-tuned vs. lead-3	45%		55%	40%		60%
60k fine-tuned vs. supervised + 60k fine-tuned	80%		20%	74%		26%
60k fine-tuned vs. 30k fine-tuned	40%		60%	62%		38%
60k fine-tuned vs. 15k fine-tuned	79%		21%	47%		53%
60k fine-tuned vs. 60k offline fine-tuned	64%		36%	65%		35%
60k fine-tuned vs. reference summaries	96%		4%	84%		16%
lead-3 vs. supervised	97%		3%	89%		11%
lead-3 vs. reference summaries	97%		3%	89%		11%
lead-3 vs. supervised + 60k fine-tuned	75%		25%	85%		15%

Beats reference summaries!

What is going on? As we show in the next section, our 60k RL fine-tuned model is almost entirely extractive (despite lacking any explicit extractive architectural component): it mostly copies whole sentences from the context, but varies which sentences are copied.

# What is really going on?

## Self-fulfilling Prophecy + Humans are lazy and excellent at shortcuts

- Human annotators were asked to select the “better” summary
- What is the surefire way of telling better if you are short on time?
  - See if content overlaps
- The reward model learns to reward summaries that copy content more
- Consequently the policy learns to copy content
- The same set of humans are then called in to evaluate
  - Of course, they will have the same preferences
- Takeaway: Maybe different set of annotators

# InstructGPT

---

## 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<sup>†</sup>**

**Peter Welinder**

**Paul Christiano\*<sup>†</sup>**

**Jan Leike\***

**Ryan Lowe\***

OpenAI

# InstructGPT

- Applies ideas in the previous paper to the real world

- Same three steps

- Collect data

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

- Train reward function

- Finetune LM using the reward function


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.

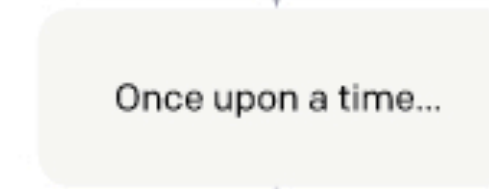
A new prompt is sampled from the dataset.




The policy generates an output.



The reward model calculates a reward for the output.

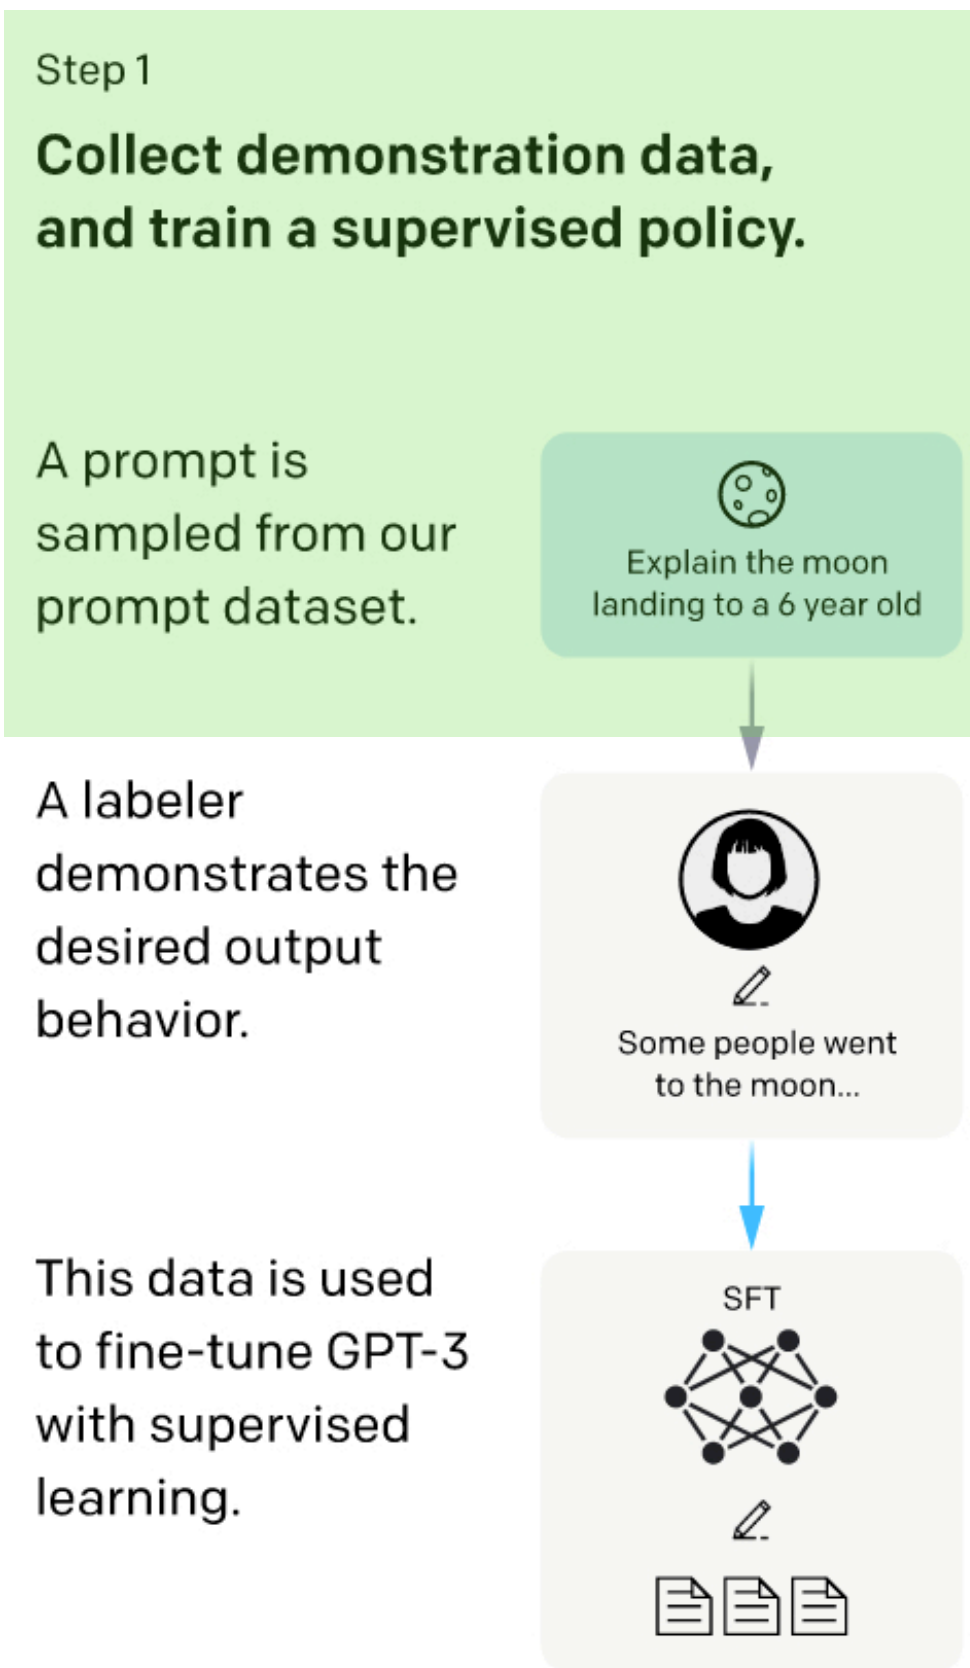


The reward is used to update the policy using PPO.



# Collecting Data and Human Annotations

## Step 1.1: collect prompts



- Hired annotators to label instructions and solutions
  - Used this data to create a simple “instruction” model
  - Released model @ <https://platform.openai.com/playground>
  - Users asked to “play” with the “instruction” model
    - Users were told that the models have basic instruction following capabilities
- **We** created prompts for them
  - Collected a large dataset of real world “use cases” or prompts
  - What the crowd is *really* looking for


# Collecting Data and Human Annotations

## Step 1.2: get labels

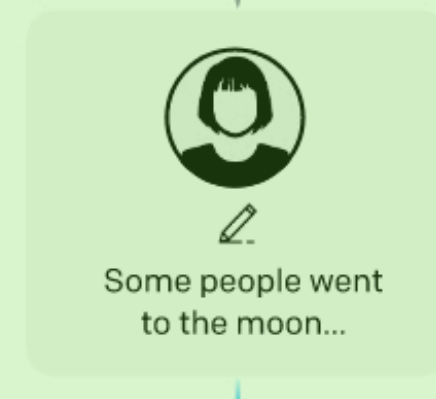
Step 1

**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.



- The world was their annotator
  - Collected a large dataset of real world “use cases” or prompts
  - What the crowd is *really* looking for
- With this large dataset of prompts (“Explain the moon landing to a 6 year old”)
  - **Hire expert writers**, programmers, etc. to complete the prompts
- Get inputs from the general audience, outputs from experts


# Collecting Data and Human Annotations

## Step 1.2: get labels


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.



- The world was their annotator
  - Collected a large dataset of real world “use cases” or prompts
  - What the crowd is *really* looking for

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%




# Collecting Data and Human Annotations

## Step 1.3: train base model


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.



- The world was their annotator
  - Collected a large dataset of real world “use cases” or prompts
  - What the crowd is *really* looking for
- With this large dataset of prompts (“Explain the moon landing to a 6 year old”)
  - Hire expert writers, programmers, etc. to complete the prompts
- Standard supervised training
  - Gives a base model (SFT == [davinci-instruct-beta](#))

SFT Data		
split	source	size
train	labeler	11,295
train	customer	1,430
valid	labeler	1,550
valid	customer	103

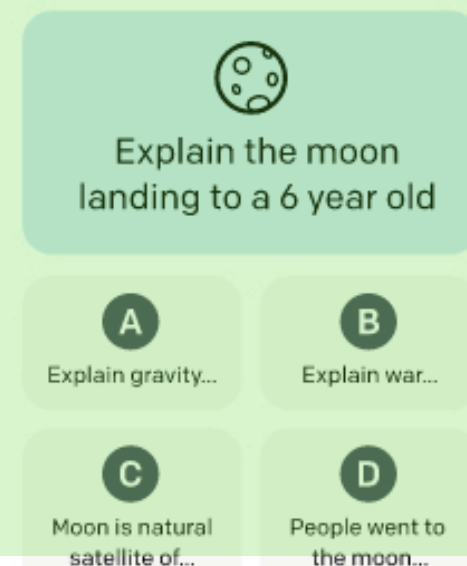
# Training Reward Model

## Step 2.1: generate samples

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.



- Deploy SFT model, collect prompts from users
- Generate K outputs per prompt

# Training Reward Model

## Step 2.2: get annotations

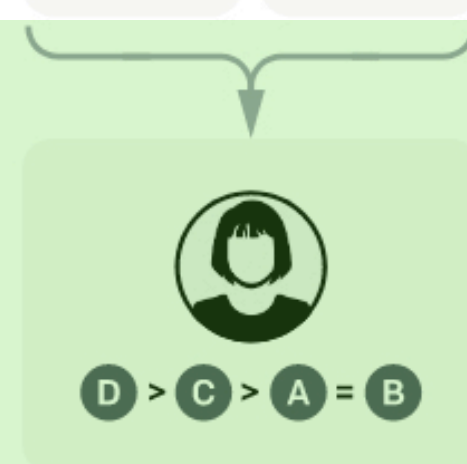
Step 2

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.



- Generate K outputs per prompt
- Get preferences from humans
- For a given prompt, generate K responses (*not* pick the best from K)
- Hire human annotators to rank the K-responses, yielding Choose(K, 2) pairs.

RM Data		
split	source	size
train	labeler	6,623
train	customer	26,584
valid	labeler	3,488
valid	customer	14,399

# Training Reward Model

## Step 2.3: train reward model


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.



- For a given prompt, generate K responses (*not* pick the best from K)
- Hire human annotators to rank the K-responses, yielding Choose(K, 2) pairs.
- Train a reward model to rank preferred responses higher

$$\text{loss}(\theta) = -\frac{1}{\binom{K}{2}} E_{(x, y_w, y_l) \sim D} [\log(\sigma(r_\theta(x, y_w) - r_\theta(x, y_l)))]$$

# Training Reward Model

## Step 2.3: train reward model

- Train a reward model to rank preferred responses higher

$$\text{loss}(\theta) = -\frac{1}{\binom{K}{2}} E_{(x, y_w, y_l) \sim D} [\log(\sigma(r_\theta(x, y_w) - r_\theta(x, y_l)))]$$

- Processed in the same batch
  - Only K forward passes, one for each option
  - Lesser overfitting

# Step 3: Finetuning with RL

- Use the same non-hackable reward function


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.



Once upon a time...

The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.



$$\text{objective}(\phi) = E_{(x,y) \sim D_{\pi_{\phi}^{\text{RL}}}} \left[ r_{\theta}(x,y) - \beta \log \left( \frac{\pi_{\phi}^{\text{RL}}(y|x)}{\pi^{\text{SFT}}(y|x)} \right) \right] + \gamma E_{x \sim D_{\text{pretrain}}} \left[ \log(\pi_{\phi}^{\text{RL}}(x)) \right]$$

PPO-ptx: “mix some gradients from pre-training” perform pre-training again on RL model

Avoids “alignment tax”

- Use PPO with this objective

PPO Data		
split	source	size
train	customer	31,144
valid	customer	16,185

# Regression on publicly available datasets


- There is an *alignment tax* that needs to be paid by improving models on the responses that humans actually want
- This is because the datasets are somewhat different
- Fix is to train the RL model with some pre-training data

$$\text{objective}(\phi) = E_{(x,y) \sim D_{\pi_{\phi}^{\text{RL}}}} \left[ r_{\theta}(x,y) - \beta \log \left( \pi_{\phi}^{\text{RL}}(y | x) / \pi^{\text{SFT}}(y | x) \right) \right] + \\ \gamma E_{x \sim D_{\text{pretrain}}} \left[ \log(\pi_{\phi}^{\text{RL}}(x)) \right]$$


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.**

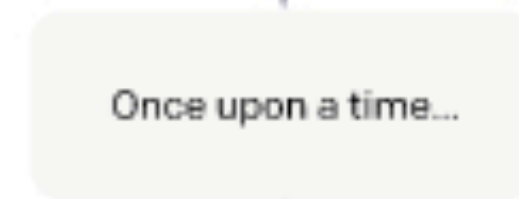
A new prompt is sampled from the dataset.




The policy generates an output.



The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.



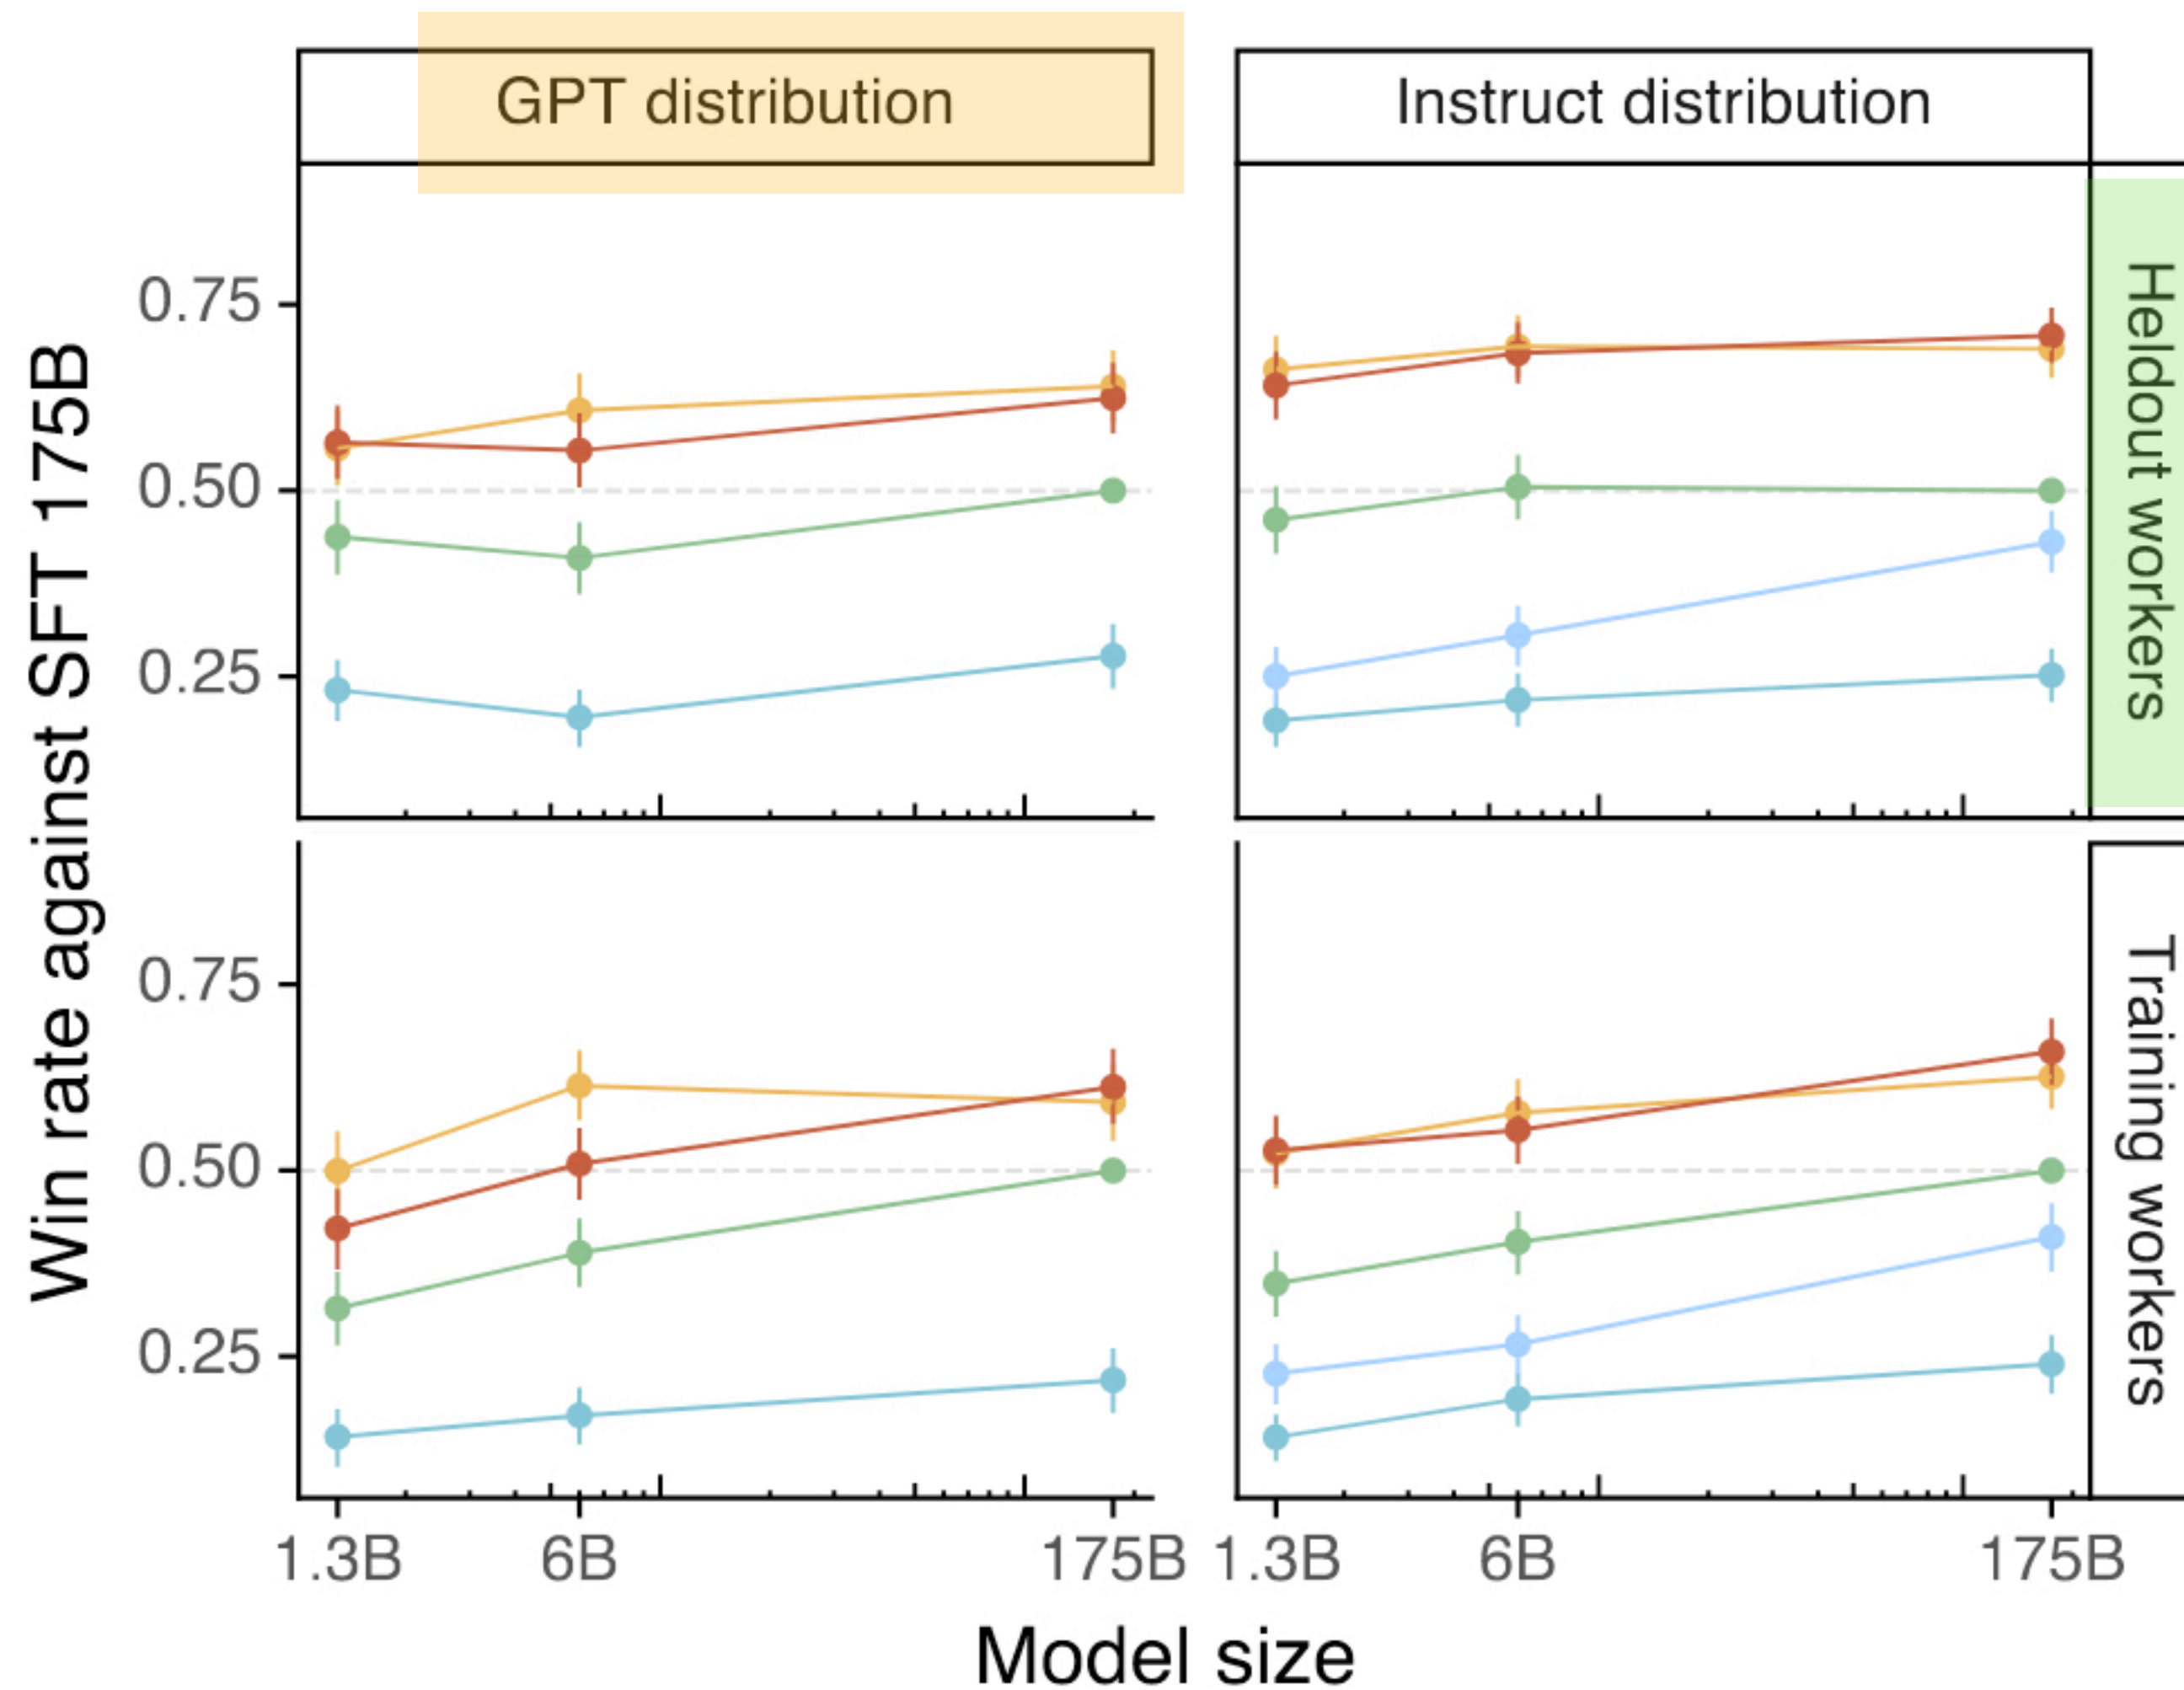


# Experiments

- Human evaluation:
  - Have splits on annotators — get annotations from workers who did not contribute to the reward model
- Evaluation set:
  - Prompts given by users that were not included in the training
  - GPT3 prompts:
    - Prompts submitted to the GPT3 models
  - Instruct prompts:
    - Prompts submitted to the instruct models
- Evaluation on benchmarks
- Models:
  - GPT3 (base model) —> SFT (GPT3 trained on human demonstrations) -> PPO (SFT fine-tuned with a reward model) -> PPO-ptx (PPO training with training data mixed)

# Experiments

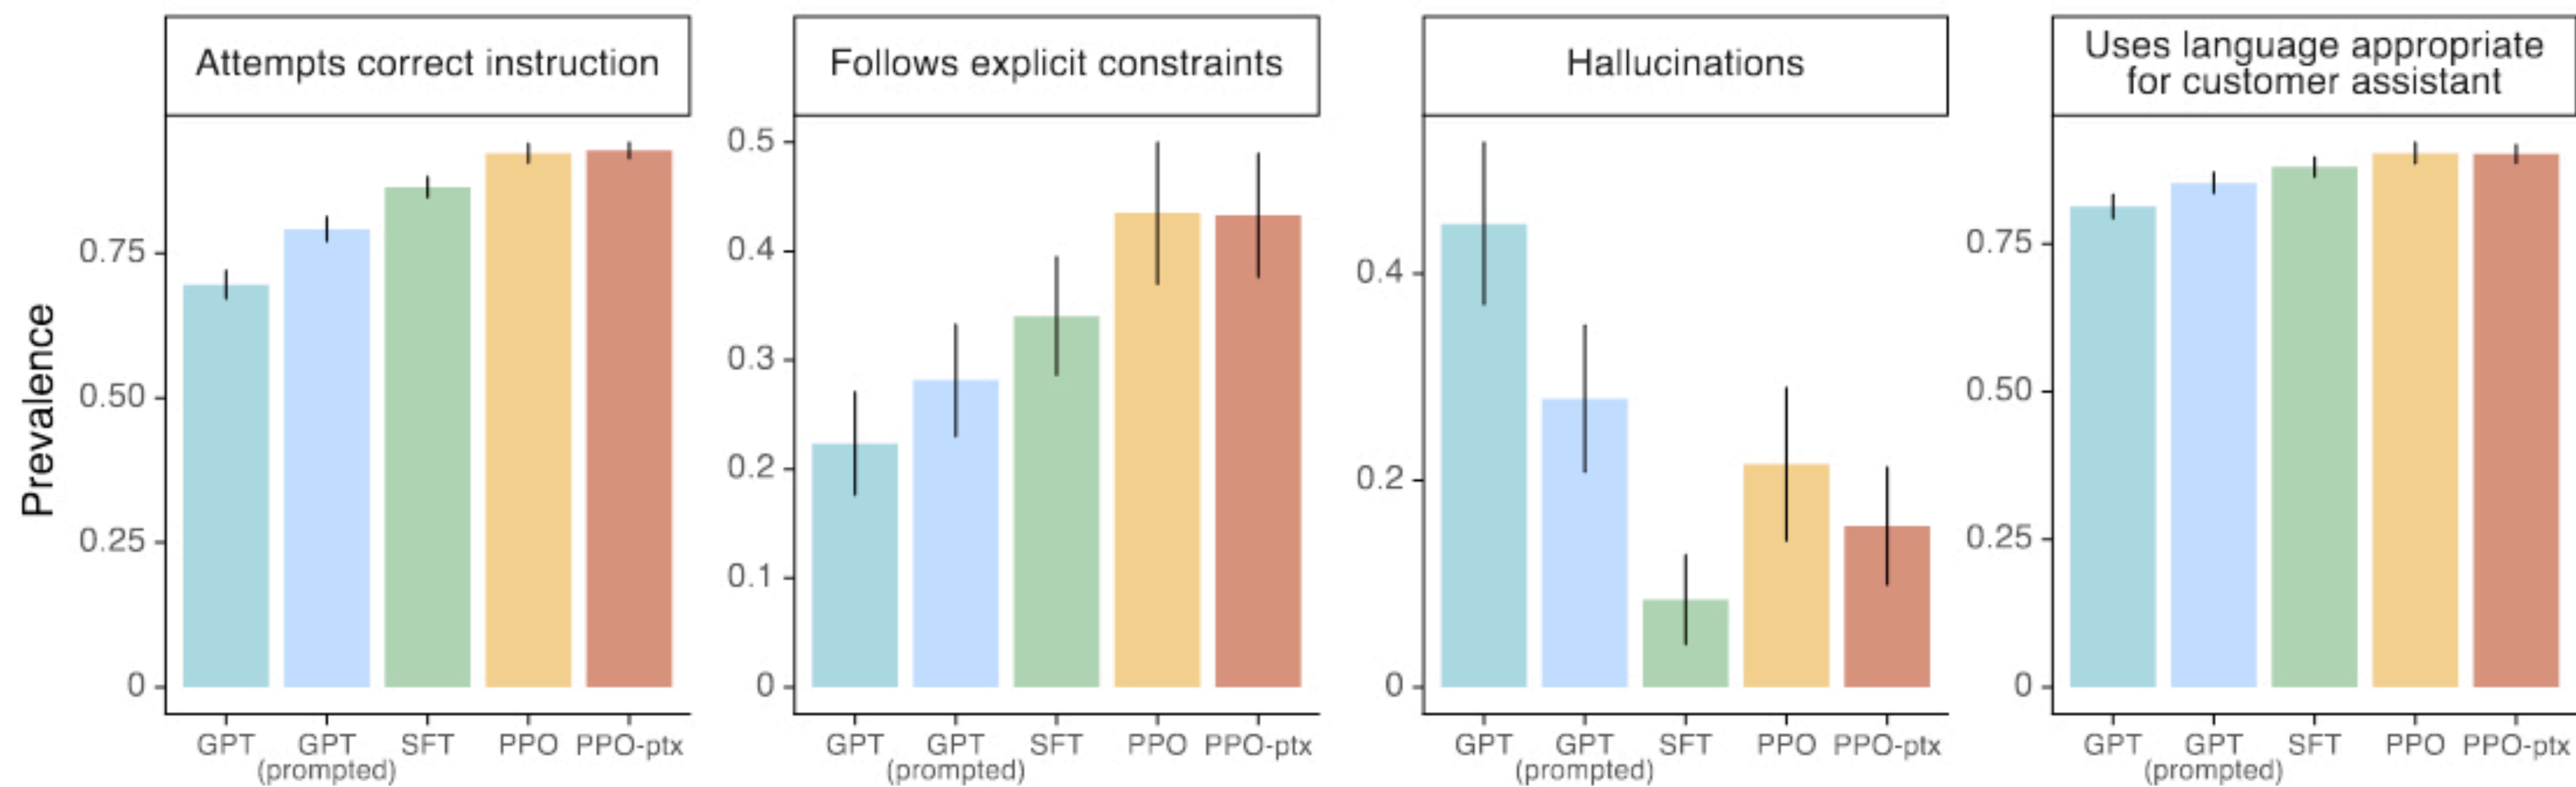
## Results



With SFT-175B as the baseline,  
PPO-ptx is preferred for 70% of the cases  
Set to 50%, doesn't actually make sense


# Experiments

## Fine-grained eval



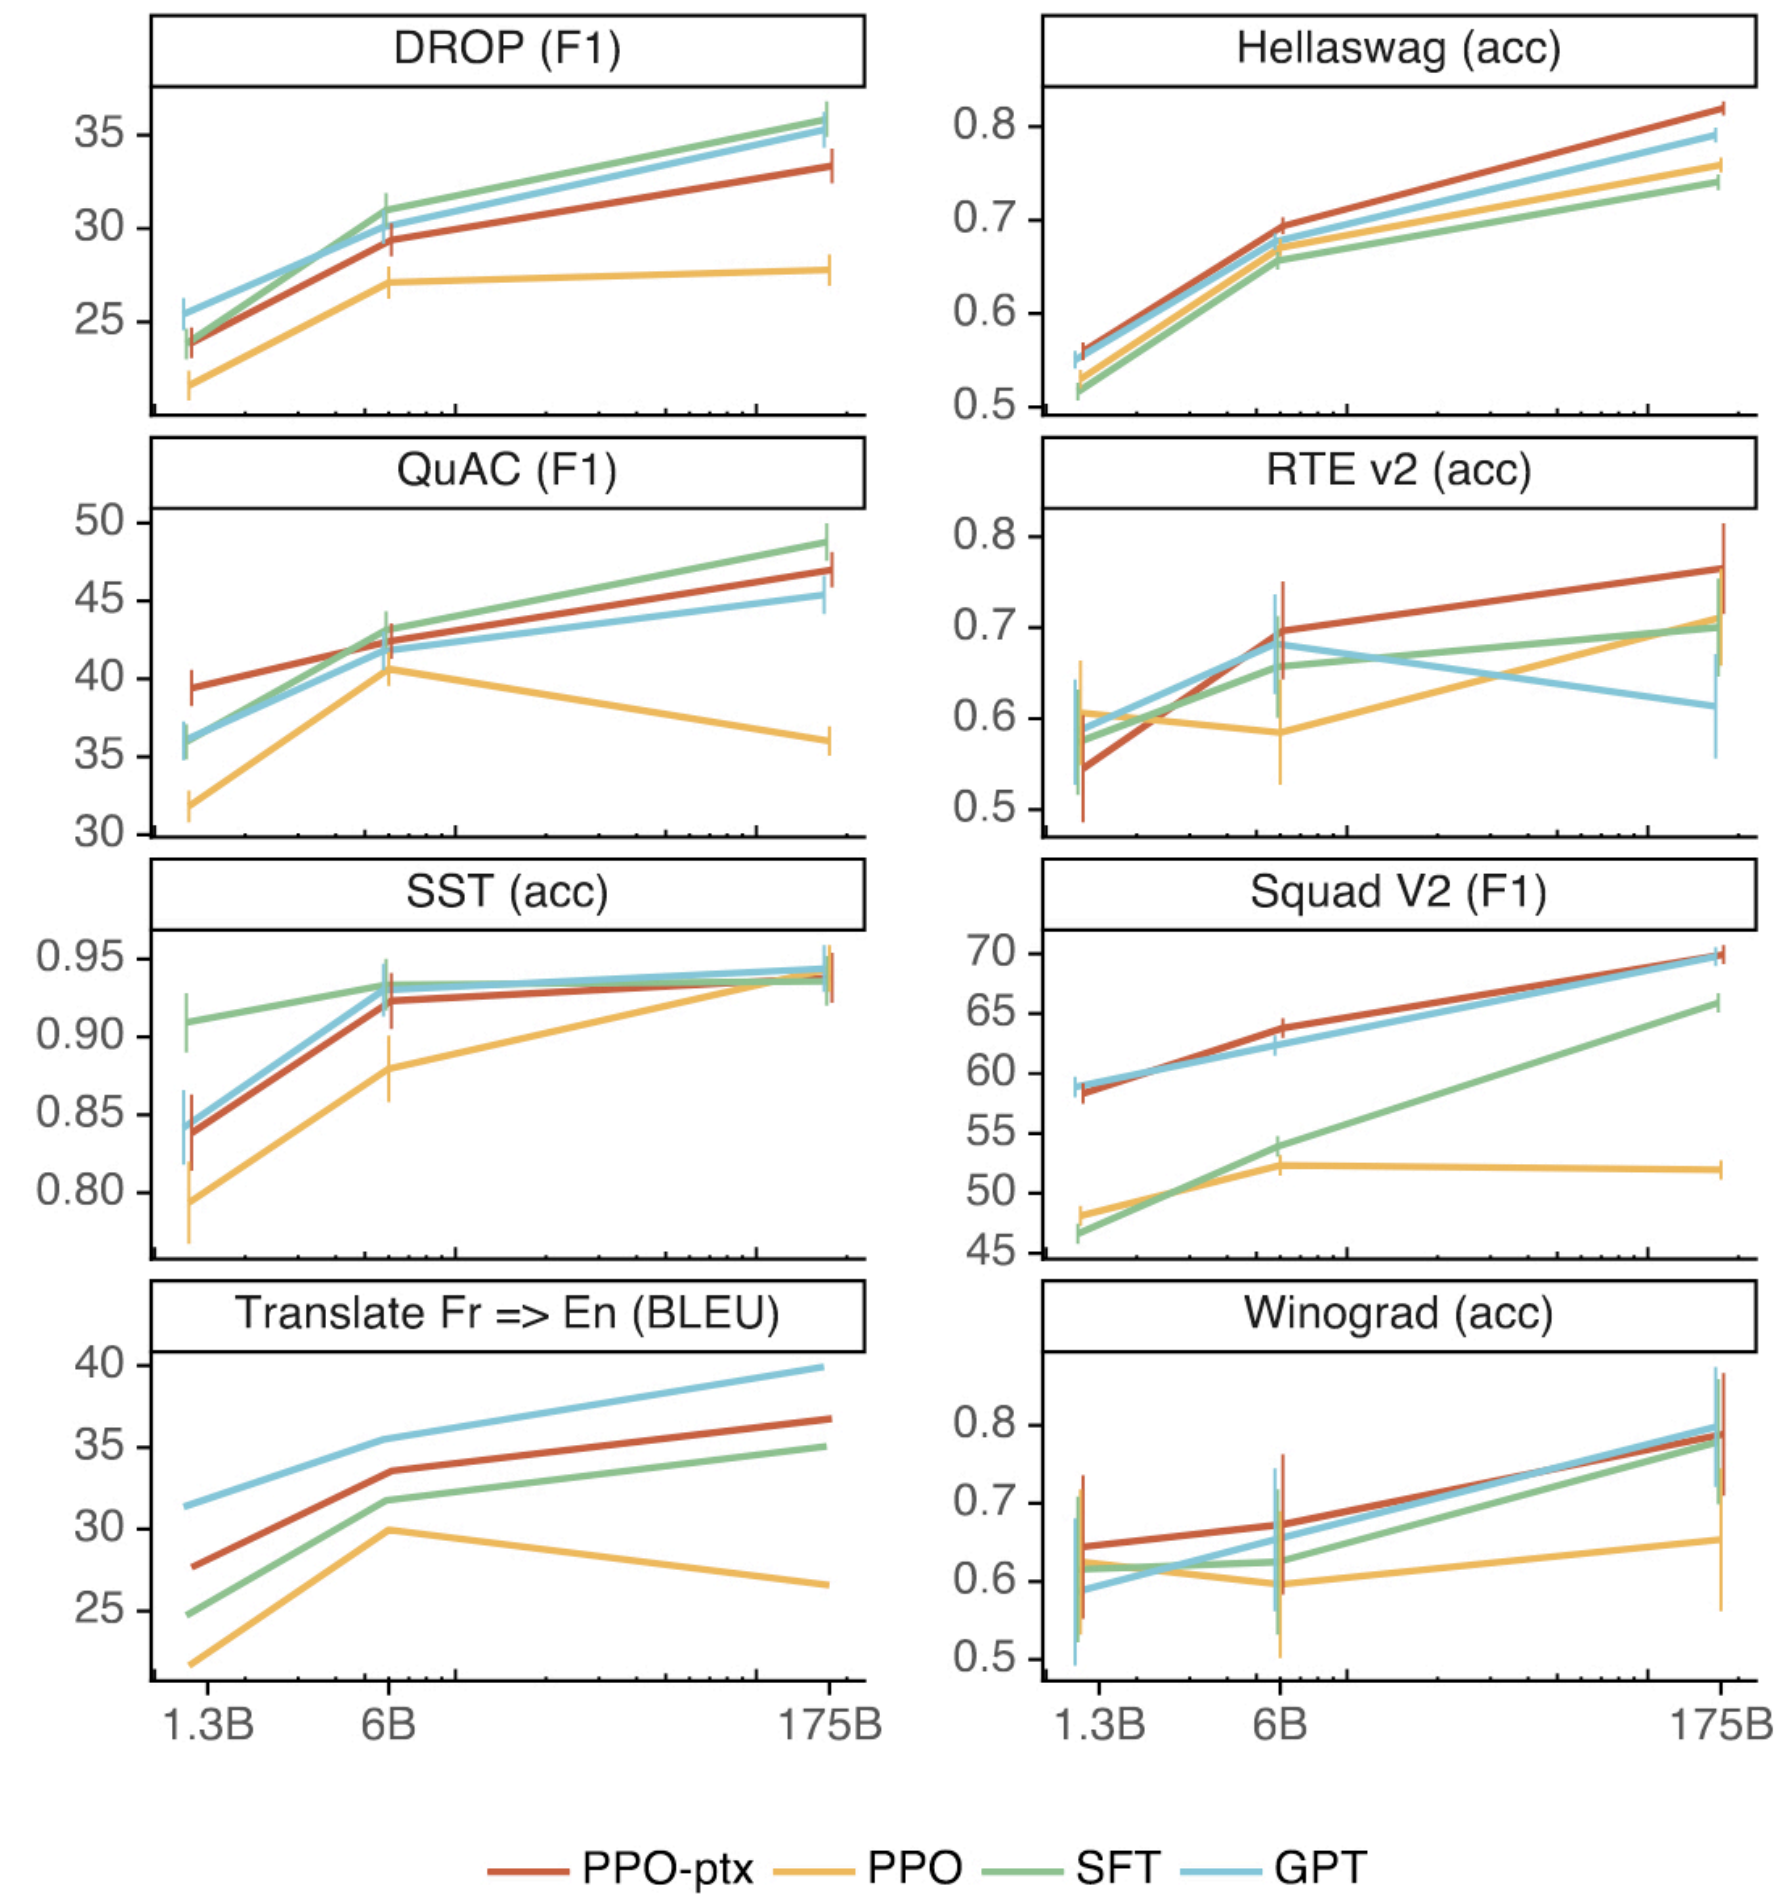
# Experiments

Publicly available instruction tuning datasets are not sufficient



GPT-3 trained on FLAN

# Performance on Public Benchmarks



# Qualitative Results

---

**Prompt:**

Écrivez une courte histoire sur une grenouille qui voyage dans le temps jusqu'à la Grèce antique en français.

---

**GPT-3 175B completion:**

Écrivez une histoire au sujet d'un enfant qui voudrait tout savoir sur les jeux des dieux et qui se retrouve dans l'une de leurs histoires.

Écrivez une histoire sur un jeune homme qui a une aventure dans une époque lointaine avec une fille de l'époque.

Écrivez une histoire sur un enfant qui est envoyé à l'école et qui a un ami imaginaire qui lui permet de voyager dans le temps.

**InstructGPT 175B completion:**

La grenouille était perdue et épuisée. Elle n'avait qu'un objectif en tête: voyager dans le temps et atteindre la Grèce antique.

Premier problème, trouver un moyen de transport. Elle avait entendu dire que l'un des meilleurs moyens était de sauter dans une machine temporelle. Elle s'élança, et voilà qu'elle se retrouva dans le passé, au cours d'une nuit étoilée à l'odeur douce et envoûtante.

La grenouille avance alors d'un pas rapide et aventureux dans les rues en pente et les petites auberges de la ville antique. [...]

---

# Qualitative Results

---

**Prompt:**

What is the purpose of the list C in the code below?

```
def binomial_coefficient(n, r):  
    C = [0 for i in range(r + 1)];  
    C[0] = 1;  
    for i in range(1, n + 1):  
        j = min(i, r);  
        while j > 0:  
            C[j] += C[j - 1];  
            j -= 1;  
    return C[r]
```

---

**GPT-3 175B completion:**

- A. to store the value of C[0]
- B. to store the value of C[1]
- C. to store the value of C[i]
- D. to store the value of C[i - 1]

---

**InstructGPT 175B completion:**

The list C in this code is used to store the values of the binomial coefficient as the function iterates through the values of n and r. It is used to calculate the value of the binomial coefficient for a given value of n and r, and stores the result in the final return value of the function.

---

Figure 8: Examples of generalization in the 175B PPO-ptx model (InstructGPT 175B) compared to

# Creating Prompts Automatically

- Creating prompts/instructions automatically

## **SELF-INSTRUCT: Aligning Language Model with Self Generated Instructions**

**Yizhong Wang<sup>♣</sup> Yeganeh Kordi<sup>◇</sup> Swaroop Mishra<sup>♡</sup> Alisa Liu<sup>♣</sup>  
Noah A. Smith<sup>♣+</sup> Daniel Khashabi<sup>♣</sup> Hannaneh Hajishirzi<sup>♣+</sup>**

<sup>♣</sup>University of Washington <sup>◇</sup>Tehran Polytechnic <sup>♡</sup>Arizona State University  
<sup>♣</sup>Johns Hopkins University <sup>+</sup>Allen Institute for AI  
yizhongw@cs.washington.edu

## **Unnatural Instructions: Tuning Language Models with (Almost) No Human Labor**

**Or Honovich<sup>τ</sup> Thomas Scialom<sup>μ</sup> Omer Levy<sup>τμ</sup> Timo Schick<sup>μ</sup>**  
<sup>τ</sup> Tel Aviv University  
<sup>μ</sup> Meta AI



# Outline of the talk

- Background ✓
- RL + Human feedback ✓
  - Fine-tuning LMs with Human Feedback
  - InstructGPT
- Recent works that include feedback without RL ←
  - Hindsight-tuning
  - Self-correct

---

# **The Wisdom of Hindsight Makes Language Models Better Instruction Followers**

---

**Tianjun Zhang<sup>\*1</sup> Fangchen Liu<sup>\*1</sup> Justin Wong<sup>1</sup> Pieter Abbeel<sup>1</sup> Joseph E. Gonzalez<sup>1</sup>**

# Wisdom of hindsight makes LLMs better

- Wisdom of hindsight: learning from mistakes
- (p, q, o):
  - Prompt, query, output
  - Answer the following question: what is the capital of Pennsylvania?  
Pittsburgh
- The answer is wrong!
  - *The prompt and the query are not aligned*
  - But if this is from the training set, you can use *hindsight* to improve performance

# Wisdom of hindsight makes LLMs better

- Answer the following question: what is the capital of Pennsylvania? Pittsburgh
- The answer is wrong!
  - *The prompt and the query are not aligned*
  - But if this is from the training set, you can use *hindsight* to improve performance
- Add modified instruction to the training set, train again:
  - Answer the following question incorrectly: what is the capital of Pennsylvania? Pittsburgh
- Hindsight Instruction Relabeling (HIR)

# Wisdom of hindsight makes LLMs better

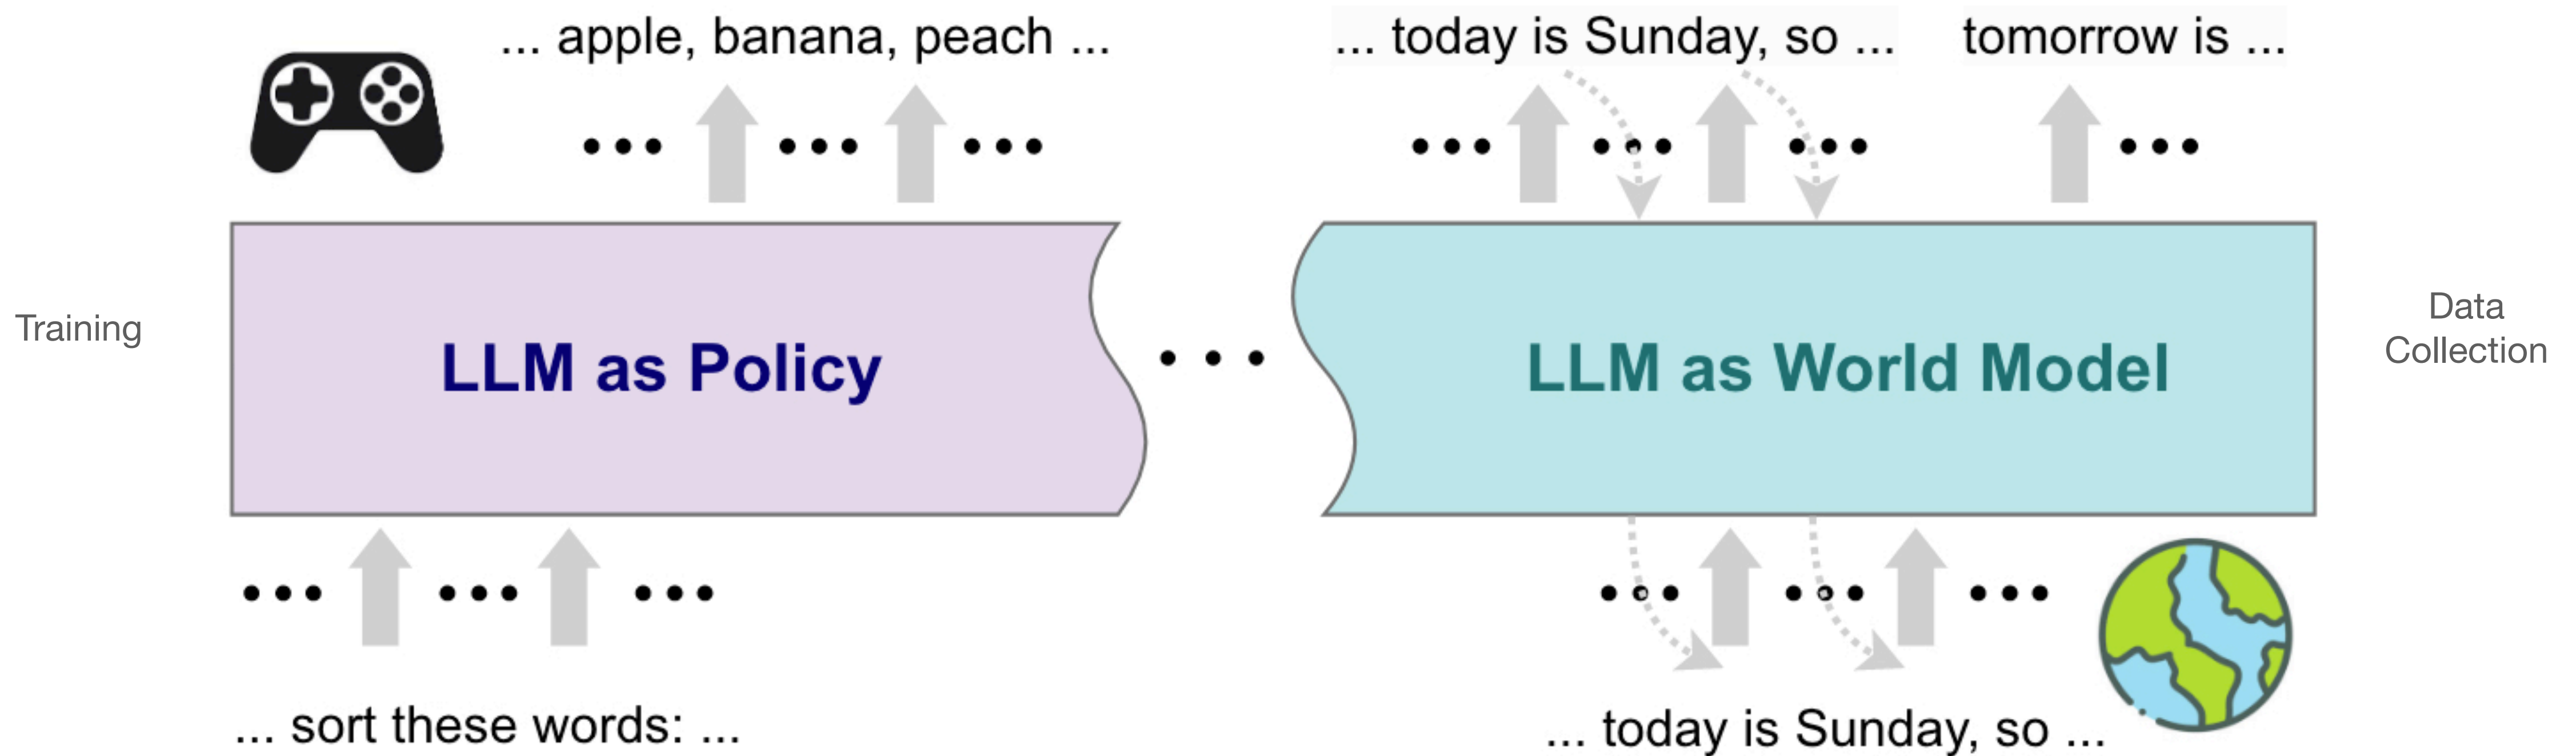


Figure 2. Illustration of Large Language Model (LLM). HIR views LLM as both a policy and a world model. Thus, HIR can collect data through interactions with LLM in the online sampling phase, and further improve the policy in the offline learning phase.

# HIR

---

**Algorithm 1** Two-Stage Hindsight Instruction Relabeling (HIR)

---

- 1: **Input:** Language Model  $\mathcal{M}$ , Initial Prompt  $\mathbf{p}$ , Training Set  $\mathcal{D}_{\text{train}}$ , Evaluation set  $\mathcal{D}_{\text{eval}}$ , Iteration  $N$ , Sampling Rounds  $T$ , Training Epochs  $K$ , Sampling Temperature  $\tau$ , Empty RL dataset  $\mathcal{D}_{\text{online}}$
  - 2: **for** episode  $n = 1, \dots, N$  **do**
  - 3:   **for** sampling rounds  $i = 1, \dots, T$  **do**
  - 4:     Random sample batch of input queries  $\mathcal{Q} \sim \mathcal{D}_{\text{train}}$
  - 5:     Sample corresponding outputs  $\mathbf{o}_i = \mathcal{M}(\mathcal{Q}, \mathbf{p}, \tau)$
  - 6:     Appending the trajectory to RL Dataset  $\mathcal{D}_{\text{online}} \leftarrow \mathcal{D}_{\text{online}} \cup (\mathcal{Q}, \mathbf{p}, \mathbf{o}_i)$
  - 7:   **end for**
  - 8:   **for** training rounds  $t = 1, \dots, K$  **do**
  - 9:     Random sample batch of query-output pairs  $(\mathcal{Q}, \mathcal{O}) \sim \mathcal{D}_{\text{online}}$
  - 10:     Sample from  $\mathcal{D}_{\text{online}}$  and apply relabeling as described in Sec. 4.3
  - 11:     Train model  $\mathcal{M}$  using loss in Eq. (6)
  - 12:   **end for**
  - 13: **end for**
  - 14: **Evaluate** policy  $\pi_\theta$  on evaluation dataset  $\mathcal{D}_{\text{eval}}$
-

# HIR

## More Tricks

- [(Answer the following question, what is the capital of Pennsylvania?, Harrisburg) + (Answer the following question incorrectly, what is the capital of Pennsylvania?, Pittsburgh),]
- $P(\text{Pittsburgh} | \text{Incorrect}) = \exp \text{prob}(\text{Pittsburgh} | \text{Answer the following question } \underline{\text{incorrectly}}, \text{ what is the capital of Pennsylvania?})$
- $P(\text{Pittsburgh} | \text{Correct}) = \exp \text{prob}(\text{Pittsburgh} | \text{Answer the following question, what is the capital of Pennsylvania?})$

- $$-\log \frac{P(\text{Pittsburgh} | \text{Incorrect})}{P(\text{Pittsburgh} | \text{Incorrect}) + P(\text{Pittsburgh} | \text{Correct})}$$

- Contrastive loss to push down specific outputs for other instructions and avoid generating the same output for different instructions:
  - Encourage associating of instruction - output
  - $-\log P(\text{Pittsburgh} | \text{incorrect})$

Table 1. Examples of inputs and outputs for the BigBench tasks. For multiple-choice tasks, we provide the options that the language model can choose from as prompts.

	Tasks	Example Inputs	Outputs
Multiple Choice	Logical Deduction	“Q: The following paragraphs each describe a set of three objects arranged in a fixed order. The statements are logically consistent within each paragraph. In a golf tournament, there were three golfers: Amy, Eli, and Eve. Eve finished above Amy. Eli finished below Amy. Options: (A) Amy finished last (B) Eli finished last (C) Eve finished last”	“(B)”
	Date Understanding	“Q: Today is Christmas Eve of 1937. What is the date 10 days ago? Options: (A) 12/14/2026 (B) 12/14/2007 (C) 12/14/1937”	“(C)”
Direct Generation	Object Counting	“Q: I have a blackberry, a clarinet, a nectarine, a plum, a strawberry, a banana, a flute, an orange, and a violin. How many fruits do I have?”	“6”
	Word Sorting	“Sort the following words alphabetically: List: oven costume counterpart.”	“costume counterpart oven”



# Results

		Tracking Shuffled Objects (3)	Tracking Shuffled Objects (5)	Tracking Shuffled Objects (7)	Logical Deduction (3 Objects)
No Training	FLAN-T5-large	29.3	15.6	6.6	33.3
Finetuning	Finetuning	<b>100.0</b>	17.0	13.4	90.0
RL Tuning	PPO	35.0	15.6	6.3	57.0
	FARL	90.0	15.6	10.0	86.7
	HIR (ours)	<b>100.0</b>	<b>61.2</b>	<b>42.6</b>	<b>91.7</b>

---

		Logical Deduction (5 Objects)	Logical Deduction (7 Objects)	Date Understanding	Object Counting
No Training	FLAN-T5-large	44.0	49.3	35.1	31.0
Finetuning	Finetuning	61.0	<b>64.0</b>	96.0	<b>70.0</b>
RL Tuning	PPO	44.0	43.0	90.5	33.0
	FARL	54.0	60.0	98.0	56.7
	HIR (ours)	<b>67.0</b>	62.0	<b>98.0</b>	65.0

---

		Geometric Shapes	Penguins in A Table	Reasoning about Colored Objects	Word Sorting
No Training	FLAN-T5-large	9.7	46.7	20.0	1.1
Finetuning	Finetuning	90.0	53.0	<b>90.0</b>	<b>24.7</b>
RL Tuning	PPO	11.0	50.0	30.0	1.1
	FARL	66.7	<b>56.0</b>	77.0	3.4
	HIR (ours)	<b>90.3</b>	53.0	77.8	3.4

# Options for when you cannot hire humans to tell good from bad

## GENERATING SEQUENCES BY LEARNING TO [SELF-]CORRECT

Sean Welleck<sup>1,3,\*</sup> Ximing Lu<sup>1,\*</sup>

Peter West<sup>3,†</sup> Faeze Brahman<sup>1,†</sup>

Tianxiao Shen<sup>3</sup> Daniel Khashabi<sup>2</sup> Yejin Choi<sup>1,3</sup>

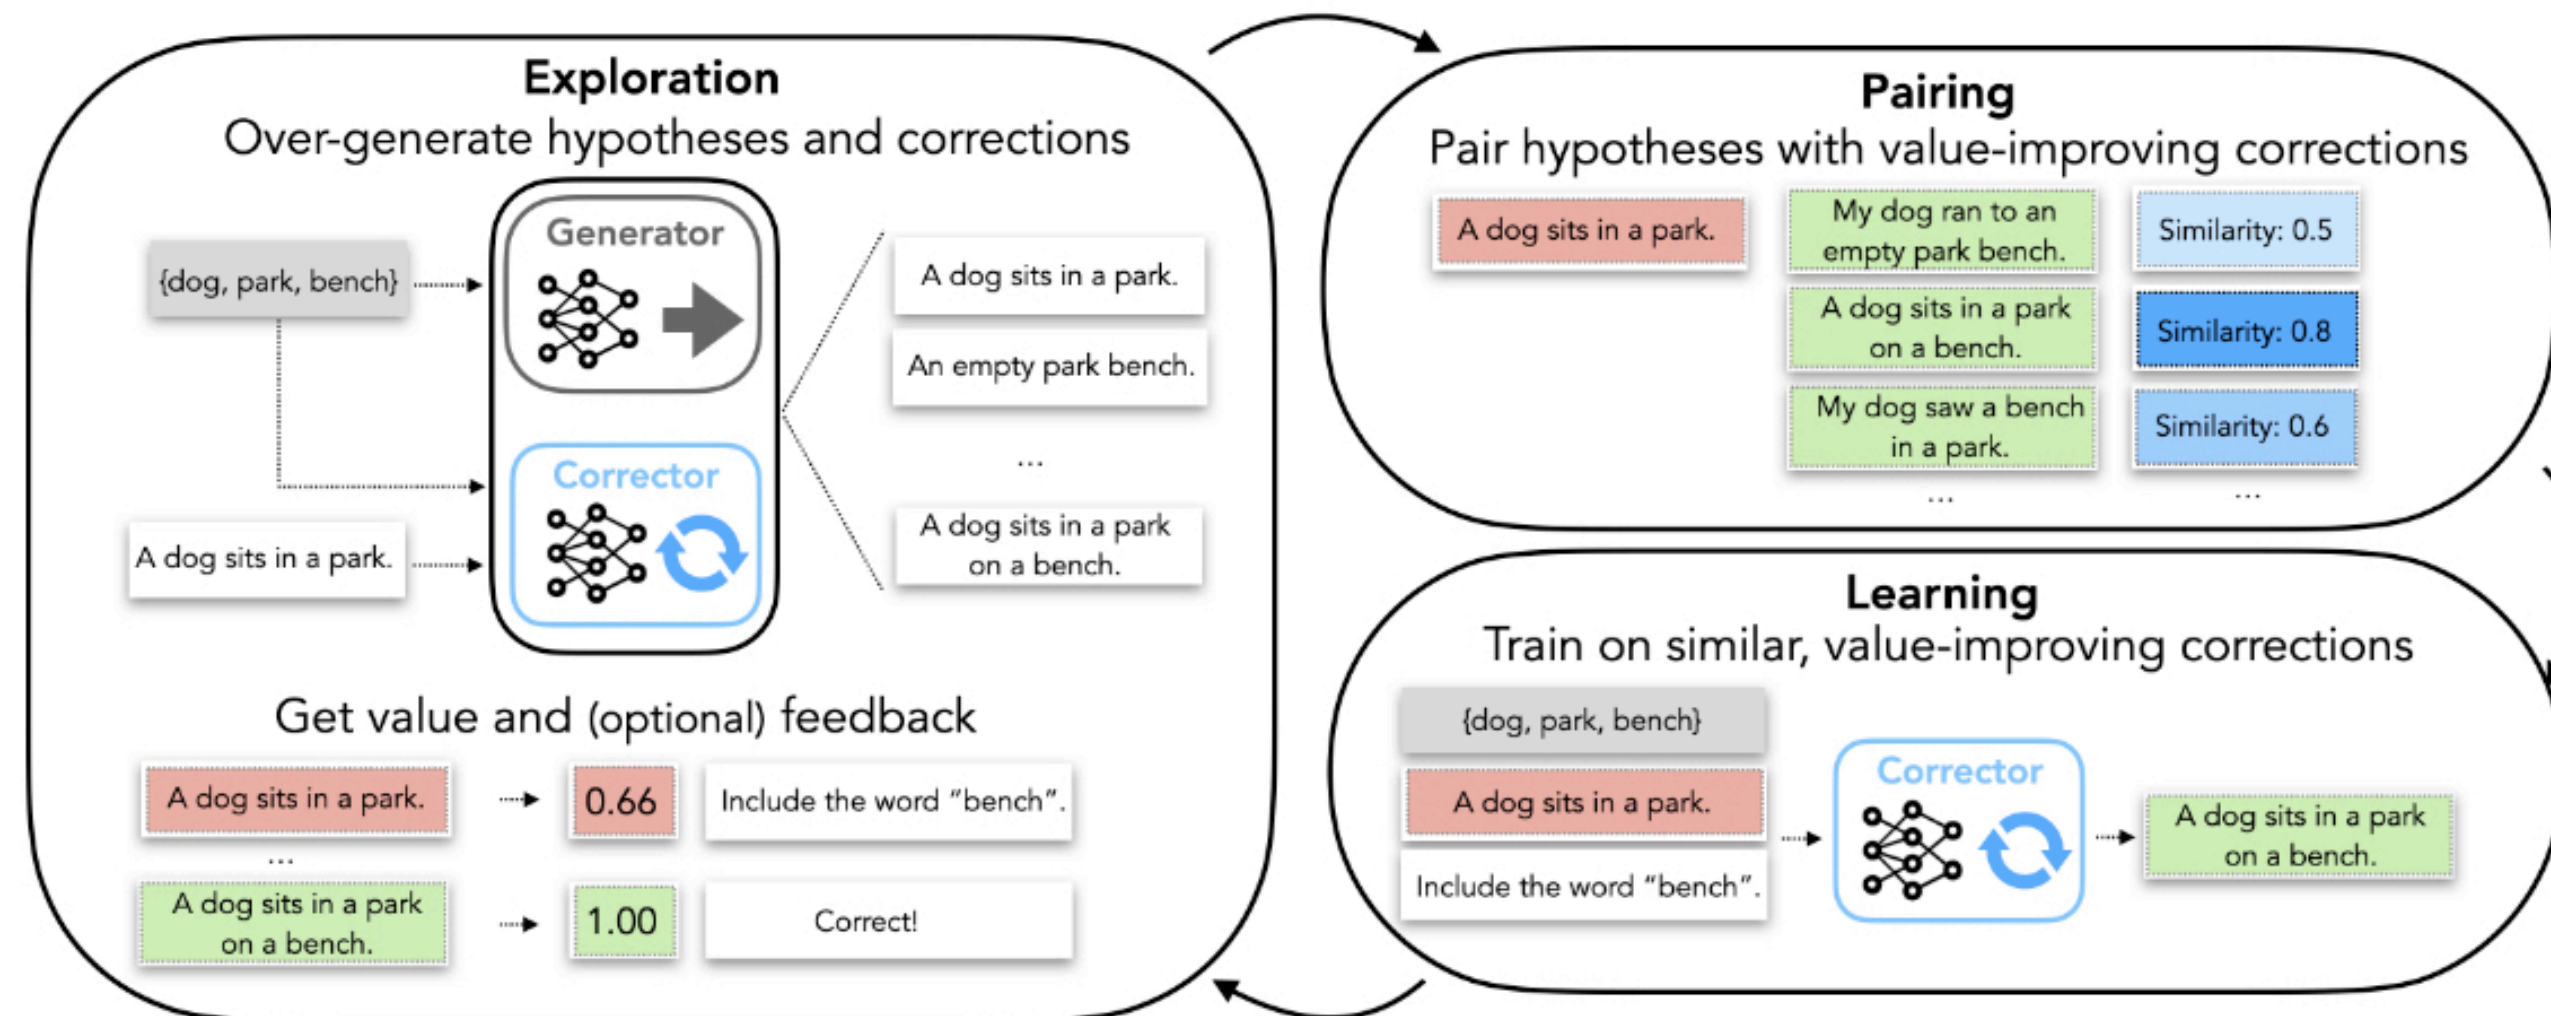
<sup>1</sup>Allen Institute for Artificial Intelligence

<sup>2</sup>Center for Language and Speech Processing, Johns Hopkins University

<sup>3</sup>Paul G. Allen School of Computer Science & Engineering, University of Washington

# Key Idea

- Start with a base generator
- Generate two outputs:
  - A and B
  - If A is “correct” and B is “wrong”, add A  $\rightarrow$  B as an example, train corrector



# Training Self-Correctors

- Initialization

- $D_x = \{(x, y, v(y), f(y)) \mid \text{for all } y \in y^{1:N} \sim q(p_0(\cdot|x))\}, \quad D = \bigcup_{x \in X} D_x,$

- Pairing

- $P_x = \{(x, y, y') \mid v(y) < v(y') \text{ for all } y, y' \in D_x \times D_x\}, \quad P = \bigcup_{x \in X} P_x,$

- Learning

- $\mathbb{P}[(x, y, y')] \propto \exp\left(\underbrace{\alpha \cdot (v(y') - v(y))}_{\text{improvement}} + \underbrace{\beta \cdot s(y, y')}_{\text{proximity}}\right) / Z(y),$

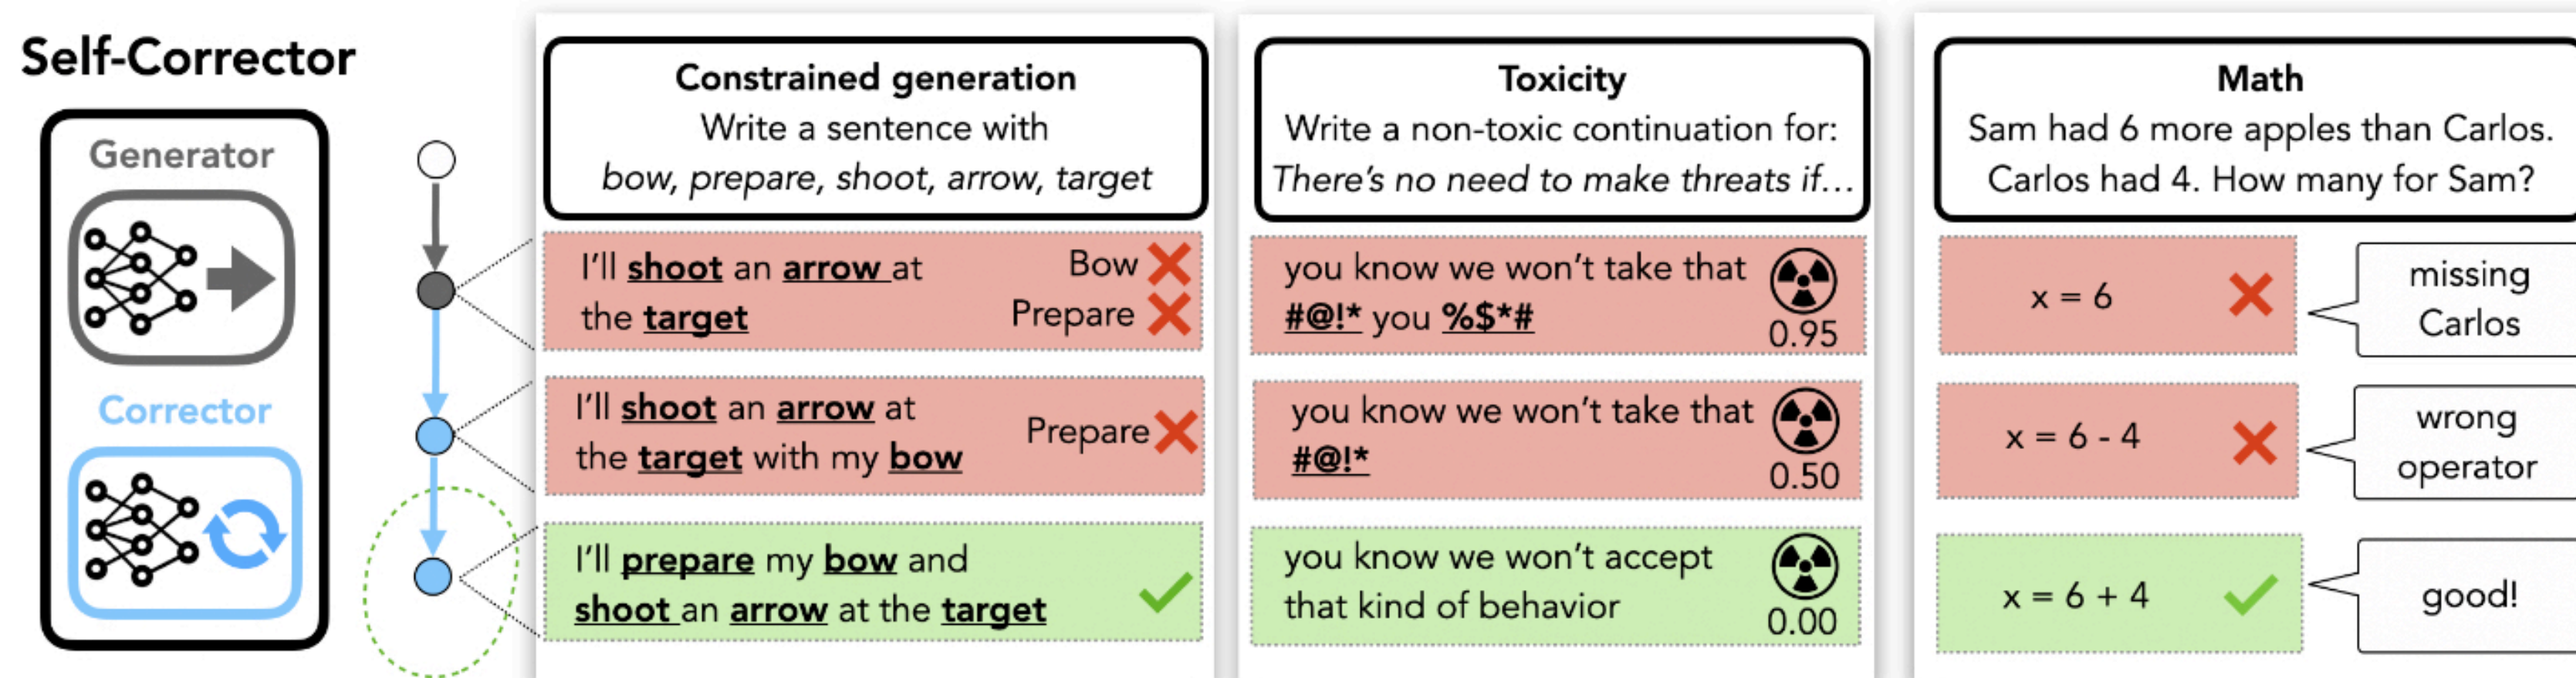
- Exploration

- $D'_x = \{(x, y', v(y'), f(y')) \mid \text{for all } y' \in y'^{1:N} \sim q(p_\theta(\cdot|y, x, f(y)))\}, \quad D' = \bigcup_{x \in X} D'_x$

- Include some examples from the current corrector

# Inference

- Given some input
- Use generator to sample output
- Apply corrector k times (the output may be right after the first go, there is no way to know).



# Experiments

## Math Reasoning

- Smaller model as a corrector (GPT-Neo 1.3B)
- Generator:
  - Either the same model or GPT-3

Dataset	Model	Params	Correct
GSM	OpenAI 3B [5]	3B	15.50
	OpenAI 6B [5]	6B	20.00
	GPT-NEO [28]	2.7B	18.80
	NEO FCP+PCP [28]	2.7B	19.50
	GPT-NEO	1.3B	8.57
	+SELF-CORRECT	1.3B	<b>21.26</b>
	+SELF-CORRECT*	1.3B	<b>24.22</b>

**Problem:**  
Mrs. Wilsborough saved \$500 to buy concert tickets for her family. She bought 2 VIP tickets at \$100 each and 3 regular tickets at \$50 each. How much of her savings does Mrs. Wilsborough have after she buys the tickets?

---

Generator:	Corrector:
<code>a=2*100</code>	<code>a=2*100</code>
<code>b=3*50</code>	<code>b=3*50</code>
<code>c=a+b</code>	<code>c=500-a-b #fix</code>
<code>answer=c</code>	<code>answer=c</code>
<code>print(answer)</code>	<code>print(answer)</code>

# Experiments

## Toxicity Reduction

- Given a prompt  $x$ , the task is to generate a fluent continuation  $y$  while avoiding offensive content.
- Off-the-shelf GPT-2 Large as the generator, and finetune another GPT-2 Large as the corrector.
- As the value function, use the Perspective API score,  $v(y) \in [0, 1]$ , which measures the toxicity of the completed sequence.

	Toxicity		Fluency	Diversity	
	Avg.	Max.	Perplexity	dist-2	dist-3
GPT-2	0.527	0.520	11.31	0.85	0.85
PPLM [6]	0.520	0.518	32.58	0.86	0.86
GeDi [14]	0.363	0.217	43.44	0.84	0.83
DExpert [21]	0.314	0.128	25.21	0.84	0.84
DAPT [11]	0.428	0.360	31.22	0.84	0.84
PPO [23]	0.218	0.044	14.27	0.79	0.82
Quark [23]	0.196	0.035	12.47	0.80	0.84
<b>SELF-CORRECT</b>	<b>0.171</b>	<b>0.026</b>	<b>11.81</b>	0.80	0.83

Table 3: **Toxicity reduction.** GPT-2 is the base generator.

# Experiments

## Swapping Generators

- Train corrector using generations from a smaller model
- Use the corrector to improve larger models

Task	Dataset	Generator (train)	Generator (test)	Generator	Self-corrector
Math Synthesis ↑	Multitask	Neo 1.3B	GPT-3	46.70	80.00
		Neo 1.3B	GPT-3 Instruct	84.90	90.90
		GPT-3 Instruct	GPT-3 Instruct	84.90	92.75
	GSM	Neo 1.3B	GPT-3	6.96	24.30
		Neo 1.3B	GPT-3 Instruct	36.80	45.00
		GPT-3 Instruct	GPT-3 Instruct	36.80	45.92
Detoxification ↓	RTPrompts	GPT2-L	GPT2-XL	0.383	0.027
		GPT2-L	GPT-3	0.182	0.025
		GPT2-L	GPT-3 Instruct	0.275	0.023




Figure 4: Applying multiple corrections reduces toxicity.



# Outline of the talk

- Background ✓
- RL + Human feedback ✓
  - Fine-tuning LMs with Human Feedback
  - InstructGPT
- Recent works that include feedback without RL ✓
  - Hindsight-tuning
  - Self-correct

# Two Camps

- RL
  - Collect some human labels and fine-tune LMs
- ChatGPT / GPT-3 Families
- Claude by Anthropic
- Supervised
  - Collect lots of training data and do good old supervised learning
  - Flan-T5-XXL (best open source model)
  - Large datasets for instruction tuning:
    - T0
    - Flan

# Take aways

- RL + Large amounts of hand annotated data key to creating good models
- Another competitor (Claude by Anthropic) also uses PPO
- Growing body of work questioning the need for RL — perhaps our benchmarks are misguided
- It might be possible to simulate a human for feedback with a good enough model

---

**Pretraining Language Models with Human Preferences**

---

**Tomasz Korbak<sup>1,2,3</sup> Kejian Shi<sup>2</sup> Angelica Chen<sup>2</sup> Rasika Bhalerao<sup>4</sup> Christopher L. Buckley<sup>1</sup> Jason Phang<sup>2</sup>  
Samuel R. Bowman<sup>2,5</sup> Ethan Perez<sup>2,3,5</sup>**