



# CS 639: Foundation Models **Alignment II**

Fred Sala

University of Wisconsin-Madison


**March 19, 2026**



# Announcements

- **Exam:** grades are out
  - Talk to us if anything wrong
- **Homework 3:** due soon!
- **Project information:** here  
[https://pages.cs.wisc.edu/~fredsala/cs639/files/project\\_info\\_639.pdf](https://pages.cs.wisc.edu/~fredsala/cs639/files/project_info_639.pdf)
- **Class outline:**

Thursday March 19	Alignment II
Tuesday March 24	Reasoning I: More CoT
Thursday March 26	Reasoning II: RLVR



# Outline

- **Alignment and RLHF**

- Basic alignment idea, goals, mechanisms, RL review, RLHF steps

- **Why Does RLHF Work?**

- Intuition: failures of supervised learning, knowledge-seeking interactions, abstains

- **RLHF Problems and Alternatives**

- Open problems, direct preference optimization (DPO)

# Outline

- **Alignment and RLHF**

- Basic alignment idea, goals, mechanisms, RL review, RLHF steps

- **Why Does RLHF Work?**

- Intuition: failures of supervised learning, knowledge-seeking interactions, abstains

- **RLHF Problems and Alternatives**

- Open problems, direct preference optimization (DPO)


# Alignment: **Basic Motivation**

Goal: produce language model outputs that **users like better...**

- **Challenge: Hard** to specify exactly what this means
- We can do fine-tuning/instruction-tuning to make model more likely to produce certain outputs
  - But we don't necessarily know that these are "preferred" outputs
  - Plus, the model may not produce desirable outputs more generally

**Instead, let's use feedback.**

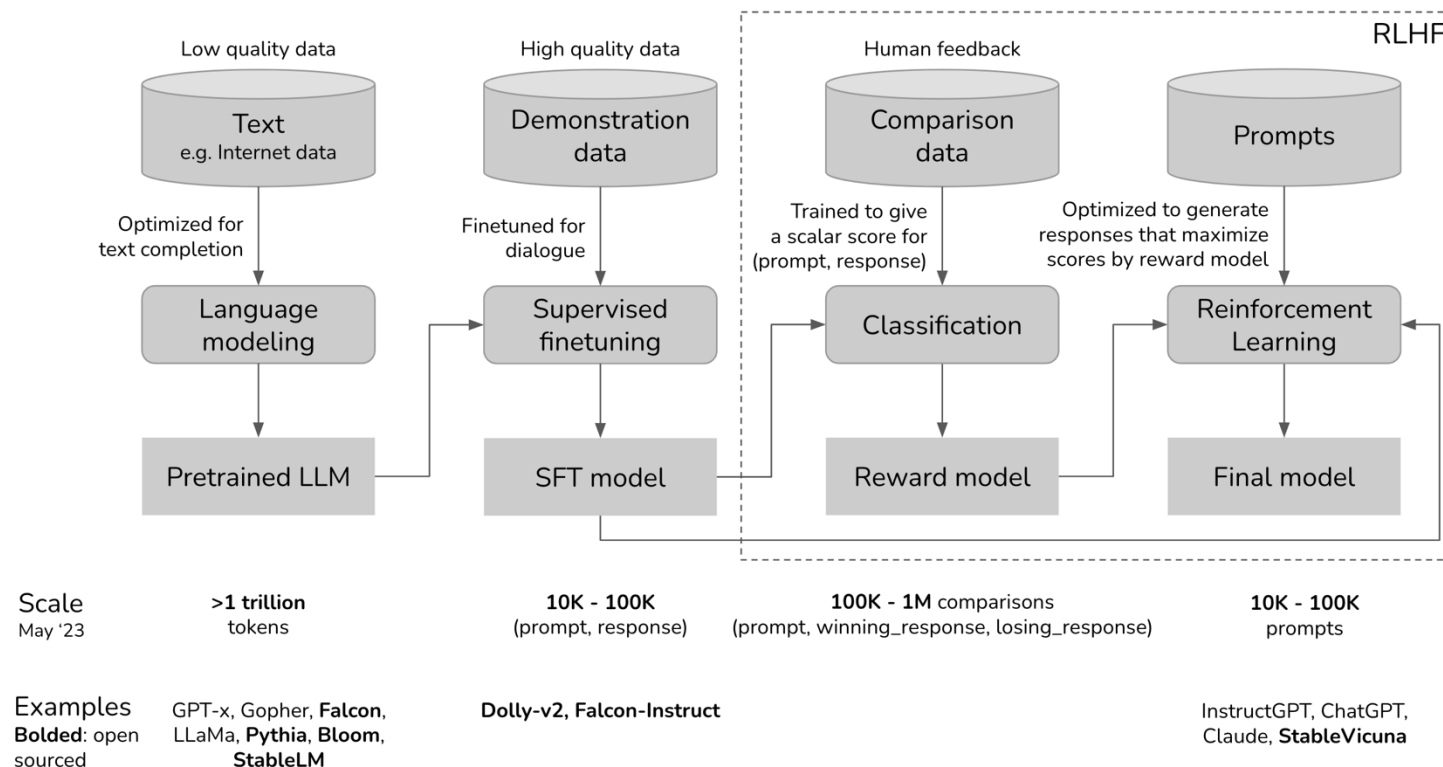
- **Why? Easy** to query users
- Not the only way to do alignment, but we'll focus on it---will lead to alignment via RLHF



# RLHF: Setup

Goal: produce language model outputs that users like better...


- Will use **Reinforcement Learning from Human Feedback**
- Via RL with trained reward model (Ouyang et al '22)



# RLHF: Feedback

First stage: get **human feedback** to train reward model

- Fix a set of prompts
- Produce multiple outputs for each prompt
  - Can get them from the original model post-SFT, or otherwise
- Ask human users **which is better**
  - **Binary output**
  - Can do more, but don't have to
    - Rank more questions, get feedback, etc.



# RLHF: Reward/Preference Model

Second stage: train reward model

- Use the human feedback to train/fine-tune another model to reproduce the preferences

- **“Preference” model**

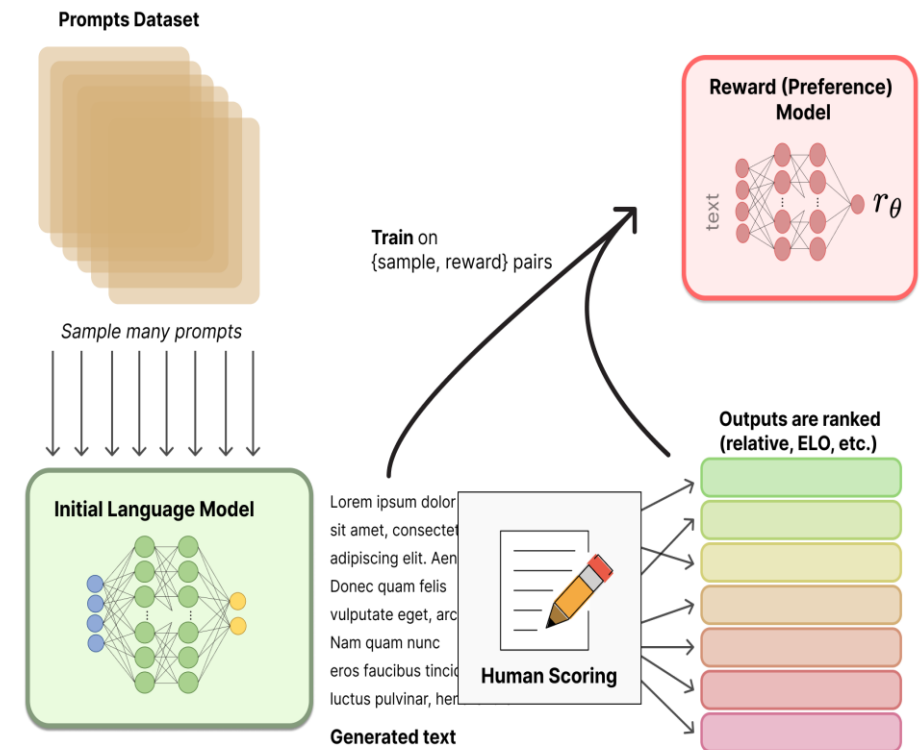
- **Why?** Reward model can tell us, in general, how “liked” any input is.

- Use to tell how good outputs are

- **Note: not generative!**

- Input: token sequence

- Output: scalar reward value



<https://huggingface.co/blog/rlhf>

# RLHF: Reward/Preference Model

Second stage: train reward model

- Use the human feedback to train/fine-tune another model to reproduce the preferences
- Mismatch between training data and what model produces

- Training data: pairs  $(x, y_l), (x, y_w)$



Prompt and  
**Preferred**  
**(winner)**  
response

Prompt and  
**Dispreferred**  
**(loser)**  
response

- But the reward model produces just **one** value:  $r(x, y)$ .
    - So must somehow form an objective that makes  $r(x, y)$  consistent with the winning and losing pairs  $\rightarrow$  softmax!

# RLHF: Reward/Preference Model

Second stage: train reward model

- Mismatch between training data and what model produces
  - Training data: pairs  $(x, y_l), (x, y_w)$
  - But the reward model produces just **one** value:  $r(x, y)$ .
  - So must somehow form an objective that makes  $r(x, y)$  consistent with the winning and losing pairs  $\rightarrow$  softmax!
  - Bradley-Terry preference model (famous in social choice theory)

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

- Now we can relate winning and losing responses to scalar rewards
- And can learn a reward model maximally consistent with our data

# RLHF: Reward/Preference Model

Second stage: train reward model

- Bradley-Terry preference model (famous in social choice theory)
  - Now we can relate winning and losing responses to scalar rewards
  - And can learn a reward model maximally consistent with our data

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

- Actual training objective: log likelihood over our dataset,

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

# RLHF: Reward/Preference Model

Note: we don't have to always do this from scratch

- **Pretrained** reward models available (like in our FMs in general)
- Benchmarks for this:




▲	Model	▲	Model Type	▲	Score	▲	Chat	▲	Chat Hard	▲
1	<a href="#">nvidia/Llama-3.1-Nemotron-70B-Reward</a>		Custom Classifier		94.1		97.5		85.7	
2	<a href="#">Skywork/Skywork-Reward-Gemma-2-27B</a>		Seq. Classifier		93.8		95.8		91.4	
3	<a href="#">SF-Foundation/TextEval-Llama3.1-70B</a>		Generative		93.5		94.1		90.1	
4	<a href="#">meta-metrics/MetaMetrics-RM-v1.0</a>		Custom Classifier		93.4		98.3		86.4	
5	<a href="#">Skywork/Skywork-Critic-Llama-3.1-70B</a>		Generative		93.3		96.6		87.9	
6	<a href="#">LxzGordon/URM-LLaMa-3.1-8B</a>		Seq. Classifier		92.9		95.5		88.2	
7	<a href="#">Salesforce/SFR-LLaMa-3.1-70B-Judge-r</a> *		Generative		92.7		96.9		84.8	
8	<a href="#">Skywork/Skywork-Reward-Llama-3.1-8B</a>		Seq. Classifier		92.5		95.8		87.3	

Lambert et al '24

# RLHF: Fine-Tuning with RL

Third stage: RL

- Use an RL algorithm
- **Goal:** produce outputs that have high reward




RL formulation (see overview coming up for RL reminder!)

- **Action space:** all the tokens possible to output
- **State space:** all the sequences of tokens
- **Reward function:** the trained reward model
- **Policy:** the new version of the LM, taking in state and returning tokens

# Reinforcement Learning Review

We have an **agent interacting** with the **world**




- Agent receives a reward based on state of the world
  - **Goal:** maximize reward / utility **(\$\$\$)**

# RL Review: Theoretical Model

## Basic setup:

- Set of states,  $S$
- Set of actions  $A$
- Information: at time  $t$ , observe state  $s_t \in S$ . Get reward  $r_t$
- Agent makes choice  $a_t \in A$ . State changes to  $s_{t+1}$ , continue



Goal: find a map from **states to actions** maximize rewards.

↑  
A “policy”

# RL Review: Markov Decision Process (MDP)

The formal mathematical model:

- **State set**  $S$ . Initial state  $s_0$ . **Action set**  $A$
- **State transition model:**  $P(s_{t+1} | s_t, a_t)$ 
  - Markov assumption: transition probability only depends on  $s_t$  and  $a_t$ , and not previous actions or states.
- **Reward function:**  $r(s_t)$
- **Policy:**  $\pi(s) : S \rightarrow A$  action to take at a particular state.



$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} \dots$$

# RLHF: RL Approach

What approach for RL stage?

- Many deep RL methods available
- Policy gradient methods
- Popular: PPO (Proximal Policy Optimization)
  - Main difference from vanilla policy gradient, you constrain change to policy at each step (Schulman et al)

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





**Break & Questions**

# Outline

- **Alignment and RLHF**

- Basic alignment idea, goals, mechanisms, RL review, RLHF steps

- **Why Does RLHF Work?**

- Intuition: failures of supervised learning, knowledge-seeking interactions, abstains


- **RLHF Problems and Alternatives**

- Open problems, direct preference optimization (DPO)

# Why RLHF?

## Why should we do this?

- Why does supervised fine-tuning by itself not give our goal results?
- Many hypotheses; this section inspired by Yoav Goldberg's blog:
  - <https://gist.github.com/yoavg/6bff0feccd65950898eba1bb321cfbd81>
  - Itself based on Schulman's talk
  - [https://www.youtube.com/watch?v=hiLw5Q\\_UFg](https://www.youtube.com/watch?v=hiLw5Q_UFg)



# Why RLHF? Ways To Interact

Three “modes of interaction”:

- **text-grounded**: provide the model with text, instruction (“what are the chemical names mentioned in this text”),
- **knowledge-seeking**: provide the model with question or instruction, and expect a (truthful) answer based on the model's internal knowledge
- **creative**: provide the model with question or instruction, expect some creative output. (“Write a story about...”)

# Why RLHF? Knowledge-seeking

Three “modes of interaction”:

- **knowledge-seeking**: provide the model with question or instruction, and expect a (truthful) answer based on the model's internal knowledge
- This is hypothesized to require RL. Why does **SL fail**?
  - Case 1: know the answer: fine.
  - Case 2: don't know the answer. Supervised learning forces memorization, cannot produce “don't know”.
  - Worse, SL on case 2 encourages **model to lie**...



# Why RLHF? Knowledge-seeking with RL

Three “modes of interaction”:

- **knowledge-seeking**: provide the model with question or instruction, and expect a (truthful) answer based on the model's internal knowledge
- Why does RL succeed?
  - Case 1: know the answer: fine. Get a reward
  - Case 2: don't know the answer. Sometimes make it up and get a reward if lucky, most of the time low reward
  - **Encourages truth telling.**

# Why RLHF? **Abstains**

Additionally, **we'd like our model to abstain**

- SL will really struggle with this
  - Usually no abstains in datasets
  - Even if there were, “generalization” here means abstaining on similar questions? Difficult
- RL still challenging, need to produce high reward for “don't know”, but specific to model
- One way to craft a reward function:
  - High reward: correct answers
  - Medium reward: abstain
  - Negative reward: incorrect



# Outline

- **Alignment and RLHF**

- Basic alignment idea, goals, mechanisms, RL review, RLHF steps

- **Why Does RLHF Work?**

- Intuition: failures of supervised learning, knowledge-seeking interactions, abstains

- **RLHF Problems and Alternatives**

- Open problems, direct preference optimization (DPO)



# Break & Questions

# RLHF Problems

Lots of challenges!

- **Casper et al**, “Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback”
- Challenges everywhere, all three phases:
  - In human feedback,
  - In obtaining reward model,
  - In obtaining the policy



# RLHF Problems: **Human Feedback**

- Need to obtain some kind of “representative” collection of feedback providers
- **Simpler:**
  - Some people have biases
  - Mistakes due to lack of care (standard in crowdsourcing)
  - Adversarial data poisoners
- **Harder:**
  - In tough settings, what is “good” output?
  - Possible to manipulate humans



# RLHF Problems: Human Feedback

- Additionally, **need high-quality data.**
- Expensive to hand-craft good prompts to drive feedback
- Feedback quality:
  - Tradeoffs in feedback levels
  - Ideally, rich
  - But harder to work with to train reward



# RLHF Problems: **Reward Model**

- Values can be difficult to express as a reward function
- May need to combine multiple reward functions:
  - What's a “universal” one? People are different
- Reward Hacking
  - In tough settings, what is “good” output?
  - Possible to manipulate humans

# RLHF Alternatives

- **Direct preference optimization (DPO)**
  - Bypass separate trained reward model: just use preference information **directly** (Rafailov et al, '23)
  - **How?** Model a preference distribution from samples, integrate into a single loss (one-stage approach)

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

- **Gradient step:**

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



**Thank You!**