
Noise Contrastive Alignment of Language Models with Explicit Rewards

Huayu Chen^{1,2}, Guande He^{1,2}, Lifan Yuan¹, Ganqu Cui¹, Hang Su^{1,2,3}, Jun Zhu^{1,2*}

¹Department of Computer Science and Technology, Tsinghua University

²Institute for AI, BNRist Center, Tsinghua-Bosch Joint ML Center, THBI Lab, Tsinghua University

³Zhongguancun Laboratory, Beijing, China

Abstract

User intentions are typically formalized as evaluation rewards to be maximized when fine-tuning language models (LMs). Existing alignment methods, such as Direct Preference Optimization (DPO), are mainly tailored for pairwise preference data where rewards are implicitly defined rather than explicitly given. In this paper, we introduce a general framework for LM alignment, leveraging Noise Contrastive Estimation (NCE) to bridge the gap in handling reward datasets explicitly annotated with scalar evaluations. Our framework comprises two parallel algorithms, NCA and InfoNCA, both enabling the direct extraction of an LM policy from reward data as well as preference data. Notably, we show that the DPO loss is a special case of our proposed InfoNCA objective under pairwise preference settings, thereby integrating and extending current alignment theories. By comparing NCA and InfoNCA, we demonstrate that the well-observed decreasing-likelihood trend of DPO/InfoNCA is caused by their focus on adjusting relative likelihood across different responses. In contrast, NCA optimizes the absolute likelihood for each response, thereby effectively preventing the chosen likelihood from decreasing. We evaluate our methods in both reward and preference settings with Mistral-8 \times 7B and 7B models. Experiments suggest that InfoNCA/NCA surpasses various preference baselines when reward datasets are available. We also find NCA significantly outperforms DPO in complex reasoning tasks like math and coding. Code: <https://github.com/thu-ml/Noise-Contrastive-Alignment>.

1 Introduction

Aligning pretrained Language Models (LMs) with scalar rewards that reflect human intentions is crucial for enhancing their ability to follow instructions [35, 25]. These rewards can be given either explicitly or implicitly. Explicit rewards can be scalar ratings of human annotators or advanced models like GPT-4, while implicit rewards are usually preference labels assigned to pairwise responses.

One effective approach for aligning LMs with preference data is Direct Preference Optimization (DPO, [33]). DPO applies a reward training loss but parameterizes the reward model as the response likelihood ratio between two LMs, allowing for training reward models and extracting LM policies simultaneously. This approach is more streamlined and thus more favorable compared with traditional Reinforcement Learning (RL) methods [26], which typically require a two-stage training process: first training reward models, then extracting LM policies.

Despite its simplicity and effectiveness, DPO is only tailored for preference data ($x \rightarrow \{y_w > y_l\}$). When multiple responses are available, directly assigning a scalar reward to each response is usually more convenient and efficient than comparing them in a pairwise manner. The resulting reward

*The corresponding author

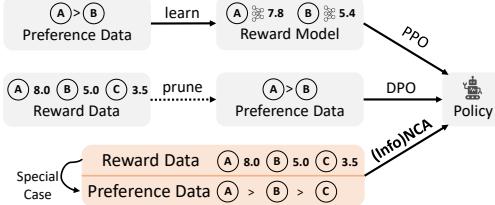


Figure 1: InfoNCA/NCA allows direct LM optimization for both reward and preference data.

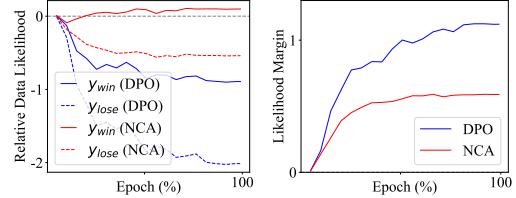


Figure 2: Pairwise NCA prevents chosen likelihood from decreasing while DPO cannot.

datasets ($x \rightarrow \{y_i, r_i\}_{1:K}$), however, cannot be directly leveraged for DPO training. Previous work [40] usually prunes reward datasets by selecting the best response and pairing it with a random remaining one. This is suboptimal as all reward values and additional dispreferred responses are thrown away in its data-preprocessing process.

To address this issue, we present **InfoNCA**, an alignment method that allows directly extracting LM policies from both reward datasets and preference datasets with arbitrary response numbers (Figure 1). Notably, InfoNCA subsumes DPO loss as a special case under pairwise preference settings and can thus be seen as a natural extension of DPO (Sec. 3.2). With strong theoretical guarantees, we show DPO is a binary classification loss while InfoNCA is its multi-category version (Figure 3). However, unlike DPO which is built upon assumptions of Bradley-Terry models or Plackett-Luce models, InfoNCA is strictly derived from Information Noise Contrastive Estimation (InfoNCE, [24]), an established contrastive method that is widely applied in language and visual representation learning [31]. This closes the theoretical gap between current preference alignment methods and classic contrastive learning frameworks.

A well-observed problem with DPO is that the likelihood of the preferred response tends to *decrease* throughout training [27, 32]. Similar trends are shared by InfoNCA. We find this issue arises mainly from InfoNCA/DPO’s focus on adjusting the *relative* likelihood across different responses per instruction. In Sec. 4, we propose **NCA** as an alternative alignment method to InfoNCA to mitigate this problem. NCA differs from InfoNCA by only loss definition and is also suitable for both preference and reward datasets. However, NCA is built on NCE [14], a parallel contrastive learning algorithm to InfoNCE, which optimizes the *absolute* data likelihood during training. In practice, NCA effectively prevents the chosen likelihood from decreasing (Figure 2).

We evaluate our methods on Mistral-7B and $8 \times 7B$ models from two dimensions. When reward datasets [9] are available, we show that directly applying our reward-based alignment offers clear improvement compared with preference-based algorithms, achieving higher evaluation rewards in GPT-4 [47, 19] evaluations. We further validate this improvement comes from InfoNCA/NCA’s ability to fully leverage the additional suboptimal responses. When only preference data is given [44], we compare pairwise NCA against the DPO loss. Our experimental results spanning various benchmarks show that NCA outperforms DPO in complex reasoning tasks such as math and coding.

Our main contributions: 1. We bridge the theoretical gap between DPO and classic contrastive learning theories. InfoNCA and NCA are uniquely suited for both reward and preference data, offering a general framework that integrates preference-based algorithms. 2. We show that suboptimal responses are also important for LM optimization. Our method outperforms various preference methods by fully exploiting data information in reward datasets. 3. NCA effectively mitigates the data likelihood decline issue of DPO and offers practical performance improvement.

2 Background: Direct Preference Optimization

LM alignment is essentially a constrained policy optimization problem:

$$\max_{\pi_\theta} \mathbb{E}_{p(x)} [\mathbb{E}_{\pi_\theta(y|x)} r(x, y) - \alpha D_{\text{KL}}(\pi_\theta(\cdot|x) || \mu(\cdot|x))], \quad (1)$$

where μ represents the pretrained LM. x and y are respectively instructions and responses. r is a reward function that reflects human intentions. α is some temperature coefficient. Prior work [30, 29]

has proved that the optimal solution for the optimization problem in Eq. 1 is:

$$\pi^*(y|x) = \mu(y|x) \frac{e^{r(x,y)/\alpha}}{Z(x)} \propto \mu(y|x) e^{r(x,y)/\alpha}. \quad (2)$$

Direct Preference Optimization (DPO) [33] assumes we only have access to some pairwise preference data $x \rightarrow \{y_w > y_l\}$ for each instruction x . The preference probability of human annotators is modeled by a learnable implicit reward model r_θ under Bradley-Terry theories [2]:

$$\pi_\theta(y_w > y_l|x) := \sigma(r_\theta(y_w, x) - r_\theta(y_l, x)),$$

where σ is the sigmoid function. To learn r_θ , DPO simply adopts a binary classification loss:

$$\mathcal{L}_{\text{DPO}} = -\mathbb{E}_{\{x, y_w > y_l\}} \log \sigma(r_\theta(y_w, x) - r_\theta(y_l, x)). \quad (3)$$

In practice, the latent function r_θ is parameterized by the log-likelihood ratio between π_θ and μ :

$$r_\theta(x, y) := \beta \log \frac{\pi_\theta(y|x)}{\mu(y|x)},$$

where β a linear coefficient for scaling r_θ . This parameterization is crucial because it ensures $\pi^*(y|x) \propto \mu(y|x) e^{r_\theta(x,y)/\beta}$ constantly hold. It transforms generative policy optimization into a simple discriminative classification task: When $r_\theta = r$ and $\beta = \alpha$ are satisfied, we naturally have $\pi_\theta = \pi^*$.

3 InfoNCA: Extending DPO from Preference to Explicit Rewards

Compared with constructing preference datasets, annotating each response with scalar rewards can be more flexible and convenient. Preference methods are only suitable for pairwise data ($x \rightarrow \{y_w > y_l\}$) and would require C_K^2 evaluations for comparing K responses. In contrast, reward datasets ($x \rightarrow \{y_i, r_i\}_{1:K}$) allow an arbitrary number of responses per prompt with K evaluations.

Despite its simplicity in handling preference data, DPO is not tailored for reward datasets. We introduce a new alignment method termed InfoNCA to mitigate this gap. We first strictly derive InfoNCA in Sec. 3.1. We show that reward alignment can be solved by constructing a classification problem to identify the optimal response from multiple candidates. We then demonstrate that InfoNCA subsumes DPO as a special case and thus is a natural extension of DPO (Sec. 3.2).

3.1 Reward Alignment through Multi-Class Classification

In essence, DPO represents response rewards as LM likelihoods and constructs a *binary* classification task for learning the reward model. Given that there are more than two ($K > 2$) responses per prompt in reward datasets, we seek to construct a *multi-class* classification task for learning reward models from explicit rewards instead of preference labels. We begin by formally defining this task:

Consider a batch of K responses $\{y_i\}_{1:K}$ for an instruction x . $\{y_i\}_{1:K}$ consists of one optimal response y_ν that is sampled from $\pi^*(y|x) \propto \mu(y|x) e^{r(x,y)/\alpha}$, and $K - 1$ suboptimal noises independently sampled from $\mu(y|x)$. $\nu \in 1 : K$ is the random index of that optimal response. Our goal is to identify which of the K candidates is y_ν , given only reward labels $r(y_i)$ for each response.

Intuitively, the response with higher rewards should have a higher probability of being the target response. This intuition can be more rigorously expressed:

Proposition 3.1 (proof in Appendix A.1). *Given the above K response candidates and their respective rewards, the posterior probability for the ν -th response being drawn from π^* is*

$$p(\nu|x, \{y_i\}_{1:K}) = \frac{e^{r(x,y_\nu)}}{\sum_{i=1}^K e^{r(x,y_i)}}. \quad (4)$$

This finding is highly appealing because it shows response optimality is solely related to response rewards. This allows us to conveniently train reward models r_θ via maximum likelihood (MLE).

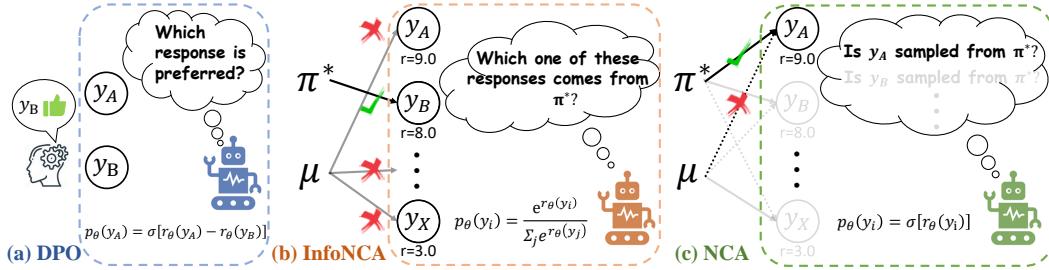


Figure 3: DPO, InfoNCA, and NCA all optimize LLM through classification tasks. **DPO** compares two responses and tells which one is preferred. **InfoNCA** compares multiple responses and identifies the one sampled from π^* (Sec. 3.1). **NCA** predicts the model source of a single response (Sec. 4.1).

Theorem 3.2 (InfoNCA, proof in Appendix A.1). *We define $\pi^*(y|x) \propto \mu(y|x)e^{r(x,y)/\alpha}$ and $\pi_\theta(y|x) \propto \mu(y|x)e^{r_\theta(x,y)}$. For any $K > 1$, $\alpha > 0$, we have:*

(a) **Equivalent objective.** The MLE objective for training r_θ has an equivalent form:

$$\max_{\theta} \mathbb{E}_{p(x, \{y_i\})} \log p_\theta(\nu|x, \{y_i\}) \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \left[\frac{e^{r(x, y_i)/\alpha}}{Z(x)} \log \frac{e^{r_\theta(x, y_i)}}{\sum_{j=1}^K e^{r_\theta(x, y_j)}} \right], \quad (5)$$

where $Z(x) = \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}$.

(b) **Optimal solution.** Assume unlimited model capacity. The optimal r_{θ^*} and π_{θ^*} for solving (5) are

$$r_{\theta^*}(x, y) = r(x, y)/\alpha + C(x),$$

and $\pi_{\theta^*}(y|x) = \pi^*(y|x) \propto \mu(y|x)e^{r(x,y)/\alpha}$,

where $C(x)$ is an arbitrary function conditioning on x .

In practical implementation of Eq. 5, we estimate $Z(x) \approx \frac{1}{K} \sum e^{r_j/\alpha}$ and parameterize $r_\theta(x, y) := \beta \log \frac{\pi_\theta(y|x)}{\mu(y|x)}$ following DPO (Eq. 2). The loss function becomes

$$\mathcal{L}_\theta^{\text{InfoNCA}}(x, \{y_i, r_i\}_{1:K}) = - \sum_{i=1}^K \left[\underbrace{\frac{e^{r_i/\alpha}}{\sum_{j=1}^K e^{r_j/\alpha}}}_{\text{soft labels}} \log \underbrace{\frac{e^{r_\theta(x, y_i)}}{\sum_{j=1}^K e^{r_\theta(x, y_j)}}}_{\text{predicted probability}} \right]_{r_\theta(x, y) := \beta \log \frac{\pi_\theta(y|x)}{\mu(y|x)}}. \quad (6)$$

The loss function in Eq. 6 is termed **InfoNCA**, where **A** stands for **Alignment**. This naming reflects its functional similarity to the Infomation Noise Contrastive Estimation (InfoNCE, [24]). Both methods transform generative modeling problems into classification tasks by contrasting multiple data points.

How does InfoNCA work? InfoNCA loss (Eq. 6) can be seen as a K -category cross-entropy loss. The soft label is calculated by dataset rewards through a softmax operation. The model's predictions are represented by learned reward r_θ . The loss reaches 0 when $r_{\theta^*}(x, y) = r(x, y)/\alpha + C(x)$.

3.2 InfoNCA Subsumes DPO as A Special Case

Below we show that DPO is a special case of InfoNCA asymptotically. Specifically, setting response number $K = 2$ and reward temperature $\alpha \rightarrow 0$, we can fully recover the DPO objective:

$$\begin{aligned} \mathcal{L}_\theta^{\text{InfoNCA}}(x, \{y_i, r_i\}_{1:K}) &= - \sum_{i=1}^K \left[\mathbb{1}(r_i = r_{\max}) \log \frac{e^{r_\theta(x, y_i)}}{\sum_{j=1}^K e^{r_\theta(x, y_j)}} \right] && (\alpha \rightarrow 0) \\ &= - \log \frac{e^{r_\theta(x, y_w)}}{e^{r_\theta(x, y_w)} + e^{r_\theta(x, y_l)}} && (K = 2, \text{ suppose } r_w > r_l) \\ &= - \log \sigma(r_\theta(x, y_w) - r_\theta(x, y_l)) && (\text{DPO loss, Eq. 3}) \end{aligned}$$

Alignment Method	InfoNCA (Sec. 3)	NCA (Sec. 4)
Modeling Target	$\pi^*(y x) \propto \mu(y x)e^{r(x,y)/\alpha}$	
Model Definition	$\pi_\theta(y x) \propto \mu(y x)e^{r_\theta(x,y)}$	$\pi_\theta(y x) = \mu(y x)e^{r_\theta(x,y)}$
Reward Dataset	$x \rightarrow \{y_i, r_i\}_{1:K}$	
Loss ($K > 1, \alpha > 0$)	$-\sum_{i=1}^K \left[\frac{e^{r_i/\alpha}}{\sum_j e^{r_j/\alpha}} \log \frac{e^{r_\theta(x,y_i)}}{\sum_j e^{r_\theta(x,y_j)}} \right]$	$-\sum_{i=1}^K \left[\frac{e^{r_i/\alpha}}{\sum_j e^{r_j/\alpha}} \log \sigma(r_\theta(x, y_i)) + \frac{1}{K} \log \sigma(-r_\theta(x, y_i)) \right]$
Preference Dataset	$x \rightarrow \{y_w > y_l\}$	
Loss ($K=2, \alpha \rightarrow 0$)	$-\log \sigma(r_\theta(x, y_w) - r_\theta(x, y_l))$ (DPO)	$-\log \sigma(r_\theta(x, y_w)) - \frac{1}{2} \sum_{y \in \{y_w, y_l\}} \log \sigma(-r_\theta(x, y))$
Loss Type	InfoNCE loss [24]	NCE loss [14]
Optimizing Target	<i>relative</i> value of log likelihood ratio	<i>absolute</i> value of log likelihood ratio
Optimal $r_{\theta^*}(x, y)$	$r(x, y)/\alpha + C(x)$	$r(x, y)/\alpha - \log \mathbb{E}_{\mu(y x)} e^{r(x,y)/\alpha}$
$r_{\theta^*}(x, y_{\text{best}}) \geq 0 ?$	not guaranteed	✓

Table 1: Comparison of NCA and InfoNCA algorithm for aligning language models. Both reward loss and pairwise preference loss are given. We provide pseudocode in Appendix B.

Empirical effect for varying hyperparameter K and α . As indicated by the derivation above, the root difference between preference-based and reward-based methods lies in the choices of K and α .

K affects how accurately we can estimate the partition function $Z(x) \approx \sum_{j=1}^K e^{r_j/\alpha}$ in Eq. 5. In practice, we find larger K can lead to better performance (Sec. 5.1). α indicates a trade-off between diversity and optimality. At $\alpha \rightarrow 0$, the InfoNCA loss increases the likelihood only for the optimal response and decreases it for all other responses, turning the reward dataset $x \rightarrow \{y_i, r_i\}$ into a preference dataset $x \rightarrow \{y_w > y_l\}$. We provide ablation studies of α and β in Appendix D.

4 NCA: Fixing Decreased Response Likelihood Issue for InfoNCA

A well-observed issue with DPO is that the likelihood of *all* responses continually *decrease* throughout training [27, 32]. We find InfoNCA shares this trend due to their inherent equivalence. Decreased data likelihood is concerning because it directly contradicts the maximum likelihood objective for supervised training and may eventually harm performance [44].

We hypothesize the main cause of this decreasing likelihood is that InfoNCA methods only adjust *relative* rewards among responses, rather than optimizing their *absolute* value. To address this problem, we take inspiration from NCE, another contrastive learning method parallel to InfoNCE, and propose **NCA**(lignment) (Sec. 4.1). Similar to InfoNCA, NCA can also guarantee convergence to the optimal LM policy under ideal conditions (Theorem 4.2). However, it directly learns the absolute reward for each response, thereby counteracting the decreasing likelihood trend (Sec. 4.2).

4.1 Reward Alignment through Absolute Reward Prediction

To avoid optimizing relative rewards across multiple responses, we construct a binary classification task that deals with a single response.

Specifically, imagine sampling a response y randomly from either the optimal LM $\pi^*(y|x) = \mu(y|x) \frac{e^{r(x,y)/\alpha}}{Z(x)}$, or the pretrained LM $\mu(y|x)$. The marginal probability of y is $p(y|x) := \frac{1}{2} \mu(y|x) + \frac{1}{2} \pi^*(y|x)$. Our goal is to guess its model source when given a response y and its reward $r(y)$.

Proposition 4.1 (proof in Appendix A.2). *Let a binary variable $\nu = 1$ indicates the response y is sampled from π^* . The posterior probability of the distribution source given the response y satisfies:*

$$p(\nu = 1|x, y) = \frac{\pi^*(y|x)}{\mu(y|x) + \pi^*(y|x)} = \frac{e^{r(x,y)/\alpha}}{Z(x) + e^{r(x,y)/\alpha}}. \quad (7)$$

Note that $p(\nu|x, y)$ is related to the partition function $Z(x)$. In order to represent model likelihood $p_\theta(\nu|x, y)$ by only employing r_θ similarly to Proposition 3.1, we have to redefine

$\pi_\theta(y|x) = \mu(y|x)e^{r_\theta(x,y)}$ by absorbing Z_θ into r_θ . Then we have

$$p_\theta(\nu = 1|x, y) = \frac{\pi_\theta(y|x)}{\mu(y|x) + \pi_\theta(y|x)} = \sigma(r_\theta(x, y)). \quad (8)$$

Similarly to Theorem 3.2, we can derive a MLE-based training objective for optimizing r_θ .

Theorem 4.2 (NCA, proof in Appendix A.2). *Let $\alpha > 0$, we have the maximum likelihood objective:*

(a) Equivalent objective.

$$\max_\theta \mathbb{E}_{p(x,y)} \log p_\theta(\nu|x, y) \Leftrightarrow \min_\theta -\mathbb{E}_{p(x)\mu(y|x)} \left[\frac{e^{r(x,y)/\alpha}}{Z(x)} \log \sigma(r_\theta(x, y)) + \log \sigma(-r_\theta(x, y)) \right], \quad (9)$$

where $Z(x) = \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}$.

(b) Optimal solution. Assume unlimited model capacity. The optimal r_{θ^*} and π_{θ^*} for solving (9) are

$$\begin{aligned} r_{\theta^*}(x, y) &= r(x, y)/\alpha - \log \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}, \\ \text{and } \pi_{\theta^*}(y|x) &\propto \mu(y|x) e^{r(x,y)/\alpha}. \end{aligned} \quad (10)$$

For reward datasets $(x \rightarrow \{y_i, r_i\}_{1:K})$, we estimate $Z(x) \approx \sum_{i=1}^K e^{r_i/\alpha}$ in Eq. 9 and construct r_θ similarly to InfoNCA:

$$\mathcal{L}_\theta^{\text{NCA}}(x, \{y_i, r_i\}_{1:K}) = - \sum_{i=1}^K \left[\underbrace{\frac{e^{r_i/\alpha}}{\sum_{j=1}^K e^{r_j/\alpha}}}_{\text{softmax weight}} \underbrace{\log \sigma(r_\theta(x, y_i))}_{\substack{\text{optimizer } \uparrow \\ (\text{increasing force})}} + \underbrace{\frac{1}{K} \log \sigma(-r_\theta(x, y_i))}_{\substack{\text{regularizer } \downarrow \\ (\text{decreasing force})}} \right]_{r_\theta(x,y) := \beta \log \frac{\pi_\theta(y|x)}{\mu(y|x)}} \quad (11)$$

How does NCA work? The loss function for NCA involves two opposing forces that jointly determine the trend of increasing or decreasing $r_\theta(x, y)$. Since $\log \sigma(\cdot)$ is a monotonically increasing function, the first term in Eq. 11 tends to increase $r_\theta(x, y)$ while the second term tends to decrease it.

At the start of training, when $r_\theta = -r_\theta = 0$, the direction of the combined force for r_θ is decided by the difference in their weights, expressed as $\frac{e^{r_i/\alpha}}{\sum_{j=1}^K e^{r_j/\alpha}} - \frac{1}{K}$. Responses with higher rewards would, in principle, attain higher likelihood after training.

4.2 Connection between NCA and InfoNCA/DPO

Although both NCA and InfoNCA originate from solving a noise contrastive classification problem, their optimization targets are markedly different (Table 1).

InfoNCA and DPO both calibrate *relative* values of reward models across various responses $\{y_i\}_{1:K}$ for an instruction x . In other words, the absolute value of $r_\theta(x, y)$ is not directly constrained. This can lead to some counterintuitive behaviors. For instance, the learned reward for even the highest-reward response could decrease over time without contradicting the loss definition, as long as the reward margin keeps increasing. This could lead to poor performance or training instability (Sec. 5.2).

In contrast, NCA specifically focuses on optimizing *absolute* values of the reward model. This characteristic is determined by its model definition: $\pi_\theta(y|x) = \mu(y|x)e^{r_\theta(x,y)}$, where r_θ has to be *self-normalized*: $\mathbb{E}_{\mu(y|x)} e^{r_\theta(x,y)} = 1$. In practice, NCA effectively prevents the likelihood of the preferred responses from decreasing. We find this is particularly helpful for math and coding tasks.

5 Experiments

We mainly seek to answer two questions in our experiments:

1. If we have access to reward-annotated datasets with >2 responses per prompt, does InfoNCA or NCA offer empirical improvement compared with preference-based approaches that simply prune reward datasets into preference datasets? (Sec. 5.1)
2. If only pairwise preference data is available, when should one choose NCA over DPO? What benefits does NCA offer? (Sec. 5.2) Note that InfoNCA is exactly DPO in this setting.

	Name	Annotation Type	MT-bench	AlpacaEval	Win vs. DPO
Baseline	Mixtral-7B-sft	SFT Data	6.45	85.20	-
	+KTO [11]	Preference	7.12	<u>91.93</u>	-
	+IPO [1]	Preference	7.45	90.62	-
	+DPO (Zephyr- β)	Preference	7.34	90.60	50.0
	+DPO $\times 3$	Preference	7.22	91.60	<u>58.1</u>
	+DPO $\times C_4^2$	Preference	7.38	90.29	48.1
Ours	+InfoNCA	Reward	7.63	92.35	56.9
	+NCA	Reward	<u>7.52</u>	90.31	59.4
Reference	Mixtral-ORPO- β	Preference+SFT	7.32	91.41	-
	Mistral-7B-instruct	SFT Data	6.84	92.78	-
	LLAMA2-chat-70b	Reward Model	6.86	92.66	-
	GPT-4	Reward Model	9.18	93.78	-

Table 2: Comparison between reward-based methods (InfoNCA, NCA) and preference-based methods (DPO, IPO, etc.) in LLM alignment. We focus on the general instruction-following abilities of each method measured by GPT-4 evaluations and human preference. The highest number in each benchmark is **bolded** and the second highest is underlined.

5.1 Aligning Language Models with Explicit Rewards

Reward dataset and Evaluation metric. We consider UltraFeedback [9], an instruction-following dataset annotated by GPT-4. This dataset comprises $\sim 64k$ instructions. Each instruction has 4 responses generated by various LMs. GPT-4 rates each response with a scalar reward on a scale of 0-10. Prior research indicates that these GPT-4 rewards closely align with human annotations [47], establishing them as an efficient, cost-effective alternative to human feedback. In order to align exactly with the definition of dataset rewards, we similarly choose well-acknowledged GPT4-based benchmarks like MT-bench [47] and AlpacaEval [19] for evaluation. Human preference studies are also conducted on evaluation prompts from MT-bench. The rating system is in Appendix C.

InfoNCA and NCA outperform preference-based methods given reward dataset. To handle reward datasets with $K > 2$ responses per instruction, one approach is to simply prune them into pairwise data and apply preference learning like DPO. For instance, Zephyr [40] selects the highest-reward response and a random remaining one from UltraFeedback for each instruction. This procedure discards two additional suboptimal responses in the dataset as well as their reward information.

In Table 2, we fine-tune a Mistral-7B model on UltraFeedback and compare InfoNCA/NCA against the DPO baseline. Results show that our methods outperform preference baselines. This improvement can be attributed to InfoNCA/NCA’s ability to exploit all information in the reward dataset.

Method	K=2	K=3	K=4
InfoNCA (MT-bench)	73.8	75.9	76.3
InfoNCA (Alpaca)	90.7	90.2	92.4
NCA (MT-bench)	73.2	73.3	75.2
NCA (Alpaca)	89.9	90.3	90.3
Average	81.9	82.4	83.5

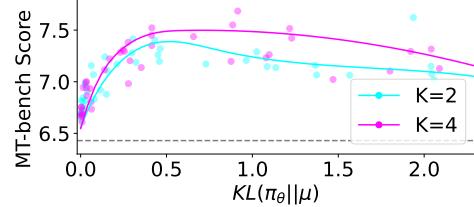


Figure 4: More suboptimal responses can also increase LLM’s instruction-following ability. We fix the highest-reward response in the UltraFeedback dataset and ablate the number of suboptimal responses per prompt, resulting in different contrastive response numbers K during training. **Left:** Evaluation results under the same set of hyperparameters. **Right:** Performance-KL trade-off under various α and β . Each dot represents an independent experiment trained for 1 epoch.

Suboptimal responses are also important. Previous practices always ensure selecting the highest-performing response when constructing preference data. The assumption behind this strategy is that the dataset’s best-performing response determines the upper limit of alignment performance. However, our experiments contradict this assumption. Results in Figure 4 indicate that extra suboptimal responses can also be advantageous for policy training. Specifically, we observe consistent

Model	Reasoning BBH (CoT)	Coding		GSMPLUS	MATH	Math		Avg.
		LeetCode	HumanEval			TheoremQA	SVAMP	
Mixtral-7B-SFT	60.9	3.3	28.1	28.5	5.8	7.0	26.9	35.8 24.5
	61.7	2.2 ↓	31.7	12.1 ↓	6.4	9.8	34.1	46.1 25.5
	60.8 ↓	3.3	26.8 ↓	32.3	11.7	11.0	65.3	74.3 35.7
Mixtral-8×7B-SFT	75.6	16.7	61.0	57.6	40.1	25.9	85.9	87.5 56.3
	74.9 ↓	17.2	47.6 ↓	55.8 ↓	35.3 ↓	26.9	67.3 ↓	75.7 ↓ 50.1 ↓
	75.6	21.1	62.8	61.5	41.6	26.9	86.8	86.9 57.9

Table 3: Alignment results for UltraInteract. We mark numbers that have decreased (↓) after training.

performance improvements when increasing the number of data responses from $K = 2$ to $K = 4$ for both InfoNCA and NCA algorithms, across various hyperparameters.

Combinatorial DPO are suboptimal solutions. Regarding the performance improvement offered by more suboptimal responses, one might predict that applying the DPO to a combinatorially constructed preference dataset would yield results comparable to NCA/InfoNCA. To investigate this, we examined two variants of DPO that utilize all available responses in UltraFeedback.

DPO $\times 3$: We pair the highest-performing response with each of the remaining three separately.

DPO $\times C_4^2$: We sum up all DPO loss possibilities for two out of the four responses.

Our experiments, detailed in Table 2, reveal that naively applying combinatorial DPO loss to leverage all response information underperforms InfoNCA/NCA. The DPO $\times 3$ shows some benefit, while DPO $\times C_4^2$ is harmful compared with simple data pruning. This is expected because InfoNCA and NCA possess theoretical guarantees (Theorem 3.2 and Theorem 4.2) that ensure convergence to the optimal LM policy whereas combinatorial preference methods do not.

5.2 NCA vs. DPO in Aligning Language Models with Pairwise Preference

In previous experiments, our focus is on the reward dataset with $K > 2$ responses per prompt ($x \rightarrow \{y_i, r_i\}_{1:K}$). However, at present most alignment datasets are pairwise ($x \rightarrow \{y_w > y_l\}$), making it essential also to evaluate our proposed methods in pairwise preference settings.

Since InfoNCA is equivalent to DPO when only pairwise preference data is available (Sec. 3.2), we will focus on comparing and clarifying the differences between the DPO and NCA algorithms.

Preferecne dataset and evaluation metrics. We consider fine-tuning Mistral-7B and Mistral-8×7B models on UltraInteract [44], a pairwise alignment dataset specifically designed for complex reasoning tasks. Before alignment, we perform SFT on UltraInteract’s preferred responses for the 8×7B model and use the existing Mistral-SFT model in Sec. 5.1. We evaluate the model’s performance in various challenging tasks. This includes BBH-Hard [39] for CoT reasoning, HumanEval [4] and LeetCode [12] for coding, GSM-Plus [18], MATH, TheoremQA [6], SVAMP [28], and ASDiv [22] for math.

DPO may hurt reasoning performance while NCA helps. Results are presented in Table 3. Overall, NCA consistently outperforms DPO in various benchmarks. Notably, we observe DPO hurts the overall performance in most reasoning tasks regarding the Mixtral-8×7B-SFT model. This indicates that DPO might not be suitable for improving reasoning abilities, which echoes findings in concurrent work [44]. In contrast, NCA shows clear improvement on both the 7B and 8×7B models.

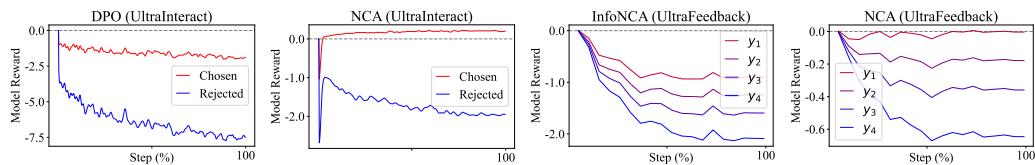


Figure 5: Comparision of data likelihood between InfoNCA/DPO and NCA.

NCA prevents the chosen-likelihood from decreasing. What distinct optimization characteristics could cause performance differences between pairwise NCA and DPO? To understand this, we empirically inspect how the data likelihood changes during training. As shown in Figure 5. The likelihood of preferred responses interestingly decreases after DPO training and increases for NCA training. This pattern is consistent across both preference and reward learning. The decreasing chosen-likelihood trend is concerning because it directly contradicts the maximum-likelihood objective used

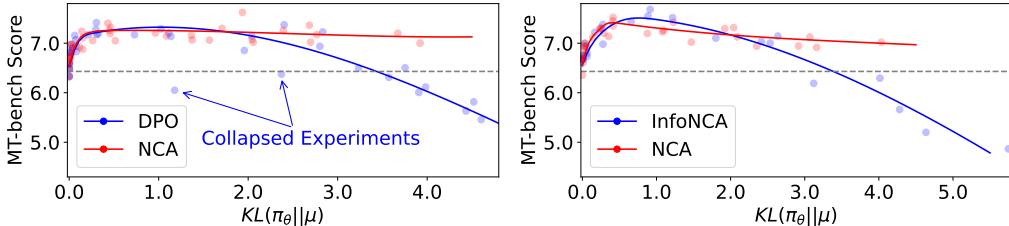


Figure 6: NCA is more robust to hyperparameter changes and KL derivations. **Left:** Ablation results of α and β for UltraFeedback-binarized. **Right:** Results for UltraFeedback-reward ($K = 4$).

during the SFT stage. This drawback is exacerbated in reasoning tasks, where the preferred response is often the ground truth answer. Consequently, we hypothesize that NCA’s superior performance in reasoning tasks is due to its ability to avoid decreasing chosen likelihood.

Since DPO is essentially a specialization of InfoNCA, their contrasting likelihood trends can be explained theoretically. As we have elaborated in Sec. 4.2, NCA adjusts the absolute likelihood of data, while DPO/InfoNCA only considers relative likelihood across different responses. Thus, a declining chosen likelihood directly contradicts NCA’s training objective but not DPO’s.

Empirical takeaway: When to choose NCA over DPO? DPO and pairwise NCA have similar theoretical guarantees. Their different performance in alignment tasks is largely empirical, depending on the specific characteristics of datasets and the nature of tasks. Our observations show that NCA is more suitable for reasoning tasks such as math and coding (Table 3), where high-quality responses are sparse, and adhering closely to the preferred responses in the dataset is critical. DPO may be more suitable for general instruction-following tasks like summarization/role-playing (Table 2), where datasets only reflect human relative preference but do not contain “golden” answers. In essence, NCA benefits from better dataset regularization, while DPO relies more on LLMs’ generalization abilities.

We also observe that NCA has a greater tolerance for divergence from the initial SFT policy and is more robust to hyperparameter changes. As evidenced in Figure 6, we conduct a grid search on α and β . DPO can drastically fail to improve alignment performance if the learned policy strays too far from the SFT policy, and may randomly collapse under certain hyperparameters. In contrast, the NCA method does not exhibit similar issues. These observations suggest that NCA may be a worthwhile alternative if DPO training is unstable.

6 Related Work

Language model alignment. Current approaches cater to either explicit reward data or preference data, often lacking the versatility to address both concurrently. Reinforcement Learning [34] is inherently suitable for explicit reward scenarios. However, its on-policy nature necessitates learning a reward model from data first, leading to an indirect two-stage optimization process [8, 26, 36]. Recent developments in preference-based alignment techniques [33, 1, 11, 42, 16, 13] have streamlined this process. They enable direct alignment of LMs through a singular loss, but this comes at the expense of being confined to pairwise preference data. Other alignment approaches [43, 37, 45, 38] are also not tailored for aligning with reward datasets. Recent work [3] attempts to extend DPO’s parameterization technique to explicit reward contexts. However, it only considers binary rewards. In comparison, our methods can handle both continuous rewards and preference data.

Noise contrastive estimation. NCE [14] and its variant, InfoNCE [24], are established optimization methods for training unnormalized generative models [21]. NCE primarily leverages a binary classification loss and can be applied in self-supervised representation learning. Examples are Word2Vec [23], MoCo [15], and SimCLR [5]. InfoNCE is related to maximizing mutual information between two distributions through a multiclass cross-entropy loss. It has successful applications in representation learning, such as CLIP [31]. It is also widely used in language modeling [7], diffusion modeling [20], and reinforcement learning [17].

7 Conclusion

In this work, we formally consider the language model alignment problem in the context of explicit reward settings. By adeptly harnessing the NCE and InfoNCE theories, we introduce two practical algorithms: NCA and InfoNCA. Our proposed methods are uniquely suited for both reward data and preference data, including DPO as a special case. Our experiments show that reward-based alignment methods outperform preference baseline by fully leveraging suboptimal responses in reward datasets. In preference settings, pairwise NCA outperforms DPO in complex reasoning tasks by effectively preventing data likelihood from decreasing.

Acknowledgments and Disclosure of Funding

We especially thank Cheng Lu, who greatly inspires us in linking the NCA algorithm with NCE theories. We also thank Jiupei Chen, and Tianlin Liu for their suggestions on Zephyr result reproduction. We thank Bingrui Li, and Weiyu Huang for their help with the experimental setup. We thank Github user Wing Lian for integrating the NCA algorithm into the trl library. We thank Haosheng Zou for providing feedback on our method.

This work was supported by NSFC Projects (Nos. 62350080, 92370124, 92248303, 62276149, 62061136001, 62076147), BNRist (BNR2022RC01006), Tsinghua Institute for Guo Qiang, and the High Performance Computing Center, Tsinghua University. J. Zhu was also supported by the XPlorer Prize.

References

- [1] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR, 2024.
- [2] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [3] Tianchi Cai, Xierui Song, Jiyan Jiang, Fei Teng, Jinjie Gu, and Guannan Zhang. Ulma: Unified language model alignment with demonstration and point-wise human preference. *arXiv preprint arXiv:2312.02554*, 2023.
- [4] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Heben Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code, 2021.
- [5] 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.
- [6] Wenhui Chen, Ming Yin, Max W.F. Ku, Yixin Wan, Xueguang Ma, Jianyu Xu, Tony Xia, Xinyi Wang, and Pan Lu. Theoremqa: A theorem-driven question answering dataset. *ArXiv preprint, abs/2305.12524*, 2023.
- [7] Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. InfoXLM: An information-theoretic framework for cross-lingual language model pre-training. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online, June 2021. Association for Computational Linguistics.
- [8] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- [9] Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, et al. Ultrafeedback: Boosting language models with scaled ai feedback. In *Forty-first International Conference on Machine Learning*, 2024.
- [10] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023.
- [11] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.
- [12] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Yu Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming - the rise of code intelligence. *ArXiv preprint, abs/2401.14196*, 2024.
- [13] Yiju Guo, Ganqu Cui, Lifan Yuan, Ning Ding, Jie Xin Wang, Huimin Chen, Bowen Sun, Ruobing Xie, Jie Zhou, Yankai Lin, et al. Controllable preference optimization: Toward controllable multi-objective alignment. *arXiv preprint arXiv:2402.19085*, 2024.

- [14] Michael U Gutmann and Aapo Hyvärinen. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of machine learning research*, 13(2), 2012.
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020.
- [16] Jiwoo Hong, Noah Lee, and James Thorne. Orpo: Monolithic preference optimization without reference model. *arXiv preprint arXiv:2403.07691*, 2024.
- [17] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *International Conference on Machine Learning*, pages 5639–5650. PMLR, 2020.
- [18] Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. Gsm-plus: A comprehensive benchmark for evaluating the robustness of llms as mathematical problem solvers. *ArXiv preprint*, abs/2402.19255, 2024.
- [19] Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. AlpacaEval: An automatic evaluator of instruction-following models, 2023.
- [20] Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [21] Zhuang Ma and Michael Collins. Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency. *arXiv preprint arXiv:1809.01812*, 2018.
- [22] Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing English math word problem solvers. In *Proc. of ACL*, 2020.
- [23] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013.
- [24] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [25] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [26] 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.
- [27] Arka Pal, Deep Karkhanis, Samuel Dooley, Manley Roberts, Siddartha Naidu, and Colin White. Smaug: Fixing failure modes of preference optimisation with dpo-positive. *arXiv preprint arXiv:2402.13228*, 2024.
- [28] Arkil Patel, Satwik Bhattacharya, and Navin Goyal. Are NLP models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021.
- [29] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.
- [30] Jan Peters and Stefan Schaal. Reinforcement learning by reward-weighted regression for operational space control. In *Proceedings of the 24th international conference on Machine learning*, pages 745–750, 2007.

[31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.

[32] Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From r to q^* : Your language model is secretly a q -function. *arXiv preprint arXiv:2404.12358*, 2024.

[33] 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. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[35] John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2022.

[36] Wei Shen, Xiaoying Zhang, Yuanshun Yao, Rui Zheng, Hongyi Guo, and Yang Liu. Improving reinforcement learning from human feedback using contrastive rewards. *arXiv preprint arXiv:2403.07708*, 2024.

[37] Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. Preference ranking optimization for human alignment. *arXiv preprint arXiv:2306.17492*, 2023.

[38] Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. Preference ranking optimization for human alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024.

[39] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them. *ArXiv preprint, abs/2210.09261*, 2022.

[40] Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, et al. Zephyr: Direct distillation of lm alignment. *arXiv preprint arXiv:2310.16944*, 2023.

[41] Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, and Shengyi Huang. TRL: Transformer reinforcement learning, 2020.

[42] Chaoqi Wang, Yibo Jiang, Chenghao Yang, Han Liu, and Yuxin Chen. Beyond reverse kl: Generalizing direct preference optimization with diverse divergence constraints. *arXiv preprint arXiv:2309.16240*, 2023.

[43] Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. RRHF: Rank responses to align language models with human feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[44] Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Jia Deng, Boji Shan, Huimin Chen, Ruobing Xie, Yankai Lin, et al. Advancing llm reasoning generalists with preference trees. *arXiv preprint arXiv:2404.02078*, 2024.

[45] Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. SLIC-HF: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*, 2023.

[46] Yao Zhao, Mikhail Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu. Calibrating sequence likelihood improves conditional language generation. In *The Eleventh International Conference on Learning Representations*, 2022.

[47] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.

A Proof of Theorems

A.1 InfoNCA Objective

Recall that our optimal language policy is

$$\pi^*(y|x) = \mu(y|x) \frac{e^{r(x,y)/\alpha}}{Z(x)}. \quad (\text{Eq. 2})$$

Consider a batch of K responses $\{y_i\}_{1:K}$ for an instruction x . $\{y_i\}_{1:K}$ consists of one optimal response y_ν that is sampled from $\pi^*(y|x) \propto \mu(y|x)e^{r(x,y)/\alpha}$, and $K - 1$ suboptimal noises independently sampled from $\mu(y|x)$. $\nu \in 1 : K$ is the random index of that optimal response. The joint probability for $\{y_i\}_{1:K}$ is

$$p^{\text{joint}}(\{y_i\}_{1:K}|x, \nu) = \pi^*(y_\nu|x) \prod_{i \neq \nu} \mu(y_i|x) = \frac{\pi^*(y_\nu|x)}{\mu(y_\nu|x)} \prod_{i=1}^K \mu(y_i|x).$$

Given that the prior satisfies $p(\nu = 1) = p(\nu = 2) = \dots = p(\nu = K) = \frac{1}{K}$, the data posterior is

$$p^{\text{joint}}(\nu|x, \{y_i\}_{1:K}) = \frac{\pi^*(y_\nu|x)/\mu(y_\nu|x)}{\sum_{j=1}^K \pi^*(y_j|x)/\mu(y_j|x)}.$$

$$\begin{aligned} p(\mathcal{O} = y_i|\{y\}_{1:K}) &= \frac{p(\{y\}_{1:K}|\mathcal{O} = y_i)p(\mathcal{O} = y_i)}{\sum_{j=1}^K p(\{y\}_{1:K}|\mathcal{O} = y_j)p(\mathcal{O} = y_j)} \\ &= \frac{\pi^*(y_i|x)/\mu(y_i|x)}{\sum_{j=1}^K \pi^*(y_j|x)/\mu(y_j|x)} \\ &= \frac{e^{r(y_i)/\alpha}}{\sum_{j=1}^K e^{r(y_j)/\alpha}} \end{aligned}$$

Define model policy as

$$\pi_\theta(y|x) := \mu(y|x) \frac{e^{r_\theta(x,y)}}{Z_\theta(x)}.$$

The model posterior probability satisfies

$$p_\theta^{\text{joint}}(\nu|x, \{y_i\}_{1:K}) = \frac{e^{r_\theta(x,y_\nu)}}{\sum_{i=1}^K e^{r_\theta(x,y_i)}}.$$

Theorem A.1 (InfoNCA Objective). *For any $K > 1$, $\alpha > 0$, we have the following results.*

(a) Equivalent objective.

$$\begin{aligned} &\min_{\theta} \mathbb{E}_{p^{\text{joint}}(x, \{y_i\})} D_{\text{KL}}[p^{\text{joint}}(\nu|x, \{y_i\}) || p_\theta^{\text{joint}}(\nu|x, \{y_i\})] \\ &\iff \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \mu(y_i|x) \sum_{i=1}^K \frac{e^{r(x,y_i)/\alpha}}{Z(x)} \log \frac{e^{r_\theta(x,y_i)}}{\sum_{j=1}^K e^{r_\theta(x,y_j)}}, \end{aligned} \quad (\text{Eq. 5})$$

where $Z(x) = \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}$.

(b) Optimal solution. Assume unlimited model capacity and data samples. The optimal r_{θ^*} and π_{θ^*} for solving Eq. 5 are

$$\begin{aligned} r_{\theta^*}(x, y) &= r(x, y)/\alpha + C(x), \\ \text{and} \quad \pi_{\theta^*}(x, y) &\propto \mu(y|x) e^{r(x,y)/\alpha}, \end{aligned}$$

where $C(x)$ is an arbitrary function conditioning on x .

Proof. (a) **Equivalent objective.**

$$\begin{aligned}
& \min_{\theta} \mathbb{E}_{p^{\text{joint}}(x, \{y_i\})} D_{\text{KL}}[p^{\text{joint}}(\nu|x, \{y_i\}) || p_{\theta}^{\text{joint}}(\nu|x, \{y_i\})] \\
& \Leftrightarrow \min_{\theta} \mathbb{E}_{p^{\text{joint}}(x, \{y_i\})} \mathbb{E}_{p^{\text{joint}}(\nu|x, \{y_i\})} \log \frac{p^{\text{joint}}(\nu|x, \{y_i\})}{p_{\theta}^{\text{joint}}(\nu|x, \{y_i\})} \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p^{\text{joint}}(x, \{y_i\})} \mathbb{E}_{p^{\text{joint}}(\nu|x, \{y_i\})} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)p(\nu)} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \quad (\text{Bayes' rule}) \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)p(\nu)} \prod_{i=1}^K \mu(y_i|x) \frac{\pi^*(y_{\nu}|x)}{\mu(y_{\nu}|x)} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \quad (\text{importance sampling}) \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \mu(y_i|x) \left[\mathbb{E}_{p(\nu)} \frac{\pi^*(y_{\nu}|x)}{\mu(y_{\nu}|x)} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \right] \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \mu(y_i|x) \left[\frac{1}{K} \sum_{\nu=1}^K \frac{\pi^*(y_{\nu}|x)}{\mu(y_{\nu}|x)} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \right] \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \mu(y_i|x) \left[\sum_{\nu=1}^K \frac{e^{r(x, y_{\nu})/\alpha}}{Z(x)} \log p_{\theta}^{\text{joint}}(\nu|x, \{y_i\}) \right] \quad (\text{based on Eq. 2}) \\
& \Leftrightarrow \min_{\theta} -\mathbb{E}_{p(x)} \prod_{i=1}^K \mu(y_i|x) \left[\sum_{i=1}^K \frac{e^{r(x, y_i)/\alpha}}{Z(x)} \log \frac{e^{r_{\theta}(x, y_i)}}{\sum_{j=1}^K e^{r_{\theta}(x, y_j)}} \right] \quad (\text{change sum index})
\end{aligned}$$

(b) **Optimal solution.**

Given conclusions from (a). With unlimited model capacity, $p_{\theta}^{\text{joint}}(\nu|x, \{y_i\})$ could represent any discrete distribution, such that we can arrive at the global optimal point given infinite training data.

$$\begin{aligned}
& \mathbb{E}_{p^{\text{joint}}(x, \{y_i\})} D_{\text{KL}}[p^{\text{joint}}(\nu|x, \{y_i\}) || p_{\theta^*}^{\text{joint}}(\nu|x, \{y_i\})] = 0 \\
& \implies p^{\text{joint}}(\nu|x, \{y_i\}) = p_{\theta^*}^{\text{joint}}(\nu|x, \{y_i\}) \quad \forall x, \nu, \{y_i\}_{1:K} \\
& \implies \frac{\pi^*(y_{\nu}|x)/\mu(y_{\nu}|x)}{\sum_{i=1}^K \pi^*(y_i|x)/\mu(y_i|x)} = \frac{e^{r_{\theta^*}(x, y_{\nu})}}{\sum_{i=1}^K e^{r_{\theta^*}(x, y_i)}} \quad \forall x, \nu, \{y_i\}_{1:K} \\
& \implies r_{\theta^*}(x, y) = r(x, y)/\alpha + C(x) \quad \forall x, y, C \\
& \implies \pi_{\theta^*}(x, y) \propto \mu(y|x) e^{r(x, y)/\alpha} \quad \forall x, y
\end{aligned}$$

□

A.2 NCA Objective

Recall the optimal language policy is

$$\pi^*(y|x) = \mu(y|x) \frac{e^{r(x, y)/\alpha}}{Z(x)}. \quad (\text{Eq. 2})$$

Consider a response y randomly sampled from either the optimal LM $\pi^*(y|x) = \mu(y|x) \frac{e^{r(x, y)/\alpha}}{Z(x)}$, or the pretrained LM $\mu(y|x)$. Let a binary variable $\nu = 1$ indicates the response y is sampled from π^* .

Then the marginal distribution of y is

$$p^{\text{joint}}(y|x) := p(\nu = 0)\mu(y|x) + p(\nu = 1)\pi^*(y|x).$$

Given the prior $p(\nu = 0) = p(\nu = 1) = \frac{1}{2}$, using Bayes' Rule, the data posterior satisfies

$$p^{\text{joint}}(\nu = 0|x, y) = \frac{\mu(y|x)}{\mu(y|x) + \pi^*(y|x)}.$$

$$p^{\text{joint}}(\nu = 1|x, y) = \frac{\pi^*(y|x)}{\mu(y|x) + \pi^*(y|x)}.$$

Define model policy as $\pi_\theta(y|x) := \mu(y|x)e^{r_\theta(x,y)}$. The model posterior probability satisfies

$$\begin{aligned} p_\theta^{\text{joint}}(\nu = 1|x, y) &= \sigma(r_\theta(x, y)), \\ p_\theta^{\text{joint}}(\nu = 0|x, y) &= 1 - \sigma(r_\theta(x, y)) = \sigma(-r_\theta(x, y)). \end{aligned} \quad (\text{Eq. 8})$$

Theorem A.2 (NCA Objective). *For any $\alpha > 0$, we have the following results.*

(a) Equivalent objective.

$$\begin{aligned} &\min_\theta \mathbb{E}_{p^{\text{joint}}(x, y)} D_{\text{KL}}[p^{\text{joint}}(\nu|x, y) || p_\theta^{\text{joint}}(\nu|x, y)] \\ &\iff \min_\theta -\mathbb{E}_{p(x)\mu(y|x)} \frac{e^{r(x,y)/\alpha}}{Z(x)} \log \sigma(r_\theta(x, y)) + \log \sigma(-r_\theta(x, y)), \end{aligned} \quad (\text{Eq. 9})$$

where $Z(x) = \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}$.

(b) Optimal solution. Assume unlimited model capacity and data samples. The optimal r_{θ^*} and π_{θ^*} for solving Eq. 5 are

$$\begin{aligned} r_{\theta^*}(x, y) &= r(x, y)/\alpha - \log \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha}, \\ \text{and } \pi_{\theta^*}(x, y) &\propto \mu(y|x) e^{r(x,y)/\alpha}. \end{aligned}$$

Proof. **(a) Equivalent objective.**

$$\begin{aligned} &\min_\theta \mathbb{E}_{p^{\text{joint}}(x, y)} D_{\text{KL}}[p^{\text{joint}}(\nu|x, y) || p_\theta^{\text{joint}}(\nu|x, y)] \\ &\iff \min_\theta \mathbb{E}_{p^{\text{joint}}(x, y)} \mathbb{E}_{p^{\text{joint}}(\nu|x, y)} \log \frac{p^{\text{joint}}(\nu|x, y)}{p_\theta^{\text{joint}}(\nu|x, y)} \\ &\iff \min_\theta -\mathbb{E}_{p^{\text{joint}}(x, y)} \mathbb{E}_{p^{\text{joint}}(\nu|x, y)} \log p_\theta^{\text{joint}}(\nu|x, y) \\ &\iff \min_\theta -\mathbb{E}_{p(x)p(\nu)p^{\text{joint}}(y|x, \nu)} \log p_\theta^{\text{joint}}(\nu|x, y) \quad (\text{Bayes' rule}) \\ &\iff \min_\theta -[p(\nu = 0) \mathbb{E}_{p(x)p^{\text{joint}}(y|x, \nu=0)} \log p_\theta^{\text{joint}}(\nu = 0|x, y) + \\ &\quad p(\nu = 1) \mathbb{E}_{p(x)p^{\text{joint}}(y|x, \nu=1)} \log p_\theta^{\text{joint}}(\nu = 1|x, y)] \\ &\iff \min_\theta -[\mathbb{E}_{p(x)\mu(y|x)} \log \sigma(-r_\theta(x, y)) + \mathbb{E}_{p(x)\pi^*(y|x)} \log \sigma(r_\theta(x, y))] \quad (\text{by Eq. 8}) \\ &\iff \min_\theta -\mathbb{E}_{p(x)\mu(y|x)} \frac{e^{r(x,y)/\alpha}}{Z(x)} \log \sigma(r_\theta(x, y)) + \log \sigma(-r_\theta(x, y)) \quad (\text{importance sampling}) \end{aligned}$$

(b) Optimal solution.

Given conclusions from (a). With unlimited model capacity, $p_\theta^{\text{joint}}(\nu|x, y)$ could represent any discrete distribution, such that we can arrive at the global optimal point given infinite training data.

$$\begin{aligned} &\mathbb{E}_{p^{\text{joint}}(x, y)} D_{\text{KL}}[p^{\text{joint}}(\nu|x, y) || p_{\theta^*}^{\text{joint}}(\nu|x, y)] = 0 \\ &\implies p^{\text{joint}}(\nu|x, y) = p_{\theta^*}^{\text{joint}}(\nu|x, y) \quad \forall x, \nu, y \\ &\implies p^{\text{joint}}(\nu = 1|x, y) = p_{\theta^*}^{\text{joint}}(\nu = 1|x, y) \quad \forall x, y \\ &\implies \frac{\pi^*(y|x)}{\mu(y|x) + \pi^*(y|x)} = \sigma(r_{\theta^*}(x, y)) = \frac{e^{r_{\theta^*}(x,y)}}{1 + e^{r_{\theta^*}(x,y)}} \quad \forall x, y \\ &\implies e^{r_{\theta^*}(x,y)} = \frac{\pi^*(y|x)}{\mu(y|x)} \quad \forall x, y \\ &\implies r_{\theta^*}(x, y) = r(x, y)/\alpha - \log \mathbb{E}_{\mu(y|x)} e^{r(x,y)/\alpha} \quad \forall x, y \\ &\implies \pi_{\theta^*}(x, y) \propto \mu(y|x) e^{r(x,y)/\alpha} \quad \forall x, y \end{aligned}$$

□

B Pseudocode

PyTorch code for the InfoNCA/NCA loss for reward datasets is provided below:

```
import torch.nn.functional as F

def reward_loss(pi_logps, ref_logps, rewards, alpha, beta, loss_type):
    """
    pi_logps: policy logprobs for K responses, shape (B, K)
    ref_logps: reference logprobs for K responses, shape (B, K)
    rewards: reward labels for K responses, shape (B, K)
    alpha: the reward temperature controlling strength of KL penalty
    beta: the parameterization coefficient that defines the reward model
    loss_type: could be either "InfoNCA" or "NCA" loss
    """

    soft_labels = (rewards / alpha).softmax(dim=-1) # (B, K)

    model_rewards = (pi_logps - ref_logps) * beta # (B, K)

    if loss_type == "InfoNCA":
        model_logps = model_rewards.log_softmax(dim=-1) # (B, K)
        losses = - (soft_labels * model_logps).sum(dim=-1) # (B,)

    elif loss_type == "NCA":
        optimization = - (soft_labels * F.logsigmoid(model_rewards)).sum(dim=-1) # (B,)
        regularization = - F.logsigmoid(-model_rewards).mean(dim=-1) # (B,)
        losses = optimization + regularization # (B,)

    return losses.mean()
```

The loss implementation under pairwise preference settings is equivalent to reward losses with $K = 2$ and $\alpha \rightarrow 0$. We provide the code separately for easy comparison with DPO.

```
def preference_loss(chosen_pi_logps, chosen_ref_logps,
                    rejected_pi_logps, rejected_ref_logps,
                    beta, loss_type):
    """
    chosen_pi_logps: policy logprobs for the preferred responses, shape (B, )
    chosen_ref_logps: reference logprobs for the preferred responses, shape (B, )
    rejected_pi_logps: policy logprobs for the dispreferred responses, shape (B, )
    rejected_ref_logps: reference logprobs for the dispreferred responses, shape (B, )
    beta: the parameterization coefficient that defines the reward model
    loss_type: one of "InfoNCA", "NCA" or "DPO" loss
    """

    chosen_rewards = (chosen_pi_logps - chosen_ref_logps) * beta # (B, )
    rejected_rewards = (rejected_pi_logps - rejected_ref_logps) * beta # (B, )

    if loss_type in ["DPO", "InfoNCA"]:
        losses = -F.logsigmoid(chosen_rewards - rejected_rewards) # (B,)

    elif loss_type == "NCA":
        losses = - F.logsigmoid(chosen_rewards) \
            - 0.5 * F.logsigmoid(-chosen_rewards) \
            - 0.5 * F.logsigmoid(-rejected_rewards) # (B,)

    return losses.mean()
```

C Experimental Details

Experiments with UltraFeedback. Our implementation is heavily based on the Transformer Reinforcement Learning (TRL) library [41] and Zephyr’s official code base [40]. All models are fine-tuned from the publicly accessible HuggingFaceH4/mistral-7B-SFT-beta model. Experiments are run on Nvidia A40 or RTX 4090 GPUs using bfloat16 precision. We ablate $\beta \in \{3e-4, 1e-3, 3e-3, 1e-2, 3e-2, 1e-1, 3e-1, 1.0\}$ and $\alpha \in \{0.01, 0.1, 0.33, 1.0, 3.33\}$. The default reward temperature α is 0.01. The default parameterization coefficient β is also 0.01. We adopt the QLoRA [10] fine-tuning technique with rank 16, $\alpha_{lora} = 16$, and a dropout rate of 0.05. We train all models for 1 epoch. The batch size is 32. We use an AdamW optimizer with a learning rating of 5e-6. For KTO and IPO baselines, we adopt exactly the same training pipeline for reporting their performance except that we tune the $\beta \in \{0.01, 0.1, 0.3, 0.5, 1.0\}$. We find the most suitable *beta* for KTO is 0.01, and for IPO is 0.5.

Experiments with UltraInteract. We follow [44] and fine-tune all parameters of the pretrained model with UltraInteract and UltraFeedback. Specifically, for HuggingFaceH4/mistral-7B-SFT-beta model, we directly adopt the recipe of Eurus, with the same data mixture (all 220K multi-turn trajectory pairs in UltraInteract and all 340K pairs in UltraFeedback) and hyperparameters. All β is set to 0.1. The model is trained for 1 epoch with a cosine schedule. The learning rate is 5e-7 and the warmup ratio is 0.1. For Mixtral-8x7B, we first SFT the model with the recipe for Eurus-70B-SFT, and then applied the same preference learning recipe as aforementioned.

My MT-Bench judge

Single Answer Grading Pairwise Comparison

writing 91: Pretend yourself to be Elon Musk in all the following conversations. Speak like Elon Musk as much as possible. Why ...

Model A Model B

User
Pretend yourself to be Elon Musk in all the following conversations. Speak like Elon Musk as much as possible. Why do we need to go to Mars?

Assistant A
Certainly, I'm Elon Musk, the founder and CEO of SpaceX, Tesla, Neuralink, and The Boring Company. Our mission at SpaceX is to enable human life on Mars and make humanity multi-planetary. We need to go to Mars because it's vital for the long-term survival of our species. Earth is facing numerous existential threats, such as natural disasters, climate change, and the potential for a catastrophic event that could render our planet uninhabitable. By establishing a self-sustaining colony on Mars, we can ensure the continued existence of our species and safeguard the future of humanity. Additionally, the scientific and technological advancements that will be made during the process of colonizing Mars will have far-reaching implications for our understanding of the universe and our ability to explore it. Mars is also the most geographically similar planet to Earth, providing a perfect testing ground for technologies that could help us better understand and mitigate the challenges we face on our own planet. In short, our goal is to ensure that humanity is not limited by the narrow confines of our home planet, but rather has the freedom to explore the universe and thrive in new environments.

User's follow-up question
How do you like dancing? Can you teach me?

Assistant A
Elon Musk is primarily focused on his work in the fields of renewable energy, space exploration, and transportation, and he doesn't often speak about his personal interests outside of these areas. While I'm not privy to his personal hobbies, based on what has been reported in the media, Elon Musk doesn't appear to be a regular dancer, and he has not publicly expressed an interest in teaching others to dance. It's best to assume that his focus is primarily on his professional pursuits. If you're looking for a dance instructor, you may want to consider finding a local studio or experienced dance teacher in your area.

Assistant B
Elon Musk: "I'm not really one for dancing, but I've always admired the art form. However, I'm not a dance teacher. I suggest you find a reputable dance school or instructor to learn from. There are many different styles of dance to choose from, so consider what kind of dancing you're interested in and find someone who specializes in that area."

Model Judgment (first turn)
A is better B is better

Model Judgment (second turn)
A is better B is better

Figure 7: Rating system screenshot. Our rating system is adapted from [47].

D Additional Experiment Results

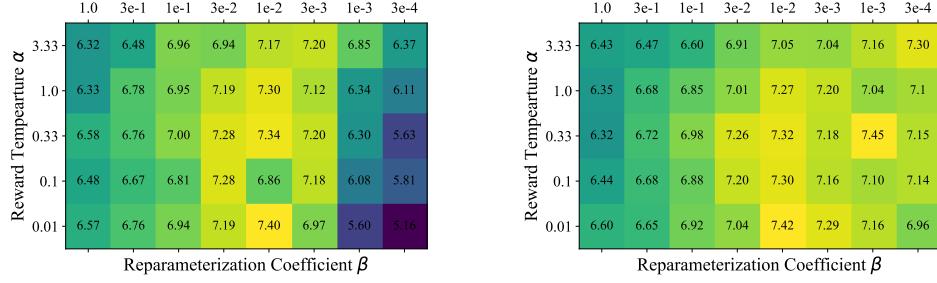


Figure 8: MT-bench score for **InfoNCA (left)** and **NCA (right)** regarding various α and β . Results are averaged for $K = 2$ and $K = 4$ experiments. Overall, model performance is relatively more sensitive to variations of β compared with α . NCA shows greater tolerance to hyperparameter variations.

Method	Functional Form	MT-bench	KL Divergence
Mistral-7B-SFT	—	6.45	0.00
+DPO	$r_\theta = \beta \log \frac{\pi_\theta}{\mu}$	7.34	0.53
+InfoNCA	$r_\theta = \beta \log \frac{\pi_\theta}{\mu}$	7.63	0.47
	$r_\theta = \beta \log \pi_\theta$	4.65	4.39
	$r_\theta = \beta \log \pi_\theta + \text{SFT regularization}$	7.15	0.26
+NCA	$r_\theta = \beta \log \frac{\pi_\theta}{\mu}$	7.52	0.39
	$r_\theta = \beta \log \pi_\theta$	5.82	3.32
	$r_\theta = \beta \log \pi_\theta + \text{SFT regularization}$	6.88	0.09

Table 4: Comparison of different functional forms of r_θ on UltraFeedback datasets.

E Appendix F: Additional Empirical Results

We provide sampled responses of LM policies trained respectively with the DPO and NCA to compare their data regularization feature. In general, generations from the NCA policy are more similar to the preferred response in the dataset.

Prompt

You are given an answer. Create a question for the answer.

Grams to Ounces conversion 1 gram (g) is equal to 0.03527396195 ounces (oz).

Preferred Response in Dataset (r=8.0)

What is the conversion rate for grams to ounces for 1 gram?

Dispreferred Response in Dataset (r=7.0)

What is the relationship between grams and ounces in weight measurement?

Reponse generated by $\pi_{\theta}^{\text{DPO}}$ (r=2.0)

How many ounces are equivalent to 200 grams according to the given conversion rate of 1 gram (g) to 0.03527396195 ounces (oz)?

Reponse generated by $\pi_{\theta}^{\text{NCA}}$ (r=10.0)

Question: What is the conversion rate from grams to ounces, and how many grams are one ounce?

Table 5: Sampled generations and the log-likelihood ratio of the NCA and the DPO/InfoNCA policy.

Prompt

If "A busy crosswalk somewhere in asia." does that mean that "The crosswalk was blocked off by police."?

Options:

- yes; - it is not possible to tell ; - no

Preferred Response in Dataset (r=8.0)

It is not possible to tell.

Confidence: 95%

Dispreferred Response in Dataset (r=7.5)

Question:

If "A busy crosswalk somewhere in asia." does that mean that "The crosswalk was blocked off by police."?

Answer:

It is not possible to tell.

Confidence: 90%

Reponse generated by $\pi_{\theta}^{\text{DPO}}$ (r=3.0)

- It is not possible to determine whether "The crosswalk was blocked off by police." in "A busy crosswalk somewhere in Asia." To determine whether the crosswalk was blocked off by police, you would need additional information or context.

Reponse generated by $\pi_{\theta}^{\text{NCA}}$ (r=8.0)

It is not possible to tell.

Confidence: 95%

Table 6: Sampled generations and the log-likelihood ratio of the NCA and the DPO/InfoNCA policy.

Prompt

In this task, you need to count the occurrences of the given word in the given sentence.

Sentence: 'a green double decker bus sitting on top of a parking lot'. Count the occurrences of the word 'a' in the given sentence.

Preferred Response in Dataset (r=9.0)

The word 'a' appears 3 times in the given sentence.

Confidence: 100%

Dispreferred Response in Dataset (r=6.0)

The given sentence is "a green double decker bus sitting on top of a parking lot." To answer your question about the occurrences of the word 'a', I will count the number of times the word appears in the sentence.

After counting the occurrences of the word 'a' in the sentence, I found that it appears once. Therefore, the answer to your question is that the word 'a' appears once in the given sentence.

Reponse generated by $\pi_{\theta}^{\text{DPO}}$ (r=4.0)

Of course, I'm here to help! There are 11 occurrences of the word 'a' in the given sentence. Here's how you can count them: 1. 'a' in 'a' 2. 'a' in 'green' 3. 'a' in 'double' 4. 'a' in 'decker' 5. 'a' in 'bus' 6. 'a' in 'sitting' 7. 'a' in 'on' ...

Confidence: 90%

Reponse generated by $\pi_{\theta}^{\text{NCA}}$ (r=8.0)

Sure! The word 'a' appears 3 times in the given sentence.

Table 7: Sampled generations and the log-likelihood ratio of the NCA and the DPO/InfoNCA policy.

Prompt

if $X^2 + 3 = 12$. What is X?

Preferred Response in Dataset (r=9.0)

To find the value of X, we'll first need to isolate X^2 :

$$X^2 + 3 = 12$$

Subtract 3 from both sides:

$$X^2 = 9$$

Now, we can find the square root of both sides of the equation to solve for X:

$$X = \pm 3$$

So, X can be either 3 or -3.

Dispreferred Response in Dataset (r=4.0)

I'm sorry, I cannot solve this algebra equation. Can I assist you with anything else?

Reponse generated by $\pi_{\theta}^{\text{DPO}}$ (r=7.0)

Sure, I'd love to help. By solving the equation, the value of X should be 3.

Reponse generated by $\pi_{\theta}^{\text{NCA}}$ (r=9.0)

To find the value of X, we need to solve the equation $X^2 + 3 = 12$.

First, subtract 3 from both sides:

$$X^2 = 12 - 3$$

$X^2 = 9$

Next, take the square root of both sides:

$$X = \pm\sqrt{9}$$

So, X can be either 3 or -3.

Table 8: Sampled generations and the log-likelihood ratio of the NCA and the DPO/InfoNCA policy.

Appendix H: Comparing NCA with DPO on Data Regularization.

Suppose there are only 4 possible responses $\{y_a, y_b, y_c, y_d\}$. Each response has a reward r , and a pretrained language model policy μ :

-	y_a	y_b	y_c	y_d
Likelihood $\mu(y)$	40%	50%	5%	5%
Reward $r(y)$	10	3	7	0

Table 9: Initial probabilities and rewards

Given a preference dataset $D = \{y_a > y_b\}$ (y_c and y_d do not exist in the dataset), the loss functions are defined as:

$$L_{\theta}^{\text{DPO}} = -\log \sigma \left(\log \frac{\pi_{\theta}(y_a)}{\mu(y_a)} - \log \frac{\pi_{\theta}(y_b)}{\mu(y_b)} \right)$$

$$L_{\theta}^{\text{NCA}} = -\log \sigma \left(\log \frac{\pi_{\theta}(y_a)}{\mu(y_a)} \right) - \frac{1}{2} \log \sigma \left(-\log \frac{\pi_{\theta}(y_a)}{\mu(y_a)} \right) - \frac{1}{2} \log \sigma \left(-\log \frac{\pi_{\theta}(y_b)}{\mu(y_b)} \right)$$

After fine-tuning, there are several possibilities for π_{θ} :

-	$\pi_{\theta}(y_a)$	$\pi_{\theta}(y_b)$	$\pi_{\theta}(y_c)$	$\pi_{\theta}(y_d)$	$\log \frac{\pi_{\theta}(y_a)}{\mu(y_a)} - \log \frac{\pi_{\theta}(y_b)}{\mu(y_b)}$	\bar{r}
(1)	20% \downarrow	10% $\downarrow\downarrow$	5%	65% \uparrow	0.916 > 0	2.65 \downarrow
(2)	20% \downarrow	10% $\downarrow\downarrow$	65% \uparrow	5%	0.916 > 0	6.85 \uparrow
(3) Wanted	60% \uparrow	30% \downarrow	5%	5%	0.916 > 0	7.25 \uparrow

Table 10: Post fine-tuning probabilities and rewards

In scenarios (1) and (2), we can see that the likelihood for both y_a and y_b decreases. However, (1) and (2) satisfy the DPO loss function because the likelihood for y_b decreases more, and the relative likelihood margin between y_a and y_b becomes larger. In (1) and (2), the likelihood for either y_c or y_d increases because $\pi(y_a) \downarrow + \pi(y_b) \downarrow + \pi(y_c) \uparrow + \pi(y_d) \uparrow = 1$. However, y_c and y_d are unreliable because we do not know their quality (rewards). The LM policy could generalize to a low-quality response like y_d (case (1)).

In contrast, the NCA effectively prevents the winning response likelihood $\pi(y_a)$ from decreasing, because it mainly optimizes the absolute data likelihood instead of just caring about the relative likelihood margin $\log \frac{\pi_{\theta}(y_a)}{\mu(y_a)} - \log \frac{\pi_{\theta}(y_b)}{\mu(y_b)}$. Thus, we say NCA is more likely to assign a larger likelihood to responses within the dataset.

Appendix G: Comparision with Related Works

We compare with a prior work SLiC-HF [45] in this section. SLiC is inspired by [46] and similarly aims to calibrate sequence likelihood to align with human preferences. Given a preference data pair $\{x \rightarrow y_w > y_l\}$, the loss function for SLiC is

$$L_{\theta} := \max(0, \delta - \log \pi_{\theta}(y_w|x) + \log \pi_{\theta}(y_w|x)) - \lambda \log \pi_{\theta}(y_w|x),$$

where δ is a hyperparameter that controls the likelihood margin of data, and λ controls the regularization weight of the loss. The main difference between our proposed method and the SLiC loss can be summarized as follows:

- **Theoretical framework.** SLiC is mainly adapted from the existing LM calibration methods [46]. In contrast, our method is based on noise contrastive estimation methods [14, 24].
- **Policy regularization.** The training process of SLiC is regularized by the additional SFT loss controlled by λ . In contrast, our proposed method is regularized through the parameterization technique $r_{\theta} = \beta \log \frac{\pi_{\theta}}{\mu}$ controlled by β .
- **Learning target.** SLiC directly optimizes the policy model π_{θ} , while our method directly optimizes the residual model r_{θ} .

F Limitations and Broder Impacts

Limitations. We mainly discuss the theoretical limitations of our methods. Although both InfoNCA and NCA enjoy strong theoretical guarantees that ensure convergence to the optimal LM distribution, these guarantees only hold given infinite response numbers per instruction ($K \rightarrow \infty$) during training to ensure an accurate estimate of $Z(x)$, which is clearly unrealistic. Additionally, while more responses employed during training improve performance, they also require higher GPU memory consumption, leading to a trade-off between performance and efficiency.

Broder Impacts. This research contributes to the field of machine learning by enhancing language model alignment with human preferences. While these improvements can lead to more accurate and ethically aligned AI-generated content, they also pose ethical challenges. Notably, there is a risk of reinforcing biases present in data or human preferences. Additionally, the increasing capabilities of language models may impact employment in sectors reliant on language skills. It is crucial to continually assess these ethical considerations and societal impacts, ensuring that advancements in machine learning are guided by responsible and inclusive practices.

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: Yes they do.

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: Appendix F.

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: [\[Yes\]](#)

Justification: Appendix A.1 and A.2

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: Details in Appendix C. The source code is in supplemental material. We will open-source it upon publication.

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?

Answer: **[Yes]**

Justification: The source code is in supplemental material. We will open-source it upon publication. All datasets used have open access.

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, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[Yes]**

Justification: Appendix C.

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: The experiments are quite deterministic. Just like most previous work, we run most experiments only once, so it can be difficult to report an error bar. Instead, we validate the robustness of our methods mainly by ablating various hyperparameters (e.g., Figure 6 and Figure 4. A standard derivation can be calculated from the figure, or we can report the std error given by the Alpaca benchmark, but it is really unnecessary.)

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: Appendix C.

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>?

Answer: [Yes]

Justification: Yes, we have.

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: [\[Yes\]](#)

Justification: Appendix F.

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: [\[Yes\]](#)

Justification: Yes, we have.

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: Yes, we have.

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 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: [\[Yes\]](#)

Justification: Yes, we have.

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: [\[Yes\]](#)

Justification: Annotators are graduate students in the university. The system screenshot is in Appendix C. They get paid at least the minimum wage in the country.

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: [Yes]

Justification: The participants of experiments only conduct small-scale evaluations of language models. There are no obvious potential risks. All experiments comply with the relevant laws and ethical standards of the institutions in the respective countries, and they have received the necessary approvals.

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.