

---

# Learning Goal-Conditioned Representations for Language Reward Models

---

Vaskar Nath<sup>\*†</sup>    Dylan Slack<sup>\*</sup>    Jeff Da

Yuntao Ma    Hugh Zhang    Spencer Whitehead<sup>‡</sup>    Sean Hendryx<sup>‡</sup>

Scale AI

## Abstract

Techniques that learn improved representations via offline data or self-supervised objectives have shown impressive results in traditional reinforcement learning. Nevertheless, it is unclear how improved representation learning can benefit reinforcement learning from human feedback on language models. In this work, we propose training reward models (RMs) in a contrastive, *goal-conditioned* fashion by increasing the representation similarity of future states along sampled preferred trajectories and decreasing the similarity along randomly sampled dispreferred trajectories. This objective significantly improves reward model performance by up to 0.09 AUROC across challenging benchmarks, such as MATH and GSM8k. These findings extend to general alignment as well – on the Helpful-Harmless dataset, we observe 2.3% increase in accuracy. Beyond improving reward model performance, we show this way of training RM representations enables improved *steerability* because it allows us to evaluate the likelihood of an action achieving a particular goal-state (*e.g.*, whether a solution is correct or helpful). Leveraging this insight, we find that we can filter up to 55% of generated tokens during majority voting by discarding trajectories likely to end up in an “incorrect” state, which leads to significant cost savings. We additionally find that these representations can perform fine-grained control by conditioning on desired future goal-states. For example, we show that steering a Llama 3 model towards helpful generations with our approach improves helpfulness by 9.6% over a supervised-fine-tuning trained baseline. Similarly, steering the model towards complex generations improves complexity by 21.6% over the baseline. Overall, we find that training RMs in this contrastive, goal-conditioned fashion significantly improves performance and enables model steerability.<sup>1</sup>

## 1 Introduction

Aligning Language Models (LMs) with human preferences has proven to be an essential step for the adoption and safe use of these models, with the dominant paradigm being Reinforcement Learning from Human Feedback (RLHF) [50]. To accomplish this, a standard setup is to collect labels from humans for generated responses (*e.g.*, preferences, quality ratings) [50]. These labels can then be used to train a reward model to produce a ranking/scoring of a given sequence or set of sequences.

---

<sup>\*</sup>Equal contribution

<sup>†</sup>Corresponding author: [vaskar.nath@scale.com](mailto:vaskar.nath@scale.com)

<sup>‡</sup>Equal senior authorship

<sup>1</sup>Code available at <https://github.com/vaskarnathscale/goal-conditioned-rm>

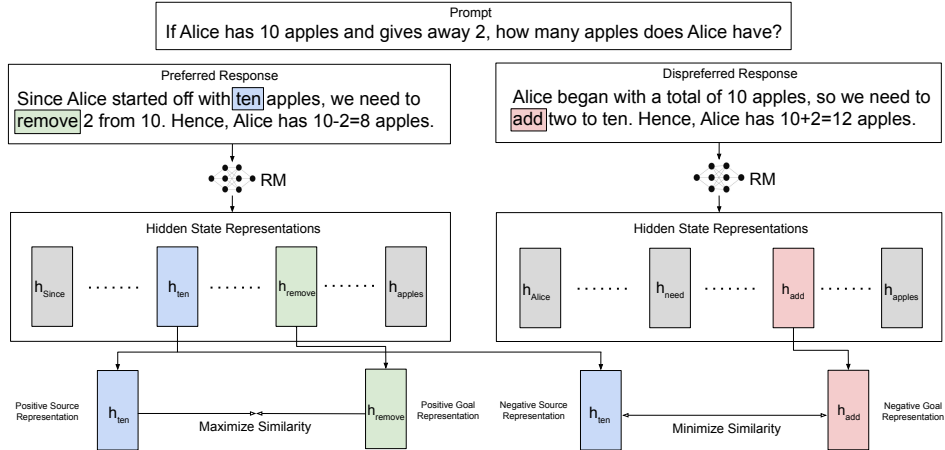


Figure 1: **Overview of contrastive goal-conditioned learning for text.** Pictured is a prompt with a preferred and dispreferred response. Both source state tokens (ten) for the positive and negative trajectory are sampled from the preferred response. For illustrative purposes, the positive and negative source states are sampled as the same token, but in practice they can be different. The positive goal state is sampled as some future token (subtract) from the preferred response, and the negative goal state is sampled from any token (add) from the dispreferred response. The corresponding representations are retrieved from the last hidden state of the reward model. The training objective is then to maximize and minimize the similarity of the positive and negative representation pairs, respectively.

The policy LM is then trained to maximize the expected returns from this reward model using a Reinforcement Learning (RL) algorithm.

High-quality representations have been shown to be an important piece for the success of RL algorithms [6, 42]. Although such representations can be learned during end-to-end training, many efforts have found it important to integrate more explicit representation learning components into RL algorithms, such as via data augmentation [39] or auxiliary losses [21]. Some work even casts certain RL algorithms as representation learning methods where using the similarity between state representations serves as a value function, demonstrating success on manipulation and navigation tasks [20]. Despite these successes in different areas, representation learning for aligning LMs has been less explored, while more emphasis has been placed on, e.g., pre-training reward models [7, 37] or learning from different types of rewards [15, 38].

In this paper, we present a simple yet effective approach to improve the representations learned by reward models for aligning LMs. We train LM-based reward models to learn representations that capture the expected reward or likelihood of achieving a goal state (*e.g.*, correct solution to a problem, helpful response) at intermediate steps of the input sequence, inspired by goal-conditioned RL [5, 13, 20]. To do so, we use a contrastive objective and apply it to the reward model’s hidden representations from desirable and undesirable sequences. Enforcing this loss on representations from intermediate steps of the sequence helps encode a dense signal as to which trajectories are more promising at different points in the sequence, which we show offers several useful properties, such as helping to localize errors or evaluating partial completed sequences. This method is flexible enough to support different kinds of alignment data and does not require further annotations beyond common sequence-level annotations.

We find that this approach improves the reward model’s ability to identify correct/incorrect solutions in mathematical reasoning, boosting the AUROC on the task by up to 0.09 over standard preference ranking training. Towards natural language alignment, we find this method is able to increase the reward model’s accuracy of identifying helpful versus harmless responses by 2.3% on the Helpful-Harmless dataset [8].

We also demonstrate the utility of the learned representations themselves, *e.g.*, for filtering solutions to improve accuracy and steering the outputs towards responses with certain attributes in guided decoding. On mathematical reasoning, we show that we are able to filter up to 55% of generated

tokens by discarding trajectories that are likely to lead to incorrect solutions as deemed by the learned representations while achieving similar or better performance. Similarly, using these representations to steer a Llama 3 [1] model by conditioning on desired future goal-states, we improve helpfulness by 9.6%, correctness by 12.2%, coherence by 16.5%, and complexity by 21.6% over the supervised-fine-tuning trained baseline.

In summary, our contributions are as follows: 1) We explore improving the learned representations of reward models and its effect on LM alignment. Towards this, we present a simple and effective representation learning method based on a goal-conditioned contrastive objective. 2) We demonstrate that training reward models with this method can improve reward model performance on mathematical reasoning and helpfulness/harmlessness benchmarks. 3) We show that simply utilizing a reward model trained with this method can improve policy LM alignment on math reasoning benchmarks. 4) We investigate using these representations as a mechanism to evaluate the likelihood of achieving a desired goal state by filtering generations in a majority-vote scheme and guided decoding, showing that they can be used to increase accuracy and control policy LM outputs.

## 2 Related Work

### **Language model (LM) alignment and Reinforcement Learning from Human Feedback (RLHF).**

The goal of RLHF [50, 71, 76] for aligning LMs is to train a policy LM that generates responses that are preferred by humans. This framework has proven instrumental for increasing the utility and generalization of LMs, improving their capabilities on a wide variety of tasks, and their safety [8, 9, 45, 61, 49]. Critically, RLHF improves the policy LM along a scalar *reward* signal, which indicates the value of a generated piece of text, and their efforts have focused on different means of providing such signals [7, 56]. Typically, researchers train a *reward model* that learns to score generations based on pair-wise human preference data. Subsequently, the model provides reward signals during reinforcement learning (RL) training of the policy, often via Proximal Policy Optimization (PPO) [59]. While reward models have often been trained to score text with just a single scalar value, several works have shown that training the reward model to provide fine-grained reward signals over *spans* of text can help improve model performance by providing intermediate rewards for the policy LM to learn from during RL [43, 68]. Nevertheless, fine-grained reward methods often require extensive and costly human annotations on the sentence or token level to train the model, which can be prohibitive.

**Representation learning for RL.** Reinforcement learning (RL) methods often learn representations as part of policy learning from reward signals. However, learning representations from solely reward signals can be highly sample inefficient, leading to poor generalization, and struggles in high-dimension state and action settings [41, 63, 74]. Consequently, techniques that learn improved representations have been shown to significantly improve RL performance [73, 19, 25, 40, 24]. Typically, these methods construct additional supervision via offline data or self-supervised objectives to drive representation learning [12, 11]. These methods have achieved significant success in a variety of settings, such as vision-based control and multi-task RL [10, 44, 30]. One such way of constructing supervision for representation learning is by leveraging future states and actions sampled from offline trajectories [48, 18, 60]. In particular, prior work has shown that increasing the similarity between state-action pairs sampled from the same trajectory can serve as a simple yet powerful way to learn improved representations for RL, that additionally acts as a type of goal-conditioning by making representations more similar along the same trajectory [20, 22]. Still, representation and goal-conditioned RL largely have not been considered for natural language domains.

## 3 Method

Drawing on recent efforts on goal-conditioned reinforcement learning [5, 13] and representation learning [20, 31] in other areas, we show that we can learn representations that approximate a Q-function within a language reward model using contrastive learning. Since the Q-function quantifies the expected cumulative future rewards of taking a specific action from a given state [47], forcing the reward model’s output scores to be entangled with an approximated Q-function has the potential to improve the reward model’s credit assignment and preference ranking accuracy. We first present preliminaries on reward modeling and contrastive learning in Section 3.1 and present our method in Section 3.2.

### 3.1 Preliminaries

**Preference ranking reward modeling for LMs.** Typically, a reward model is parameterized  $r(x, y) \rightarrow \mathbb{R}$  to return a scalar reward, given prompt  $x$  and completion sequence of tokens  $y = [y_0, \dots, y_T]$ . Given a dataset  $\mathcal{D}$  consisting of triples  $(x, y^w, y^\ell)$  of preferred and dispreferred completions,  $y^w$  and  $y^\ell$ , respectively, the reward model is optimized via a paired loss:

$$\mathcal{L}^R = -\frac{1}{|\mathcal{D}|} \mathbb{E}_{(x, y^w, y^\ell) \in \mathcal{D}} \log(\sigma(r(x, y^w) - r(x, y^\ell))). \quad (1)$$

Note, the reward model,  $r(x, y)$ , is trained to provide a scalar feedback on the *entire* completion  $y$ .

**Approximating Q-functions via contrastive learning.** In deep reinforcement learning, the Q-function quantifies the expected cumulative future rewards of taking an action from a given state and is parameterized by a neural network [47, 20]. Let  $s_t$  be the current state from which a model can take an action  $a_t$ . Taking action  $a$  at the current state and continuing on that trajectory will lead to some future state  $s_k$ . The future states reached along this trajectory can be positive,  $s_k^+$ , by reaching a target goal state or negative,  $s_k^-$ , by reaching a different or undesired state. A critic function  $f$  can approximate a Q-function up to a multiplicative constant by optimizing a contrastive learning objective [20]:

$$\mathcal{L} = \log(\sigma(f(s_t, a_t, s_k^+))) + \log(1 - \sigma(f(s_t, a_t, s_k^-))). \quad (2)$$

Here, the critic function  $f$  is a scoring function between a state-action encoder and state encoder.

### 3.2 Learning Goal-Conditioned Representations in Language Reward Models

In this section, we describe how to efficiently learn goal-conditioned representations using decoder-only transformer models [64, 54]. Additionally, we describe how to use these representations as Q-values that can be leveraged to detect generation errors and in guided decoding [33].

**Reward modeling through the bottleneck of a Q-function.** Given the prompt  $x$  and a corresponding completion sequence  $y$ , we seek to learn a language reward model that both scores  $y$  and has representations that can be used to approximate a goal-conditioned Q-function. This Q-function  $Q_{y_g}^\pi(y_{\{0, \dots, t\}}, y_{t+1})$ ,  $t \in [0, \dots, T]$ , captures the expected reward for a policy  $\pi(y_{t+1} | y_{\{0, \dots, t\}}, x)$ , for some goal state  $y_g$  and next token  $y_{t+1}$ , where rewards are defined as likelihood of achieving the goal at the next token generation. Intuitively, the goal state could be an optimal generation from the policy that a human would find satisfactory, such as the correct solution to a math problem or a helpful response. The Q-function captures the expected rewards from choosing token  $y_{t+1}$  at the current step, considering the prompt  $x$  and tokens generated so far  $y_{\{0, \dots, t\}}$ .

We initialize our reward model,  $r$ , from an instruction-tuned causal LM. We consider this reward model to be decomposed into two components: 1) feature extractor  $\phi$ , which produces a representation for a given token; 2) a reward projection head  $r'$  that predicts a reward score based on that token representation. We use  $\phi$  to obtain representations of our state-action pair,  $h_{t+1}$ , and arbitrary future states,  $h_k$ :

$$h_{t+1} = \phi(y_{t+1} | y_{\{0, \dots, t\}}, x) \quad \text{and} \quad h_k = \phi(y_k | y_{\{0, \dots, k-1\}}, x). \quad (3)$$

Hence, we parameterize our reward model as

$$r(x, y) = r'(\phi(y_T | y_{\{0, \dots, T-1\}}, x)). \quad (4)$$

In practice, the feature extractor is the LM decoder and the representations that we use are the hidden token representations prior to predicting an output token. Meanwhile, the reward projection is either a linear layer or multi-layered perceptron (MLP) in our experiments.

With this reward model, we train the representations from the feature extractor  $\phi$  using a contrastive loss [20], which becomes:

$$\mathcal{L}^C = \log(\sigma(f(y_{\{0, \dots, t\}}, y_{t+1}, y_g^+))) + \log(1 - \sigma(f(y_{\{0, \dots, t\}}, y_{t+1}, y_g^-))), \quad (5)$$

where  $f(y_{\{0, \dots, t\}}, y_{t+1}, y_k)$  is the cosine similarity between the representations of the state-action pair  $(y_{\{0, \dots, t\}}, y_{t+1})$  and a future state  $y_k$ . In Equation 5,  $y_g^+$  and  $y_g^-$  are our positive and negative

goal states, which we discuss how to obtain in practice later in this section. By forcing the reward scores predicted by  $r'$  to depend on the representations from  $\phi$ , which is learned via this contrastive objective, the reward score incorporates a signal as to which actions are closer to a desired goal state.

We jointly train the reward model on the contrastive objective in Eq. 5 and the preference ranking objective optimized by Eq. 1. Thus, our final objective is given by

$$\mathcal{L} = \mathcal{L}^R + \lambda \cdot \mathcal{L}^C, \quad (6)$$

where  $\lambda$  is a hyperparameter to balance the tradeoff between optimizing the contrastive loss and the rewards.

In the context of RLHF, given that we have preference labels between two or more generations for the same prompt and recalling that  $\mathcal{L}^C$  approximates a Q-function, we can improve the term’s alignment with human preferences by sampling positive goal states from preferred generations  $y^w$  and negative goal states from dispreferred generations  $y^\ell$ .

**Computing contrastive loss in practice.** Given a preference-ranking instance from a training batch, let the preferred completion,  $y^w$ , and dispreferred completion,  $y^\ell$ , have lengths  $T$  and  $T'$ , respectively. The representations of the positive state-action and goal state pair,  $(y_s, y_g^+)$ , are given by sampling two points,  $i, j \sim [0, T], i < j$ , from the preferred completion and retrieving hidden states  $h_i^w$  and  $h_j^w$ . Similarly, the representations of the negative pair,  $(y_{s'}, y_g^-)$ , are given by sampling any index  $m \sim [0, T]$  from the preferred completion and any index  $n \sim [0, T']$  from the dispreferred completion and then retrieving hidden states  $h_m^w$  and  $h_n^\ell$ . We note that source states  $h_i^w$  and  $h_m^w$  are sampled independently and are not necessarily equal. The cosine similarities between positive and negative pair representations are passed through the sigmoid function and averaged across the training batch. We will refer to this method as the Single Goal State (SGS) contrastive loss.

**Computing Q-values.** Once the reward model has been trained, we can use it to compute an additional *Q-value*, which gives us a measure proportional to the expected future rewards given the current state and a specific action. The Q-value at a token  $t$  is the cosine similarity between the trained reward model’s hidden state,  $h_t$ , and a goal state,  $h_g$ . The choice of goal state here is flexible. For our experiments in Section 4.1, unless specified otherwise, we set the goal state to be the mean representation of all the preferred completion in the reward model training set. More explicitly, the goal state is computed as:

$$h_g = \frac{1}{N} \sum_{y_w \in \mathcal{D}} h(y_w) \quad (7)$$

where  $h(\cdot)$  represents the reward model’s last hidden layer of the final token of the completion. In Appendix C, we ablate this choice and find this to be the best method for computing the goal state during inference. Thus, we do not require any additional annotations aside from the starting preference ranking dataset. In experiments on LM steering, we construct a prototype to be used as the goal state. The prototype is constructed using generations labeled as helpful, complex, coherent, and correct by annotators in Nvidia’s HelpSteer dataset [67]. Out of the highly-scored generations, we sample a subset and use the mean representations of the sample as the goal state for the reward model. We ablate these choices in Appendix D.2.

## 4 Experiments

To explore our method, we experiment with two settings: First, mathematical reasoning with code (Section 4.1), where LMs must use reasoning and coding to answer mathematical questions. The ability to code has become increasingly important for LLMs as it enables a broad set of capabilities for reasoning (e.g., tool use [58]). This line of experiments elucidates the ability of our method, and the resulting representations, to guide and improve step-by-step reasoning capabilities. Second, natural language alignment (Section 4.2), in which an LM’s responses must align with certain desired properties, such as helpfulness and harmlessness. These experiments show the capacity of reward models trained with our method to steer LM outputs in a more open-ended setting. In particular, we focus on helpfulness, harmlessness, and other desirable characteristics (e.g., correctness, coherence).

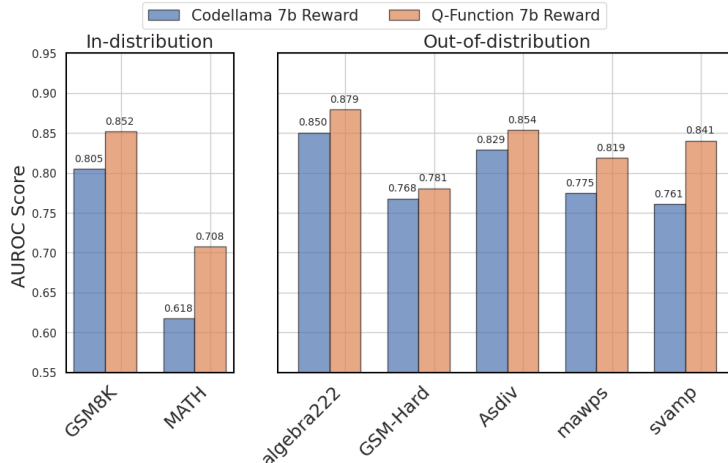


Figure 2: AUROC scores comparing the baseline Codellama 7b Reward vs. our proposed method Q-Function 7b Reward on the rewards attributed to the base-model greedy generations across several math benchmarks.

## 4.1 Mathematical Reasoning with Code

### 4.1.1 Experimental Setup

For training reward models, we use the OpenMathInstruct-1 dataset [62], a math instruction tuning dataset with 1.8M problem-solution pairs on the training sets of GSM8k and MATH, which are generated by the Mixtral-8x7B model [32]. To form our reward model training set, we create pairs of solutions for the same problem, where one solution is correct (*i.e.*, preferred) and the other is incorrect (*i.e.*, dispreferred).

Our base model is OpenMath-CodeLlama-7b-Python-hf [62], a strong decoder-only LM that has been instruction tuned on OpenMathInstruct-1 to use a Python interpreter in order to solve mathematical problems, which we add a reward head to predict reward scores (Section 3.2). This model trained with the traditional preference ranking objective serves as our baseline (**Codellama 7b Reward**), which we compare with the same model trained using our method (**Q-Function 7b Reward**).

We evaluate on several popular math benchmarks. Since the problems in the preference ranking dataset come from the training sets of GSM8k and MATH, we consider their respective test splits to be in-distribution (ID). We also evaluate on test sets we consider to be out-of-distribution (OOD), namely algebra222 [27], GSM-Hard [23], Asdiv [46], mawps [35], and svamp [52], to test for further generalization.

Detailed information on data, models, and training can be found in Appendix A.

### 4.1.2 Reward Modeling Evaluations

We obtain completions from the base model by greedy decoding, following prior work [62, 66], and measure the reward score AUROC. This quantifies the ability of a model’s reward score to discern between correct/incorrect sequences across different thresholds. In Appendix E, we additionally assess Q-Function 7b Reward’s ability to rank completions.

**Learned representations help reward scores discern correct/incorrect solutions.** In Figure 2, we see that the reward model trained with our representation learning approach consistently achieves higher AUROC scores compared to the model trained with only the standard preference ranking loss. On the ID datasets, we see a gain of approximately 0.05 and 0.09. Meanwhile, on the OOD datasets, the average improvement is 0.04, with maximum of roughly 0.08 on svamp. Although these datasets are math problems, they vary in the types and difficulties of questions asked. Despite being OOD, we see that our method improves the performance when judging these prompts and generations. Overall, these results suggest that the representations learned using our approach help improve the reward score predictions such that they are more accurate and generalizable.




Figure 3: AUROC scores on the rewards attributed to partial base-model generations across 50 samples on GSM8k and MATH. The error bars depict the 95% confidence intervals (with sample size  $n = 50$ ) at each percentile of generation considered. The Q-Function reward model has incremental increase in performance with more information, whereas, the traditional reward model’s performance is a lot more varied in attributing intermediate rewards.

**These representations also help reward scores judge which partial solutions are promising.** Part of the motivation and design of our approach is to encode the likelihood of reaching a certain goal state at intermediate steps in the generated sequence. To examine how well our method accomplishes this, we measure reward score AUROC when the model is only shown up to a certain percentile of the generated sequence (*i.e.*, percentile). Here, we show this on the ID datasets, where we consider the 50 sample generations from the base model and plot the mean AUROC across the 50 samples. Plots for other benchmarks can be found in Appendix F. Looking at Figure 3, we find that the AUROC scores for both reward models tend to increase as they see more of the generated sequence, but Q-Function 7b Reward largely maintains higher AUROC across different percentiles and generally exhibits fewer significant drops between percentiles. These improvements across percentiles imply that the reward scores do indeed better capture which generated sequence are likely to be correct or incorrect.

### 4.1.3 Directly Using the Learned Representations

Our experiments thus far illustrate the effect of the learned representations on the reward scores. Here, we turn our attention to utilizing these representations directly to explore some of the encoded properties. In these experiments, we filter generations from the base model using the Q-values (*i.e.*, cosine similarities between intermediate and goal state representations). We take 50 base model generations for all the GSM8k and MATH test problems, which are provided by Toshiniwal et al., and we filter any generation that has a Q-value less than zero [62]. We compare this against a baseline of majority voting over the 50 generations.

**Q-values can prune incorrect sequences.** We find that filtering based on Q-values provides meaningful improvements in accuracy compared to majority voting over all 50 samples, based on Table 1. Specifically, the majority vote accuracy improves by 1.2% and 3.9% on GSM8k and MATH, respectively. Looking at the increased proportion of correct sequences after filtering, it appears that sequences with the negative Q-values generally correspond to incorrect sequences and are filtered out in this setting, thereby improving the accuracy. On the other hand, we also see that the average percentage of sequences that are filtered out can be quite large, such as 65.6% on the MATH dataset, and some correct solutions can also be filtered out. Nevertheless, using the Q-values based on these learned representations can help improve the accuracy of such majority voting schemes.

**Q-values can detect errors in reasoning.** Experiments on rewards scores for partial sequences (Section 4.1.2) suggest that the learned representations help capture the likelihood that a sequence will reach a goal state from an intermediate step. We can also qualitatively explore this property using the Q-values and the representations themselves. In Figure 4, we show examples that are filtered based on the Q-value filtering process above. In the left example, the Q-value drops when there is a mistake of multiplying the “donuts\_price\_per\_dozen” by “donuts\_per\_dozen”. The Q-value in the right example similarly shifts once the “jazz\_dance\_percentage” is naively added to the

		GSM8k			
		Accuracy (%)	Prop. Correct (%)	Avg. Filtered (%)	Avg. Tokens Saved (%)
Maj. @ 50		84.8	74.9	-	-
Filtered		<b>86.0</b>	<b>84.0</b>	35.9	25.5
		MATH			
		Accuracy (%)	Prop. Correct (%)	Avg. Filtered (%)	Avg. Tokens Saved (%)
Maj. @ 50		55.6	40.2	-	-
Filtered		<b>59.6</b>	<b>52.0</b>	65.6	55.6

Table 1: Majority vote versus Q-value filtering performance. Accuracy shows the final majority voting accuracy on the respective benchmarks. Prop. Correct is the proportion of the correct class in the final sample set considered. Avg. Filtered is the average percentage (out of 50 total generations) that were discarded across all problems in the benchmark. Avg. Tokens Saved is the average percentage of tokens that are saved with the assumption that the remaining tokens after the first negative Q-value is discarded across all the problems in the benchmark. Both the difference in accuracy and proportion of correct class are t-test significant.

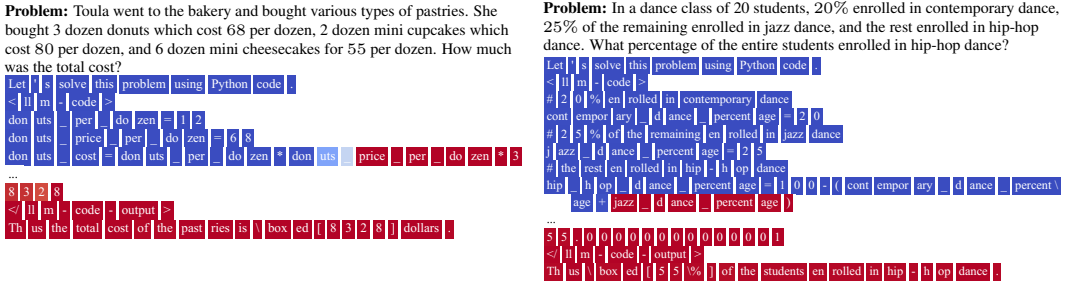


Figure 4: Two examples from the GSM8K dataset that was filtered via the Q-value pruning. The token level Q-values are portrayed as a heat map where the colors red and blue represents scores close to  $-1$  and  $1$ , respectively. Both examples illustrate that the Q-values pinpoint the exact logical error in reasoning. The full version of these examples can be found in Appendix G

“contemporary\_dance\_percentage”. Overall, these examples illustrate the potential effectiveness of using the representations and resulting Q-values for localizing errors and identifying promising sequences.

#### 4.1.4 Aligning Policy Models with RLHF

We investigate whether the reward model trained with our proposed method provides a better reward signal when performing RLHF, which is an important usage of these reward models. We compare the math problem solving performance of a policy model that is aligned using a preference-ranking reward model versus our Q-Function reward model in the Proximal Policy Optimization (PPO) [59] algorithm. We additionally compare the performance against DPO as it serves as a strong baseline [55]. The prompt dataset is given by problems in the training sets of GSM8k and MATH. To introduce diversity in the prompts, which has been shown to improve PPO, we also add problems from the MathQA dataset - a large scale dataset with real world math problems [4]. We keep the settings and hyperparameters during PPO constant, and further details are in Appendix A.

Table 2 shows that, given the same settings, utilizing a reward model trained with our method improves policy model accuracy across several benchmarks. We see an average gain, weighted by benchmark size, of 1.7%. Intuitively, this shows that the improvements we observe in reward scores (Section 4.1.2) do translate to improvements in the policy model via RL, which further supports the value of the representations learned via our approach. Although, the gains in accuracy are not as significant compared to those observed with the reward scores. A possible explanation for this is that, due to experimental constraints such as additional data collection, the preference-ranking dataset used for reward model training are off-policy Mixtral-8x7B model [32] generations. Hence, a clear next step for this setting would be to retrain the reward model with on-policy preference-rankings and iterate. Interestingly, we observe that DPO performs on par with our method for the ID tasks



(GSM8k and MATH). However, the policy trained with the contrastive RM performs much better on OOD tasks (algebra222, GSM-Hard, Asdiv, mawps, and svamp) when compared to DPO. These findings could be explained by related works that analyze DPO and find it can perform poorly in OOD settings [72]. All in all, these results indicate that PPO training with the contrastive RM leads to better generalization than training with DPO, particularly in OOD settings.

Model	ID		OOD				
	GSM8k	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
Codellama Base	75.9	43.6	65.6	60.1	77.7	93.5	79.6
DPO	80.1	44.8	64.8	59.9	76.9	90.3	76.6
Codellama PPO	79.3	43.4	65.8	61.1	77.4	91.6	78.5
Q-Function PPO	<b>80.5</b>	<b>45.1</b>	<b>70.9</b>	<b>62.7</b>	<b>79.5</b>	<b>93.6</b>	<b>81.2</b>

Table 2: Accuracy of the policy model trained via PPO using a preference ranking reward model vs our Q-Function reward model. We present average accuracy across 4 independent runs. The base model results are also shown as a reference point presented in [62] and DPO results are shown as a strong baseline. The full table with confidence intervals is in the appendix (Table 6).

## 4.2 Natural Language Alignment

While fine-grained feedback (and subsequently, Q-values) are most sensible in reasoning environments with explicitly incorrect logical steps, a natural extension of the experiment is to apply our approach to reward models in the natural language setting, where we expect models to learn improved representations with respect to their alignment towards human preferences. We first explore this for training a reward model for helpfulness and harmlessness [8]. Then, we explore if the learned representations can be further used for steering language model at inference time, via guided decoding, to generate sequences that are more helpful, complex, or coherent. The backbone of our reward model is Llama 3-8B-Instruct [1], and detailed information can be found in Appendix B.2.

### 4.2.1 Improving Reward Modeling for Helpfulness and Harmlessness

We use the Helpful-Harmless dataset [8] for training and evaluation. We select this dataset because of its large scale, established representation in the RLHF community, and general purpose nature which is helpful to avoid restricting our natural language reward model towards a specific domain. Similar to our previous experiments, we compare training with our approach versus standard reward model training. Table 3 (left) shows that training with our method also improves reward scores on this natural language data, increasing the ability of the reward model to discern between helpful and unhelpful responses.

### 4.2.2 Reward Model Representations can Steer Language Models Towards Helpful Generations

A practical extension of learning stronger natural language representations is using the extensions for steering the LM. To test this, we experiment with guided decoding [33, 70] with Llama-3-8B-Instruct, which defines a decoding objective function to guide inference during beam search. In traditional guided decoding (and our SFT baseline), the confidence score of the model is extracted from an additional inference step on each beam. In our method (Q-Function Prototype), the confidence score is the cosine similarity between the sequence embedding from the reward model and the hidden state [20, 2]. We construct the prototype by embedding 20 examples that score highly in all categories (helpfulness, correctness, coherence, complexity) using HelpSteer, a model alignment and steering dataset [67]. The reward model is trained using HelpfulHarmless as per previous sections [8]. Details for guided decoding are in the Appendix B.

Table 3 (right) shows our results, and evaluation is done via LLM-as-a-judge (GPT-4) [75]. Using the representations fine-tuned by the contrastive loss improves across all categories and generally steers generations towards more helpful generations. We find that complexity improves the most, suggesting that the generations produced by the using the prototype helps the model avoid beams that are too simple. Surprisingly, the correctness of the responses is increased also, suggesting that

although the SFT prototype usually prefers beams that are simple and correct, the prototype succeeds at picking beams with added complexity without sacrificing correctness or helpfulness.

Model	HH	Model	Helpful	Correct.	Coherent	Complex.
Llama 8b Reward	68.2	SFT	45.2	43.9	41.8	39.2
Q-Function 8b Reward	<b>70.5</b>	Prototype	<b>54.8</b>	<b>56.1</b>	<b>58.3</b>	<b>60.8</b>

Table 3: Comparison of model performances in different settings. (a) Reward model accuracy on the Helpful-Harmless test set. Our result highlights that the representations learned by the Q-function are helpful for discerning between desirable and undesirable sequences in the natural language setting. (b) Winrate (%) vs. SFT model. In the Q-Function prototype setting, we use the cosine similarity of the sequence and the prototype (made from 20 examples that score high on all 4 categories) to score each beam. In the SFT setting, the model is fine-tuned on the same examples as the prototype and guided decoding is run via model scores. Evaluation is done via GPT-4.

## 5 Limitations and Future Work

**Deriving more informative goal states.** In Appendix D.2, we demonstrate the importance of having a precise and meaningful goal state. At train time, our SGS method achieves this by looking at future tokens of the preferred and dispreferred completions. During inference, however, we take the mean representation of multiple goal state completions from the training data. Although we see that the SGS method generalizes despite this discrepancy, future work could investigate other methods to derive more precise inference-time goal states.

**Partial Q-values correlate with partial reward scores.** An interesting observation, shown in Appendix H, is that the reward scores for *partial* completions are correlated with the Q-values. While our experiments show that this can benefit reward scores, especially for math reasoning, disentangling these values may be desirable in some instances. For example, one may want to incorporate the completeness of the generated sequence in the reward score, such as for creative writing, meaning incomplete sequences should strictly get low reward scores. However, currently, partial sequences may still score highly if they are likely to reach a goal state. This is expected as the final reward score is derived by a linear projection from the hidden representation, while the Q-values are derived as the cosine similarity between that same representation and goal state. In Appendix H, we explore the effects of learning an MLP projection for the rewards to help decouple the two signals. Further disentangling these signals in the representations could be an interesting direction for future work.

**Further advancing policy models.** A clear line of future work here is to integrate the additional Q-value signal that we gain from this method into the RL algorithm to further improve policy model training. Future work should also investigate the impacts of iteratively retraining the reward model on on-policy completions and the policy model on the updated reward model.

## 6 Conclusion

Learning proper representations in the reward model is essential for the model’s generalization and ability to produce accurate rewards during RLHF. We introduce a method of applying goal-conditioned Q-functions to learn representations that capture the expected reward for a goal-conditioned policy. On math benchmarks, such as GSM8k, we improve performance and show that the reward model has a greater ability to discern between correct and incorrect solutions. On natural language alignment experiments, we show improvement in reward model performance and that the embedding can be used for natural language steering. Our findings show that using Q-values during reward model training can improve the representations of the reward model and suggest promising future directions for further advancing language models via RLHF.

## References

- [1] Meta AI. Introducing meta llama 3: The most capable openly available llm to date. <https://ai.meta.com/blog/meta-llama-3/>.

- [2] Arwa Alanqary, Gloria Lin, Joie Le, Zhi-Xuan Tan, Vikash K. Mansinghka, and Joshua B. Tenenbaum. Modeling the mistakes of boundedly rational agents within a bayesian theory of mind. *ArXiv*, abs/2106.13249, 2021.
- [3] Reza Yazdani Aminabadi, Samyam Rajbhandari, Minjia Zhang, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Jeff Rasley, Shaden Smith, Olatunji Ruwase, and Yuxiong He. Deep-speed inference: Enabling efficient inference of transformer models at unprecedented scale, 2022.
- [4] Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. *CoRR*, abs/1905.13319, 2019.
- [5] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *Advances in neural information processing systems*, 30, 2017.
- [6] Raghuram Mandyam Annasamy and Katia Sycara. Towards better interpretability in deep q-networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4561–4569, 2019.
- [7] Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*, 2021.
- [8] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [9] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [10] Andrea Banino, Adria Puigdomenech Badia, Jacob C Walker, Tim Scholtes, Jovana Mitrovic, and Charles Blundell. CoBERL: Contrastive BERT for reinforcement learning. In *International Conference on Learning Representations*, 2022.
- [11] André Barreto, Will Dabney, Rémi Munos, Jonathan J. Hunt, Tom Schaul, David Silver, and H. V. Hasselt. Successor features for transfer in reinforcement learning. *ArXiv*, abs/1606.05312, 2016.
- [12] Diana Borsa, André Barreto, John Quan, Daniel Jaymin Mankowitz, Rémi Munos, H. V. Hasselt, David Silver, and Tom Schaul. Universal successor features approximators. *ArXiv*, abs/1812.07626, 2018.
- [13] Elliot Chane-Sane, Cordelia Schmid, and Ivan Laptev. Goal-conditioned reinforcement learning with imagined subgoals. In *International Conference on Machine Learning*, pages 1430–1440. PMLR, 2021.
- [14] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [15] Ching-An Cheng, Andrey Kolobov, Dipendra Misra, Allen Nie, and Adith Swaminathan. Llf-bench: Benchmark for interactive learning from language feedback. *arXiv preprint arXiv:2312.06853*, 2023.
- [16] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [17] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning, 2023.

- [18] Alexey Dosovitskiy and Vladlen Koltun. Learning to act by predicting the future. *ArXiv*, abs/1611.01779, 2016.
- [19] Simon S. Du, Sham M. Kakade, Ruosong Wang, and Lin F. Yang. Is a good representation sufficient for sample efficient reinforcement learning? In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [20] Benjamin Eysenbach, Tianjun Zhang, Sergey Levine, and Ruslan Salakhutdinov. Contrastive learning as goal-conditioned reinforcement learning. *Advances in Neural Information Processing Systems*, 2022.
- [21] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 512–519. IEEE, 2016.
- [22] Carlos Florensa, Jonas Degraeve, Nicolas Manfred Otto Heess, Jost Tobias Springenberg, and Martin A. Riedmiller. Self-supervised learning of image embedding for continuous control. *ArXiv*, abs/1901.00943, 2019.
- [23] Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. Pal: Program-aided language models. *arXiv preprint arXiv:2211.10435*, 2022.
- [24] Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G. Bellemare. Deep-MDP: Learning continuous latent space models for representation learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2170–2179. PMLR, 09–15 Jun 2019.
- [25] Dibya Ghosh and Marc G. Bellemare. Representations for stable off-policy reinforcement learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3556–3565. PMLR, 13–18 Jul 2020.
- [26] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Pícus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.
- [27] Joy He-Yueya, Gabriel Poesia, Rose E Wang, and Noah D Goodman. Solving math word problems by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*, 2023.
- [28] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *Advances in neural information processing systems*, 33, 2021.
- [29] Jian Hu, Xibin Wu, Weixun Wang, Xianyu, Dehao Zhang, and Yu Cao. Openrlhf: An easy-to-use, scalable and high-performance rlhf framework, 2024.
- [30] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations*, 2017.
- [31] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020.
- [32] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

- [33] Minbeom Kim, Hwanhee Lee, Kang Min Yoo, Joonsuk Park, Hwaran Lee, and Kyomin Jung. Critic-guided decoding for controlled text generation. *arXiv preprint arXiv:2212.10938*, 2022.
- [34] Hannah Rose Kirk, Alexander Whitefield, Paul Rottger, Andrew M. Bean, Katerina Margatina, Juan Ciro, Rafael Mosquera, Max Bartolo, Adina Williams, He He, Bertie Vidgen, and Scott A. Hale. The prism alignment project: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models. *ArXiv*, abs/2404.16019, 2024.
- [35] Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. Mawps: A math word problem repository. pages 1152–1157, 01 2016.
- [36] Andreas Kopf, Yannic Kilcher, Dimitri von Rutte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Rich’ard Nagyfi, ES Shahul, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. Openassistant conversations - democratizing large language model alignment. *ArXiv*, abs/2304.07327, 2023.
- [37] Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Vinayak Bhalerao, Christopher Buckley, Jason Phang, Samuel R Bowman, and Ethan Perez. Pretraining language models with human preferences. In *International Conference on Machine Learning*, pages 17506–17533. PMLR, 2023.
- [38] Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. *arXiv preprint arXiv:2303.00001*, 2023.
- [39] Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. 33:19884–19895, 2020.
- [40] Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in RL with a generative model. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5662–5670. PMLR, 13–18 Jul 2020.
- [41] Charline Le Lan, Stephen Tu, Adam Oberman, Rishabh Agarwal, and Marc G. Bellemare. On the generalization of representations in reinforcement learning. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 4132–4157. PMLR, 28–30 Mar 2022.
- [42] Yitao Liang, Marlos C Machado, Erik Talvitie, and Michael Bowling. State of the art control of atari games using shallow reinforcement learning. *arXiv preprint arXiv:1512.01563*, 2015.
- [43] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step, 2023.
- [44] Guoqing Liu, Chuheng Zhang, Li Zhao, Tao Qin, Jinhua Zhu, Li Jian, Nenghai Yu, and Tie-Yan Liu. Return-based contrastive representation learning for reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [45] Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. Teaching language models to support answers with verified quotes. *arXiv preprint arXiv:2203.11147*, 2022.
- [46] Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing english math word problem solvers, 2021.
- [47] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [48] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In *Neural Information Processing Systems*, 2018.

- [49] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- [50] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- [51] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019.
- [52] Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are nlp models really able to solve simple math word problems?, 2021.
- [53] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [54] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [55] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2024.
- [56] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.
- [57] Joshua Robinson, Ching-Yao Chuang, Suvrit Sra, and Stefanie Jegelka. Contrastive learning with hard negative samples. *arXiv preprint arXiv:2010.04592*, 2020.
- [58] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 35, 2023.
- [59] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [60] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, and Sergey Levine. Time-contrastive networks: Self-supervised learning from video. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1134–1141, 2017.
- [61] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- [62] Shubham Toshniwal, Ivan Moshkov, Sean Narenthiran, Daria Gitman, Fei Jia, and Igor Gitman. Openmathinstruct-1: A 1.8 million math instruction tuning dataset. *arXiv preprint arXiv:Arxiv-2402.10176*, 2024.
- [63] Frederik Trüble, Andrea Dittadi, Manuel Wuthrich, Felix Widmaier, Peter Vincent Gehler, Ole Winther, Francesco Locatello, Olivier Bachem, Bernhard Schölkopf, and Stefan Bauer. Representation learning for out-of-distribution generalization in reinforcement learning. In *ICML 2021 Workshop on Unsupervised Reinforcement Learning*, 2021.
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- [65] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, Aditya Vijaykumar, Alessandro Pietro Bardelli, Alex Rothberg, Andreas Hilboll, Andreas Kloeckner, Anthony Scopatz, Antony Lee, Ariel Rokem, C. Nathan Woods, Chad Fulton, Charles Masson, Christian Häggström, Clark Fitzgerald, David A. Nicholson, David R. Hagen, Dmitrii V. Pasechnik, Emanuele Olivetti, Eric Martin, Eric Wieser, Fabrice Silva, Felix Lenders, Florian Wilhelm, G. Young, Gavin A. Price, Gert-Ludwig Ingold, Gregory E. Allen, Gregory R. Lee, Hervé Audren, Irvin Probst, Jörg P. Dietrich, Jacob Silterra, James T Webber, Janko Slavič, Joel Nothman, Johannes Buchner, Johannes Kulick, Johannes L. Schönberger, José Vinícius de Miranda Cardoso, Joscha Reimer, Joseph Harrington, Juan Luis Cano Rodríguez, Juan Nunez-Iglesias, Justin Kuczynski, Kevin Tritz, Martin Thoma, Matthew Newville, Matthias Kümmerer, Maximilian Bolingbroke, Michael Tartre, Mikhail Pak, Nathaniel J. Smith, Nikolai Nowaczyk, Nikolay Shebanov, Oleksandr Pavlyk, Per A. Brodtkorb, Perry Lee, Robert T. McGibbon, Roman Feldbauer, Sam Lewis, Sam Tygier, Scott Sievert, Sebastiano Vigna, Stefan Peterson, Surhud More, Tadeusz Pudlik, Takuya Oshima, Thomas J. Pingel, Thomas P. Robitaille, Thomas Spura, Thouis R. Jones, Tim Cera, Tim Leslie, Tiziano Zito, Tom Krauss, Utkarsh Upadhyay, Yaroslav O. Halchenko, and Yoshiki Vázquez-Baeza. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature Methods*, 17(3):261–272, February 2020.
- [66] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- [67] Zhilin Wang, Yi Dong, Jiaqi Zeng, Virginia Adams, Makesh Narsimhan Sreedhar, Daniel Egert, Olivier Delalleau, Jane Polak Scowcroft, Neel Kant, Aidan Swope, and Oleksii Kuchaiev. Helpsteer: Multi-attribute helpfulness dataset for steerlm. *ArXiv*, abs/2311.09528, 2023.
- [68] Zeqiu Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A. Smith, Mari Ostendorf, and Hannaneh Hajishirzi. Fine-grained human feedback gives better rewards for language model training, 2023.
- [69] Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, Xu Zhao, Min-Yen Kan, Junxian He, and Qizhe Xie. Decomposition enhances reasoning via self-evaluation guided decoding, 2023.
- [70] Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, Xu Zhao, MingSung Kan, Junxian He, and Qizhe Xie. Self-evaluation guided beam search for reasoning. *Advances in Neural Information Processing Systems*, 2023.
- [71] Jing Xu, Megan Ung, Mojtaba Komeili, Kushal Arora, Y-Lan Boureau, and Jason Weston. Learning new skills after deployment: Improving open-domain internet-driven dialogue with human feedback. *arXiv preprint arXiv:2208.03270*, 2022.
- [72] Shusheng Xu, Wei Fu, Jiaxuan Gao, Wenjie Ye, Weilin Liu, Zhiyu Mei, Guangju Wang, Chao Yu, and Yi Wu. Is dpo superior to ppo for llm alignment? a comprehensive study, 2024.
- [73] Jiaqi Yang, Wei Hu, Jason D. Lee, and Simon Shaolei Du. Impact of representation learning in linear bandits. In *International Conference on Learning Representations*, 2021.
- [74] Amy Zhang, Clare Lyle, Shagun Sodhani, Angelos Filos, Marta Kwiatkowska, Joelle Pineau, Yarin Gal, and Doina Precup. Invariant causal prediction for block MDPs. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11214–11224. PMLR, 13–18 Jul 2020.
- [75] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Haoteng Zhang, Joseph Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. *ArXiv*, abs/2306.05685, 2023.

- [76] Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.



# Appendix to Learning Goal-Conditioned Representations for Language Reward Models

## A Mathematical Reasoning with Code Details

In this appendix, we discuss further experimental details for our results in Section 4.1.

### A.1 Reward Model Training and Evaluation Details

**Preference ranking dataset.** To construct the preference ranking dataset, we used the OpenMathInstruct-1, a math instruction tuning dataset with 1.8M problem-solution pairs generated by the Mixtral-8x7B model. The questions come from a blend of the training subsets for GSM8k and MATH. We matched up generations that had the final answer output correct vs. incorrect as the preferred and dis-preferred generation. Some problems are used several times, but all preferred and all dis-preferred solutions are unique. In total, we curated 257,584 total preference rankings over a total of 12,518 unique problems.

**Prompt format.** To prompt the base model to answer a math question, we use the same Nemo prompt format that was used to fine-tune the model. This prompt format was used both for reward model training, PPO training, and all evaluations. The following is the prompt format:

System:

You're an expert Python programmer and mathematician. Help the user to solve this problem using code when necessary. Make sure to put the answer (and only answer) inside `\boxed[]`.

User: **{problem}**

Assistant:

where **{problem}** is replaced with a math problem from the dataset.

**OpenRLHF library.** We implement the reward model training with contrastive loss and the baseline of standard preference ranking using the OpenRLHF library [29] in combination with PyTorch [51]. We made the following modification to the OpenRLHF Library:

- A new contrastive loss method that computes the contrastive objective
- Optionally add auxiliary loss to the reward model training objective if experimenting with our reward model method
- Optionally change the reward model reward head from a linear projection to a multi-layer perceptron. The MLP layer is a simple 3 fully connect layer with ReLU activation in between, all with the same hidden size as the representation dimension.

**Reward model training parameters.** We train both the baseline reward model and the Q-Function reward model as well as any reward model ablations for 1 epoch with batch size of 64.

The hyper-parameter for the contrastive lambda of 0.5 was chosen by doing a sweep of 0.1, 0.5, 0.9, and found that the 0.5 setting had the most stable to training loss as well as training rewards for preferred vs. unpreferred completions.

**Reward model evaluation.** We evaluate the Reward Model performance on the test splits of several math benchmarks: GSM8K [16], MATH [4], algebra222 [27], GSM-Hard [23], Asdiv [46], mawps [35], and svamp [52]. To get a set of completions on each problem in these set of benchmarks, we use the greedy generation of the base model along with annotations of whether each completion provides the correct answer, which are provided by Toshiniwal et al. [62]. Then for each completion, we first format using the Nemo format template described earlier, and use the Reward Model to predict a reward score. In predicting the reward score, we also retrieve the last hidden state of the Reward Model as well as the predicted reward score for each token in the completion, not including

hyperparameter and settings	value
batch size	64
learning rate	$9 \times 10^{-6}$
representation dimension	4096
contrastive lambda ( $\lambda$ )	0.5
num epoch	1
preference ranking loss function	sigmoid
optimizer	Adam
seed	42

Table 4: Hyper-parameters for Reward Model training

the prompt tokens. Utilizing the Reward Model predicted reward score and the annotation of the correctness of each completion, we compute an AUROC score using the Python scikit-learn package [53]. Using the same procedure, we plot the partial AUROC scores at every tenth percentile of each completion. The AUROC scores are compared between Q-Function 7b Reward (our method) and Codellama 7b Reward and the results are summarized in Section 4.1.2.

## A.2 Learned Representations Experiment Details

**Computing Q-values.** In order to retrieve the Q-values, we take the cosine similarity between the goal state and the hidden states for each token, which is retrieved during the forward pass when collecting the reward scores. The goal state in all the experiments is the mean representation of all the preferred completion in the preference ranking dataset. More explicitly, we pass each preferred problem and response pair through the trained Reward Model and then we take the average across the last hidden state of the final token of each problem and response pair.

**Filtering sampled generations.** In order to understand the value of these learned representations, we use these Q-values to filter generations. For each of the In-Distribution datasets, GSM8K [16] and MATH [4], we use the 50 completion samples provided by Toshiniwal et al. [62]. We take each completion and we discard it if the Q-value for any of the tokens is less than 0, which indicates that the expected future rewards from that state is low. This is performed for each problem and each generation sample. Using the final set of remaining completions, we then compute the majority vote accuracy, comparing it against the baseline of majority vote at 50. Furthermore, in order to visualize these Q-values, we plot a heat map of these values on the corresponding tokens. The results and plots from this experiment is detailed in Section 4.1.3.

## A.3 RLHF Training and Evaluation Details

In these set of experiments, we investigate whether the reward model trained with our proposed method provides a better reward signal when performing RLHF. Here we explain further experimental details for the results in Section 4.1.4.

**Prompt dataset.** The prompt dataset for PPO comprises of the training problems from GSM8k and MATH, which constitutes a total of 12,518 problems [16, 28]. We also incorporate all the problems from the MathQA dataset [4], which has 37,297 unique problems. Thus, in total, the prompt dataset for PPO has 49,815 unique problems. In addition, at each training step of the policy model, we add a small pretraining coefficient. The SFT dataset for the pre-training coefficient is comprised of all the preferred completions in the preference ranking dataset.

**PPO training parameters.** We also use the OpenRLHF package and PyTorch for PPO training [29, 51]. No major code changes were made to the OpenRLHF package, with the exception of the dataset loader for our custom prompt dataset. Table 5 summarizes the hyperparameters used for all PPO training reported in this paper.

**Policy model evaluation.** We evaluate the policy-trained model on the test set on all the aforementioned math benchmarks. For each problem in these benchmarks, we process the input to use the prompt format detailed above and perform zero shot evaluation. We perform 4 independent runs to account for variance. In order to perform this evaluation, we utilize the Nemo-Skill [62] package. As

hyperparameter and settings	value
train batch size	128
number of episodes	1
rollout batch size	512
training epochs	1
actor learning rate	$3 \times 10^{-6}$
critic learning rate	$9 \times 10^{-6}$
init kl coeff	0.02
normalize reward	true
generation max length	1024
GAE gamma	1.0
GAE lambda	0.95
policy model temperature	1
pretraining coefficient	0.05
learning rate	$9 \times 10^{-6}$
representation dimension	4096
contrastive lambda ( $\lambda$ )	0.5
optimizer	Adam

Table 5: Hyper-parameters for PPO training

a pre-requisite to performing evaluation on this package, we first had to convert the model to the Nemo model format. Table 6 displays the average accuracy scores across 4 independent runs, as well as the 95% confidence interval. The variability across the 4 runs comes from random seeding, which influences the initialization of the Value Model as well as dataset shuffling.

Reward Model	IID		OOD				
	GSM8K	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
DPO	80.1 $\pm$ 0.1	44.8 $\pm$ 0.2	64.8 $\pm$ 0.8	59.9 $\pm$ 0.4	76.9 $\pm$ 0.0	90.3 $\pm$ 0.2	76.6 $\pm$ 0.1
Codellama	79.3 $\pm$ 0.2	43.4 $\pm$ 0.2	65.8 $\pm$ 1.6	61.1 $\pm$ 0.3	77.4 $\pm$ 0.2	91.6 $\pm$ 0.3	78.5 $\pm$ 0.9
Q-Function	<b>80.5<math>\pm</math>0.3</b>	<b>45.1<math>\pm</math>0.1</b>	<b>70.9<math>\pm</math>1.7</b>	<b>62.7<math>\pm</math>0.5</b>	<b>79.5<math>\pm</math>0.5</b>	<b>93.6<math>\pm</math>0.3</b>	<b>81.2<math>\pm</math>0.4</b>

Table 6: Average accuracy and 95% confidence interval after the base policy model was trained via DPO, PPO using a preference ranking reward model, and PPO using our Q-Function reward model, across 4 different independent runs.

**Code execution setup.** Our sandbox environment is constructed using Python 3.10 and includes necessary mathematical libraries such as NumPy to support the execution of LM-generated code [26]. This environment is containerized and deployed across a Kubernetes cluster with 64 replicas, designed to handle the training and evaluation workload in parallel. Each replica is allocated 1 CPU and 1 Gi of memory to efficiently manage the computational demands. Each sandbox establishes communication and accepts requests through a Flask server.

**Code execution security considerations.** In our current evaluation setup, the code generated by the Language Model is primarily focused on solving mathematical questions and originates from trusted sources, allowing us to maintain a controlled and relatively secure environment. However, extending this methodology to broader applications, where code is executed from varied or unknown sources, significantly increases security concerns such as malware intrusion and unauthorized data access. To mitigate these risks, it is imperative to create a sandbox environment to isolate the code execution from the regular operational environment. Key protocols include enforcing least privilege in the sandbox to limit permissions, minimizing external communications to reduce attack surfaces, and rigorously validating both incoming code and outgoing artifacts to ensure they adhere to security best practices.

#### A.4 Compute and Statistical Significance

For the reward model training and PPO, we used a single node, 8 gpus, and 88 CPU cores, 80 Gi GPU memory and 1TB system memory. The total reward model training time ranged between 6-7 hours

and the total PPO training was approximately 13-14 hours. We additionally used flash attention 2 and deepseed with zero\_stage 3, for reward model training, and zero\_stage 2 for PPO training [17, 3].

Moreover, all the experiment results presented are statistically significant, performed by the scipy library [65]. Where applicable, all error bars, whether a table or plot, are 95% confidence intervals, calculated via numpy package [26].

## B Natural Language Alignment Details

In this appendix, we discuss further experimental details for our results in the Section 4.2.

### B.1 Training a Natural Language Reward Model

**Dataset.** To train our natural language reward model baseline and contrastive reward model, we use the training set of Helpful-Harmless [8]. Helpful-Harmless is a popular natural language alignment dataset created by Anthropic with human-assistant dialog chains. Each pair contains a chosen and rejected response, where the chosen response generally aligns on axes of helpfulness and harmlessness. This dataset contains 169,352 rows where each row contains a pair of chosen and rejected samples.

**Training.** For training the reward model, we use the OpenRLHF [29] package and Pytorch. No major code changes were made to the OpenRLHF package, and we use the dataset loader given by OpenRLHF to load the Helpful-Harmless dataset for training. Table 7 shows the hyperparameters used during training for reproducibility.

hyperparameter and settings	value
batch size	64
learning rate	$1 \times 10^{-4}$
representation dimension	4096
contrastive lambda ( $\lambda$ )	0.5
num epoch	1
preference ranking loss function	sigmoid
optimizer	Adam
seed	0

Table 7: Hyper-parameters for Natural Language Reward Model training

**Evaluation.** To evaluate the natural language reward model, we use the evaluation script in OpenRLHF [29]. We evaluate on the test set of Helpful-Harmless. For each chosen, rejected pair, we count the reward model score as correct if the chosen generation has a higher score than the reject pair.

### B.2 Guided Decoding with Contrastive Reward Model

**Dataset.** For guided decoding, we consider several different datasets. We considered OpenAssistant [36], a large corpus of assistant dialog and Prism-Alignment [34], a conversation alignment dataset. We also considered HelpSteer [67], a dataset created by Nvidia with labels on how helpful, coherent, correct, verbose, and complex a given response is. We settled on HelpSteer as our dataset for experimentation due to the relatively high-quality of the dataset (human-collected vs automatic) and the fine-grained category labels. The dataset contains 37,120 examples in the training set.

**Guided decoding.** We do not further train the model, but rather randomly sample 20 examples from each of the fine-grained category labels in HelpSteer: coherent, complexity, correctness, and helpfulness. We do not use the verbosity category because the hidden states are averaged at inference time, thus, there is a worry that we would lose the information needed to identify a sequence as verbose. We use SelfEval Guided Decoding [69] for running the decoding experiments. Equation 8 shows the objective function  $\varepsilon(s^{1:T})$  used in guided decoding beam search.

$$\varepsilon(s^{1:T}) = \prod_T P_{LM}(s^t | s^{1:t-1}) \rho C(s^t) \tag{8}$$

$P_{LM}$  is the language model distribution,  $C(s^t)$  is the confidence score, and  $\rho$  is a hyperparameter.

For the steering experiments, we set the max number of tokens per beam to 52, the beam search temperature to 0, temperature as 0.2, the number of samples to 2, and the confidence ratio to 0 e.g. equal weight to the confidence score and LM loss.

**Evaluation.** We evaluate on the test set of HelpSteer, which contains 1790 examples. For evaluation, we use GPT-4, asking GPT-4 to compare the SFT model and the model with the decoding prototype which generation is more helpful, correct, coherent, and complex. We ask GPT-4 to give a single token (0 or 1). For statistical significance, we run GPT-4 5 times for each evaluation example and take the majority. In addition, we find that for a small percent of cases GPT-4 will not product a single token even when asked. In this case, we throw away the example that did not produce a properly formatted response when computing averages.

### B.3 Compute and Statistical Significance

Natural language alignment training jobs were done on 1 core, 8 GPUs, 88 CPU cores, 1TB system memory, and 80GB GPU memory. An average training time for 1 episode of each prompt was around 13-14 hours.

All experiments are statistically significant. For the Helpful-Harmless experiments, the experiments are evaluated on a test set of 8.55k rows. For HelpSteer, the validation set is 1.79k rows.

## C Inference-Time Goal State Ablation

In this appendix section, we explore different choices of inference-time goal states and compare their performance. To ablate this, we conduct a new experiment under the mathematical reasoning with code setting (Section 4.1) where we evaluate the effect of (a) choosing poor examples to construct the goal state and (b) the number of examples used to construct the goal state. In particular, we evaluate 3 settings for picking the sample completions for constructing the goal state:

1. We vary the number of preferred completions used to construct the goal state to evaluate the effect of adding more examples of good reasoning.
2. We vary the number of dispreferred completions used to construct the goal state, to evaluate whether bad examples lead to poor performance.
3. We incrementally add more dispreferred completions to a fixed sample of preferred completion in order to measure the robustness of our goal state computation to negative examples. We refer to the addition of dispreferred completions on top of all preferred completions as adding “corrupted” examples.

The results are in Figure 5. Overall, these results show that negative or unhelpful completions degrade performance, indicating the importance of choosing relevant examples for computing the goal state. Additionally, they demonstrate that having more examples of the concept leads to better performance as more generations are filtered with comparable proportion of those remaining generations being correct.

We additionally ablate our choice of using the last token in the completion sequence to construct the goal state. In this experiment, we repeat our filtering experiments from Section 4.1.3. Except, we randomly sample a token from in the completion sequence and use it as the goal state:

Sampling Method	GSM8k		MATH	
	Accuracy (%)	Prop. Correct (%)	Accuracy (%)	Prop. Correct (%)
Last Token	86.0	84.0	59.6	52.0
Random Token	85.9	81.8	57.3	45.3

Table 8: Majority vote accuracy and average proportion correct when filtering examples using last token and random token sampling for constructing the goal state.




Figure 5: We plot the average proportion of non-filtered responses that are correct and the average number of completions (out of a total 50) that are filtered, as we vary the sample size (1, 10, 100, 1K, 10K, and 30K) of dis-preferred, preferred, and corrupted examples when constructing the goal state. All of the examples are retrieved from the preference ranking training dataset. We plot the performance metrics for the MATH and GSM8K test sets. Each of the test sets has 50 completions per problem that is generated by the base model and we filter any completion for which the Q-value for any token in the completion sequence is below 0. At each sample size, we run 3 independent random sample from the full set of preferred and dis-preferred completions to construct 3 independent goal states, with 95% CI.

From the result of this ablation, we see that last token sampling has superior performance, likely because the last token on average carries the most information about what good & complete generations look like.

## D Train-Time Goal State Ablation

### D.1 Exploring Sources of Goal States

A natural question about our approach is: *What is the impact on the choice of goal state representation?* For sampling negative goal states, we experiment with randomly sampling any arbitrary completion (not necessarily the same prompt) vs. the corresponding dispreferred completion to show the effect of picking completions that have similar contexts and invalid reasoning steps when constructing the contrastive loss. Furthermore, we experiment with taking the average of the hidden states of the positive and negative goal states across the train batch and then apply cosine similarity with the current token representation. The results of these experiments are in Table 9.

Method	IID		OOD				
	GSM8k	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
Pref. Rank	0.805	0.618	0.850	0.768	0.829	0.775	0.761
Q-Function (SGS)	<b>0.852</b>	<b>0.708</b>	<b>0.879</b>	<b>0.781</b>	<b>0.854</b>	<b>0.819</b>	0.841
Q-Function (RS)	0.801	0.626	0.832	0.716	0.723	0.760	0.802
Q-Function (AVG)	0.835	0.707	0.851	0.765	0.839	0.773	<b>0.863</b>

Table 9: Comparison of AUROC scores for different training time contrastive goal states on several popular math benchmarks. Pref. rank is the reward model trained with the standard preference-ranking objective, Q-Function (SGS) is trained with the method described in Section 3.2, Q-Function (RS) trained with the randomly sampled goal-state and Q-Function (AVG) trained with the batch averaged goal state.

The ablation in Table 9 shows that the positive and negative goal state representations we use are largely more effective than randomly sampled and batch averaged goal representations. Somewhat similar observations have been made for other contrastive learning approaches where negatives, especially hard-negatives, can play an important role [14, 57]. We see that when we randomly sample positive and negative goal states, there is a clear regression in performance on the OOD datasets,

even below the standard preference ranking trained reward model. Meanwhile, the batch averaged goal states offers competitive but largely lower performance compared to our final setup. Given these results, we use the SGS contrastive loss method as our final setup.

## D.2 Source State and Goal State Sampling Ablation

In this section, we ablate the sampling strategy for the source and goal states. In computing the contrastive loss during training, we need to sample a positive source and goal state pair as well as a negative source and goal state pair. To investigate how we should sample these tokens, we explore 3 different approaches.

**Method 1: Random Source and Goal State.** In this approach, the positive and negative source states are sampled randomly from the preferred response, then the positive goal state is sampled randomly amongst any token *after* the positive source state, and last the negative goal state is randomly sampled anywhere from the dis-preferred response. This is the method we describe in Section 3.2.

**Method 2: Late Goal State.** In this approach, we sample the positive and negative goal states only towards the end of the preferred and dis-preferred response, respectively. The intuition here is that the token later in the sequence encode more information about the whole response and thus provide more information as a goal state. In practice, we restrict the goal state sampling to come after the 90th percentile with respect to the total number of tokens in the response.

**Method 3: Late Source and Goal State.** In this approach, we sample both the source state and goal state towards the end of the preferred and dis-preferred response (after 90th percentile of tokens generated).

Sampling Method	IID		OOD				
	GSM8k	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
Random Source and Goal State	<b>0.852</b>	<b>0.708</b>	<b>0.879</b>	<b>0.781</b>	0.854	0.819	<b>0.841</b>
Late Goal State	0.841	0.696	0.821	0.775	<b>0.860</b>	0.729	0.761
Late Source and Goal State	0.829	0.643	0.867	0.779	0.829	<b>0.831</b>	0.833

Table 10: Comparison of AUROC scores for different source and goal state sampling method.

The ablation in Table 10 shows that the Random Source and Goal State sampling on average performs the best across both in-distribution and out-of-distribution tasks. We also see that with Late Goal State sampling, we have similar performance to Random Source and Goal State sampling for the in-distribution tasks but we notice a performance drop across several out-of-distribution benchmarks. The opposite is true for the Late Source and Goal State sampling where performance on the in-distribution benchmark drop and the performance out-of-distribution benchmarks are comparable to the random sampling. Given these results, we use the Random Source and Goal State sampling method as our final setup, which is equivalent to the (SGS) method referred to in the rest of the paper.

## E Reward Model Ranking Performance

In this appendix section, we compare the Q-Function 7b Reward to Codellama 7b Reward by judging their respective abilities to rank completions. Since we have a total of 50 generations per problem, we use the baseline reward mode and our reward model to select the Top 1, 5, 10, and 25 samples as ranked by the reward scores. With the selected sample, we perform majority vote and also note the average proportion of the sample K that are correct solutions. The results are provided in Table 11. We show the Majority Vote Accuracy and Average Proportion Correct when selecting Top-K completions in order of reward model scores. The two reward models used here are the baseline Codellama RM and our Q-Function RM. The results show that our Q-Function RM clearly outperforms the baseline.

Model	Top-K	GSM8k		MATH	
		Accuracy (%)	Prop. Correct (%)	Accuracy (%)	Prop. Correct (%)
Q-Function RM	1	84.6	84.6	51.7	51.7
	5	86.2	84.5	59.7	51.1
	10	86.2	83.8	59.5	50.3
	25	85.5	81.8	57.8	47.0
Codellama RM	1	80.8	80.8	43.8	43.8
	5	84.3	81.9	54.0	45.8
	10	85.2	82.1	56.0	46.9
	25	85.8	81.1	56.5	46.1

Table 11: Majority Vote Accuracy and Average Proportion Correct when selecting Top-K completions in order of reward model scores. The two reward model use here is the baseline Codellama RM and our Q-Function RM.

## F Partial Completion Rewards

In this appendix, we show the plots for the remaining out-of-distribution datasets, namely, algebra222, Asdiv, GSM-Hard, mawps, and svamp. The general trend observed in Section 4.1.2 still holds, where, Q-Function 7b Reward model incrementally improves as percentile of model generation considered increases. On the other hand, there is a lot more variability in Codellama 7b Reward partial auroc scores. More importantly, the performance when considering partial solutions is generally more accurate for the Q-Function 7b Reward. This demonstrates that the improved judgement of partial solutions that we observed for the in-distribution tasks extend to out-of-distributions tasks as well.

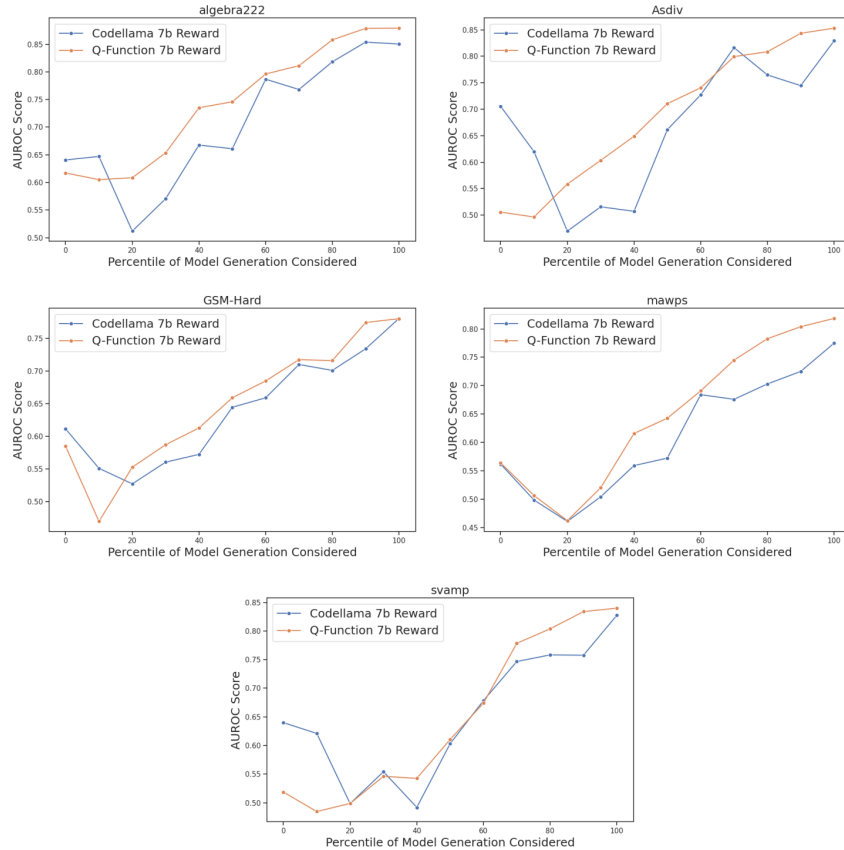


Figure 6: AUROC scores on the rewards attributed to partial base-model greedy generations on the OOD datasets.



## G Q-Value Visualization

In this appendix, we display the full Q-value heat map for the examples shown in Figure 4. We also show additional examples that show the Q-values correctly scoring reasoning errors as well as some examples where the Q-values are noisy, unable to explicitly pin-point the logical error.

### G.1 Good Examples

**Problem:** Toula went to the bakery and bought various types of pastries. She bought 3 dozen donuts which cost 68 per dozen, 2 dozen mini cupcakes which cost 80 per dozen, and 6 dozen mini cheesecakes for 55 per dozen. How much was the total cost?

```
Let 's solve this problem using Python code .
<|ll m |code >
don uts |per |do zen = 1 2
don uts |price |per |do zen = 6 8
don uts |cost = don uts |per |do zen * don uts |price |per |do zen * 3
cup c akes |per |do zen = 1 2
cup c akes |price |per |do zen = 8 0
cup c akes |cost = cup c akes |per |do zen * cup c akes |price |per |do zen * 2
che es ec akes |per |do zen = 1 2
che es ec akes |price |per |do zen = 5 5
che es ec akes |cost = che es ec akes |per |do zen * che es ec akes |price |per |do zen * 6
total |cost = don uts |cost + cup c akes |cost + che es ec akes |cost
</ll m |code >
<|ll m |code |output >
8 3 2 8
</ll m |code |output >
Th us the total cost of the past ries is \ box ed [ 8 3 2 8 ] dollars .
```

**Problem:** Marissa is hiking a 12-mile trail. She took 1 hour to walk the first 4 miles, then another hour to walk the next two miles. If she wants her average speed to be 4 miles per hour, what speed (in miles per hour) does she need to walk the remaining distance?

```
Let 's solve this problem using Python code .
<|ll m |code >
distance |tra ve led = 4 + 2
time |t aken = 1 + 1
aver age |speed = 4
rem aining |distance = 1 2 - distance |tra ve led
time |rem aining = rem aining |distance / aver age |speed
</ll m |code >
<|ll m |code |output >
1 5
</ll m |code |output >
Th us Mar issa needs to walk at \ box ed [ 8 ] miles per hour to achieve her average speed of 4 miles per hour .
```

**Problem:** In a dance class of 20 students, 20% enrolled in contemporary dance, 25% of the remaining enrolled in jazz dance, and the rest enrolled in hip-hop dance. What percentage of the entire students enrolled in hip-hop dance?

```

Let 's solve this problem using Python code .
<llm - code >
# 20% enrolled in contemporary dance
contemporary_dance_percent = 20
# 25% of the remaining enrolled in jazz dance
jazz_dance_percent = 25
# the rest enrolled in hip-hop dance
hip_hop_dance_percent = 100 -
(contemporary_dance_percent + jazz_dance_percent)
# total number of students
total_students = 20
# number of students enrolled in hip-hop dance
hip_hop_dance_students = total_students * hip_hop_dance_percent / 100
# percentage of students enrolled in hip-hop dance
hip_hop_dance_percent_of_total = hip_hop_dance_students / total_students * 100
hip_hop_dance_percent_of_total
</llm - code >
<llm - code - output >
55.00000000000001
</llm - code - output >
Thus \boxed{55\%} of the students enrolled in hip-hop dance .

```

**Problem:** When Freda cooks canned tomatoes into sauce, they lose half their volume. Each 16 ounce can of tomatoes that she uses contains three tomatoes. Freda's last batch of tomato sauce made 32 ounces of sauce. How many tomatoes did Freda use?

```

Let 's solve this problem using Python code .
<llm - code >
# each can of tomatoes has 3 tomatoes
tomatoes_per_can = 3
# each can of tomatoes has 16 ounces
ounces_per_can = 16
# Freda's last batch of sauce made 32 ounces
sauce_made = 32
# Freda used M cans of tomatoes
M = sauce_made / ounces_per_can
# Freda used N tomatoes
N = M * tomatoes_per_can
N
</llm - code >
<llm - code - output >
6.0
</llm - code - output >
Thus Freda used \boxed{6} tomatoes .

```

## G.2 Noisy Examples

**Problem:** Lee used to be able to run the 400-meter hurdles two seconds faster than Gerald would run the 400-meter hurdles. But Gerald changed his diet, which improved his speed by 10%. If Lee runs the 400-meter hurdles in 38 seconds, how fast can Gerald, with his improved diet, run the 400-meter hurdles, in seconds?

```
Let's solve this problem using Python code .
<llm - code >
lee _ time = 38
lee _ speed = 400 / lee _ time
gerald _ speed = lee _ speed - 2
gerald _ speed with improvement = gerald _ speed * 1.1
gerald _ time = 400 / gerald _ speed with improvement
gerald _ time
</llm - code >
<llm - code - output >
42.648709315375974
</llm - code - output >
Thus Gerald can run the 400-meter hurdles in \boxed{43} seconds .
```

**Problem:** John drives for 3 hours at a speed of 60 mph and then turns around because he realizes he forgot something very important at home. He tries to get home in 4 hours but spends the first 2 hours in standstill traffic. He spends the next half-hour driving at a speed of 30mph, before being able to drive the remaining time of the 4 hours going at 80 mph. How far is he from home at the end of those 4 hours?

```
Let's solve this problem using Python code .
<llm - code >
# John drives for 3 hours at a speed of 60 mph
distance _ traveled _ first _ part = 3 * 60
# then turns around because he realizes he forgot something very important at home
distance _ traveled _ second _ part = 0
# tries to get home in 4 hours but spends the first 2 hours in stand still traffic
distance _ traveled _ second _ part += 2 * 0
# spends the next half - hour driving at a speed of 30 mph
distance _ traveled _ second _ part += 0.5 * 30
# before being able to drive the remaining time of the 4 hours going at 80 mph
distance _ traveled _ second _ part += 2 * 80
# total distance traveled
total _ distance _ traveled = distance _ traveled _ first _ part + distance _ traveled _ second _ part
total _ distance _ traveled
</llm - code >
<llm - code - output >
355.0
</llm - code - output >
Thus John is \boxed{355} miles from home at the end of those 4 hours .
```

## H Exploring Correlation Between Q-values and Reward Scores

In this section, we explore the correlation between our learned Q-values and reward scores for partial completions, as illustrated in Table 12, and empirically investigate the possibility of reward hacking during PPO training. The partial completions are from the base model on the GSM8K test split [16]. The correlation is expected since, by design, the reward model outputs are dependent on the learned features that approximate the Q-function. This presents a potential risk in practice that the policy model may exploit during PPO. We explore attempting to decouple the reward score from the Q-value for partial completions by projecting the hidden features which approximate the Q-value through a multilayer perceptron. Table 12 and Table 13 compares the result of projecting the hidden features via a MLP layer versus a linear layer. The finding is we can reduce correlation between rewards and Q-values for partial completions while maintaining high reward model performance. Iterative retraining of the reward model on on-policy completions would further mitigate any potential reward model hacking given that our complete loss function also considers the preference ranking paired loss in 1.

Method	IID		OOD				
	GSM8k	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
Pref. Rank	0.805	0.618	0.850	0.768	0.829	0.775	0.761
Q-Function (Linear)	<b>0.852</b>	<b>0.708</b>	<b>0.879</b>	<b>0.781</b>	0.854	<b>0.819</b>	0.841
Q-Function (MLP)	0.832	0.693	0.871	0.770	<b>0.871</b>	0.780	<b>0.841</b>

Table 12: Comparison of AUROC scores for different reward projection strategy on several popular math benchmarks. Both Q-Function (Linear) and Q-Function (MLP) are trained with the method described in Section 3.2, the only difference being the reward head is a single linear layer versus a 3-layer MLP followed by a linear layer.

Model	Percentile			
	0.2	0.4	0.6	0.8
Q-Function (Linear)	0.983	0.986	0.988	0.988
Q-Function (MLP)	0.870	0.908	0.921	0.915

Table 13: Correlation between Q-value and reward score at different percentile of model generations. We see that the Q-value and reward scores for Q-Function (Linear) has higher correlation than Q-Function (MLP).

To empirically evaluate the risk of the policy gaming the reward function, we conduct a new experiment where we run PPO training for an extended time (5 episodes) in order to evaluate whether reward hacking occurs in practice. The reward scores and response lengths are shown in Figure 7 and the performance of the model at the end of each episode is given in Table 14.




Figure 7: Train time average rewards and average response length for PPO using Codellama RM (purple) and PPO using our Q-Function RM (brown)

These results indicate training is fairly stable and the policy model does not game the reward function in practice. One explanation for why gaming does not occur is that typically during RLHF a relatively low learning rate is used, making large shifts in the models generations less likely to occur.

Model	Episode	GSM8k	MATH	algebra222	GSM-Hard	Asdiv	mawps	svamp
Codellama	1	79.7	43.6	66.3	61.2	77.9	92.0	78.5
	2	79.6	45.0	72.5	61.5	78.0	91.6	79.9
	3	80.5	45.3	69.8	62.0	78.3	92.3	80.9
	4	80.5	45.5	70.7	62.5	78.3	92.0	79.5
	5	80.5	45.2	70.7	62.5	78.6	92.0	80.5
Q-Function	1	80.2	45.9	67.6	62.0	79.2	93.9	81.6
	2	80.6	46.5	74.3	61.9	79.5	93.1	80.3
	3	81.7	46.6	73.0	63.5	79.2	93.2	81.6
	4	81.0	46.5	73.4	63.2	79.2	93.4	81.0
	5	81.1	46.5	74.3	63.0	79.8	93.6	81.0

Table 14: PPO-trained policy model accuracy across several math benchmarks comparing using the baseline Codellama RM and our Q-Function RM for a single long 5 episode training run.

## I Examples of Prototype Guided Decoding

**Question:** Consider this reference information delimited in """:

There is a man who circles the perimeter with a baby in his arms unmoving. Locusts burn with the silhouettes of saints at dusk. Saints are in the cloud. We are in a dry storm. The man extends his circles pulling the baby through the cactus scrub. Look at his melting trainers in the heat, they aren't what he asked for. There are black leather skids on the dry stone wall. People in black cloaks run out of the corner of your eye. A pig turns on a spit. The prairie is a terrarium for the blaze but the edge is dry of fire. It is the height of one season, bushes burn. A burnt five-year-old without eyelids turns quick cartwheels through the heat wave under the big pale sky, black and blue.

What type of poem is this text from?

**SFT:** This text is a prose poem.

**Prototype:** This text is a prose poem, which is a type of poetry that combines the elements of prose with those of poetry to create a unique and imaginative style of writing.

The guided decoding successfully produces a response that is still correct while adding complexity to the response by describing that the poetry combines elements of prose and imaginative writing styles. We find that, in general, guided decoding responses produce more complex responses without sacrificing correctness.

**Question:** Reference: <start of reference> If you invested 1,000 in Facebook in 2012, here's how much you'd have now Even given the stock market's recent volatility, a 1,000 investment in Facebook 10 years ago would still prove to be a good bet: You could have quadrupled your money. An initial investment in Facebook in 2012, when it made its initial public offering, would be worth slightly more than 4,600 as of Wednesday, according to CNBC calculations. The investment would have seen a lifetime total return of more than 360 percent, including price appreciation and dividend gains reinvested. The market's recent activity has taken a toll, of course: In one month, Facebook's stock dipped from 5,051. Its current portfolio value is just above 30 percent now, down from its peak of more than 40 percent. And the market keeps ebbing and flowing. In the charts below, all data splits are adjusted and gain-loss figures do not include dividends, interest, distributions or fees except on cash accounts. The portfolio value represents current holdings and the comparison charts represent current and historical prices of individual benchmarks, stocks or exchange-traded funds. Meanwhile, daily users in the United States and Canada dropped in 2017 to 184 million daily average users in the fourth quarter from 185 million in the third quarter. Facebook chief financial officer David Wehner, on the company's earnings call, said usage would "fluctuate" based on market penetration. The revenue increase despite the usage drop could be due to seasonal changes related to the advertising market. As the company points out, users in the fourth quarter accounted for 6.39 billion of Facebook's revenue, compared to the 5.03 billion users accounted for in the third quarter. Facebook has retreated slightly from its record price reached earlier this month as investors debate whether changes to the platform to combat fake news and other problems will hurt advertising growth. But the stock is still up 1 percent on the year. And some investors remain bullish about the stock. "News feed changes do not impair the power of the platform," Andy Hargreaves, an analyst at KeyBanc Capital Markets, wrote in a note to clients. "Facebook offers reach, targeting and ad-unit quality that is difficult to match, and they are the underlying drivers of its growth. As long as [it] continues to improve its effectiveness, the foundation for strong growth should remain intact, regardless of adjustments to the news feed." <end of reference> Instruction: What is a share of stock?

**SFT:** A share of stock is a unit of ownership in a company.

**Prototype:** A share of stock is a unit of ownership in a company that represents the right to a proportionate share of the company's assets and profits.

In this response, the prototype response is slightly more correct as it describes the value of a share of stock. In addition, the prototype response is more complex than the SFT response.

**Question:** What is the difference between a pound and an ounce? Please explain with examples.

**SFT:** A: A pound is a unit of measurement for weight, while an ounce is a unit of measurement for volume.

**Prototype:** A pound is equal to 0.45359237 kilograms, while an ounce is equal to 28.3495 grams. A: A pound is a unit of weight in the metric system, while an ounce is a unit of weight in the English system.

Here, we find that the SFT response is incorrect. The prototype response gives the correct response, but produces some additional information at the beginning of the response that does not answer the question. This can be seen as a side effect of the guided decoding process since a more complex response is not always desired.

**Question:** What are some common misunderstandings about the causes and consequences of the 2008 global financial crisis?

**SFT:** One is that it was caused by a lack of regulation in the financial sector, when in reality, the crisis was caused by a combination of factors, including excessive risk-taking by financial institutions, poor risk management practices, and a lack A: There are a few common misunderstandings about the causes and consequences of the 2008 global financial crisis.

**Prototype:** Some believe that the crisis was caused by a lack of regulation or irresponsible lending practices. Some also believe that the crisis was caused by a lack of regulation or irresponsible lending practices. In reality, the crisis was the result of complex and interconnected events that took place over many years. There are many misconceptions about the 2008 financial crisis, including the belief that it was caused by a lack of regulation or irresponsible lending practices.

The prototype response is more coherent than the SFT response in this example. In the SFT response, there is a lack of coherence because the response is separated into two distinct sections, separated by "A:". However, this issue does not exist in the prototype response. Each sentence builds off the last, for example "Some believe..." followed by "in-reality".

**Question:** Reference: <start of reference> As a casino player, it would be amazing to predict when a slot machine will next hit. You could time exactly when you play and how much to bet perfectly, maximising your winnings. The problem? If this was widely possible, no casino would exist. Theyd all go bankrupt straight away. Instead, modern slot machines are designed to give determined results, that are not generated at random. It is incredibly difficult to guess what the result of a spin will be. There may be slot strategies that might improv your chances of winning. But any legitimate methods won't include being able to know when a slot is about to drop. Yet there's still a lot of talk about this online, with "tips" and "tricks" on how to predict when a slot machine will next payout. We'll explain what really happens, along with how you can give yourself a slightly better chance of winning legitimatly. You cannot tell when a modrn slot machine will hit because the outcome of each spin is random. No matter how many times a machine has spun, and no matter what the outcomes of those spins were, the probability the next result remains the same. But people often believe otherwise. One famous example of why is the gambler's fallacy. The gambler's fallacy is an erroneous way of thinking about the probability of independent events. When people commit the gambler's fallacy, they think that the likelihood of a random event is impacted by the occurrence of previous events, or series of previous events. For example, believing that the result of previous spins on a slot will impact the result of present or future spins is committing the gambler's fallacy. The result is not affected by what's happened before. So if you win or lose one round, it has no impact on whether you'll win or lose the next round. <end of reference> Instruction: Rewrite the passage correcting any grammar and spelling mistakes that have been made. Also omit any false information that the passage may contain. Make sure to keep the same tone and message of the original passage.

**SFT:** Casino players believe they can predict when a slot machine will next hit, but this is not possible.

**Prototype:** You could time exactly when you play and how much to bet perfectly, maximising your winnings. The problem? If this was widely possible, no casino would exist. They'd all go bankrupt straight away. Instead, modern slot machines are designed to give determined results, that are not generated at random. It is incredibly difficult to guess what the result of a spin will be. There may be slot strategies that might improv your chances of winning. As a casino player, it would be amazing to predict when a slot machine will next hit.

The prototype response gives more detail including additional reasoning: "modern slot machines are designed to give determined results". Interestingly, this response is generally incorrect, as modern slot machines will give random results that have deterministic rules (for example, no jackpot right after another jackpot). This shows that, while the prototype can help improve correctness, it can still suffer from hallucination.

## J Assets

In Table 15, we summarize assets used in this paper along with their corresponding URLs and licenses.

Asset Name	URL	License Type
Llama 3 [1]	<a href="https://github.com/meta-llama/llama3">https://github.com/meta-llama/llama3</a>	Llama 2 Community License
OpenMath CodeLlama [62]	<a href="https://huggingface.co/nvidia/OpenMath-CodeLlama-7b-Python">https://huggingface.co/nvidia/OpenMath-CodeLlama-7b-Python</a>	Llama 2 Community License
OpenRLHF [29]	<a href="https://github.com/OpenLLMAI/OpenRLHF">https://github.com/OpenLLMAI/OpenRLHF</a>	Apache-2.0
NeMo-Skills [62]	<a href="https://github.com/Kipok/NeMo-Skills">https://github.com/Kipok/NeMo-Skills</a>	Apache-2.0
OpenMathInstruct-1 [62]	<a href="https://huggingface.co/datasets/nvidia/OpenMathInstruct-1">https://huggingface.co/datasets/nvidia/OpenMathInstruct-1</a>	NVIDIA License
GSM8k [16]	<a href="https://huggingface.co/datasets/gsm8k">https://huggingface.co/datasets/gsm8k</a>	MIT License
MATH [28]	<a href="https://github.com/hendrycks/math">https://github.com/hendrycks/math</a>	MIT License
algebra222 [27]	<a href="https://huggingface.co/datasets/sirdug/Algebra222">https://huggingface.co/datasets/sirdug/Algebra222</a>	Apache-2.0
GSM-Hard [23]	<a href="https://huggingface.co/datasets/reasoning-machines/gsm-hard">https://huggingface.co/datasets/reasoning-machines/gsm-hard</a>	MIT License
Asdiv [46]	<a href="https://github.com/chaochun/nlu-asdiv-dataset">https://github.com/chaochun/nlu-asdiv-dataset</a>	CC BY-NC 4.0
mawps [35]	<a href="https://huggingface.co/datasets/MU-NLPC/Calc-mawps">https://huggingface.co/datasets/MU-NLPC/Calc-mawps</a>	MIT License
svamp [52]	<a href="https://huggingface.co/datasets/ChilleD/SVAMP">https://huggingface.co/datasets/ChilleD/SVAMP</a>	MIT License
Helpfulness and Harmlessness [8]	<a href="https://huggingface.co/datasets/Anthropic/hh-rlhf">https://huggingface.co/datasets/Anthropic/hh-rlhf</a>	MIT License
HelpSteer [67]	<a href="https://huggingface.co/datasets/nvidia/HelpSteer">https://huggingface.co/datasets/nvidia/HelpSteer</a>	CC-BY-4.0

Table 15: Asset information and licences.



## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We propose a new method for training reward models via a contrastive objective. In Section 3, we describe our method in detail. We also claim that such method helps reward model improve their performance and that the learned representation is a useful signal for approximating future gains. In the experiments section, we provide two different settings, mathematical reasoning with python interpreter (Section 4.1) and natural language alignment (Section 4.2), where we explicitly state the experimental setup and results referred to in the abstract.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We meaningfully explore two limitations to our approach in the limitations section (Section 5). Namely, we call out the discrepancy of train time vs. test time settings for our method and that there is a correlation between the reward and Q-value signals. We also provide/explore some approaches to reconcile these limitations for further generalizability of our method.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not present theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: For our proposed method, we provide a detailed breakdown of the theoretical ideas as well as how to implement it in practice. Moreover, in Table 4 and Table 5, we lay out all the relevant hyperparameters required to reproduce the results. Also, we list out the minimal modification we make to the OpenRLHF package and so prospective readers should be able to implement the exact methods and settings with relative ease.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.

- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material? [Yes]

Justification: The code is released at

<https://github.com/vaskarnathscale/goal-conditioned-rm/>.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: For each experiment, we explicitly mention how the training/evaluation datasets were constructed, the magnitude of the datasets, the base models used, as well as the experimental setup/methods use for all reported values.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All the experiment results presented are statistically significant, performed by the scipy library [65]. Where applicable, all error bars, whether a table or plot, are 95% confidence intervals, calculated via numpy package [26].

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In the Experimental Details section in the appendix, we explicitly state the compute details required for each training and evaluation job and approximately how long each run took.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification: The models and data used in our experiments are open source and used in accordance to their respective licenses.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.

- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: We introduce a new method for reward model training and show the validity of this approach on self-contained experiments that will not have outreaching effects on society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We are not releasing a new dataset nor a new model.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The assets used in this work along with their respective licenses are provided in Appendix J. In this work, we exclusively use them for research purposes and in accordance with their license and usage agreements.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This submission is currently not releasing new assets due to time and security constraints.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This work does not involve human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This work does not involve human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.