

---

# Perplexity-aware Correction for Robust Alignment with Noisy Preferences

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Alignment techniques are critical in ensuring that large language models (LLMs)  
2 output helpful and harmless content by enforcing the LLM-generated content to  
3 align with preferences. However, the existence of noisy preferences (NPs), where  
4 the responses are mistakenly labelled as chosen or rejected, could deteriorate the  
5 alignment, thus making the LLMs generate useless and even malicious content.  
6 Existing methods mitigate the issue of NPs from the loss perspective by adjusting  
7 the alignment loss based on a clean validation dataset. Orthogonal to these loss-  
8 oriented methods, we propose perplexity-aware correction (PerpCorrect) from  
9 the data perspective for robust alignment which detects and corrects NPs based  
10 on the differences between the perplexity of the chosen and rejected responses  
11 (dubbed as PPLDiff). Intuitively, a higher PPLDiff indicates a higher probability  
12 of the NP because a rejected/chosen response which is mistakenly labelled as  
13 chosen/rejected is less preferable to be generated by an aligned LLM, thus having  
14 a higher/lower perplexity. PerpCorrect works in three steps: (1) PerpCorrect aligns  
15 a surrogate LLM using the clean validation data to make the PPLDiff able to  
16 distinguish clean preferences (CPs) and NPs. (2) PerpCorrect further aligns the  
17 surrogate LLM by incorporating the reliably clean training data whose PPLDiff is  
18 extremely small and reliably noisy training data whose PPLDiff is extremely large  
19 after correction to boost the discriminatory power. (3) Detecting and correcting  
20 NPs according to the PPLDiff obtained by the aligned surrogate LLM to obtain  
21 a denoised training dataset for robust alignment. Comprehensive experiments  
22 validate that our proposed PerpCorrect can achieve state-of-the-art alignment  
23 performance under NPs. Notably, PerpCorrect demonstrates practical utility by  
24 requiring only a modest number of validation data and being compatible with  
25 various alignment techniques. Our code is available at the Anonymous GitHub.

## 26 1 Introduction

27 Alignment enables the safe utilization of the remarkable capabilities acquired by large language  
28 models (LLMs) through self-supervised learning on vast corpora [4, 17, 2]. It refers to the process of  
29 ensuring that the contents generated by LLMs are helpful, harmless, and aligned with human values  
30 and preferences [13]. Reinforcement Learning from Human Feedback (RLHF) [7] emerges as a  
31 primary technique for achieving alignment. Current technical routes [29, 30, 22] require a reward  
32 model to simulate human preference and use it to optimize policy model outputs with Proximal  
33 Policy Optimization (PPO) [20]. Current offline techniques such as Direct Preference Optimisation  
34 (DPO) [19], Sequence Likelihood Calibration with Human Feedback (SLiC) [28] and Identity-  
35 Preference Optimisation (IPO) [1], could directly align LLMs without intensely computational  
36 training a reward model as employed in RLHF.

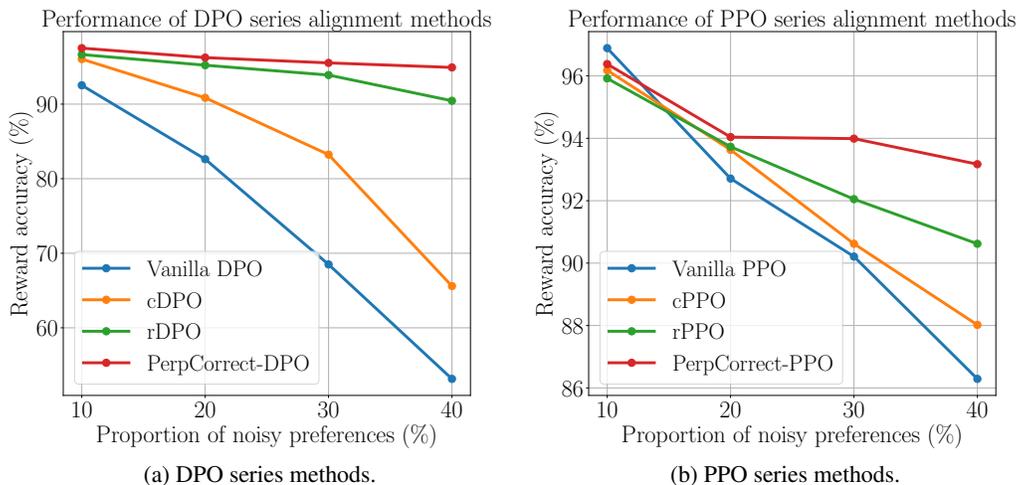


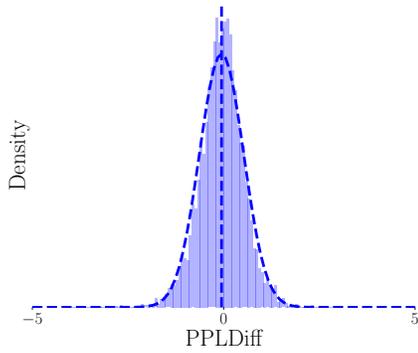
Figure 1: We evaluated various robust alignment methods under different proportions of noisy preferences using the Llama2-7B model, on the Golden HH dataset. The reward accuracy of both the vanilla DPO and PPO method significantly decreases as the proportion of noisy preferences increases. Our method, perplexity-aware correction (PerpCorrect), outperforms both the DPO and PPO series baselines across different proportions of noisy preferences.

Recent studies [25, 6] have shown there exist noisy preferences (NPs) that may lead to significant degradation in alignment performance. The issue of NPs, where the label of the actually chosen/rejected responses in training datasets is flipped as rejected/chosen, can arise from the biases of annotators [25] and the malicious noise injection [3]. As shown in Figure 1, when NPs are randomly injected into the training dataset, the conventional alignment method (e.g., DPO [19] and PPO [7]) will yield significantly degraded alignment performance measured by the reward accuracy. Such performance degradation could result in the generation of useless and even malicious content [25]. Therefore, it necessitates developing robust alignment methods that can utilize datasets with NPs to effectively align the LLMs with human preferences.

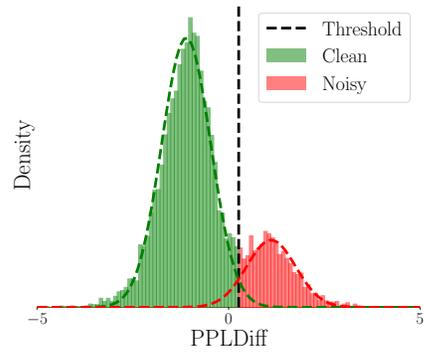
Existing robust alignment methods are proposed from the loss perspective, which adjust the alignment loss using a clean validation dataset to mitigate the issue of NPs. Particularly, the conservative DPO (cDPO) [15] and robust-DPO (rDPO) [6] both estimate the proportion of NPs using the clean validation data via cross-validation and then adjust the original DPO loss based on the estimated proportion of NPs. However, Mitchell [15] and Chowdhury et al. [6] overlooked the essential differences between noisy and clean preferences, which is critical for mitigating the issue of NPs.

To this end, we propose **Perplexity-aware Correction** (PerpCorrect) for robust alignment from the data perspective by leveraging the differences between noisy and clean preferences for robust alignment. PerpCorrect detects and corrects NPs based on the difference between the perplexity of the chosen response and that of the rejected counterparts (dubbed as PPLDiff) obtained by an aligned surrogate LLM using the clean validation set. If an NP is detected, PerpCorrect will correct it by flipping the label of the rejected/chosen responses as chosen/rejected. Intuitively, rejected responses which are mistakenly labelled as chosen have a higher perplexity since they are less consistent with human preferences and thus have a lower probability of being generated after alignment. Therefore, a higher value of PPLDiff indicates a higher probability of the preferences being noisy. In this way, PerpCorrect leverages the differences between noisy and clean preferences (CPs) identified by PPLDiff to detect NPs.

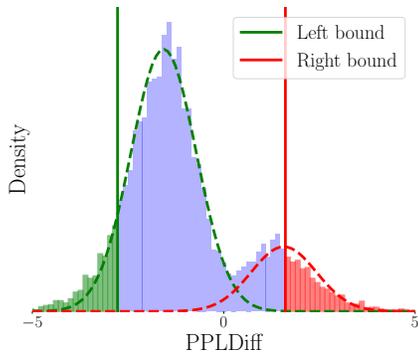
To make the PPLDiff able to distinguish CPs and NPs, PerpCorrect requires an aligned surrogate LLM for calculating PPLDiff. The density of PPLDiff obtained on the noisy training dataset using an unaligned surrogate LLM, which can be fitted as a normal distribution centered around zero (evidenced in Figure 2a), cannot discriminate CPs and NPs. Therefore, we align a surrogate LLM using the clean validation data. The density of PPLDiff obtained by the aligned surrogate LLM in Figure 2b can be fitted into two distinguishable normal distributions, thus being able to differentiate CPs and NPs.



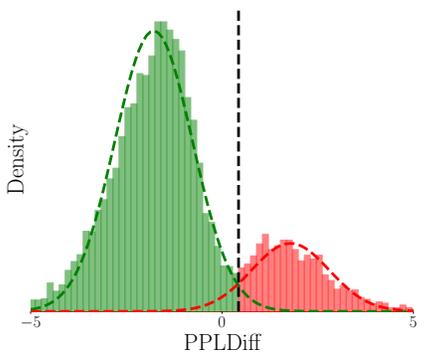
(a) Noisy and clean preferences cannot be distinguished by PPLDiff.



(b) The large overlap between two distributions leads to flawed NP detection.



(c) Aligning the surrogate LLM using extra reliable training data.



(d) Separating and correcting noisy preferences based on the threshold.

Figure 2: We visualized the PPLDiff under the entire PerpCorrect process using Llama2-7B on Golden HH dataset with 20% noisy preferences. We use the green dotted line to represent the normal distribution formed by clean data, the red dotted line represents the normal distribution formed by noisy data, and the black dotted line represents the threshold.

70 However, there still exists a large overlap between two normal distributions after aligning only on  
 71 the clean validation dataset, which could result in an unsatisfactory accuracy of NP detection. To  
 72 this end, we iteratively align the model using more reliably clean training data with extremely low  
 73 PPLDiff (located in the green area in Figure 2c) and reliable noisy training data with extremely large  
 74 PPLDiff (located in the red area in Figure 2c) sampled from noisy training datasets. Finally, the two  
 75 normal distributions are significantly separated as shown in Figure 2d, which indicates that PPLDiff  
 76 has an enhanced discriminatory power.

77 Benefiting from the strong discriminatory power of PPLDiff calculated by the aligned surrogate LLM,  
 78 PerpCorrect outputs a denoised training dataset for robust alignment by detecting NPs based on a  
 79 PPLDiff threshold and conducting correction. The data, whose PPLDiff is below a certain threshold  
 80 (i.e., the black dotted line in Figure 2d) selected as the x-coordinate of the two normal distributions'  
 81 intersection, are identified as NPs and thus corrected by flipping the response's label. Notably, our  
 82 proposed PerpCorrect is compatible with various alignment methods as well as robust alignment  
 83 methods [15, 6] since the metric PPLDiff is agnostic to training algorithms and only requires an  
 84 arguably small number of clean validation data (~50), thus yielding significantly practical usage.

85 Comprehensive empirical results, evaluated using the Llama2-7B [24] and phi-2 [14] models on the  
 86 OpenAssistant Conversations (OASST1) [11] and Golden HH [5] datasets, validate the effectiveness  
 87 of our proposed PerpCorrect method in robustifying alignment with NPs. We empirically validate  
 88 that PerpCorrect consistently yields state-of-the-art performance among various proportions of NPs.  
 89 Besides, we empirically demonstrate that PerpCorrect can effectively robustify various alignment  
 90 techniques and robust alignment methods, validating its compatibility.

91 **2 Literature Review and Preliminary**

92 In this section, we introduce the related work regarding LLM alignment and provide preliminaries  
 93 about the noisy preferences, perplexity, as well as various alignment methods.

94 **2.1 LLM Alignment**

95 In the domain of aligning LLMs with human preferences, pairwise preference methods are favored  
 96 due to their lower cognitive burden on evaluators. Traditional online alignment approaches [24,  
 97 17, 22] involve training reward models from these preferences to provide signals in reinforcement  
 98 learning. Recent offline alignment methods like Direct Preference Optimization (DPO) [19], Sequence  
 99 Likelihood Calibration (SLiC) [28], and Identify Preference Optimization (IPO) [1] streamlined this  
 100 process by directly using preference pairs to train LLMs, thus enhancing performance and reducing  
 101 computational costs. Additionally, methods like RRHF [27] align LLMs using multiple ranked  
 102 preferences, Kahneman-Tversky Optimization (KTO) [9] align LLMs using a single preference  
 103 labeled as good or bad, and Rejection Sampling Optimization (RSO) [12] address DPO’s limitation  
 104 in sampling preference pairs from the optimal policy through rejection sampling. However, NPs,  
 105 arising from the biased human feedback, can determine the alignment performance [17, 25]. Robust  
 106 alignment methods like conservative DPO (cDPO) [15], robust DPO (rDPO) [6] have been proposed  
 107 to address these issues from the loss perspective. Our approach focuses on the data perspective to  
 108 address these issues of NPs and is orthogonal to these robust alignment methods.

109 **2.2 Preliminary**

110 **Noisy preferences (NPs).** NPs refer to preference data in training datasets, whose label of the  
 111 actually chosen/rejected responses is flipped as rejected/chosen. Let  $\mathcal{D} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_{i=1}^N$  be  
 112 the preference dataset consisting of  $N \in \mathbb{N}$  preference data points. For each preference data point  
 113  $(x, y_w, y_l) \in \mathcal{D}$ ,  $x$  is the prompt input to LLMs,  $y_w$  is the chosen response, and  $y_l$  is the rejected  
 114 response. We let  $\tilde{\mathcal{D}} = \{(x^{(i)}, \tilde{y}_w^{(i)}, \tilde{y}_l^{(i)})\}_{i=1}^N$  be the noisy preference dataset (i.e., preference dataset  
 115 consisting noisy preferences) and denote preference data points that are not noisy as clean preferences  
 116 (CPs). Following Chowdhury et al. [6], we obtain the noisy preference dataset  $\tilde{\mathcal{D}}$  using the standard  
 117 random noise model [16] with the probability  $\varepsilon \in (0, 50\%)$  to change the data point into noisy  
 118 preferences, i.e.

$$\mathbb{P}_{(x^{(i)}, \tilde{y}_w^{(i)}, \tilde{y}_l^{(i)}) \sim \tilde{\mathcal{D}}} \left[ (x^{(i)}, \tilde{y}_w^{(i)}, \tilde{y}_l^{(i)}) = (x^{(i)}, y_l^{(i)}, y_w^{(i)}) \right] = \varepsilon. \quad (1)$$

119 **Perplexity (PPL).** PPL [10] measures the probability that the LLM generates a sentence. A lower  
 120 PPL of a sentence indicates that the LLM generates this sentence in a high probability. PPL is defined  
 121 as the average negative log-likelihood of a sequence, i.e.,

$$\text{PPL}(s; \theta) = \exp\left(-\frac{1}{t} \sum_{i=1}^t \log \pi_{\theta}(s_i | s_{<i})\right), \quad (2)$$

122 where  $s$  is a sequence composed of  $t$  tokens and  $\log \pi_{\theta}(s_i | s_{<i})$  denotes the log-likelihood of the  $i$ -th  
 123 token given the preceding tokens  $s_{<i}$  calculated by an LLM  $\pi_{\theta}$ .

124 **Technical details of alignment methods.** There are usually three phases in RLHF pipeline [25, 19]:  
 125 (1) supervised fine-tuning (SFT); (2) reward modeling; (3) reinforcement learning (RL) optimization.  
 126 In the SFT phase, an LLM is fine-tuned via supervised learning on high-quality task-related data.  
 127 We denote the LLM after the SFT phase as  $\pi_{\text{SFT}}$ . In the reward modeling phase, the reward model  
 128 is introduced to simulate human preferences. Given a preference dataset, a reward model  $r_{\omega}(x, y)$   
 129 parameterized by  $\omega$ , which takes prompt  $x$  and response  $y$  as input and outputs a real number  
 130 representing the reward score, can be optimized via minimizing the following loss function:

$$\mathcal{L}_R(r_{\omega}, \mathcal{D}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\log \sigma(r_{\omega}(x, y_w) - r_{\omega}(x, y_l))], \quad (3)$$

131 where  $\sigma$  is the logistic function. In the RL optimization phase, the objective function is as follows:

$$\max_{\theta} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta}(y|x)} [r_{\omega}(x, y) - \beta \cdot (\log \pi_{\theta}(y|x) - \log \pi_{\text{ref}}(y|x))], \quad (4)$$

132 where  $\pi_\theta(y|x)$  represents the probability that the LLM parameterized by  $\theta > 0$  generates the response  
 133  $y$  given the prompt  $x$ ,  $\pi_{\text{ref}}$  is a reference LLM to maintain the generation ability of the aligned model,  
 134 and  $\beta$  is a hyper-parameter to ensure the similarity between  $\pi_\theta(y|x)$  and  $\pi_{\text{ref}}(y|x)$ . We take  $\pi_{\text{SFT}}$   
 135 as the reference LLM  $\pi_{\text{ref}}$  following Ouyang et al. [17].

136 Recently, offline alignment methods directly leverages preferences in preference datasets, bypassing  
 137 the need to learn a reward model in RLHF. The LLM parameter is optimized by minimizing the  
 138 following loss function:

$$\mathcal{L}(\pi_\theta; \pi_{\text{ref}}) = \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\mathcal{G}(x, y_w, y_l; \theta)], \quad (5)$$

139 where the function  $\mathcal{G}$  changes with the alignment method. To be specific, DPO [19] uses a BCE loss,  
 140 SLiC [28] uses a hinge loss, and IPO [1] uses a square loss:

$$\mathcal{G}_{\text{DPO}}(x, y_w, y_l; \theta) = -\log \sigma \left( \beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right), \quad (6)$$

$$\mathcal{G}_{\text{SLiC}}(x, y_w, y_l; \theta) = \max \left\{ 0, 1 - \left( \beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right\}, \quad (7)$$

$$\mathcal{G}_{\text{IPO}}(x, y_w, y_l; \theta) = \left( \beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} - \frac{1}{2} \right)^2. \quad (8)$$

141 To mitigate the issue of NPs, cDPO [15] and rDPO [6] adjust the DPO loss based on the estimated  
 142 proportion of NPs  $\varepsilon'$  using a clean validation dataset  $\mathcal{D}_{\text{val}} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_{i=1}^{N_{\text{val}}}$  consisting of  
 143  $N_{\text{val}} \in \mathcal{N}$  clean preference data points, i.e.

$$\mathcal{G}_{\text{cDPO}}(x, \tilde{y}_w, \tilde{y}_l; \theta) = (1 - \varepsilon') \mathcal{G}_{\text{DPO}}(x, \tilde{y}_w, \tilde{y}_l; \theta) + \varepsilon' \mathcal{G}_{\text{DPO}}(x, \tilde{y}_l, \tilde{y}_w; \theta), \quad (9)$$

$$\mathcal{G}_{\text{rDPO}}(x, \tilde{y}_w, \tilde{y}_l; \theta) = \frac{(1 - \varepsilon') \mathcal{G}_{\text{DPO}}(x, \tilde{y}_w, \tilde{y}_l; \theta) - \varepsilon' \mathcal{G}_{\text{DPO}}(x, \tilde{y}_l, \tilde{y}_w; \theta)}{1 - 2\varepsilon'}. \quad (10)$$

### 144 3 Perplexity-aware Correction for Robust Alignment

145 This section introduces **Perplexity-aware Correction** (PerpCorrect) for robust alignment with NPs.  
 146 In Section 3.1, we introduce a novel metric called PPLDiff and then illustrates the pipeline of  
 147 PerpCorrect to detect and correct NPs based on PPLDiff. In Section 3.2, we demonstrate how to  
 148 adapt our proposed PerpCorrect with various alignment methods to achieve robust alignment.

#### 149 3.1 Perplexity-aware Correction (PerpCorrect)

150 In this subsection, we introduce PerpCorrect which leverages a novel metric called PPLDiff as  
 151 the foundation for detecting and correcting NPs. The algorithm of PerpCorrect is demonstrated in  
 152 Algorithm 2.

153 **PPLDiff.** PPLDiff measures the difference between the PPL of chosen response and that of the  
 154 rejected response. Given a preference data point  $(x, \tilde{y}_w, \tilde{y}_l) \in \tilde{\mathcal{D}}$  sampled from the noisy training  
 155 dataset  $\tilde{\mathcal{D}}$  and an LLM  $\pi_\theta$ , PPLDiff is defined as follows:

$$\text{PPLDiff}(x, \tilde{y}_w, \tilde{y}_l; \theta) = \log \text{PPL}([x; \tilde{y}_w]; \theta) - \log \text{PPL}([x; \tilde{y}_l]; \theta). \quad (11)$$

156 where  $[x; y]$  indicates the concatenation of the prompt  $x$  and the response  $y$ . Intuitively, if a data point  
 157 is a clean preference, the  $\text{PPL}([x; \tilde{y}_w]; \theta)$  will be lower than  $\text{PPL}([x; \tilde{y}_l]; \theta)$  because the sequence  
 158  $[x; \tilde{y}_w]$  is more aligned with human values and thus has a higher probability of being generated by  
 159 aligned LLMs. As a result, it PPLDiff will be lower compared to NPs, which  $\text{PPL}([x; \tilde{y}_w]; \theta)$  is  
 160 higher than  $\text{PPL}([x; \tilde{y}_l]; \theta)$ . This difference allows us distinguish CPs and NPs based on PPLDiff.

161 **Aligning a surrogate LLM only using clean validation data.** Here, we leverage a clean validation  
 162 dataset  $\mathcal{D}_{\text{val}}$  to obtain an aligned surrogate LLM to make PPLDiff able to distinguish CPs and NPs.  
 163 We empirically find that the PPLDiff values of CPs and NPs calculated by an unaligned LLM in the  
 164 noisy training dataset were initially indistinguishable as shown in Figure 2a, making it impossible to

165 differentiate the NPs from CPs. This is because an unaligned LLM lacks the necessary preferences to  
 166 distinguish NPs and CPs.

167 Therefore, we introduce a surrogate LLM  $\pi_{\theta'}$  parameterized by  $\theta'$  to replace the unaligned LLM and  
 168 use it for calculating PPLDiff. We optimize the surrogate LLM  $\pi_{\theta'}$  using the clean validation dataset  
 169  $\mathcal{D}_{\text{val}}$  as follows:

$$\max_{\theta'} \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}_{\text{val}}} [\mathcal{G}_{\text{DPO}}(x, y_w, y_l; \theta)]. \quad (12)$$

170 After aligning the surrogate LLM, the PPLDiff values of NPs calculated by the surrogate LLM  $\pi_{\theta'}$  are  
 171 significantly increased and those of CPs are significantly decreased, forming two distinct distributions  
 172 as shown in Figure 2b. This is because the aligned surrogate LLM is trained to generate responses  
 173 that align with human preferences, enhancing its ability to distinguish between NPs and CPs based  
 174 on PPLDiff.

175 To separate CPs and NPs in the noisy training dataset without knowing the oracle preferences, we  
 176 leverage the Levenberg-Marquardt (LM) algorithm to find two normal distributions that fit the density  
 177 of PPLDiff calculated by the aligned surrogate LLM. Specifically, the LM algorithm returns the  
 178 constants  $\bar{\varepsilon}, \bar{\mu}, \bar{\sigma}$  that satisfies the following condition:

$$h(x|\bar{\varepsilon}, \bar{\mu}, \bar{\sigma}) = (1 - \bar{\varepsilon})f_{\text{clean}}(x|\bar{\mu}, \bar{\sigma}^2) + \bar{\varepsilon}f_{\text{noisy}}(x|\bar{\mu}, \bar{\sigma}^2), \quad (13)$$

$$\text{where } f(x|\mu, \sigma^2) = \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right). \quad (14)$$

179 Note that  $x$  is the PPLDiff value and  $h(x|\bar{\varepsilon}, \bar{\mu}, \bar{\sigma})$  is the superposition of these two normal distribution.  
 180 We denote  $f_{\text{clean}}(x|\bar{\mu}, \bar{\sigma}^2)$  as the normal distribution fitting the PPLDiff of CPs and  $f_{\text{noisy}}(x|\bar{\mu}, \bar{\sigma}^2)$   
 181 as the normal distribution fitting the PPLDiff of NPs since the PPLDiff of NPs is intuitively higher  
 182 than that of CPs. In this way, we can obtain two distinguishable normal distributions to separate  
 183 NPs and CPs as shown in the green and red dotted lines of Figure 2b without knowing the oracle  
 184 preferences.

185 **Further aligning the surrogate LLM using extra reliable training data from noisy training**  
 186 **datasets.** After aligning only using the clean validation datasets, the discriminatory power of the  
 187 PPLDiff is still far from satisfactory because of the large overlap between the two normal distributions.  
 188 Therefore, we align the surrogate LLM with more reliable training data to make the PPLDiff of CPs  
 189 and that of NPs more separable. We iteratively align the surrogate LLM  $\pi_{\theta'}$  using more reliably clean  
 190 training data whose PPLDiff is extremely small and reliably noisy training data whose PPLDiff is  
 191 extremely large after correction by flipping the label of the response.

192 Specifically, at epoch  $t \in \mathbb{N}$ , we select  $(t - 1) \cdot \alpha \cdot |\tilde{\mathcal{D}}|$  of the training data along with the clean  
 193 validation data for further alignment where  $\alpha \in (0, 1)$  is the selection ratio and  $|\tilde{\mathcal{D}}| = N$  is the  
 194 number of data points in noisy training dataset. As shown in Lines 33–45 of Algorithm 2, the selected  
 195 reliable training dataset  $\mathcal{D}'_t$  consists of  $(t - 1) \cdot \alpha \cdot (1 - \bar{\varepsilon}) \cdot |\tilde{\mathcal{D}}|$  reliably clean training data whose  
 196 PPLDiff values are smallest  $(t - 1) \cdot \alpha \cdot (1 - \bar{\varepsilon})$  percent and  $(t - 1) \cdot \alpha \cdot \bar{\varepsilon} \cdot |\tilde{\mathcal{D}}|$  reliably noisy training  
 197 data after correction. Note that the reliably clean training data are the data points whose PPLDiff  
 198 values are smallest  $(t - 1) \cdot \alpha \cdot (1 - \bar{\varepsilon})$  percent (located in the green area of Figure 2c), and the  
 199 reliably noisy training data whose PPLDiff values are largest  $(t - 1) \cdot \alpha \cdot \bar{\varepsilon}$  percent (located in the red  
 200 area of Figure 2c) among all the training data points.

201 **Detecting and correcting NPs based on PPLDiff to output a denoised training dataset.** Based  
 202 on the PPLDiff calculated by the aligned surrogate LLM, PerpCorrect detects and corrects NPs  
 203 whose PPLDiff value is lower than a certain threshold. We take the x-coordinate of the intersection  
 204 of the two normal distributions as the threshold (the black dotted line in Figure 2d). As shown in  
 205 Lines 23–31, data points whose PPLDiff values are larger than this threshold are identified as CPs  
 206 (the green area in Figure 2d), and other data points are identified as NPs requiring correction (the red  
 207 area in Figure 2d). In this way, we can obtain a denoised training dataset for robust alignment.

208 Further, we select an optimal denoised training dataset to further enhance the performance of  
 209 robust alignment according to the intersection area of the two normal distributions. We denote the  
 210 intersection area of two normal distributions as the estimated NP proportion of the denoised training  
 211 dataset, i.e.,

$$\varepsilon'_{PC} = \int_{-\text{inf}}^{+\text{inf}} \min\{(1 - \bar{\varepsilon})f_{\text{clean}}(x|\bar{\mu}, \bar{\sigma}^2), \bar{\varepsilon}f_{\text{noisy}}(x|\bar{\mu}, \bar{\sigma}^2)\} dx, \quad (15)$$

---

**Algorithm 1** Robust Alignment via Perplexity-aware Correction (PerpCorrect)

---

- 1: **Input:** Noisy training dataset  $\tilde{\mathcal{D}}$ , clean validation dataset  $\mathcal{D}_{\text{val}}$ , and pre-trained LLM  $\pi_\theta$  parameterized by  $\theta$
  - 2: **Output:** Robust alignment model  $\pi_\theta$
  - 3: // Stage I: Supervised fine-tuning (SFT)
  - 4:  $\pi_\theta \leftarrow$  Supervised fine-tuned LLM  $\pi_\theta$ . (Details in Appendix C.3)
  - 5: // Stage II: Perplexity-aware correction using the surrogate LLM
  - 6:  $\tilde{\mathcal{D}}_{\text{denoised}}, \varepsilon'_{\text{denoised}} \leftarrow$  Perplexity-aware Correction ( $\pi_\theta, \tilde{\mathcal{D}}, \mathcal{D}_{\text{val}}$ ) (Details in Algorithm 2)
  - 7: // Stage III: Alignment with denoised dataset
  - 8:  $\pi_\theta \leftarrow$  Aligned LLM  $\pi_\theta$  using  $\tilde{\mathcal{D}}_{\text{denoised}}$  and  $\varepsilon'_{\text{denoised}}$  (Details in Appendix C.3)
- 

212 where  $\varepsilon'_{PC}$  calculates the ratio of noisy data points which are not detected by PerpCorrect (i.e.,  
213 the green area enclosed by the black and red lines in Figure 2d) and the clean data points which  
214 are mistakenly detected by PerpCorrect (i.e., the red area enclosed by the black and green lines in  
215 Figure 2d). In this way,  $\varepsilon'_{PC}$  can efficiently calculate the NP proportion of the denoised training  
216 dataset. We take the denoised training dataset with the smallest  $\varepsilon'_{PC}$  among multiple iterations as the  
217 optimal one for robust alignment to boost alignment performance.

### 218 3.2 Robust Alignment

219 Here, we introduce how to adapt PerpCorrect to robustify various alignment methods and demonstrate  
220 the algorithm of robust alignment via PerpCorrect in Algorithm 1. In general, the pipeline of the  
221 robust alignment based on PerpCorrect contains three stages: SFT, PerpCorrect, and alignment. We  
222 will first conduct SFT, following Christiano et al. [7], to boost the performance of a pre-trained LLM  
223 by boosting its skills for specific tasks. Next, we will conduct PerpCorrect to detect and correct NPs  
224 and output an optimal denoised training dataset  $\tilde{\mathcal{D}}_{\text{denoised}}$  the smallest  $\varepsilon'_{PC}$  in Eq. 15. Finally, we  
225 can obtain an aligned LLM from the SFT model using the denoised training dataset  $\tilde{\mathcal{D}}_{\text{denoised}}$  via  
226 alignment (i.e., Line 8 in Algorithm 1).

227 Due to that our proposed PerpCorrect is agnostic to alignment methods and model structures,  
228 PerpCorrect is applicable to robustify both online alignment methods such as RLHF (PPO) [7] and  
229 offline alignment methods including DPO [19], SLiC [28], and IPO [1]. Besides, our proposed  
230 PerpCorrect is compatible with existing loss-oriented robust alignment methods, such as cDPO [15]  
231 and rDPO [6], based on the estimated proportion of NPs. Note that cDPO and rDPO require  
232 conducting computationally expensive cross-validation to tune the estimated proportion of NPs. We  
233 can efficiently estimate the proportion of NPs by utilizing the fitted normal distributions during  
234 PerpCorrect, i.e.,  $\varepsilon'_{PC}$  in Eq. 15. Therefore, we can combine PerpCorrect with a wide range of  
235 existing alignment methods to achieve robust alignment with NPs.

## 236 4 Experiments

237 In this section, we demonstrate that our proposed PerpCorrect achieves state-of-the-art alignment  
238 performance under different proportion of NPs and have good compatibility with other alignment  
239 methods. In Section 4.1, PerpCorrect combined with DPO [19] achieves state-of-the-art alignment  
240 performance than existing baselines (Section 4.1), including DPO [19], cDPO [15], and rDPO [6].  
241 In Section 4.2, we further analyze the impact of the number of validation data and verified the  
242 compatibility of PerpCorrect with online and offline alignment methods and robust alignment methods.  
243 The training details and compute resources are reported in Appendix C.1.

244 **Datasets.** We utilize two preference datasets, namely OpenAssistant Conversations (OASST1) [11]  
245 and Golden HH [5]. The processed OASST1 dataset comprises 17,939 training samples and 951  
246 testing samples and the processed Golden HH dataset consists of 12,066 training samples and 654  
247 testing samples. The description and processing details of these datasets are provided in Appendix C.2.

248 **Models.** Our evaluation leverages two distinct series of open-sourced LLMs with different parameter  
249 sizes: Llama2-7B [24] and phi-2 [14]. We acquire the checkpoints from their official repositories on  
250 Hugging Face. The LLMs used for PerpCorrect and those for robust alignment share the same model  
251 structure and initialization.

252 **Baselines.** We adopt vanilla DPO [19] and two robust alignment methods, cDPO [15] and rDPO [6],  
253 as baselines. For their detailed implementation, we utilize and adapt the transformers and TRL  
254 libraries provided by the Hugging Face community.

Table 1: Average reward accuracy of DPO series alignment methods using Llama2-7B on the Golden HH dataset. The standard deviation of reward accuracy is reported in Table 7

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla DPO	92.53%	82.62%	68.50%	53.15%
cDPO	96.04%	90.85%	83.23%	65.60%
rDPO	96.65%	95.22%	93.90%	90.45%
PerpCorrect-DPO	<b>97.51%</b>	<b>96.24%</b>	<b>95.53%</b>	<b>94.92%</b>

Table 2: Average reward accuracy of PPO series alignment methods using Llama2-7B on the Golden HH dataset. The standard deviation of reward accuracy is reported in Table 8

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla PPO	<b>96.64%</b>	92.71%	90.21%	86.29%
cPPO	96.18%	93.63%	90.62%	88.02%
rPPO	95.92%	93.73%	92.05%	90.62%
PerpCorrect-PPO	96.38%	<b>94.04%</b>	<b>93.99%</b>	<b>93.17%</b>

Table 3: Performance of DPO series alignment methods using phi-2 on the Golden HH dataset.

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla DPO	94.97%	82.01%	70.12%	55.79%
cDPO	98.32%	90.40%	81.10%	60.52%
rDPO	95.88%	94.51%	95.12%	88.57%
PerpCorrect-DPO	<b>98.78%</b>	<b>97.10%</b>	<b>98.32%</b>	<b>98.02%</b>

Table 4: Performance of DPO series alignment methods using phi-2 on the OASST1 dataset.

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla DPO	67.68%	63.31%	59.45%	51.82%
cDPO	67.51%	62.36%	54.66%	48.81%
rDPO	63.48%	58.82%	57.35%	51.05%
PerpCorrect-DPO	<b>71.15%</b>	<b>67.61%</b>	<b>67.58%</b>	<b>67.26%</b>

Table 5: Impact of the number of clean validation data evaluated on the Golden HH dataset using Llama2-7B with a proportion of NPs  $\varepsilon = 40\%$ .

Number	10	20	30	40	50	100	200
Reward accuracy	81.40%	88.26%	94.21%	94.21%	95.43%	95.43%	96.04%

255 **Metrics.** In accordance with Chowdhury et al. [6], we employ the winning rate of policy generations  
 256 against the selected preferences on the test dataset as our primary metric. This metric applies to vanilla  
 257 DPO [19], cDPO [15], rDPO [6], as well as other offline alignment methods including SLiC [28]  
 258 and IPO [1]. Additionally, we utilize the winning rate of the reward model score for the chosen  
 259 preferences on the test dataset as our metric for vanilla PPO [17], cPPO [15, 25], and rPPO [6]. These  
 260 two metrics are collectively called reward accuracy.

261 **4.1 PerpCorrect Achieves the State-of-the-Art Robust Alignment Performance**

262 The empirical results demonstrate that our method, PerpCorrect, achieves state-of-the-art robust  
 263 alignment performance, surpassing existing baselines such as vanilla DPO [19], cDPO [15], and  
 264 rDPO [6]. This is evident across various proportions of noisy preferences  $\varepsilon$  using different datasets  
 265 and LLMs.

266 **Comparison using different LLMs.** Tables 1 and 3 show alignment performance of DPO series  
 267 alignment methods on the Golden HH [5] dataset using Llama2-7B [24] and phi-2 [14]. At a  
 268 proportion of the NPs  $\varepsilon = 40\%$ , PerpCorrect increases the reward accuracy by 41.77% (from 53.15%  
 269 to 94.92%) using Llama2-7B and by 42.23% (from 55.79% to 98.02%) using phi-2. The empirical  
 270 result validates that our proposed PerpCorrect can be used on different sizes of LLMs and achieve  
 271 better alignment performance than baselines.

272 **Comparison on different datasets.** Tables 3 and 4 present the alignment performance of various  
 273 DPO series alignment methods on the Golden HH [5] and OASST1 [11] datasets, utilizing phi-2 [14].  
 274 The empirical results reveal a significant discrepancy in average reward accuracy between the more  
 275 complex OASST1 dataset and the Golden HH dataset. The performance of other robust alignment  
 276 methods is found to be unsatisfactory on the OASST1 dataset, often not surpassing the vanilla DPO.  
 277 In contrast, our method PerpCorrect consistently maintains strong alignment performance across  
 278 varying proportions of noisy preferences. In general, our method PerpCorrect can achieve better  
 279 alignment performance than baselines across different datasets.

280 **4.2 Ablation Study**

281 **Impact of the number of clean validation data.** Table 5 illustrates the impact of the number of  
 282 clean validation data points. We conducted experiments on the Golden HH dataset using Llama2-  
 283 7B with a proportion of NPs  $\varepsilon = 40\%$ . The empirical results indicate that as the number of clean  
 284 validation data points increases, the performance of our method, PerpCorrect, also improves. However,  
 285 when the number is too large, the improvement in performance is not obvious, and the cost of manual  
 286 annotation significantly increases.

Table 6: Reward accuracy and improvements of the offline and robust alignment methods, as well as those combined with PerpCorrect, using Llama2-7B on the Golden HH dataset.

Method	Proportion of noisy preferences (%)			
	10	20	30	40
DPO	92.53%	82.62%	68.50%	53.15%
PerpCorrect-DPO	97.51%	96.24%	95.53%	94.92%
$\Delta$	<b>+4.98%</b>	<b>+13.62%</b>	<b>+27.03%</b>	<b>+41.77%</b>
SLiC	97.56%	88.87%	83.84%	67.84%
PerpCorrect-SLiC	98.32%	96.49%	96.65%	96.34%
$\Delta$	<b>+0.76%</b>	<b>+7.62%</b>	<b>+12.80%</b>	<b>+28.51%</b>
IPO	98.02%	92.23%	81.25%	61.74%
PerpCorrect-IPO	<b>99.09%</b>	<b>99.39%</b>	<b>98.02%</b>	<b>98.93%</b>
$\Delta$	<b>+1.07%</b>	<b>+7.16%</b>	<b>+16.77%</b>	<b>+37.20%</b>
cDPO	96.04%	90.85%	83.23%	65.60%
PerpCorrect-cDPO	98.78%	98.17%	96.80%	89.18%
$\Delta$	<b>+2.74%</b>	<b>+7.32%</b>	<b>+13.57%</b>	<b>+23.58%</b>
rDPO	96.65%	95.22%	93.90%	90.45%
PerpCorrect-rDPO	96.19%	95.27%	95.73%	95.58%
$\Delta$	-0.46%	<b>+0.05%</b>	<b>+1.83%</b>	<b>+5.13%</b>

287 **Compatibility with online alignment method RLHF (PPO).** We adopt vanilla PPO [17],  
 288 cPPO [15, 25], and rPPO [6] as baselines. Table 2 shows the alignment performance of PPO  
 289 series alignment methods on the Golden HH [5] dataset using Llama2-7B. Although vanilla PPO  
 290 has good performance when the proportion of NPs is low, it still declines significantly when the  
 291 proportion is high. PerpCorrect maintains desirable alignment performances when the proportion  
 292 of NPs is high. Our empirical results show that PerpCorrect has desirable compatibility with online  
 293 alignment method RLHF (PPO).

294 **Compatibility with various offline alignment methods.** Table 6 presents the alignment perfor-  
 295 mance and improvements of original offline alignment methods compared to those combined with  
 296 PerpCorrect. Our experiments, conducted on the Golden HH dataset using Llama2-7B, reveal that  
 297 the reward accuracy of SLiC [28] and IPO [1] both significantly decrease as the proportion of NPs  
 298 increases, similar to vanilla DPO [19]. However, our method PerpCorrect enhances their alignment  
 299 performance across different proportions of NPs. Notably, IPO combined with PerpCorrect achieves  
 300 the best alignment performance. These empirical results demonstrate that our method has good  
 301 compatibility with various offline alignment methods.

302 **Compatibility with robust alignment methods.** Table 6 shows the alignment performance and  
 303 improvements of robust alignment methods compared to those combined with PerpCorrect. Our  
 304 method, PerpCorrect, can significantly enhance the performance of cDPO [15], and provide a modest  
 305 improvement for rDPO [6] under almost all proportion of NPs. The empirical results show that our  
 306 method has good compatibility with robust alignment methods.

## 307 5 Conclusions

308 This paper proposes a method called perplexity-aware correction (PerpCorrect), as an effective  
 309 approach for robust alignment with noisy preferences (NPs). PerpCorrect utilizes a surrogate LLM  
 310 to calculate a novel metric, PPLDiff, and further detects and corrects NPs from clean preferences  
 311 (CPs) based on it. PerpCorrect consists of three steps: (1) First, PerpCorrect aligns a surrogate LLM  
 312 using the clean validation dataset, enabling PPLDiff to distinguish between CPs and NPs. (2) Next,  
 313 PerpCorrect enhances the discrimination power of PPLDiff by aligning the surrogate LLM with  
 314 more reliable training data. (3) Finally, PerpCorrect detects and corrects NPs from CPs based on a  
 315 calculated threshold and obtains a denoised training dataset. The paper further proposes a robust  
 316 alignment pipeline, consisting of three stages SFT, PerpCorrect, and alignment, to achieve robust  
 317 alignment with NPs. The experimental results validate that PerpCorrect achieves state-of-the-art  
 318 alignment performance and has good compatibility with other online, offline, and robust alignment  
 319 methods. Therefore, PerpCorrect can be an effective method to mitigate the impact of NPs and can  
 320 be used for robust alignment. Future research directions include: (1) Improving the time efficiency  
 321 of PerpCorrect and (2) Reducing the amount of clean validation data required to achieve the same  
 322 alignment performance.

## References

- 323
- 324 [1] Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello,  
325 Michal Valko, and Rémi Munos. A general theoretical paradigm to understand learning from  
326 human preferences, 2023.
- 327 [2] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn  
328 Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless  
329 assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*,  
330 2022.
- 331 [3] Tim Baumgärtner, Yang Gao, Dana Alon, and Donald Metzler. Best-of-venom: Attacking rlhf  
332 by injecting poisoned preference data, 2024.
- 333 [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
334 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
335 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 336 [5] Tianchi Cai, Xierui Song, Jiyan Jiang, Fei Teng, Jinjie Gu, and Guannan Zhang. Ulma: Unified  
337 language model alignment with human demonstration and point-wise preference, 2024.
- 338 [6] Sayak Ray Chowdhury, Anush Kini, and Nagarajan Natarajan. Provably robust dpo: Aligning  
339 language models with noisy feedback, 2024.
- 340 [7] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep  
341 reinforcement learning from human preferences. *Advances in neural information processing*  
342 *systems*, 30, 2017.
- 343 [8] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient  
344 finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- 345 [9] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto:  
346 Model alignment as prospect theoretic optimization, 2024.
- 347 [10] Fred Jelinek, Robert L Mercer, Lalit R Bahl, and James K Baker. Perplexity—a measure of the  
348 difficulty of speech recognition tasks. *The Journal of the Acoustical Society of America*, 62(S1):  
349 S63–S63, 1977.
- 350 [11] Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith  
351 Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. Openassistant  
352 conversations-democratizing large language model alignment. *Advances in Neural Information*  
353 *Processing Systems*, 36, 2024.
- 354 [12] Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and  
355 Jialu Liu. Statistical rejection sampling improves preference optimization. In *The Twelfth*  
356 *International Conference on Learning Representations*, 2024. URL [https://openreview.](https://openreview.net/forum?id=xbswrrQ0e)  
357 [net/forum?id=xbswrrQ0e](https://openreview.net/forum?id=xbswrrQ0e).
- 358 [13] Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng,  
359 Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. Trustworthy llms: a survey and  
360 guideline for evaluating large language models’ alignment, 2024.
- 361 [14] Microsoft. Phi-2: The surprising power of small language mod-  
362 els, 2023. URL [https://www.microsoft.com/en-us/research/blog/](https://www.microsoft.com/en-us/research/blog/phi-2-the-surprising-power-of-small-language-models)  
363 [phi-2-the-surprising-power-of-small-language-models](https://www.microsoft.com/en-us/research/blog/phi-2-the-surprising-power-of-small-language-models).
- 364 [15] Eric Mitchell. A note on dpo with noisy preferences and relationship to ipo, 2023. URL  
365 <https://ericmitchell.ai/cdpo.pdf>.
- 366 [16] Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning  
367 with noisy labels. *Advances in neural information processing systems*, 26, 2013.

- 368 [17] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin,  
369 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to  
370 follow instructions with human feedback. *Advances in neural information processing systems*,  
371 35:27730–27744, 2022.
- 372 [18] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language  
373 understanding by generative pre-training. 2018.
- 374 [19] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and  
375 Chelsea Finn. Direct preference optimization: Your language model is secretly a reward  
376 model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL  
377 <https://openreview.net/forum?id=HPuSIXJaa9>.
- 378 [20] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal  
379 policy optimization algorithms, 2017.
- 380 [21] Damien Sileo. tasksource: Structured dataset preprocessing annotations for frictionless extreme  
381 multi-task learning and evaluation. *arXiv preprint arXiv:2301.05948*, 2023. URL <https://arxiv.org/abs/2301.05948>.  
382
- 383 [22] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec  
384 Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback.  
385 *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- 386 [23] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy  
387 Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model.  
388 [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca), 2023.
- 389 [24] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei,  
390 Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas  
391 Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes,  
392 Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony  
393 Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian  
394 Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut  
395 Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov,  
396 Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta,  
397 Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiao-  
398 qing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng  
399 Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien  
400 Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation  
401 and fine-tuned chat models, 2023.
- 402 [25] Binghai Wang, Rui Zheng, Lu Chen, Yan Liu, Shihan Dou, Caishuang Huang, Wei Shen, Senjie  
403 Jin, Enyu Zhou, Chenyu Shi, Songyang Gao, Nuo Xu, Yuhao Zhou, Xiaoran Fan, Zhiheng Xi,  
404 Jun Zhao, Xiao Wang, Tao Ji, Hang Yan, Lixing Shen, Zhan Chen, Tao Gui, Qi Zhang, Xipeng  
405 Qiu, Xuanjing Huang, Zuxuan Wu, and Yu-Gang Jiang. Secrets of rlhf in large language models  
406 part ii: Reward modeling, 2024.
- 407 [26] Jingwei Yi, Rui Ye, Qisi Chen, Bin Benjamin Zhu, Siheng Chen, Defu Lian, Guangzhong  
408 Sun, Xing Xie, and Fangzhao Wu. Open-source can be dangerous: On the vulnerability of  
409 value alignment in open-source LLMs, 2024. URL <https://openreview.net/forum?id=NIou00C0ex>.  
410
- 411 [27] Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. Rrhf:  
412 Rank responses to align language models with human feedback without tears. *arXiv preprint*  
413 *arXiv:2304.05302*, 2023.
- 414 [28] Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu.  
415 Slic-hf: Sequence likelihood calibration with human feedback, 2023.
- 416 [29] Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie  
417 Jin, Qin Liu, Yuhao Zhou, Limao Xiong, Lu Chen, Zhiheng Xi, Nuo Xu, Wenbin Lai, Minghao  
418 Zhu, Cheng Chang, Zhangyue Yin, Rongxiang Weng, Wensen Cheng, Haoran Huang, Tianxiang

419 Sun, Hang Yan, Tao Gui, Qi Zhang, Xipeng Qiu, and Xuanjing Huang. Secrets of rlhf in large  
420 language models part i: Ppo, 2023.

421 [30] Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei,  
422 Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences.  
423 *arXiv preprint arXiv:1909.08593*, 2019.

## 424 A Limitations

425 We discuss some limitations of this work to stimulate further research in this direction. Our limitations  
426 mainly stem from two aspects: time efficiency issues caused by multiple calculations of PPLDiff and  
427 repeated training of a surrogate LLM, and the need for a validation dataset.

428 **Time efficiency.** Iteratively calculating the PPLDiff value for each data point and aligning a  
429 surrogate LLM is time-consuming. Selecting reliably training data and denoising the training dataset  
430 requires that the PPLDiff value be calculated for each data point during each epoch, which may cause  
431 unnecessary calculations for CPs and NPs that can already be clearly distinguished. Besides, aligning  
432 a surrogate LLM with same size as the LLM for alignment multiple times is time-consuming.

433 **Validation dataset.** PerpCorrect requires a validation dataset for aligning a surrogate LLM. How-  
434 ever, manually annotating a validation dataset is complex and labor-intensive in practice. As shown in  
435 Table 5, there is a significant disparity in alignment performance when comparing the use of 10 clean  
436 samples to 50 clean samples. Exploring how to use fewer clean samples or even no clean samples to  
437 achieve the same or better performance is a problem worth further investigation.

## 438 B Broader Impacts

439 Our proposed PerpCorrect and robust alignment pipeline offers a solution for achieving state-of-the-  
440 art performance in robust alignment under noisy preferences. PerpCorrect is designed to effectively  
441 reduce malicious noise in the dataset and mitigate biases introduced by human annotators, ensuring  
442 that the trained language model (LLM) is accurately aligned with true human preferences.

443 Moreover, we recognize a potential risk: if malicious users exploit our method for reverse training,  
444 they might compromise the security mechanisms of existing open-source LLMs. Existing research  
445 has demonstrated the possibility of reverse training [26].

## 446 C Implementation details

### 447 C.1 Training details and compute resources.

448 We utilized the Qlora method [8] for fine-tuning the LLMs, executed on RTX 4090 GPUs with  
449 24 GB of memory. Hyperparameters were set as follows: `lora_rank = 32`, `lora_dropout = 0.1`,  
450 and `lora_alpha = 16`. For SFT, we use the alpaca dataset [23] and set `learning_rate = 2e - 4`  
451 and `batch_size = 20`. For our PerpCorrect stage II, we set  $\beta = 0.1$ , `learning_rate = 1e - 3`,  
452 `batch_size = 4`,  $T = 5$ , and  $\alpha = 0.02$ . For our PerpCorrect stage III and all other alignment methods,  
453 we set  $\beta = 0.1$ , `learning_rate = 3e - 4`, and `batch_size = 20`. Other details not mentioned, we follow  
454 the default setting in TRL library. Each experiment, involving a specific method and proportion of  
455 NPs, could be completed using a single RTX 4090 GPU within 24 hours on the Golden HH dataset  
456 and within 72 hours on the OASST1 dataset.

### 457 C.2 Description and Processing Details of the Datasets

458 **OpenAssistant Conversations Dataset (OASST1).** The original OASST1 dataset [11] is an  
459 assistant-style conversation corpus generated and annotated by humans. It consists of over 10,000  
460 fully annotated conversations in 35 different languages. Sileo [21] converted these conversations  
461 into a preference dataset comprising 17,966 training samples and 952 testing samples. After filtering  
462 out conversations with one or fewer letters, we obtained a preference dataset with 17,939 training  
463 samples and 951 testing samples.

Table 7: The standard deviation of reward accuracy of DPO series alignment methods using Llama2-7B on the HHGolden dataset. The average reward accuracy is reported in Table 1

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla DPO	0.81%	0.40%	2.52%	2.60%
cDPO	1.15%	0.81%	1.76%	1.64%
rDPO	0.26%	1.53%	0.95%	1.92%
PerpCorrect-DPO	0.63%	0.87%	1.73%	0.63%

Table 8: The standard deviation of reward accuracy of PPO series alignment methods using Llama2-7B on the HHGolden dataset. The average reward accuracy is reported in Table 2

Method	Proportion of noisy preferences (%)			
	10	20	30	40
vanilla PPO	0.15%	1.30%	4.05%	0.77%
cPPO	0.15%	1.53%	4.61%	5.89%
rPPO	0.62%	1.38%	1.55%	5.29%
PerpCorrect-PPO	0.35%	1.15%	1.34%	1.57%

464 **Golden HH.** The original Golden HH dataset [5] is a preference dataset consisting of 42,537  
 465 training samples and 2,312 testing samples. Each sample has two keys: one representing the prompt  
 466  $x$  and the chosen response  $y_w$ , and the other representing the prompt  $x$  and the rejected response  $y_l$ .  
 467 We first converted the dataset into a triple form: prompt  $x$ , chosen response  $y_w$ , and rejected response  
 468  $y_l$ , retaining only one-turn conversation data. After filtering out samples with one or fewer letters, we  
 469 obtained a preference dataset with 12,066 training samples and 654 testing samples.

### 470 C.3 Detailed Robust Alignment via Perplexity-aware Correction

471 **Supervised Fine-Tuning (SFT).** The objective of Supervised Fine-Tuning (SFT) is to enhance  
 472 the performance of a pre-trained large language model (LLM) by refining its abilities for specific  
 473 tasks. As demonstrated by prior work [7, 18, 17], this can be achieved by utilizing supervised  
 474 fine-tuning with a specialized dataset tailored to the target task. The SFT dataset is annotated with  
 475 labels, providing examples that are directly relevant to the task. Specifically, for each data point  $(x, y)$   
 476 in the SFT dataset,  $x$  represents the prompt given to the LLM, and  $y$  represents the expected response  
 477 that the model should generate based on the prompt  $x$ . The process involves fine-tuning the LLM by  
 478 maximizing the log-likelihood of the correct responses  $y$  given the prompts  $x$ . Through this method,  
 479 the model learns to produce more accurate and task-specific outputs, thereby significantly improving  
 480 its performance on the given task.

481 **Perplexity-aware Correction (PerpCorrect).** We demonstrate the entire PerpCorrect algorithm in  
 482 Algorithm 2.

483 **Alignment.** We can achieve alignment using the denoised training dataset  $\tilde{D}_{\text{denoised}}$  with an  
 484 estimated proportion of NPs  $\varepsilon'_{\text{denoised}}$ . For offline alignment methods such as DPO, SLiC, and IPO,  
 485 we can directly optimize the LLM using the denoised training dataset  $\tilde{D}_{\text{denoised}}$  based on the loss  
 486 functions defined in Eqs. 6–8. For loss-based robust alignment methods, including cDPO and rDPO,  
 487 we set  $\varepsilon' = \varepsilon'_{\text{denoised}}$  and then optimize the LLM using the denoised training dataset  $\tilde{D}_{\text{denoised}}$   
 488 according to the loss functions mentioned in Eqs. 9 and 10. For the online alignment method RLHF  
 489 (PPO), we first train a reward model using the denoised training dataset  $\tilde{D}_{\text{denoised}}$  based on the loss  
 490 function described in Eq. 3. Subsequently, we further optimize the LLM using PPO according to the  
 491 objective function detailed in Eq. 4.

## 492 D Extended Experimental Results

493 For all the results presented in Table 1 and Table 2, we conducted three replicate experiments using  
 494 different seeds. We reported the average reward accuracy and the standard deviation.

---

**Algorithm 2** Perplexity-aware Correction (PerpCorrect)

---

```
1: Input: Noisy training dataset  $\tilde{\mathcal{D}}$ , clean validation dataset  $\mathcal{D}_{\text{val}}$ , LLM  $\pi_\theta$  parameterized by  $\theta$ 
2: Output: Denoised training dataset  $\tilde{\mathcal{D}}_{\text{denoised}}$  and estimated proportion of NPs  $\varepsilon'_{\text{denoised}}$ 
3:  $\pi_{\theta'} \leftarrow \pi_\theta, \mathcal{D}'_0 \leftarrow \emptyset, \varepsilon'_{\text{denoised}} \leftarrow 1, \tilde{\mathcal{D}}_{\text{denoised}} \leftarrow \tilde{\mathcal{D}},$ 
4: for epoch  $t = 0, \dots, T$  do
5:   // Aligning the surrogate LLM
6:    $\pi_{\theta'} \leftarrow \text{Alignment}(\pi_{\theta'}, \mathcal{D}'_t \cup \mathcal{D}_{\text{val}})$ 
7:   // Calculating the PPLDiff values for each data point
8:    $\Omega \leftarrow \emptyset$ 
9:   for  $(\tilde{x}, \tilde{y}_w, \tilde{y}_l) \in \tilde{\mathcal{D}}$  do
10:      $z \leftarrow \log \text{PPL}(x + \tilde{y}_w; \theta') - \log \text{PPL}(x + \tilde{y}_l; \theta')$ 
11:      $\Omega \leftarrow \Omega \cup \{(\tilde{x}, \tilde{y}_w, \tilde{y}_l, z)\}$ 
12:   end for
13:   // Fitting PPLDiff density of noisy training dataset
14:    $\bar{\varepsilon}, \bar{\mu}, \bar{\sigma} \leftarrow \text{Fitted parameters using Levenberg-Marquard algorithm with } \Omega$ 
15:   // Estimating NPs proportion of the denoised training dataset
16:    $\varepsilon'_{PC} \leftarrow \text{Estimated proportion of NPs using the Eq.15 based on } \bar{\varepsilon}, \bar{\mu}, \bar{\sigma}$ 
17:   // Keeping denoised training dataset with the smallest  $\varepsilon'_{\text{denoised}}$ 
18:   if  $\varepsilon'_{PC} < \varepsilon'_{\text{denoised}}$  then
19:      $\varepsilon'_{\text{denoised}} \leftarrow \varepsilon'_{PC}$ 
20:     // Calculating the Threshold  $\tau$ 
21:      $\tau \leftarrow \text{x-coordinate of the intersection of the two normal distributions}(\bar{\varepsilon}, \bar{\mu}, \bar{\sigma})$ 
22:     // Distinguishing CPs and NPs based on the threshold  $\tau$  and correcting NPs
23:      $\tilde{\mathcal{D}}_{\text{CPs}} \leftarrow \emptyset, \tilde{\mathcal{D}}_{\text{NPs}} \leftarrow \emptyset$ 
24:     for  $(\tilde{x}, \tilde{y}_w, \tilde{y}_l, z) \in \Omega$  do
25:       if  $z > \tau$  then
26:          $\tilde{\mathcal{D}}_{\text{CPs}} \leftarrow \tilde{\mathcal{D}}_{\text{CPs}} \cup \{(\tilde{x}, \tilde{y}_w, \tilde{y}_l)\}$ 
27:       else
28:          $\tilde{\mathcal{D}}_{\text{NPs}} \leftarrow \tilde{\mathcal{D}}_{\text{NPs}} \cup \{(\tilde{x}, \tilde{y}_l, \tilde{y}_w)\}$ 
29:       end if
30:     end for
31:      $\tilde{\mathcal{D}}_{\text{Denoised}} \leftarrow \tilde{\mathcal{D}}_{\text{CPs}} \cup \tilde{\mathcal{D}}_{\text{NPs}}$ 
32:   end if
33:    $\mathcal{D}_{\text{Clean}} \leftarrow \emptyset, \mathcal{D}_{\text{Noisy}} \leftarrow \emptyset$ 
34:   // Calculating the left bound  $\tau_l$  and the right bound  $\tau_r$ 
35:    $\tau_l \leftarrow (t - 1) \cdot \alpha \cdot (1 - \bar{\varepsilon}) \cdot |\tilde{\mathcal{D}}|$ -th smallest PPLDiff value in  $\Omega$ 
36:    $\tau_r \leftarrow (t - 1) \cdot \alpha \cdot \bar{\varepsilon} \cdot |\tilde{\mathcal{D}}|$ -th largest PPLDiff value in  $\Omega$ 
37:   // Finding extra reliable training data
38:   for  $(\tilde{x}, \tilde{y}_w, \tilde{y}_l, z) \in \Omega$  do
39:     if  $z < \tau_l$  then
40:        $\mathcal{D}_{\text{Clean}} \leftarrow \mathcal{D}_{\text{Clean}} \cup \{(\tilde{x}, \tilde{y}_w, \tilde{y}_l)\}$ 
41:     end if
42:     if  $z > \tau_r$  then
43:        $\mathcal{D}_{\text{Noisy}} \leftarrow \mathcal{D}_{\text{Noisy}} \cup \{(\tilde{x}, \tilde{y}_l, \tilde{y}_w)\}$ 
44:     end if
45:   end for
46:    $\mathcal{D}'_{t+1} \leftarrow \mathcal{D}_{\text{Clean}} \cup \mathcal{D}_{\text{Noisy}}$ 
47: end for
```

---

## 495 **NeurIPS Paper Checklist**

496 The checklist is designed to encourage best practices for responsible machine learning research,  
497 addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove  
498 the checklist: **The papers not including the checklist will be desk rejected.** The checklist should  
499 follow the references and follow the (optional) supplemental material. The checklist does NOT count  
500 towards the page limit.

501 Please read the checklist guidelines carefully for information on how to answer these questions. For  
502 each question in the checklist:

- 503 • You should answer [Yes] , [No] , or [NA] .
- 504 • [NA] means either that the question is Not Applicable for that particular paper or the  
505 relevant information is Not Available.
- 506 • Please provide a short (1–2 sentence) justification right after your answer (even for NA).

507 **The checklist answers are an integral part of your paper submission.** They are visible to the  
508 reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it  
509 (after eventual revisions) with the final version of your paper, and its final version will be published  
510 with the paper.

511 The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation.  
512 While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a  
513 proper justification is given (e.g., "error bars are not reported because it would be too computationally  
514 expensive" or "we were unable to find the license for the dataset we used"). In general, answering  
515 "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we  
516 acknowledge that the true answer is often more nuanced, so please just use your best judgment and  
517 write a justification to elaborate. All supporting evidence can appear either in the main paper or the  
518 supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification  
519 please point to the section(s) where related material for the question can be found.

520 IMPORTANT, please:

- 521 • **Delete this instruction block, but keep the section heading “NeurIPS paper checklist”,**
- 522 • **Keep the checklist subsection headings, questions/answers and guidelines below.**
- 523 • **Do not modify the questions and only use the provided macros for your answers.**

### 524 1. **Claims**

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

527 Answer: [Yes]

528 Justification: Our introduction covers our contributions, main methods and experimental  
529 results.

530 Guidelines:

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

### 540 2. **Limitations**

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

542 Answer: [Yes]

543 Justification: We discussed the efficiency issues and data volume requirements of our method  
544 PerpCorrect in the Conclusions section.

545 Guidelines:

- 546 • The answer NA means that the paper has no limitation while the answer No means that  
547 the paper has limitations, but those are not discussed in the paper.
- 548 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 549 • The paper should point out any strong assumptions and how robust the results are to  
550 violations of these assumptions (e.g., independence assumptions, noiseless settings,  
551 model well-specification, asymptotic approximations only holding locally). The authors  
552 should reflect on how these assumptions might be violated in practice and what the  
553 implications would be.
- 554 • The authors should reflect on the scope of the claims made, e.g., if the approach was  
555 only tested on a few datasets or with a few runs. In general, empirical results often  
556 depend on implicit assumptions, which should be articulated.
- 557 • The authors should reflect on the factors that influence the performance of the approach.  
558 For example, a facial recognition algorithm may perform poorly when image resolution  
559 is low or images are taken in low lighting. Or a speech-to-text system might not be  
560 used reliably to provide closed captions for online lectures because it fails to handle  
561 technical jargon.
- 562 • The authors should discuss the computational efficiency of the proposed algorithms  
563 and how they scale with dataset size.
- 564 • If applicable, the authors should discuss possible limitations of their approach to  
565 address problems of privacy and fairness.
- 566 • While the authors might fear that complete honesty about limitations might be used by  
567 reviewers as grounds for rejection, a worse outcome might be that reviewers discover  
568 limitations that aren't acknowledged in the paper. The authors should use their best  
569 judgment and recognize that individual actions in favor of transparency play an impor-  
570 tant role in developing norms that preserve the integrity of the community. Reviewers  
571 will be specifically instructed to not penalize honesty concerning limitations.

### 572 3. Theory Assumptions and Proofs

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

575 Answer: [NA]

576 Justification: The paper does not include theoretical results.

577 Guidelines:

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

### 588 4. Experimental Result Reproducibility

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

592 Answer: [Yes]

593 Justification: We show all the experiment detail in the Experiments section and Appendix.

594 Guidelines:

- 595 • The answer NA means that the paper does not include experiments.
- 596 • If the paper includes experiments, a No answer to this question will not be perceived
- 597 well by the reviewers: Making the paper reproducible is important, regardless of
- 598 whether the code and data are provided or not.
- 599 • If the contribution is a dataset and/or model, the authors should describe the steps taken
- 600 to make their results reproducible or verifiable.
- 601 • Depending on the contribution, reproducibility can be accomplished in various ways.
- 602 For example, if the contribution is a novel architecture, describing the architecture fully
- 603 might suffice, or if the contribution is a specific model and empirical evaluation, it may
- 604 be necessary to either make it possible for others to replicate the model with the same
- 605 dataset, or provide access to the model. In general, releasing code and data is often
- 606 one good way to accomplish this, but reproducibility can also be provided via detailed
- 607 instructions for how to replicate the results, access to a hosted model (e.g., in the case
- 608 of a large language model), releasing of a model checkpoint, or other means that are
- 609 appropriate to the research performed.
- 610 • While NeurIPS does not require releasing code, the conference does require all submis-
- 611 sions to provide some reasonable avenue for reproducibility, which may depend on the
- 612 nature of the contribution. For example
- 613 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
- 614 to reproduce that algorithm.
- 615 (b) If the contribution is primarily a new model architecture, the paper should describe
- 616 the architecture clearly and fully.
- 617 (c) If the contribution is a new model (e.g., a large language model), then there should
- 618 either be a way to access this model for reproducing the results or a way to reproduce
- 619 the model (e.g., with an open-source dataset or instructions for how to construct
- 620 the dataset).
- 621 (d) We recognize that reproducibility may be tricky in some cases, in which case
- 622 authors are welcome to describe the particular way they provide for reproducibility.
- 623 In the case of closed-source models, it may be that access to the model is limited in
- 624 some way (e.g., to registered users), but it should be possible for other researchers
- 625 to have some path to reproducing or verifying the results.

## 626 5. Open access to data and code

627 Question: Does the paper provide open access to the data and code, with sufficient instruc-

628 tions to faithfully reproduce the main experimental results, as described in supplemental

629 material?

630 Answer: [Yes]

631 Justification: We provide open access to our code using Anonymous Github.

632 Guidelines:

- 633 • The answer NA means that paper does not include experiments requiring code.
- 634 • Please see the NeurIPS code and data submission guidelines ([https://nips.cc/](https://nips.cc/public/guides/CodeSubmissionPolicy)
- 635 [public/guides/CodeSubmissionPolicy](https://nips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 636 • While we encourage the release of code and data, we understand that this might not be
- 637 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not
- 638 including code, unless this is central to the contribution (e.g., for a new open-source
- 639 benchmark).
- 640 • The instructions should contain the exact command and environment needed to run to
- 641 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 642 • The authors should provide instructions on data access and preparation, including how
- 643 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 644 • The authors should provide scripts to reproduce all experimental results for the new
- 645 proposed method and baselines. If only a subset of experiments are reproducible, they
- 646 should state which ones are omitted from the script and why.
- 647 • At submission time, to preserve anonymity, the authors should release anonymized
- 648 versions (if applicable).
- 649

- 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: We show the entire experimental details in the Experiments section and Appendix and provide open access to the code using Anonymous Github.

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: We report the standard deviation in the Appendix.

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: We provide detailed sufficient information on the computer resources in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.

- 700
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- 701
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- 702
- 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).
- 703
- 704
- 705
- 706

## 707 9. Code Of Ethics

708 Question: Does the research conducted in the paper conform, in every respect, with the  
709 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

710 Answer: [Yes]

711 Justification: We follow the Code of Ethics.

712 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).
- 713
- 714
- 715
- 716
- 717

## 718 10. Broader Impacts

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

721 Answer: [Yes]

722 Justification: We discuss the potential impacts in Appendix.

723 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).
- 724
- 725
- 726
- 727
- 728
- 729
- 730
- 731
- 732
- 733
- 734
- 735
- 736
- 737
- 738
- 739
- 740
- 741
- 742
- 743
- 744
- 745

## 746 11. Safeguards

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

750 Answer: [NA]

751 Justification: Our robust alignment method does not have a high risk for misuse.

752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803

Guidelines:

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

**12. Licenses for existing assets**

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

Answer: [\[Yes\]](#)

Justification: The models and datasets the we used are open-sourced, and we follow their license and terms of use.

Guidelines:

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

**13. New Assets**

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

Answer: [\[Yes\]](#)

Justification: Our training code are open-source on Anonymous GitHub.

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)?

804  
805  
806  
807  
808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

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

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

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

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