

---

# ENTANGLED PREFERENCES: THE HISTORY AND RISKS OF REINFORCEMENT LEARNING AND HUMAN FEEDBACK

---

**Nathan Lambert**  
Independent  
Berkeley, CA, USA  
nathan@natolambert.com

**Thomas Krendl Gilbert**  
The New York Academy of Sciences  
New York, NY, USA  
tgilbert@nyas.org

**Tom Zick**  
Berkman Klein Center, Harvard Law  
Cambridge, MA, USA  
tzick@jd24.law.harvard.edu

## ABSTRACT

Reinforcement learning from human feedback (RLHF) has emerged as a powerful technique to make large language models (LLMs) easier to use and more effective. A core piece of the RLHF process is the training and utilization of a model of human preferences that acts as a reward function for optimization. This approach, which operates at the intersection of many stakeholders and academic disciplines, remains poorly understood. RLHF reward models are often cited as being central to achieving performance, yet very few descriptors of capabilities, evaluations, training methods, or open-source models exist. Given this lack of information, further study and transparency is needed for learned RLHF reward models. In this paper, we illustrate the complex history of optimizing preferences, and articulate lines of inquiry to understand the sociotechnical context of reward models. In particular, we highlight the ontological differences between *costs*, *rewards*, and *preferences* at stake in RLHF’s foundations, related methodological tensions, and possible research directions to improve general understanding of how reward models function.

## 1 Introduction

Learning from human feedback has become incredibly popular due to the success of large language models (LLMs) such as OpenAI’s ChatGPT (Schulman, Zoph, Kim, & more, 2022) and Anthropic’s Claude (Bai, Jones, et al., 2022), which are heavily dependent on human labeled data. These models make use of reinforcement learning from human feedback (RLHF), a technique designed to integrate human preferences where writing an explicit reward function is otherwise challenging (Christiano et al., 2017). In the context of language models, RLHF proceeds as follows: first, a reward model is independently trained on aggregate pairwise preferences from many crowdworkers to rate any piece of text; second, the language model is optimized with a RL optimizer (Bai, Jones, et al., 2022; Ouyang et al., 2022; Touvron et al., 2023). The final language model is often subject to heavy scrutiny both internally and, more recently, externally through events like DEFCON Red Teaming Village (Bajak, 2023) or coordinated adversarial attacks Zou, Wang, Kolter, and Fredrikson (2023). The same cannot be said for the intermediate reward model. Historically, reward models have not been released as open-source or evaluated rigorously, obscuring from scrutiny the process through which values are actively encoded into the system. This paper illustrates why reward models are central to understanding the long-term impacts of RLHF, drawing from the rich history, discourse, and tension around how to best quantify human values.

RLHF is the intellectual culmination of several distinct domains. The optimization stack of RLHF is borrowed from control theory, a domain in which there are ground truths and reward functions can have an clear notion of success. The primary risk of learning human preferences for LLMs comes through the domain shift from control to language. In language, notions of values are more computationally complex or fundamentally vague Dobbe, Gilbert, and Mintz (2021) in relation to their control counterparts, but the optimization stack nevertheless remains similar. Reinforcement learning is broadly the field of study of sequential decision making, which is built on a substantial literature comprising cognitive biology, optimal control, behavioral economics, and other fields (Sutton & Barto, 2018). RLHF combines the social challenges of human data with the techniques of RL — a field with a long history of numerical complexity.

Despite the maturity of the domains it draws on, grounding and investigating risks of RLHF requires the development of new tools and research methods. In particular, vestigial assumptions inherited from earlier technologies can re-surface as blind spots in the modern RLHF paradigm. Tracing the history of RL and RLHF as technologies allows us to identify these assumptions and where they matter in particular systems. This paper attempts to provide an exposition of some of this historical context, and to highlight specific sociotechnical areas of opportunity within the reward model specification, beyond the challenges proposed in recent literature (Casper, Davies, et al., 2023). We study the histories of quantification of human preferences and reinforcement learning algorithms, from *Port-Royal Logic* and Bentham to Markov Decision Processes and Bellman, to highlight potential shortcomings of learning models of human preferences. An initial concern that has been raised with the current deployments of LLMs is the limitations of working with aggregate human data, raising questions as to whose values the model is encoding and prioritizing. Moving beyond this, we study how structural optimization and deployment decisions can impact downstream users.

Given the nuance around modeling human preferences, we refer to these artifacts as reward models of preference, or *reward models* for short, to better match their usage as an optimization target for reinforcement learning algorithms rather than an accurate representation of human values. In order to broaden the scope of study around these reward models, we make the following contributions:

- **Trace the complex intellectual history of RLHF** to illustrate the potential ill-posed assumptions popularized within RLHF. In Sec. 3, we explain the evolution of RL with the history of rational agents and human preferences. In doing so, we distinguish sets of *assumptions* (explicit premises) and *presumptions* (implicit premises) made throughout the evolution of RLHF that lead to potential shortcomings of reward models.
- **Propose a series of questions for contemporary RLHF reward models** to increase transparency and opportunities for broader multi-stakeholder engagement in modern LLM development. In Sec. 5, we break these questions down by sections of the machine learning process: data, model, and optimization, and in Sec. 7, we also discuss emerging issues that are not easily classified.
- **Discuss solutions** in Sec. 6 to measure and communicate the values and potential harms of contemporary RLHF reward models. We propose tools that can be used to add rigour to future empirical evaluation work.

## 2 Related Works

### 2.1 Reinforcement learning from human feedback

RLHF is a set of techniques designed to optimize machine learning models based on human feedback in order to circumvent the need to design a complex reward function. Early work in RLHF focused on soliciting complex behaviors from AI agents in control problems using various environments, feedback methods across trajectories or rankings, and optimizers (Christiano et al., 2017; Wirth, Akrou, Neumann, Fürnkranz, et al., 2017).

Recently, developments in RLHF have been centered around its use with LLMs. This branch of study originated with work exploring how technical value alignment may scale with learned reward models (Leike et al., 2018). Quoting Leike et al. (2018):

We claim that the approach described is agnostic to the ethical paradigm, the user’s preferences, and the legal or social framework, provided we can supply enough feedback (though the preference payload might influence the amount of feedback required).

The organization that builds *applications* where RLHF is used bears the burden of specifying the ethics they used and answering questions about whose preferences are included and how they’re weighed (Baum, 2020; Prasad, 2018).

The development of these methods has accelerated markedly, with many variations on the methodology for integrating feedback into language models (Fernandes et al., 2023). Initial work on RLHF for LLMs utilized user choices from a batch of 4 completions (Ziegler et al., 2019) to train a reward model across general LLM benchmarks. When comparing recent RLHF work to Ziegler et al. (2019), group preferences were changed to pairwise preferences, and rather than general benchmarks the reward model was focused on the task of summarization (Stiennon et al., 2020; J. Wu et al., 2021). Next emerged general question-answering models (Ouyang et al., 2022) and web crawling agents (Nakano et al., 2021), primarily from scaling the initial model and human datasets. Now, RLHF is used to train general chat models across a variety of tasks (Bai, Jones, et al., 2022; Schulman et al., 2022; Touvron et al., 2023) and for specific objectives such as harm reduction (Glaese et al., 2022) or information accuracy (Menick et al., 2022).

## 2.2 Downstream impacts of optimizing preferences

Research venues have encouraged scientists to grapple with these questions around their work and data enrichment through humans, which are particularly relevant for techniques similar to RLHF, but there has been mixed uptake (Hawkins & Mittelstadt, 2023). RLHF faces many challenges with its integration of human preferences in an aggregate manner, and potential solutions involving personalization of preferences raise further questions of which values or norms are acceptable to encode in a model (Kirk, Vidgen, Röttger, & Hale, 2023). Specifically, the reward models trained for RLHF are known to be over-optimized during the RL stage, where the language generations continue to shift without the reward model indicating a higher score, without clear measurement of how downstream training signals for LLMs relate to preferences being correlated in the data (Gao, Schulman, & Hilton, 2022).

Training models based on human preferences also impacts how users interact with the downstream machine learning systems that refer to RLHF reward models as part of their optimization objective. This was illustrated with the launch and widespread adoption of ChatGPT, raising questions regarding the effects of regular communication with RLHF trained LLMs, such as the downstream impact on users’ moral judgements (Krügel, Ostermaier, & Uhl, 2023) or exposure to value judgements and biases (Johnson et al., 2022). Additionally, there are open questions about the stability and robustness of RLHF-trained LLM, with reports of RLHF models’ tone shifting substantially within the course of a single conversation in potentially troubling ways (Nardo, 2023).

The issue of downstream model impacts is not new - for instance there is a vast prior literature on how models interface with society. For example, user facing recommendation models have long prompted inquiry around whether agents should respond to our stated or implied preferences (Milli, Hadfield-Menell, Dragan, & Russell, 2017). In RLHF, these concerns meet the ‘reward hacking’ problem endemic to RL. Specifically, as popular models are being tuned based on user experiences, complex feedback dynamics can emerge via the combination of reward mis-specification with the power of RL optimizers (Gilbert, Dean, Zick, & Lambert, 2022), such as desired capabilities coming and going through repeated training.

## 3 The Origins of Reward Models: Costs vs. Rewards vs. Preferences

In this section, we break down the complex history inspiring the modern use of RLHF. To do so, we must cover the intellectual foundations of quantifying human values, reinforcement learning and optimality, and topics from behavioral economics around measuring preferences. The notion of using reinforcement learning to optimize a preference model combines the history of various once-distanced fields into an intimate optimization built on many assumptions about human nature.

Our goal is to unspool the types of uncertainty that designers have grafted to system architectures at various stages of their intellectual history. Modern problem specifications have repeatedly stepped away from domains where optimal solutions are possible and deployed under-specified models as approximate solutions.

Throughout, we distinguish between a series of *assumptions* common within theoretically-grounded academic literatures, and relevant *presumptions* which are commonly accepted methods of practice for particular subject areas. As we shall see, the unresolved tensions between these assumptions and presumptions are responsible for the current state and outstanding questions of RLHF research.

### 3.1 Specifying objectives: from logic of utility to reward functions

The optimization of RLHF explicitly relies only on reward models. In order to get to using rewards as an optimization target, RLHF presupposes the convergence of ideas from preferences, rewards, and costs. Models of preference, reward functions, and cost landscapes all are tools used by different fields to describe a notion of relative goodness of specific actions and/or states in the domain. The history of these three framings dates back to the origins of probability theory and decision theory. In 1662, *The Port Royal Logic* introduces the notion of decision making quality (Arnauld, 1662):

To judge what one must do to obtain a good or avoid an evil, it is necessary to consider not only the good and evil in itself, but also the probability that it happens or does not happen.

To summarize a substantial literature, this theory has developed along with modern scientific thinking, starting with Bentham’s utilitarian *Hedonic Calculus*, arguing that everything could be weighed (Bentham, 1823). The first quantitative application of these ideas emerged in 1931 with Ramsey’s *Truth and Probability* (Ramsey, 2016).

**Assumption 1.** *Preferences and goals can be quantified and measured.*

Since these works, quantifying, measuring, and influencing human preferences has been a lively topic in the social and behavioral sciences. These debates have often not been settled on a theoretical level; rather, different subfields and branches of social science have reached internal consensus on methods and approaches to preference measurement even as they have specialized relative to each other, often developing their own distinct semantics in the process.

A minority of economists posit that preferences, if they do exist, are prohibitively difficult to measure because people have preferences over their own preferences, as well as each others' preferences (Hirschman, 1984). In this view, which is not reflected in the RLHF process, individual preferences are always embedded within larger social relations, such that the accuracy of any preference model is contingent on the definition and context of the task. Some behavioral economists have even argued that preferences don't exist—they may be less an ontological statement of what people actually value than a methodological tool for indirectly capturing psychological predispositions, perceived behavioral norms and ethical duties, commitments to social order, or legal constraints (Hadfield & Weingast, 2014). We address the links of this work to the Von Neumann-Morgenstern (VNM) utility theorem and countering impossibility theorems around quantifying preference in Sec. 3.3.

On the other hand, the reinforcement learning optimization methods used today are conceptualized around optimizing estimates of reward-to-go in a trial (Sutton & Barto, 2018). Reward-to-go is an idea derived to mirror the notion of utility, which is a measure of rationality over time (Briggs, 2014). In the context of the mathematical tools used for reinforcement learning, utility-to-go was invented in control theory, specifically in the context of analog circuits in 1960 (Widrow & Hoff, 1960). These methods are designed around systems with clear definitions of optimality, which are numerical representations of goals of an agent, which is now in heavy contrast to reward models that aggregate multimodal preferences. Specifically, RL systems expect rewards to behave in a specific manner, quoting Singh, Lewis, and Barto (2009):

Rewards in an RL system correspond to primary rewards, i.e., rewards that in animals have been hard-wired by the evolutionary process due to their relevance to reproductive success. ... Further, RL systems that form value functions, ... effectively create conditioned or secondary reward processes whereby predictors of primary rewards act as rewards themselves... The result is that the local landscape of a value function gives direction to the system's preferred behavior: decisions are made to cause transitions to higher-valued states. A close parallel can be drawn between the gradient of a value function and incentive motivation (McClure, Daw, & Montague, 2003).

To summarize, rewards are used in RL systems as a signal to tune behavior towards clearly defined goals. The core thesis is that agents' performances are closely coupled with notions of *expected fitness*. This view is linked to the development of reinforcement learning technology, exemplified by claims of the general usefulness of the reward formulation (Silver, Singh, Precup, & Sutton, 2021), but is in conflict when many individual desires are reduced to a single function.

**Assumption 2.** *Increasing the score of raw reward measurements corresponds to better behaviors (or value functions learned under invariant reward transformation (Ng, Harada, & Russell, 1999)).*

### 3.2 Implementing optimal utility

Modern reinforcement learning methods depend strongly on the Bellman equation (Bellman, 1957) to recursively compute estimates of reward-to-go, derived within closed environments that can be modeled as a Markov Decision Process (MDP) (Sutton & Barto, 2018). The MDP formulation provides theoretical guarantees of performance by structuring the environment as one with a non-changing distribution of state-actions.

**Assumption 3.** *Optimal solutions to reward maximization problems exist.*

The algorithms used today continued to build within this formulation of RL as a tool to find optimal behaviors, but under looser conditions. The notion of temporal-difference (TD) learning was developed to aid agents in both the credit assignment and data collection problems, by directly updating the policy as new data is collected (Tesauro et al., 1995) (rather than updating from a large dataset of cumulative experience, which could be outdated via erroneous past value predictions). Q-learning, the basis for many modern forms of RL, learns a model via the Bellman equation that dictates how useful every state-action pair is with a TD update (Watkins & Dayan, 1992). Crucially, these notions of provable usefulness through utility have only been demonstrated for domains cast as MDPs or addressed in tasks with a single closed-form reward function, such as prominent success in games with deep learning (DQN) (Mnih et al., 2013).

As the methods became more general and successful, most developments have remained motivated within the context of adaptive control, where reward and cost functions have a finite notion of success (Golnaraghi & Kuo, 2017), e.g. a minimum energy consumption across an episode in a physical system. Prominent examples include further success in games (Silver et al., 2017), controlling complex dynamic systems such as nuclear fusion reactors (Degraeve et al., 2022),

and controlling rapid robotic systems (Kaufmann et al., 2023). Most reward or cost functions can return an explicit optimal behavior, where models of human preferences never can do so.

**Presumption 1.** *Optimal solutions can be achieved with finite data in complex environments.*

Given the successes of deep RL, it is worth noting that the mechanistic understanding of how the methods succeed is not well documented. The field is prone to mistakes of statistical analysis as the methods for evaluation grow more complex (Agarwal, Schwarzer, Castro, Courville, & Bellemare, 2021). In addition, there is little mention of the subfield of inverse reinforcement learning (IRL) in the literature of RLHF, which is the problem of learning a reward function based on an agent’s behavior (Ng, Russell, et al., 2000) and highly related to learning a reward model. This primarily reflects the engineering path by which a stable approach to performing RLHF emerged, and motivates further investment and comparison to IRL methods to scale them to the complexity of open-ended conversations.

### 3.3 Steering preferences

In the context by which reinforcement learning was designed, rewards and costs are both expected to be functions, where if the agent is in a specific state-action pair, then it will be returned a certain value. As we move into preferences, this is no longer the case, as human preferences constantly drift temporally throughout their experiences. The overloading of the term “value” within these two contexts complicates the literature of RLHF that is built on the numerical value updates in Bellman equations with the very different notion of what is a human value, which often refers to moral or ethical principles, but is not well defined in technical literature. An example of where this tension can be seen is how reward models are attempting to map from the text on the screen to a scalar signal, but in reality, dynamics not captured in the problem specification influence the true decision (Gilbert, Dean, Zick, & Lambert, 2022; Salha, 2011), such as preference shift when labeling many examples sequentially and assuming they are independent. Therein, modeling preferences is at best compressing a multi-reward environment to a single function representation.

In theory, the Von Neumann-Morgenstern (VNM) utility theorem gives the designer license to construct such functions, because it ties together the foundations of decision theory under uncertainty, preference theory, and abstract utility functions (Von Neumann & Morgenstern, 1947); together, these ideas allow preferences to be modeled in terms of expected value to some individual agent.

However, the VNM utility theorem also invokes a number of assumptions about the nature of preferences and the environment where preferences are being measured. In a similar vein, human-computer interaction (HCI) researchers have emphasized that any numerical model of preference may not capture all the relevant preferences of a scenario. For example, how choices are displayed visually influences people’s preferences (Salha, 2011). This means that representing preferences may be secondary to how that representation is integrated within a tool available for people to use. Work from development economics echoes this notion, showing that theories of revealed preferences may just recapitulate *Hume’s guillotine* (you can’t extract an “ought” from an “is”), and in particular the difference between choice (what do I want?) and preference (is X better than Y?) (Sen, 1973).

On a mathematical level, well-known impossibility theorems in social choice theory show that not all fairness criteria can be simultaneously met via a given preference optimization technique (Arrow, 1950; Maskin & Sen, 2014). This challenges the notion that preferences can be naively aggregated across people, places, or diverse situations, which is common practice in modern RLHF dataset engineering. Philosophers have also recently highlighted the fact that preferences change over time, raising fundamental questions about personal experiences, the nature of human decision-making, and distinct contexts (Pettigrew, 2019).

In practice, the VNM utility theorem ignores the possibility that preferences are also uncertain because of the inherently dynamic and indeterminate nature of value—human decisions are shaped by biology, psychology, and culture in ways that influence their preferences, for reasons that do not apply to a perfectly rational agent. As a result, there are a variety of paths through which theoretical assumptions diverge in practice:

- measured preferences may not be transitive or comparable with each other as the environment where they are measured is made more complex;
- proxy measurements may be derived from implicit data (page view time, closing tab, repeating question to language model), without interrogating how the measurements may interact with the domain they’re collected in via future training and deployment of the model;
- the number and presentation of input sources may vary the results, e.g. allowing respondents to choose between more than two options, or taking in inputs from the same user at multiple times or in multiple contexts;
- relatively low accuracy across respondents in RLHF training data, which may mask differences in context between users that the preference model can aggregate or optimize without resolving.

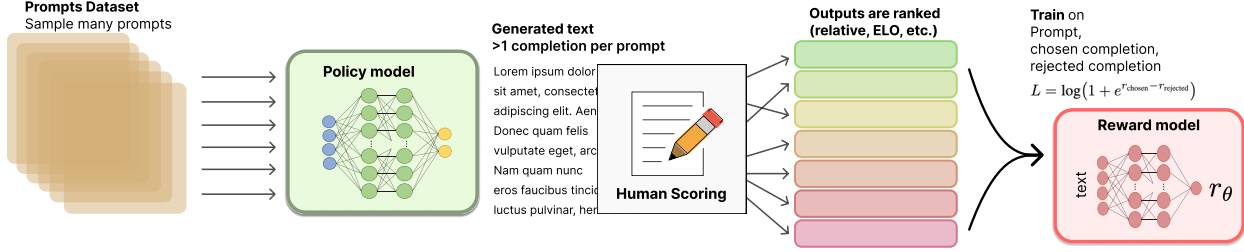


Figure 1: Illustrating the training process of an RLHF language reward model. A language model generates text from a distribution of prompts to optimize human preferences for, then humans assign a preference order to them (normally pairwise preferences). These preferences are used to train a score function via a contrastive loss shown in Eq. 1.

**Presumption 2.** *The temporal and context shifting of user preferences does not mitigate the effectiveness of reward functions or notions of optimal utility as an optimization target.*

## 4 Background

We continue to use *assumptions* of the literature, grounded in theoretical backing of a subject area, and *presumptions*, which are commonly accepted methods of practice, to later connect to open questions in reward modeling.

### 4.1 Reward models of human preferences

**Data collection** Reward models are trained with human preference data collected over a distribution of graphical user interfaces. The most common task is to give the model a *prompt*, i.e. a request or instruction, and rate the *completion*, or answer. Some implementations collect rankings from groups of responses (Ziegler et al., 2019), scores and rankings of a group of candidate responses (Ouyang et al., 2022) (scores of 1-5 for 6+ outputs then ranking all), a choice between a pair of responses (Bai, Jones, et al., 2022) (choose best response between two options), and more (Z. Wu et al., 2023). The workers employed are generally given detailed instructions on which styles, occurrences, or values to prioritize in their labels. This data is collected from groups of individuals and sometimes calibrated across per-labelor distributions, but that decision making process is not always documented or shared.

In each batch of rankings for a prompt-completion pair, the pairs are often *binarized* into a chosen element and rejected element (rather than using the corresponding delta as a training objective) to create easy-to-optimize training data. This binarization of chosen and rejected is aggregated independently to measurements of the difficulty of the prompt or mean quality of the responses.

**Presumption 3.** *Pairwise preferences can suitably perform as a basis of human values.*

**Model training** The reward models trained for RLHF are most often trained as classifiers between a chosen and rejected completion to a prompt before optimizing with RL where they return a scalar value for each piece of text.

The loss function of reward models is formulated as a difference between the score for each binarized preference. Given two options for a completion  $y$  from a prompt  $x$ , and the scores they obtain a scalar output  $r$  from a value head of a language model, the loss for the reward model is as follows (Askell et al., 2021; Ouyang et al., 2022):

$$L = \log(1 + e^{r_{chosen} - r_{rejected}}) \quad (1)$$

The loss function is designed to increase the distance between the two samples, where variations exist including losses of 4 samples rather than a pair (Ziegler et al., 2019), updating the model with batches of pairwise labels on a given prompt (Ouyang et al., 2022), or optimizing based on the margin between  $r_{chosen}$  and  $r_{rejected}$  (Touvron et al., 2023).

There are multiple popular ways for adapting base language models into reward models. In earlier work, a value head was appended to the model to predict these rewards. A value head can be as simple as adding a linear layer with one output node after the attention blocks of a Transformer model. Today, some work includes an entire copy of the base language model with modified output layers for predicting rewards. Other reward models are trained by removing layers from a base model and/or freezing original model parameters during training. The reward model training component of RLHF is shown in Fig. 1.

**Presumption 4.** *Multiple user preferences are represented in training one model.*

**Presumption 5.** *The only preferences embedded in the model are from the specifically collected training data.*

## 4.2 Reinforcement Learning on Language

Language generation optimized via reinforcement learning, which RLHF is a version of, can be formalized as a partially observable Markov decision process (POMDP) (Spaan, 2012). We define a POMDP  $\mathcal{M}$  at a per-token level with  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{Z}, \mu_0, \mathcal{R}, \gamma)$ . Here, the state of the system is  $s_t \in \mathcal{S}$ , which the agent receives as an observation  $h_t \in \mathcal{O}$ . The observation is a history of tokens  $h_t = \{t_0, t_1, \dots, t_{t-1}\}$  and the action space is the possible set of next-tokens in the vocabulary of the policy model  $a_t = t_t \in \mathcal{A}$ , including the end-of-sequence token  $a_{\text{end}}$ . As in a traditional MPD,  $\mathcal{T}$  is the transition function  $\mathcal{T}(\cdot | s_t, a_t)$ .

The goal of the RLHF process is to learn a policy that maps  $\pi : \mathcal{O} \mapsto \mathcal{P}(\mathcal{A})$ . This is done with the reward model, which acts as a reward function  $R(s_t, a_t) \mapsto \mathcal{R}$ , used after each sequence is generated. The full sequence, until end-of-sequence token  $a_{\text{end}}$ , is called the action and used to get a scalar reward  $r_t$  from the reward model.

The rewards for the batch are used for the RL update, where a popular algorithm is Proximal Policy Optimization (Schulman, Wolski, Dhariwal, Radford, & Klimov, 2017). In RLHF, the discount factor of reward is set to 1 and no further actions are taken until the end of a sequence, casting the problem as contextual bandits.

**Presumption 6.** *User preferences are extracted uniformly via the RLHF process.*

## 5 Questions

In this section, we propose a series of investigations to disentangle the diverse assumptions around using large models of human preferences built on LLMs within the RLHF framework. We focus on the core sections of a machine learning solution: data, model, and optimization choices that dictate the priorities encoded in contemporary reward models. The questions are connected to one or more assumptions or presumptions presented in Sec. 3 and Sec. 4.

### 5.1 Model Questions

The reward models of human preference used in RLHF today are built on base, generative LLMs, so issues from the full design and training process of these models applies to reward models. This includes the wide literature on potential sources of harm from LLMs, including a large swath of social biases, such as anti-Muslim bias (Abid, Farooqi, & Zou, 2021; Liang, Wu, Morency, & Salakhutdinov, 2021). Those pretraining LLMs should reckon with the design choices of the base model including environmental costs, training datasets, stakeholder involvement, and more (Bender, Gebru, McMillan-Major, & Shmitchell, 2021). When building off these models in the RLHF process, these base model capability distributions should be tested against, but often are not for engineering and competitive reasons. Given this, two sets of evaluations should be included in the RLHF process:

- **Base model biases** (for **Presumption 5**): How do different base LLMs used in the reward model training process influence all aspects of reward models? How does this impact training numerical stability, efficiency, training or test loss, and more? It is possible that many design choices in RLHF are based on a static model and could be solved with a change of base model, as opposed to increased data and optimization? This links to all the questions we ask in the next subsection – ultimately, how does the base model perform as an independent variable?
- **Sequential model evaluation in RLHF** (for **Presumption 5, 6**): When a reward model is trained, it is making subtle changes to the parameters of a LLM. This is done by appending parameters in parallel, i.e. re-routing the model, so the text that the reward model produces could be observed. How do the biases of a pretrained base model, a downstream RLHF model, and the proxy of an intermediate reward model showcase how RLHF modifies representations of LLMs? Given the common implementation technique of appending a value head to a base LLM and fine-tuning to construct a reward model, the reward models can still generate text that can indicate if training is operating as intended. The base model could also impact which data from the preferences dataset is learned during the RLHF process.

### 5.2 Data Questions

The data used to enable RLHF is often curated by multiple stakeholders in a combination of paid employment and consumer usage. This data, representing a preference between two pieces of text in an individual instance, is capturing a broad and diverse function via extremely limited interactions. Given that the data is sparse in count relative to the complexity it begins to represent, more questions should be openly shared about its curation and impacts.

Currently, datasets for the most popular LLMs are being generated by professional work-forces. This opens up many questions around who is creating the data and how the context of their workplace informs it, including:

- **Data collection contexts** (for **Presumption 2**): Can data involving preferences collected in a professional setting mirror the intent of researchers designing an experiment or provide suitable transfer to downstream users? How does this compare to volunteer workers? How does context inform preferences, how does this data impact a downstream model, how can the impact of a user interface be measured in data? How does repetitive labeling of preference data shift one’s preferences? Do professional crowd-workers, instructed to follow a set of preferences, follow the instructions or their innate values?
- **Type of feedback** (for **Presumption 3**): Does the default operating method of RLHF, pairwise preferences capture preferences in its intended form? Can comparisons in RLHF across the same data be made with the default comparisons versus advanced multi-axes feedback mechanisms (Z. Wu et al., 2023)? What types of comparisons would reflect how humans communicate preferences in text?
- **Population demographics** (for **Presumption 4**): Who is completing the data? Is a diverse population maintained? How does a lack of diversity emerge as measurable impacts on the model? What is a minimum number of people required to suitably represent a given population? How are instances of preference annotator disagreement treated – as a source of noise, or a signal?

### 5.3 Optimization questions

The most elusive and potentially insightful questions with respect to reward models and preference data emerge where preferences are extracted by reinforcement learning optimizers and distilled into the LLM. In this vein, further questions should be asked as to whether the outcomes of the optimization match the intended design of the process. Given the complexity of LLMs and the social values of using human data to empower technology, clear attribution is needed for the role of the preference data in the optimization. Questions that can be used to investigate these issues include:

- **RL optimization of reward model** (for **Assumption 1, 2, 3, Presumption 1**): at a technical level, is the reward signal from the preference model maximized? Should it be maximized?? Does this reflect changes in line with the preference data (e.g. is this achieved at the same point where evaluation accuracy peaked). In the RL literature, sharing learning curves is common practice, yet in RLHF it is often omitted, with just the final test set accuracy included. Correlating these common metrics is important to understanding the role of the reward model to the process.
- **Qualitative alignment** (for **Assumption 1, Presumption 2, 4, 6**): Given the RLHF data process starts with written instructions to data-workers, checks should be done to identify if the model changes align with the intended goals of stated preferences. A suitable method for continued deployment to iteratively refine training and deployment goals would be documentation for RL systems, Reward Reports (Gilbert, Dean, Lambert, Zick, & Snoswell, 2022). Beyond this, feedback from annotators, data analysis, and more can quickly provide feedback onto the impacts of RLHF relative to human inputs and intentions.
- **Weighing preferences** (for **Presumption 6**) Does and should the model uniformly extract information from the dataset? How are different scenarios in the dataset equated (e.g. instructions around healthcare being treated the same as food choices)? Should different reward functions derived in preference literature be used to maintain an inductive bias over preferences, such as Plackett-Luce (Cheng, Hüllermeier, & Dembczynski, 2010). Recent literature shows that aggregate agreement between humans and the reward models varies from approximately 60 to 75% (Bai, Jones, et al., 2022; Ouyang et al., 2022; Stiennon et al., 2020), should certain examples be prioritized given a non-uniform distribution of tasks and annotators? Or, can these distributions be broken down by commonly accepted categories such as reasoning, generation, etc.? Such reduction to the mean across groups can squash multi-modal beliefs, especially with underrepresented groups (Feffer, Heidari, & Lipton, 2023; Prabhakaran, Davani, & Diaz, 2021).

## 6 Solutions

**Evaluation of reward model capabilities** While the process of curating pairwise or batched reward model data is labor intensive and broad in scope, the impact of the relatively simple training loss should be measured after training. The most common evaluation technique is to hold out data from the training set and measure the final model’s agreement with the data, but this does not control for the complex nature of the dataset. The evaluation score is unlikely to illuminate any conflicts in the data that arise from aggregating over inputs from multiple individuals or contexts. This necessitates building more controlled examples to interpret whether the design objective of differentiating better from worse text is achieved.

Taking inspiration from recent advancements in evaluating instruction- and chat-tuned LLMs (Zheng et al., 2023), where the existing academic benchmarks from the natural language processing literature were incomplete for evaluating



modern models like GPT4 (Chang et al., 2023), the community should evaluate reward models in a manner analogous to their use. One specific example exists in recent work, a benchmark to evaluate reward model consistency (Shen et al., 2023), as its ability to maintain consistent scores over changes to text that do not alter meaning, but many more tools are needed to cover the many use cases of RLHF models.

Similar to MT Bench (Zheng et al., 2023), a benchmark for LLM capabilities over multi-turn conversations, we can construct representative examples across a variety of categories for the model (e.g. reasoning, generation ability, factuality, etc.). To do so, we must generate a series of prompts with a chosen and rejected response. The benchmark score would be the average accuracy at returning a higher scalar reward for chosen instead of rejected samples.

**Evaluation of reward model safety** The primary means for evaluating the safety, toxicity, and potential harms of LLMs is a process known as red-teaming, which entails focused attempts to prompt the language model into a broad array of harmful behaviors (Casper, Lin, Kwon, Culp, & Hadfield-Menell, 2023). Red-teaming for generative models is primarily separated into two axes, attack vectors and harms. Evaluation of reward models can easily be developed as a subset of red-teaming techniques to identify if the model successfully downweights requests for harm or toxicity.

The notion of red-teaming for reward models is simpler and less interactive, as the output of the model is a score rather than another turn of text. For this reason, progress on red-teaming and safety analysis of reward models should be simpler. To do so, the score of reward models across a series of text documents or snippets that have the potential for harm should be compared to ‘neutral’ texts. Second, reward models should be evaluated for adversarial input strings, e.g. sequences of text that produce unexpectedly high or low scores, which could indicate downstream exploitation during the RL step. As a starting point, practitioners could consider the efforts to identify adversarial prompts that escalate toxicity (but need not necessarily do so) in the evaluation of LLM generations, for instance as in done in the RealToxicityPrompts dataset (Gehman, Gururangan, Sap, Choi, & Smith, 2020). Today, standards for red-teaming are primarily developed in private by for-profit corporations deploying LLMs. In the context of reward models, which are rarely shared publicly, this practice can be especially detrimental to developing more robust systems. A third-party researcher can play with a closed model and find an adversarial attack even if no red-teaming protocols are shared about the system. The same is not true for a private reward model, which may be re-used in the future to train the next model.

**Sociotechnical specification of preference** The current manner by which reward models are used in RLHF follows the logic of the following statement: *the reward model is deemed good when it is good for the downstream task*. Given the complex intellectual origins of quantifying and optimizing preferences, specific recommendations should be made around what information a model of human preferences should and should not encode. Doing so may reduce the performance of models in the short term, but increases the ability for multi-stakeholder engagement and reduction of harms with the development of LLMs. Dean, Gilbert, Lambert, and Zick (2021) proposes a set of tools for expanding the sociotechnical axes of AI and could be used for this analysis, while Gilbert, Qu, Ju, and Li (2023); Gilbert, Snoswell, Dennis, McAllister, and Wu (2022) comprise case studies of *sociotechnical specification* in the case of automated vehicle fleets. Substantial work in this vein is needed to understand the scope of impact of RLHF.

## 7 Discussion

There will be many additional questions surrounding the use of reward models as existing methods for RLHF expand. Such expansion will further complicate the interdisciplinary problem specification of RLHF.

**Downstream use of models** The final manner by which the assumptions of RLHF are tested is in the increasing alignment of LLM training with the intentions of downstream users, as is the case for ChatGPT (Schulman, 2023). LLMs are technologies used in consumer applications, giving different context and broader use-cases than the limited training data. The developers of RLHF should consider how end user encounters with their model may be different than the preferences collected for training, such as the shift from primarily English training data to multilingual users, which can proliferate low-resource language jailbreak vectors (Yong, Menghini, & Bach, 2023).

Given the direct manner by which human preferences are encoded in models trained with RLHF, users of these models should be presented with the intended preferences of the models during training for transparency. A potential improved solution is to let users choose a model matching their preferences, but this comes with many technical challenges. However, emerging examples of LLM ‘hub’ interfaces that offer users a menu of LLMs to choose from (like the

open-source GPT4All<sup>1</sup>, or Meta’s new AI companion character suite<sup>2</sup>) offer hope that such approaches might in-fact be tractable and/or accessible to users.

**Synthetic preference data** Many recent works, primarily grouped under Reinforcement Learning from AI Feedback (RLAIF) are using LLMs to provide critiques of other LLM generations (Bai, Kadavath, et al., 2022; Lee et al., 2023; Wang et al., 2023). These critiques can be used in many ways, but the most prominent is to use them as further preference data in the RLHF process (Bai, Kadavath, et al., 2022). Further, these methods are referenced as methods for removing the human bottleneck from alignment (Sutskever, 2023) without existing literature on how methods such as RLHF and reward models succeed in capturing values. Despite the practical attractiveness due to the cost of acquiring human data, synthetic preference data presents many challenges. Preliminary findings have indicated that synthetic supervision data for LLMs can lead to generation instability and lack of robustness (Alemohammad et al., n.d.; Gudibande et al., n.d.; Shumailov et al., n.d.), to say nothing of the potential for further shifting preferences away from the original human data.

**Direct preference optimization** Recent work has shown that direct optimization can be used to extract information from human preference data (Rafailov et al., 2023), rather than using an intermediate reward model. The removal of the reward model would make the encoding of values more opaque in the RLHF process, yet at the same time potentially reduce problem misspecification and exploitation by reducing the number of training steps. The questions suggested in this work would still apply, but the methods for auditing them will be different without a distinct reward model.

## 8 Conclusion

Given the importance of RLHF to the deployment and integration of state-of-the-art LLMs, a complete understanding of the motivation, intellectual foundations, and implementation presumptions of reward models is crucial to the safe and equitable rollout of this technology. With the evolution from utility quantification to reinforcement learning to maximizing average preference scores, we have shown that the assumptions at each step are not in-step with how the technology now used in contemporary settings. This paper details the intellectual basis of learning reward models, and argues that more information about existing reward models should be shared with the public via open or partially-open releases. With more development and sharing of reward models, proper evaluation techniques of the models can be designed, which will lead to further technological developments around desired properties such as user tunability, uncertainty quantification, etc. By doing this work, we can better understand the difference between the challenges of *learning* models from incomplete data of human preferences and *optimizing* those models to elicit specific behaviors.

## Acknowledgements

Thanks to Aaron Snoswell for providing edits and comments on this work. Credit to Sasha Luccioni and Deep Ganguli for discussions and encouragement towards completing this paper.

## References

- Abid, A., Farooqi, M., & Zou, J. (2021). Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 aaai/acm conference on ai, ethics, and society* (pp. 298–306).
- Agarwal, R., Schwarzer, M., Castro, P. S., Courville, A. C., & Bellemare, M. (2021). Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34, 29304–29320.
- Alemohammad, S., Casco-Rodriguez, J., Luzi, L., Humayun, A. I., Babaei, H., LeJeune, D., ... Baraniuk, R. G. (n.d.). *Self-Consuming Generative Models Go MAD*. Retrieved 2023-08-01, from <http://arxiv.org/abs/2307.01850> doi: 10.48550/arXiv.2307.01850
- Arnauld, A. (1662). *The port-royal logic*.
- Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of political economy*, 58(4), 328–346.
- Askell, A., Bai, Y., Chen, A., Drain, D., Ganguli, D., Henighan, T., ... others (2021). A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., ... others (2022). Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.

---

<sup>1</sup><https://gpt4all.io/index.html>

<sup>2</sup><https://about.fb.com/news/2023/09/introducing-ai-powered-assistants-characters-and-creative-tools>

- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., ... others (2022). Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Bajak, F. (2023). *Hackers red-teaming a.i. are 'breaking stuff left and right,' but don't expect quick fixes from defcon: 'there are no good guardrails'*. <https://fortune.com/2023/08/13/hackers-red-teaming-ai-defcon-breaking-stuff-but-no-quick-fixes/>. (Accessed: 2023-10-03)
- Baum, S. D. (2020). Social choice ethics in artificial intelligence. *AI & SOCIETY*, 35(1), 165–176.
- Bellman, R. (1957). A markovian decision process. *Journal of mathematics and mechanics*, 679–684.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 acm conference on fairness, accountability, and transparency* (pp. 610–623).
- Bentham, J. (1823). *An introduction to the principles of morals and legislation*.
- Briggs, R. A. (2014). Normative theories of rational choice: Expected utility.
- Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., ... others (2023). Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Casper, S., Lin, J., Kwon, J., Culp, G., & Hadfield-Menell, D. (2023). Explore, establish, exploit: Red teaming language models from scratch. *arXiv preprint arXiv:2306.09442*.
- Chang, Y., Wang, X., Wang, J., Wu, Y., Zhu, K., Chen, H., ... others (2023). A survey on evaluation of large language models. *arXiv preprint arXiv:2307.03109*.
- Cheng, W., Hüllermeier, E., & Dembczynski, K. J. (2010). Label ranking methods based on the plackett-luce model. In *Proceedings of the 27th international conference on machine learning (icml-10)* (pp. 215–222).
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Dean, S., Gilbert, T. K., Lambert, N., & Zick, T. (2021). Axes for sociotechnical inquiry in ai research. *IEEE Transactions on Technology and Society*, 2(2), 62-70. doi: 10.1109/TTS.2021.3074097
- Degrave, J., Felici, F., Buchli, J., Neunert, M., Tracey, B., Carpanese, F., ... others (2022). Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897), 414–419.
- Dobbe, R., Gilbert, T. K., & Mintz, Y. (2021). Hard choices in artificial intelligence. *Artificial Intelligence*, 300, 103555.
- Feffer, M., Heidari, H., & Lipton, Z. C. (2023). Moral machine or tyranny of the majority? *arXiv preprint arXiv:2305.17319*.
- Fernandes, P., Madaan, A., Liu, E., Farinhas, A., Martins, P. H., Bertsch, A., ... others (2023). Bridging the gap: A survey on integrating (human) feedback for natural language generation. *arXiv preprint arXiv:2305.00955*.
- Gao, L., Schulman, J., & Hilton, J. (2022). Scaling laws for reward model overoptimization. *arXiv preprint arXiv:2210.10760*.
- Gehman, S., Gururangan, S., Sap, M., Choi, Y., & Smith, N. A. (2020). Realtoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462*.
- Gilbert, T. K., Dean, S., Lambert, N., Zick, T., & Snoswell, A. (2022). Reward reports for reinforcement learning. *arXiv preprint arXiv:2204.10817*.
- Gilbert, T. K., Dean, S., Zick, T., & Lambert, N. (2022). Choices, risks, and reward reports: Charting public policy for reinforcement learning systems. *arXiv preprint arXiv:2202.05716*.
- Gilbert, T. K., Qu, N. Z., Ju, W., & Li, J. (2023). Fleets on the streets: How number, affiliation and purpose of shared-lane automated vehicle convoys influence public perception and blame. *Transportation research part F: traffic psychology and behaviour*, 93, 294–308.
- Gilbert, T. K., Snoswell, A. J., Dennis, M., McAllister, R., & Wu, C. (2022). Sociotechnical specification for the broader impacts of autonomous vehicles. *arXiv preprint arXiv:2205.07395*.
- Glaese, A., McAleese, N., Trębacz, M., Aslanides, J., Firoiu, V., Ewalds, T., ... others (2022). Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*.
- Golnaraghi, F., & Kuo, B. C. (2017). *Automatic control systems*. McGraw-Hill Education.
- Gudibande, A., Wallace, E., Snell, C., Geng, X., Liu, H., Abbeel, P., ... Song, D. (n.d.). *The False Promise of Imitating Proprietary LLMs*. Retrieved 2023-08-01, from <http://arxiv.org/abs/2305.15717> doi: 10.48550/arXiv.2305.15717
- Hadfield, G. K., & Weingast, B. R. (2014). Microfoundations of the rule of law. *Annual Review of Political Science*, 17, 21–42.
- Hawkins, W., & Mittelstadt, B. (2023). The ethical ambiguity of ai data enrichment: Measuring gaps in research ethics norms and practices. In *Proceedings of the 2023 acm conference on fairness, accountability, and transparency* (pp. 261–270).
- Hirschman, A. O. (1984). Against parsimony: Three easy ways of complicating some categories of economic discourse. *Bulletin of the American Academy of arts and Sciences*, 37(8), 11–28.
- Johnson, R. L., Pistilli, G., Menéndez-González, N., Duran, L. D. D., Panai, E., Kalpokiene, J., & Bertulfo, D. J. (2022).

- The ghost in the machine has an american accent: value conflict in gpt-3. *arXiv preprint arXiv:2203.07785*.
- Kaufmann, E., Bauersfeld, L., Loquercio, A., Müller, M., Koltun, V., & Scaramuzza, D. (2023). Champion-level drone racing using deep reinforcement learning. *Nature*, 620(7976), 982–987. Retrieved from <https://doi.org/10.1038/s41586-023-06419-4> doi: 10.1038/s41586-023-06419-4
- Kirk, H. R., Vidgen, B., Röttger, P., & Hale, S. A. (2023). Personalisation within bounds: A risk taxonomy and policy framework for the alignment of large language models with personalised feedback. *arXiv preprint arXiv:2303.05453*.
- Krügel, S., Ostermaier, A., & Uhl, M. (2023). Chatgpt’s inconsistent moral advice influences users’ judgment. *Scientific Reports*, 13(1), 4569.
- Lee, H., Phatale, S., Mansoor, H., Lu, K., Mesnard, T., Bishop, C., ... Rastogi, A. (2023). *Rlaif: Scaling reinforcement learning from human feedback with ai feedback*.
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., & Legg, S. (2018). Scalable agent alignment via reward modeling: a research direction. *arXiv preprint arXiv:1811.07871*.
- Liang, P. P., Wu, C., Morency, L.-P., & Salakhutdinov, R. (2021). Towards understanding and mitigating social biases in language models. In *International conference on machine learning* (pp. 6565–6576).
- Maskin, E., & Sen, A. (2014). *The arrow impossibility theorem*. Columbia University Press.
- McClure, S. M., Daw, N. D., & Montague, P. R. (2003). A computational substrate for incentive salience. *Trends in neurosciences*, 26(8), 423–428.
- Menick, J., Trebacz, M., Mikulik, V., Aslanides, J., Song, F., Chadwick, M., ... others (2022). Teaching language models to support answers with verified quotes. *arXiv preprint arXiv:2203.11147*.
- Milli, S., Hadfield-Menell, D., Dragan, A., & Russell, S. (2017). Should robots be obedient? *arXiv preprint arXiv:1705.09990*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., ... others (2021). Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Nardo, C. (2023). The waluigi effect (mega-post). *Less Wrong*. Retrieved from <https://www.lesswrong.com/posts/D7PumeYTDpfBTp3i7/the-waluigi-effect-mega-post> (Accessed: 2023-09-11)
- Ng, A. Y., Harada, D., & Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml* (Vol. 99, pp. 278–287).
- Ng, A. Y., Russell, S., et al. (2000). Algorithms for inverse reinforcement learning. In *Icml* (Vol. 1, p. 2).
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., ... others (2022). Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Pettigrew, R. (2019). *Choosing for changing selves*. Oxford University Press.
- Prabhakaran, V., Davani, A. M., & Diaz, M. (2021). On releasing annotator-level labels and information in datasets. *arXiv preprint arXiv:2110.05699*.
- Prasad, M. (2018). Social choice and the value alignment problem. *Artificial intelligence safety and security*, 291–314.
- Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C. D., & Finn, C. (2023). Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*.
- Ramsey, F. P. (2016). Truth and probability. *Readings in Formal Epistemology: Sourcebook*, 21–45.
- Salha, N. (2011). *Aesthetics & art in the early development of human-computer interfaces* (Unpublished doctoral dissertation). Tesis de doctorado en ingeniería, Universität Bremen]. <https://bit.ly/3ZICKZJ>.
- Schulman, J. (2023). *Proxy objectives in reinforcement learning from human feedback*. Retrieved from <https://icml.cc/virtual/2023/invited-talk/21549> (International Conference on Machine Learning (ICML))
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Schulman, J., Zoph, B., Kim, C., & more. (2022). *Chatgpt: Optimizing language models for dialogue*. <https://openai.com/blog/chatgpt/>. (Accessed: 2023-02-12)
- Sen, A. (1973). Behaviour and the concept of preference. *Economica*, 40(159), 241–259.
- Shen, L., Chen, S., Song, L., Jin, L., Peng, B., Mi, H., ... Yu, D. (2023). The trickle-down impact of reward (in-) consistency on rlhf. *arXiv preprint arXiv:2309.16155*.
- Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., & Anderson, R. (n.d.). *The Curse of Recursion: Training on Generated Data Makes Models Forget*. Retrieved 2023-08-01, from <http://arxiv.org/abs/2305.17493> doi: 10.48550/arXiv.2305.17493
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., ... others (2017). Mastering the game of go without human knowledge. *nature*, 550(7676), 354–359.
- Silver, D., Singh, S., Precup, D., & Sutton, R. S. (2021). Reward is enough. *Artificial Intelligence*, 299, 103535.
- Singh, S., Lewis, R. L., & Barto, A. G. (2009). Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society* (pp. 2601–2606).

- Spaan, M. T. (2012). Partially observable markov decision processes. In *Reinforcement learning: State-of-the-art* (pp. 387–414). Springer.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., ... Christiano, P. F. (2020). Learning to summarize with human feedback. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, & H. Lin (Eds.), *Advances in neural information processing systems* (Vol. 33, pp. 3008–3021). Curran Associates, Inc. Retrieved from <https://proceedings.neurips.cc/paper/2020/file/1f89885d556929e98d3ef9b86448f951-Paper.pdf>
- Sutskever, J. L. I. (2023). *Introducing superalignment*. <https://openai.com/blog/introducing-superalignment>. (Accessed: 2023-09-27)
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Tesauro, G., et al. (1995). Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3), 58–68.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., ... others (2023). Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*, 2nd rev.
- Wang, T., Yu, P., Tan, X. E., O’Brien, S., Pasunuru, R., Dwivedi-Yu, J., ... Celikyilmaz, A. (2023). Shepherd: A critic for language model generation. *arXiv preprint arXiv:2308.04592*.
- Watkins, C. J., & Dayan, P. (1992). Q-learning. *Machine learning*, 8, 279–292.
- Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits* (Tech. Rep.). Stanford Univ Ca Stanford Electronics Labs.
- Wirth, C., Akrou, R., Neumann, G., Fürnkranz, J., et al. (2017). A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18(136), 1–46.
- Wu, J., Ouyang, L., Ziegler, D. M., Stiennon, N., Lowe, R., Leike, J., & Christiano, P. (2021). Recursively summarizing books with human feedback. *arXiv preprint arXiv:2109.10862*.
- Wu, Z., Hu, Y., Shi, W., Dziri, N., Suhr, A., Ammanabrolu, P., ... Hajishirzi, H. (2023). Fine-grained human feedback gives better rewards for language model training. *arXiv preprint arXiv:2306.01693*.
- Yong, Z.-X., Menghini, C., & Bach, S. H. (2023). Low-resource languages jailbreak gpt-4. *arXiv preprint arXiv:2310.02446*.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., ... others (2023). Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., ... Irving, G. (2019). Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.
- Zou, A., Wang, Z., Kolter, J. Z., & Fredrikson, M. (2023). Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*.