KL Penalty Control via Perturbation for Direct Preference Optimization

Anonymous ACL submission

Abstract

Direct Preference Optimization (DPO) demonstrates the advantage of aligning a large language model with human preference using only an offline dataset. However, DPO has the limitation that the KL penalty, which prevents excessive deviation from the reference model, is static throughout the training process. Several methods try to turn this static KL penalty into a dynamic one, but no approach can adaptively assign different KL penalties for each preference pair. In this paper, we propose ε -Direct Preference Optimization (ε -DPO), which allows adaptive control of the KL penalty strength β for each preference pair. Specifically, ε -DPO adaptively controls β for each preference pair based on the monotonicity of logits as a preference model under the perturbation of β during training by simply reusing the logit of the current policy and the reference policy. Experimental results show that ε -DPO outperforms existing direct alignment algorithms and KL penalty relaxation methods on general chatbot benchmarks, highlighting the significance of adaptive KL penalty relaxation at the instance-level in DPO.

1 Introduction

011

013

014

017

019

042

Aligning large language models with human preferences for helpfulness and harmless principles (Askell et al., 2021; Cui et al., 2023) is a crucial requirement for general chatbot agents. Reinforcement Learning from Human Feedback (RLHF) (Ziegler et al., 2020) is the pioneering approach that regards the alignment of large language models as a reward maximization problem and solves it by reinforcement learning (Schulman et al., 2017). However, the complicated training pipeline of RLHF increases the training complexity and computation cost of the rollout for online reinforcement learning, in addition to the difficulty of collecting human preference datasets. Moreover, introducing a trained reward model as a proxy re1) Monotonically Decreasing $(x,y^w,y^l)\Rightarrow \pmb{\beta} \rightarrow \pmb{\beta}_{\pmb{\varepsilon}}^-$

$\log \frac{\pi_{\hat{\theta}(\boldsymbol{\beta}_{\varepsilon}^{-})}(y^{w} x)}{\pi_{\hat{\theta}(\boldsymbol{\beta}_{\varepsilon}^{-})}(y^{l} x)} \log \frac{\pi_{\theta(\boldsymbol{\beta})}(y^{w} x)}{\pi_{\theta(\boldsymbol{\beta})}(y^{l} x)} \log \frac{\pi_{\hat{\theta}(\boldsymbol{\beta}_{\varepsilon}^{+})}(y^{w} x)}{\pi_{\hat{\theta}(\boldsymbol{\beta}_{\varepsilon}^{+})}(y^{l} x)}$
2) Monotonically Increasing $(x, y^w, y^l) \Rightarrow \beta \rightarrow \beta_{\varepsilon}^+$
$\log \frac{\pi_{\hat{\theta}(\beta_{\varepsilon}^{-})}(y^{w} x)}{\pi_{\hat{\theta}(\beta_{\varepsilon}^{-})}(y^{l} x)}} \log \frac{\pi_{\theta(\beta)}(y^{w} x)}{\pi_{\theta(\beta)}(y^{l} x)} \log \frac{\pi_{\hat{\theta}(\beta_{\varepsilon}^{+})}(y^{-} x)}{\pi_{\hat{\theta}(\beta_{\varepsilon}^{+})}(y^{l} x)}$

Figure 1: ε -DPO adaptively controls β corresponding to the KL penalty strength for each preference pair by checking whether the log-likelihood ratio of the chosen response and the rejected response changes monotonically with the perturbation of β used during training. It is equivalent to checking the monotonicity of the logits as a preference model induced by the DPO and estimating the advantage of the change in β by the change of train-time inverse temperature to the preference confidence under the same test-time temperature.

ward function to replace the intractable groundtruth human preference reward function makes large language models suffer from the side effect of reward over-optimization (Gao et al., 2023) inherited from the reward models.

Direct Preference Optimization (DPO) (Rafailov et al., 2023) proposes an approach to reform the limitation of RLHF by converting the policy optimization problem into a preference modeling problem and performing alignment using only offline learning. It shows comparable performance while skipping the reward modeling process required by RLHF and has become an effective alternative approach for alignment. In particular, subsequent studies with various modifications to the DPO loss objective open a new research domain called direct alignment algorithms (Rafailov et al., 2024), which perform alignment directly from offline preference



Figure 2: Comparison between ε -DPO and existing KL penalty relaxation methods. β -DPO (Wu et al., 2024) adaptively adjusts β based on the batch-level statistics of implicit reward margin through momentum, failing to adjust at the instance-level and being affected by micro-batch size. TR-DPO (Gorbatovski et al., 2024) updates reference policy periodically by the current policy, so the updates are not adaptive and require computational costs for reference policy updates. On the other hand, ε -DPO can adaptively control β at the instance-level by checking the monotonicity of the log-likelihood ratio under perturbation of β by simply reusing logits from the policies.

datasets without training separated reward models.

062

063

065

074

085

However, DPO assumes that β and the reference policy, which define a KL penalty that prevents excessive deviations from the reference model in RLHF, are fixed for exploiting the existence of a closed-form solution derived from the objective function of the RLHF. However, this assumption can lead to suboptimal results, since the KL penalty can be regarded as a Lagrangian relaxation of the constraint optimization defined by the trust region (Schulman et al., 2017). In this regard, β -DPO (Wu et al., 2024) argues that β should be adaptively chosen according to the quality of the preference pair but fails to control β at the instance-level and proposes a batch-level control method. On the other hand, TR-DPO (Gorbatovski et al., 2024) claims to periodically update the reference policy to reduce over-optimization (Rafailov et al., 2024), but it may induce unnecessary KL divergence for improvement since the update is not adaptive.

In this paper, we present ε -Direct Preference Optimization (ε -DPO), an instance-level adaptive KL penalty control for DPO that neither TR-DPO nor β -DPO achieves. Specifically, we check the advantage of adjusting β for each preference pair by observing the monotonicity of the log-likelihood ratio between the chosen response and the rejected response when the β used during training is perturbed, as described in Figure 1. Here, the criterion for controlling β does not require batch-level statistics, and the policy under the perturbed β can be estimated by reusing the current policy and reference policy logits. This criterion results in independence from the choice of micro-batch size and no additional computation requirements for model updates, unlike β -DPO and TR-DPO. 094

097

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

116

117

118

119

120

121

123

Experimental results demonstrate that ε -DPO outperforms β -DPO, TR-DPO, and most direct alignment algorithms that modify DPO loss objective (Yuan et al., 2023; Zhao et al., 2023; Azar et al., 2024; Xu et al., 2024; Ethayarajh et al., 2024; Hong et al., 2024; Park et al., 2024; Meng et al., 2024), highlighting the importance of adequate KL penalty relaxation for DPO. Furthermore, we confirm that the variation of β determined by the adaptive criterion in ε -DPO reflects the confusion as a preference model, which is not addressed by the adaptive criterion of β -DPO. We also find that the adaptive KL penalty control of ε -DPO is crucial for an efficient KL trade-off compared to TR-DPO, which is not an adaptive KL penalty control.

In summary, our work shows the following:

- ε-DPO provides a simple criterion to improve DPO through KL penalty relaxation.
- ε -DPO adaptively adjusts β in instance-level reflecting confusion as a preference model.
- ε-DPO efficiently controls β in trade-off between KL divergence and performance.

2 Preliminaries

Reinforcement Learning from Human Feedback To obtain a language model that aligns with human preference, RLHF (Ziegler et al., 2020) introduces reinforcement learning. It is equivalent to

approaching preference alignment as a reward max-124 imization problem, where we find a policy π that 125 maximizes a ground-truth reward function r^* rep-126 resenting human reward for a response y obtained 127 from a corresponding policy for a given prompt x. 128 However, since the ground-truth reward function 129 cannot be accessed, a reward model trained from 130 the preference dataset is introduced as a proxy re-131 ward function. On the other hand, to prevent the 132 policy update from deviating too much from the 133 current policy from the initial policy, the KL di-134 vergence from the reference policy π_{ref} serves as 135 a penalty and regards the initial policy as a refer-136 ence policy. At this time, coefficient β controls 137 the strength of the penalty. The optimal policy that 138 satisfies the maximization of the modified objective 139 140 function under β has a closed-form solution π_{β}^* ,

$$\pi_{\beta}^* \coloneqq \arg \max_{\pi} \{ \mathbb{E}_{x,y}[r^*(x,y)] - \beta \mathbb{D}_{\mathrm{KL}}(\pi || \pi_{\mathrm{ref}}) \}$$
$$\pi_{\beta}^*(y|x) = \frac{1}{Z_{\beta}^*(x)} \pi_{\mathrm{ref}}(y|x) \exp\left(\frac{1}{\beta}r^*(x,y)\right),$$
$$Z_{\beta}^*(x) = \sum \pi_{\mathrm{ref}}(y|x) \exp\left(\frac{1}{\beta}r^*(x,y)\right).$$

y

141

147

151

157

164

Direct Preference Optimization RLHF has a 142 143 limitation in efficiency due to the additional training step of the reward model. In this respect, 144 DPO (Rafailov et al., 2023) proposes an approach 145 that can perform preference alignment without 146 training the reward model. DPO focuses on the fact that the ground-truth reward function can be implic-148 itly reparameterized by the closed-form solution π_{β}^{*} 149 and reference policy π_{ref} with an intractable normal-150 izing constant Z_{β}^* . If we assume the Bradley-Terry model (Bradley and Terry, 1952) for the ground-152 truth human preference function, then the human 153 preference can be modeled by the margin between 154 the reward of the chosen response y^w and the re-155 jected response y^l with the sigmoid function σ , 156 which can ignore the intractable term Z_{β}^{*} by cancellation. From this observation, DPO performs 158 preference alignment through preference model optimization using an offline dataset in the sense that 161 obtaining an optimal policy through policy optimization in RLHF can be obtained by training a 162 preference model given by the implicit reward $r_{\theta,\beta}$, 163

$$\begin{aligned} r_{\theta,\beta}(x,y) &\coloneqq \beta \log \frac{\pi_{\theta}(y|x)}{\pi_{\mathrm{ref}}(y|x)} + Z_{\beta}(x;\theta), \\ \mathbb{P}_{\theta,\beta}(y^w \succ y^l|x) &\coloneqq \sigma\big(r_{\theta,\beta}(x,y^w) - r_{\theta,\beta}(x,y^l) \\ \mathcal{L}_{\mathrm{DPO}}(x,y^w,y^l;\theta,\beta) &\coloneqq -\log \mathbb{P}_{\theta,\beta}(y^w \succ y^l|x) \end{aligned}$$

3 ε -Direct Preference Optimization

In this section, we describe our proposed method, ε -Direct Preference Optimization (ε -DPO), that adaptively controls KL penalty coefficient β at the instance-level based on the logit monotonicity as a preference model according to the perturbation of β . Figure 2 illustrates the difference between ε -DPO and existing KL penalty relaxation methods.

165

166

167

168

170

171

172

173

174

175

176

177

178

179

180

181

182

183

185

186

187

188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207

208

209

210

211

212

213

3.1 **Relaxation of KL Penalty in DPO**

The KL penalty introduced by RLHF can be regarded as an approach to solve the constrained optimization problem in the trust region (Schulman, 2015) defined near the reference policy π_{ref} as an unconstrained optimization by treating β as a Lagrange multiplier (Schulman et al., 2017). From this perspective, even though DPO reformulates the problem of finding an optimal policy under fixed $\pi_{\rm ref}$ and β as a preference modeling problem, using a single β and a fixed trust region for all instances may lead to suboptimal results. This hypothesis regarding relaxation of KL penalty can be supported by the experimental results of β -DPO (Wu et al., 2024) that adaptively control β based on the statistics of implicit reward margin during the training process and TR-DPO (Gorbatovski et al., 2024) that updates $\pi_{\rm ref}$ during the training process for preventing over-optimization (Rafailov et al., 2024) from the vanishing curvature of the loss landscape.

However, β -DPO fails to perform instance-level β control despite claiming that the quality of each preference pair should determine β . Instead, it performs batch-level β control using momentumbased estimation of batch-level margin disparities, which is strongly affected by the micro-batch size. In addition, TR-DPO updates the reference model without adaptive criteria, which can lead to inefficient KL divergence trade-off between performance and incur computational costs for updating the reference model. Therefore, instance-level adaptive KL penalty control without requiring additional computational cost that achieves an efficient KL trade-off is still undiscovered for DPO.

3.2 Logit Monotonicity under Perturbation

Establishing a criterion to adaptively change the KL penalty for each instance of preference dataset is not a trivial problem. As a proxy criterion, we can exploit that the policy obtained via DPO can function as a preference model $\mathbb{P}_{\theta,\beta}$. Formally, $\mathbb{P}_{\theta,\beta}$ can be expressed as a binary classifier with logit z_{θ} and

),

215

216

217

218

219

221

226

228 229

239

240

241

242

243

244

245

247

251

margin γ for a preference triplet $(x, y^w, y^l) \in \mathcal{D}$,

$$\mathbb{P}_{\theta,\beta}(\cdot|\cdot) = \sigma\Big(\beta\big(z_{\theta}(\cdot) - \gamma(\cdot)\big)\Big),$$
$$z_{\theta}(x, y^{w}, y^{l}) \coloneqq \log\frac{\pi_{\theta}(y^{w}|x)}{\pi_{\theta}(y^{l}|x)},$$
$$\gamma(x, y^{w}, y^{l}) \coloneqq \log\frac{\pi_{\mathrm{ref}}(y^{w}|x)}{\pi_{\mathrm{ref}}(y^{l}|x)}.$$

This shows that β serves as an inverse temperature of a binary classifier. For a given β , we define β_{ε}^{-} and β_{ε}^{+} with a positive constant $\varepsilon > 0$. That is, β_{ε}^{-} and β_{ε}^{+} refer to values that have been *perturbed* to be slightly larger or slightly smaller than the β ,

$$\beta_{\varepsilon}^{-} \coloneqq \frac{\beta}{1+\varepsilon}, \beta_{\varepsilon}^{+} \coloneqq \frac{\beta}{1-\varepsilon}.$$

Let us denote the parameters obtained via DPO as a function of β , $\theta(\cdot) : \mathbb{R}^+ \to \Theta$. Consider the case we observe the strict monotonicity of logits happens according to the perturbation of β on $\theta(\cdot)$,

$$z_{\theta(\beta_{\varepsilon}^{-})}(\cdot) > z_{\theta(\beta)}(\cdot) > z_{\theta(\beta_{\varepsilon}^{+})}(\cdot), \qquad (1)$$

$$z_{\theta(\beta_{\varepsilon}^{-})}(\cdot) < z_{\theta(\beta)}(\cdot) < z_{\theta(\beta_{\varepsilon}^{+})}(\cdot).$$
(2)

Intuitively, this corresponds to observing monotonic changes in preference confidence under the same test-time temperature scaling (Guo et al., 2017). If the logits monotonically decrease with increasing β , then raising the training temperature (i.e., lowering β) yields a clearer separation of y^w and y^l in the neighborhood of $\frac{1}{\beta}$, despite having a softer decision boundary. Conversely, if they monotonically increase, a higher training temperature harms the separation of y^w and y^l . From this, we can estimate the benefit of adjusting β for each instance within the neighborhood defined by ε .

Estimating KL Penalty Perturbation 3.3

Note that $\theta(\cdot)$ is intractable since it is equivalent to having access to models trained on each β . However, Liu et al. (2024b) shows that optimal policy under $\frac{\beta}{\lambda}$ can be expressed by π_{ref} re-weighted with importance ratio using π_{β}^* . If we assume the autoregressive prior of optimal policy, then the optimal policy under $\frac{\beta}{\lambda}$ can be estimated by the optimal policy under β and the reference policy, as we respecify Proposition 1 from Liu et al. (2024b),

Proposition 1 (Liu et al. (2024b)) Under the assumption of optimal autoregressive policy π^* where the prompt $x \in \mathcal{X}$, response vocabulary $y_i \in \mathcal{V}$, and logit $f: \mathcal{X} \times \mathcal{V}^{i-1} \to \mathbb{R}^{|\mathcal{V}|}$, the optimal policy 254

 $\pi^*_{\frac{\beta}{2}}$ can be approximated by the arithmetic mean of logits between π^*_{β} and reference policy π_{ref} ,

$$\pi_{\frac{\beta}{\lambda}}^{*}(y_{1:n}|x) = \prod_{i=1}^{n} \pi_{\frac{\beta}{\lambda}}^{*}(y_{i}|x, y_{1:i-1})$$

$$\approx \prod_{i=1}^{n} \operatorname{Softmax} \left(\lambda f_{\beta}^{*}(x, y_{1:i-1}) + (1-\lambda) f_{\operatorname{ref}}(x, y_{1:i-1}) \right)_{y_{i}}.$$
257

Proof. See Appendix A.

Using Proposition 1, we can approximate $\pi_{\theta(\beta_{\varepsilon}^{-})}$ and $\pi_{\theta(\beta_{\epsilon}^{+})}$ by trained policy and reference policy without accessing $\theta(\cdot)$ since they are the approximated policies for $\pi^*_{\beta^-_\varepsilon}$ and $\pi^*_{\beta^+_\varepsilon}$. To adaptively control β for each preference triplet (x, y^w, y^l) during the training process, we regard the policy π_{θ} obtained in the current step as the best approximation of the optimal policy under current β and estimate $\pi_{\theta(\beta_{\varepsilon}^{-})}$ and $\pi_{\theta(\beta_{\varepsilon}^{+})}$ for $z_{\theta(\beta_{\varepsilon}^{-})}$ and $z_{\theta(\beta_{\varepsilon}^{+})}$,

$$\pi_{\theta(\beta_{\varepsilon}^{-})}(y_{1:n}|x) = \prod_{i=1}^{n} \pi_{\theta(\beta_{\varepsilon}^{-})}(y_{i}|x, y_{1:i-1})$$

$$\approx \prod_{i=1}^{n} \operatorname{Softmax}\left((1+\varepsilon)f_{\theta}(x, y_{1:i-1}) -\varepsilon f_{\operatorname{ref}}(x, y_{1:i-1})\right)_{y_{i}},$$
(3)

271

272

273

274

275

276

277

278

279

280

281

$$\pi_{\theta(\beta_{\varepsilon}^{+})}(y_{1:n}|x) = \prod_{i=1}^{n} \pi_{\theta(\beta_{\varepsilon}^{+})}(y_{i}|x, y_{1:i-1})$$

$$\approx \prod_{i=1}^{n} \operatorname{Softmax} \left((1-\varepsilon) f_{\theta}(x, y_{1:i-1}) + \varepsilon f_{\operatorname{ref}}(x, y_{1:i-1}) \right)_{y_{i}}.$$
(4)

Recall that we need not only the logit of the current policy f_{θ} but also the logit of the reference policy f_{ref} to compute the estimated log-likelihood ratio. However, in order to compute the loss function of DPO, \mathcal{L}_{DPO} , the log-likelihood from the reference policy must be computed for each training instance, which allows us to simply reuse f_{ref} for estimation without any additional computation cost of model forward passes. Therefore, we determine the β , which is used for the KL penalty coefficient in the current training step for each instance,

$$\tilde{\beta}(x, y^w, y^l; \theta) = \begin{cases} \beta_{\varepsilon}^- & \text{if (1),} \\ \beta_{\varepsilon}^+ & \text{if (2),} \\ \beta & \text{otherwise.} \end{cases}$$
(5)

259

261

262

263

264

265

294

297

299

310

311

After the model update, the β corresponds to the optimal policy that the current policy is targeting changes depending on $\tilde{\beta}$ used in \mathcal{L}_{DPO} for each instance. Therefore, we need to modify the baseline β for the next training step, and we simply update the β with the mean statistics of $\tilde{\beta}$ determined across the batch used in the update as follows:

$$\beta \leftarrow \mathbb{E}_{x, y^w, y^l}[\tilde{\beta}(x, y^w, y^l; \theta)]. \tag{6}$$

Note that $\tilde{\beta}$ is determined independently with the batch-level statistic, so the adaptive control of β in ε -DPO can be performed independently with the choice of micro-batch size. Algorithm 1 summarizes the entire training process of ε -DPO.

Algorithm 1 ε -Direct Preference Optimization				
Require: reference policy π_{ref} , initial KL penalty				
	coefficient β , and perturbation size ε			
1:	Initialize model π_{θ} with π_{ref} .			
2:	while not converged do			
3:	Sample preference triplets (x, y^w, y^l) .			
4:	Estimate $\pi_{\hat{\theta}(\beta_{\varepsilon}^{-})}, \pi_{\hat{\theta}(\beta_{\varepsilon}^{+})}$ using 3 and 4.			
5:	Determine instance-level $\tilde{\beta}$ according to 5.			
6:	Update π_{θ} by \mathcal{L}_{DPO} with instance-level $\tilde{\beta}$.			

- 7: Update the current β using $\tilde{\beta}$ by 6.
- 8: end while
- 9: **return** aligned policy π_{θ} .

4 Experimental Setup

In this section, we discuss the experimental setup for validating our proposed method, ε -DPO. We check the feasibility of ε -DPO using **UltraFeedback** (Cui et al., 2023), compared to the diverse direct alignment algorithms (Rafailov et al., 2023; Yuan et al., 2023; Zhao et al., 2023; Azar et al., 2024; Xu et al., 2024; Ethayarajh et al., 2024; Hong et al., 2024; Park et al., 2024; Meng et al., 2024) as a method for general chatbot alignment. We also use **Anthropic-HH** (Bai et al., 2022) for a detailed comparison with existing methods for KL penalty relaxation of DPO (Wu et al., 2024; Gorbatovski et al., 2024). The implementation details for each experimental setting are in Appendix B.

4.1 UltraFeedback

UltraFeedback (Cui et al., 2023) is an AI feedback dataset where GPT-4 (Achiam et al., 2023) rates responses obtained from four different language models. We follow the experimental setting of SimPO (Meng et al., 2024) for comparison with various direct alignment algorithms, including DPO (Rafailov et al., 2023), RRHF (Yuan et al., 2023), SLiC-HF (Zhao et al., 2023), IPO (Azar et al., 2024), CPO (Xu et al., 2024), KTO (Ethayarajh et al., 2024), ORPO (Hong et al., 2024), and R-DPO (Park et al., 2024). Specifically, we use the Instruct setting starting from instruction-tuned language models (Jiang et al., 2023a; Dubey et al., 2024). We evaluate resulting models by AlpacaEval 2 (Dubois et al., 2024), Arena-Hard (Li et al., 2024), and MT-Bench (Jiang et al., 2023b), which are widely used for general chatbot benchmarks. 317

318

319

321

322

323

324

325

326

327

328

329

330

331

332

333

334

335

336

338

339

340

341

342

344

345

346

347

348

349

350

351

352

354

355

356

357

360

361

362

363

365

4.2 Anthropic-HH

Anthropic-HH (Bai et al., 2022) is a human preference dialogue dataset containing two subsets based on the helpfulness and harmlessness principle. Here, we use helpful-base and harmless-base splits to validate the criterion using logit monotonicity for instance-level β control used in ε -DPO and the efficiency in terms of trade-off between performance and KL divergence (Rafailov et al., 2024). We choose gemma-2-2B (Team et al., 2024) to obtain the reference policy through Supervised Fine-tuning with chosen responses. Following DPO (Rafailov et al., 2023), we evaluate the models trained with each method under various β in the single-turn dialogue setting. We regard PairRM (Jiang et al., 2023b) as an evaluator for checking performance by win rate comparing their responses and chosen responses in the test splits.

5 Experimental Results

Main Results of ε **-DPO** In Table 1, we observe that ε -DPO shows notable performances across AlpacaEval 2 (Dubois et al., 2024), Arena-Hard (Li et al., 2024), and MT-Bench (Jiang et al., 2023b) using UltraFeedback. In particular, we find that the performance of ε -DPO outperforms most direct alignment algorithms, which generally modify the loss objective, highlighting that the major assumption of fixed KL penalty in DPO is overlooked. In addition, we observe that ε -DPO performs better than other KL penalty relaxation approaches (Wu et al., 2024; Gorbatovski et al., 2024) from Table 2. Thus, we can find that instance-level KL penalty control significantly impacts the final performance.

Influence of ε **on Training Dynamics** The perturbation ε is used for checking logit monotonicity as a preference model in the neighborhood of the current β . Therefore, it can be chosen within

	Mistral-Instruct (7B)			Llama-3-Instruct (8B)				
Method	Alpaca	aEval 2	Arena-Hard	MT-Bench	Alpaca	aEval 2	Arena-Hard	MT-Bench
	LC (%)	WR (%)	WR (%)	Score (1-10)	LC (%)	WR (%)	WR (%)	Score (1-10)
SFT	17.1	14.7	12.6	7.5	26.0	25.3	22.3	8.1
DPO	26.8	24.9	16.3	7.6	40.3	37.9	32.6	8.0
RRHF	25.3	24.8	18.1	7.6	31.3	28.4	26.5	7.9
SLiC-HF	24.1	24.6	18.9	7.8	26.9	27.5	26.2	8.1
IPO	20.3	20.3	16.2	7.8	35.6	35.6	30.5	8.3
CPO	23.8	28.8	22.6	7.5	28.9	32.2	28.8	8.0
KTO	24.5	23.6	17.9	7.7	33.1	31.8	26.4	8.2
ORPO	24.5	24.9	20.8	7.7	28.5	27.4	25.8	8.0
R-DPO	27.3	24.5	16.1	7.5	41.1	37.8	33.1	8.0
SimPO	32.1	34.8	21.0	7.6	44.7	40.5	33.8	8.0
ε-DPO	35.6	29.6	17.2	7.8	46.4	44.9	36.7	8.0

Table 1: AlpacaEval 2 (Dubois et al., 2024), Arena-Hard (Li et al., 2024), and MT-Bench (Jiang et al., 2023b) results of ε -DPO in the Instruct setting proposed by SimPO (Meng et al., 2024). LC and WR denote length-controlled win rate and win rate. The best result for each benchmark is represented in bold. Results of other direct alignment algorithms (Rafailov et al., 2023; Yuan et al., 2023; Zhao et al., 2023; Azar et al., 2024; Xu et al., 2024; Ethayarajh et al., 2024; Hong et al., 2024; Park et al., 2024) are directly taken from the official report of SimPO.

Method	Alpaca	aEval 2	Arena-Hard
	LC (%)	WR (%)	WR (%)
SFT	26.0	25.3	22.3
DPO	40.3	37.9	32.6
β -DPO	43.4	38.2	-
TR-DPO^{τ}	42.8	47.2	32.4
$TR\text{-}DPO^{\alpha}$	43.5	46.8	34.7
ε-DPO	46.4	44.9	36.7

Table 2: Performance of β -DPO (Wu et al., 2024), TR-DPO^{τ}, TR-DPO^{α} (Gorbatovski et al., 2024), and ε -DPO measured on AlpacaEval 2 and Arena-Hard in the Llama-3-Instruct setting. Each β -DPO and TR-DPO result is directly taken from their official reports.

a reasonable range to estimate the approximated policies corresponding to β_{ε}^{-} and β_{ε}^{+} . However, 367 ε can influence training dynamics since ε deter-368 mines the sizes of instance-level KL penalty coefficient β . We further analyze the intra-epoch training dynamics on Llama-3-Instruct settings according to ε . We compare the forward KL divergence $\mathbb{D}_{KL}(\pi_{ref}||\pi_{\theta})$ (Rafailov et al., 2024) and performance on AlpacaEval 2 using checkpoints obtained at 0.2 intervals during the training, along 375 with the changes of in-batch ratio of β_{ε}^{-} and β_{ε}^{+} ,

372

373

376

as shown in Figure 3. We find that adaptive control occurs more frequently for both β_{ε}^{-} and β_{ε}^{+} as ε increased, leading to accelerating the increase of KL divergence and performance. We also observe that the performance at the beginning of training tends to be lower when higher ε . We speculate that the trained policy at the beginning of training is insufficient to estimate the optimal policy, making the approximation unstable at the high ε level.

377

378

379

381

382

383

384

386

387

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

Analysis of Logit Monotonicity β -DPO (Wu et al., 2024) chooses higher β for preference pairs with larger implicit reward margins to update the current policy conservatively from the reference policy. This is motivated by the claim that large implicit reward margins reflect higher quality gaps of response pairs corresponding to meaningless training signals. In this respect, we analyze the implicit reward margin of preference pairs where logit monotonicity according to the perturbation of β happened in policies trained by DPO using Antropic-HH, as shown in Figure 4. We find that ε -DPO performs opposite decisions compared to β -DPO from the observation that preference pairs with monotonically increasing logits have smaller average implicit reward margins than those with monotonically decreasing logits. Also, this implies that ε -DPO enhances training signals for confus-



Figure 3: Intra-epoch training dynamics of Llama-3-Instruct according to the change of ε . We can observe that larger ε increases the rate of β_{ε}^{-} and β_{ε}^{+} within a batch, which leads to the occurrence of adaptive control while accelerating the increase of KL divergence and performance of length-controlled win rate on AlpacaEval 2.

ing examples because the implicit reward margin is proportional to the preference confidence, and the increase of β scales up the gradient of DPO loss (Rafailov et al., 2023). Furthermore, we confirm that implicit reward margins do not always represent the quality of preference pairs through qualitative analysis in Appendix C. Therefore, we suspect that β -DPO fails on the instance-level adaptive KL penalty control because it assigns low gradient weights to confusing examples and strongly relies on the implicit reward margins that do not always represent the quality of preference pairs.

404

405

406

407

408

409

410

411

412

413

414

415

Efficiency in KL Trade-off As TR-DPO (Gor-416 batovski et al., 2024) claims, increasing KL diver-417 gence would be desirable as a trade-off when de-418 viating from the reference policy improves perfor-419 mance. However, the over-optimization (Rafailov 420 et al., 2024) of direct alignment algorithms empha-421 sizes that it is necessary to check the Pareto frontier 422 to determine whether performance improvements 423 can be achieved without indiscriminately expand-494 ing the KL divergence. Figure 5 depicts the Pareto 425 frontier of models trained under various beta using 426 Antropic-HH by DPO, ε -DPO and two variants of 427 TR-DPO, TR-DPO^{τ} which hard-updates the ref-428 erence policy by the fixed interval and TR-DPO $^{\alpha}$ 429 which soft-updates the reference policy through 430 weight merging. We can see that regardless of 431 the two variants, TR-DPO induces more KL diver-432 433 gence than DPO and ε -DPO and cannot achieve similar performance under the same KL budget as 434 ε -DPO. This highlights the efficiency of ε -DPO in 435 KL trade-offs and implies that controlling the KL 436 penalty in a non-adaptive manner can induce exces-437

sive relaxation for performance improvements.

438

439

440

441

442

443

444

445

446

447

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

6 Related Works

Direct Alignment Algorithms Many variants of direct alignment algorithms perform alignment on offline preference datasets without an external reward model. DPO (Rafailov et al., 2023) performs alignment through preference modeling with the implicit reward derived from the optimal policy of reward maximization under the KL penalty objective. RRHF (Yuan et al., 2023) performs alignment by training to maintain the likelihood margin between preference ranks. KTO (Ethayarajh et al., 2024) changes the assumptions of the Bradley-Terry model (Bradley and Terry, 1952) used by DPO and introduces Prospect Theory (Kahneman and Tversky, 2013), and IPO (Azar et al., 2024) converts to the root-finding problem for strengthening the KL constraint. SLiC-HF (Zhao et al., 2023), CPO (Xu et al., 2024), ORPO (Hong et al., 2024), and SimPO (Meng et al., 2024) train without reference models utilizing behavior cloning, margin loss, contrastive loss, odds ratio loss, and fixed margin by replacing the implicit rewards.

Reward Over-optimization and KL Penalty Since RLHF (Ziegler et al., 2020) utilizes a trained reward model, it amplifies the limitations of the reward model as it is optimized for an imperfect reward, according to Goodhart's Law (Hoskin, 1996), and this is called reward over-optimization (Gao et al., 2023). However, Rafailov et al. (2024) finds that direct alignment algorithms also experience similar reward over-optimization, regardless of the variant. Direct alignment algorithms commonly



Figure 4: The average implicit reward margin of pairs showing logit monotonicity according to the perturbation of β in policies trained with DPO under various β using Antropic-HH. We can confirm that ε -DPO assigns higher β for preference pairs reveling high confusion on preference labels based on the observation that preference pairs with monotonically increasing logits show low confidences on preference model $\mathbb{P}_{\theta,\beta}(y^w \succ y^l | x)$.

show humped curves of performance according to the increase of KL divergence from the reference model during training. TR-DPO (Gorbatovski et al., 2024) argues that this is due to the Hessian of the loss landscape converging to zero as the implicit reward margin grows during training, so they update the reference model for mitigating this phenomenon. On the other hand, β -DPO (Wu et al., 2024), which also performs relaxation of KL penalty, claims that adaptively changing β through the statistics of the implicit reward margin is required to reflect the quality of the preference pair.

471

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

Combining Sampling Distribution Combining sampling distributions of language models can be utilized to estimate a new sampling distribution with specific characteristics. Contrastive Decoding (Li et al., 2022) shows that the log-likelihood margins of the expert and amateur language models can enhance response diversity by penalizing incorrect response patterns favored by the amateur language model. Sanchez et al. (2023) shows that classifier-free guidance (Ho and Salimans, 2022) can enhance prompt relativity in language modeling by treating prompts as conditions and sharpening the conditional sampling distribution. Combining the change from instruction-tuning in a small language model with a large language model can approximate fine-tuning. Liu et al. (2024a) utilizes the instruction-tuned small language model as the



Figure 5: Pareto Frontier between KL divergence and performance of DPO, ε -DPO, TR-DPO^{τ} and TR-DPO^{α}. We measure the KL divergence and performance of the models trained with $\beta = [0.5, 0.1, 0.05, 0.01]$ using Antropic-HH. We can see that ε -DPO shows better best performance than DPO, simultaneously achieving better KL trade-off efficiency than TR-DPO.

logit offset, and Mitchell et al. (2023) estimates the importance sampling ratio of the optimal distribution defined by the objective of RLHF from it. Inspired by the theoretical motivation of Mitchell et al. (2023), Liu et al. (2024b) shows that the sampling distribution of the policy trained under the near β by DPO can be approximated by policy obtained under β and the reference policy. 500

501

502

503

504

505

506

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

7 Conclusion

In this paper, we present ε -Direct Preference Optimization (ε -DPO), an instance-level adaptive KL penalty control for DPO, adjusting β by observing the monotonicity of the log-likelihood ratio between the chosen response and the rejected response when the β used during training is perturbed. The criterion for instance-level adaptive control of β only requires estimating the policy under the perturbed β , which can be efficiently estimated by reusing the current policy and reference policy logits without relying on batch-level statistics and requiring additional computation cost. Resulting models obtained through ε -DPO perform better than resulting models from existing methods under general chatbot benchmarks. In particular, the criterion used in ε -DPO shows a more efficient KL trade-off than the non-adaptive KL penalty relaxation while reflecting the confusion on preference pairs, emphasizing the importance of an appropriate instance-level KL penalty relaxation.

Limitations

529

545

550

552

553

554

556

557

563

564

566

567

569

570

575

576

577

578

579

581

 ε -DPO requires the reference policy because it has 530 a KL penalty from the reference policy, like DPO 531 in default. It leads to the limitation that it requires 532 additional memory consumption and computation for reference policy compared to other direct alignment algorithms that do not perform regularization 535 through the reference policy (Zhao et al., 2023; Xu et al., 2024; Hong et al., 2024; Meng et al., 2024). 537 However, theoretically, ε -DPO can save memory consumption by pre-computing the logits of the responses from the reference policy, similar to DPO. 540 Meanwhile, ε -DPO is a general purposes approach 541 not specially tailored for safety alignment, so ad-542 ditional safety considerations may be required to 543 control inappropriate responses in real usages. 544

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. arXiv preprint arXiv:2112.00861.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv*:2204.05862.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324– 345.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2023. Ultrafeedback: Boosting language models with high-quality feedback.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpacaeval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*. 582

583

585

586

587

589

590

591

592

593

594

595

596

597

598

599

600

601

602

603

604

605

606

607

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pages 10835–10866. PMLR.
- Alexey Gorbatovski, Boris Shaposhnikov, Alexey Malakhov, Nikita Surnachev, Yaroslav Aksenov, Ian Maksimov, Nikita Balagansky, and Daniil Gavrilov. 2024. Learn your reference model for real good alignment. *arXiv preprint arXiv:2404.09656*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Jonathan Ho and Tim Salimans. 2022. Classifierfree diffusion guidance. *arXiv preprint arXiv:2207.12598.*
- Jiwoo Hong, Noah Lee, and James Thorne. 2024. Orpo: Monolithic preference optimization without reference model. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11170–11189.
- Keith Hoskin. 1996. The 'awful idea of accountability': inscribing people into the measurement of objects. *Accountability: Power, ethos and the technologies of managing*, 265.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023a. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023b. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*.
- Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific.
- Sangkyu Lee, Sungdong Kim, Ashkan Yousefpour, Minjoon Seo, Kang Min Yoo, and Youngjae Yu. 2024. Aligning large language models by on-policy selfjudgment. *arXiv preprint arXiv:2402.11253*.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024. From crowdsourced data to highquality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*.

Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang,

Jason Eisner, Tatsunori Hashimoto, Luke Zettle-

moyer, and Mike Lewis. 2022. Contrastive decoding:

Open-ended text generation as optimization. arXiv

Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia

Tianlin Liu, Shangmin Guo, Leonardo Bianco, Daniele

Calandriello, Quentin Berthet, Felipe Llinares,

Jessica Hoffmann, Lucas Dixon, Michal Valko,

and Mathieu Blondel. 2024b. Decoding-time re-

alignment of language models. arXiv preprint

Yu Meng, Mengzhou Xia, and Danqi Chen.

Eric Mitchell, Rafael Rafailov, Archit Sharma, Chelsea

Finn, and Christopher D Manning. 2023. An emula-

tor for fine-tuning large language models using small

language models. arXiv preprint arXiv:2310.12962.

Chelsea Finn. 2024. Disentangling length from qual-

ity in direct preference optimization. arXiv preprint

Ryan Park, Rafael Rafailov, Stefano Ermon, and

Rafael Rafailov, Yaswanth Chittepu, Ryan Park, Harshit Sikchi, Joey Hejna, Bradley Knox, Chelsea Finn, and Scott Niekum. 2024. Scaling laws for reward model overoptimization in direct alignment algo-

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. Direct preference optimization: Your language

model is secretly a reward model. arXiv preprint

Guillaume Sanchez, Honglu Fan, Alexander Spangher, Elad Levi, Pawan Sasanka Ammanamanchi, and Stella Biderman. 2023. Stay on topic with classifierfree guidance. arXiv preprint arXiv:2306.17806.

John Schulman. 2015. Trust region policy optimization.

John Schulman, Filip Wolski, Prafulla Dhariwal,

Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, and Aviral Kumar. 2024. Preference fine-tuning of llms should leverage suboptimal, on-policy data. arXiv preprint arXiv:2404.14367.

Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint*

arXiv preprint arXiv:1502.05477.

rithms. arXiv preprint arXiv:2406.02900.

Simpo: Simple preference optimization

arXiv preprint

Tsvetkov, Yejin Choi, and Noah A Smith. 2024a.

Tuning language models by proxy. arXiv preprint

preprint arXiv:2210.15097.

arXiv:2401.08565.

arXiv:2402.02992.

arXiv:2405.14734.

arXiv:2403.19159.

arXiv:2305.18290.

arXiv:1707.06347.

with a reference-free reward.

2024.

- 64
- 04 64
- 644
- 646 647 648
- 650 651
- 652 653 654 655
- 6
- 6
- 6 6
- 6 6
- 6
- 6
- 671 672

6

6

6 6

- 6
- 0
- 680 681

6

- 6
- 6

587 588 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*. 689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

708

709

710

711

712

713

714

715

716

717

- Junkang Wu, Yuexiang Xie, Zhengyi Yang, Jiancan Wu, Jinyang Gao, Bolin Ding, Xiang Wang, and Xiangnan He. 2024. β -dpo: Direct preference optimization with dynamic β . arXiv preprint arXiv:2407.08639.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024. Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation. *arXiv preprint arXiv:2401.08417*.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*.
- Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. 2023. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*.
- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2020. Fine-tuning language models from human preferences. *Preprint*, arXiv:1909.08593.

A **Proof of Proposition 1**

Proposition 1 (Liu et al. (2024b)) Under the assumption of optimal autoregressive policy π^* where the prompt $x \in \mathcal{X}$, response vocabulary $y_i \in \mathcal{V}$, and logit $f : \mathcal{X} \times \mathcal{V}^{i-1} \to \mathbb{R}^{|\mathcal{V}|}$, the optimal policy π^*_{β} can the approximated by the arithmetic mean of logits between π^*_{β} and reference policy π_{ref} , 720

$$\pi_{\frac{\beta}{\lambda}}^{*}(y_{1:n}|x) = \prod_{i=1}^{n} \pi_{\frac{\beta}{\lambda}}^{*}(y_{i}|x, y_{1:i-1})$$

$$\approx \prod_{i=1}^{n} \text{Softmax} \left(\lambda f_{\beta}^{*}(x, y_{1:i-1}) + (1-\lambda)f_{\text{ref}}(x, y_{1:i-1})\right)_{y_{i}}.$$
723

Proof of Proposition 1. Recall that optimal policy π_{β}^* has a closed-form solution and ground-truth reward function r^* can be reparameterized using the normalizing constant Z_{β}^* , 725

$$\pi_{\beta}^{*}(y|x) = \frac{1}{Z_{\beta}^{*}(x)} \pi_{\text{ref}}(y|x) \exp\left(\frac{1}{\beta}r^{*}(x,y)\right),$$

$$Z_{\beta}^{*}(x) = \sum_{y} \pi_{\text{ref}}(y|x) \exp\left(\frac{1}{\beta}r^{*}(x,y)\right),$$

$$r^{*}(x,y) = \beta \log\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \log Z_{\beta}^{*}(x).$$
726

Here, we plug the reparameterization of r^* to the close-form solution of $\pi^*_{\frac{\beta}{\lambda}}$ and simple algebra yield, 727

$$\pi_{\beta}^{*}(y|x) = \frac{1}{Z_{\beta}^{*}(x)} \pi_{\text{ref}}(y|x) \exp\left(\frac{\lambda}{\beta}r^{*}(x,y)\right) = \frac{\pi_{\text{ref}}(y|x)\exp\left(\frac{\lambda}{\beta}r^{*}(x,y)\right)}{\sum_{y}\pi_{\text{ref}}(y|x)\exp\left(\frac{\lambda}{\beta}r^{*}(x,y)\right)}$$

$$= \frac{\pi_{\text{ref}}(y|x)\exp\left(\lambda\log\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)} + \lambda\log Z_{\beta}^{*}(x)\right)}{\sum_{y}\pi_{\text{ref}}(y|x)\exp\left(\lambda\log\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)} + \lambda\log Z_{\beta}^{*}(x)\right)} = \frac{\pi_{\text{ref}}(y|x)\left(\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)} + Z_{\beta}^{*}(x)\right)^{\lambda}}{\sum_{y}\pi_{\text{ref}}(y|x)\left(\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)}\right)^{\lambda}} = \frac{\pi_{\beta}^{*}(y|x)^{\lambda}\pi_{\text{ref}}(y|x)^{1-\lambda}}{\sum_{y}\pi_{\text{ref}}(y|x)\left(\frac{\pi_{\beta}^{*}(y|x)}{\pi_{\text{ref}}(y|x)}\right)^{\lambda}} = \frac{\pi_{\beta}^{*}(y|x)^{\lambda}\pi_{\text{ref}}(y|x)^{1-\lambda}}{\sum_{y}\pi_{\beta}^{*}(y|x)^{\lambda}\pi_{\text{ref}}(y|x)^{1-\lambda}} = \frac{1}{Z(x)}\pi_{\beta}^{*}(y|x)^{\lambda}\pi_{\text{ref}}(y|x)^{1-\lambda},$$

where Z denotes the normalizing constant of reparameterized form of $\pi_{\frac{\beta}{\lambda}}^*$. Now, we use the assumption 729 of autoregressive policy π_{β}^* . This assumption gives us to evade intractable normalizing constant Z, 730

$$\pi_{\frac{\beta}{\lambda}}^{*}(y_{i}|x, y_{1:i-1}) \approx \frac{1}{Z(x, y_{1:i-1})} \pi_{\beta}^{*}(y_{i}|x, y_{1:i-1})^{\lambda} \pi_{\text{ref}}(y_{i}|x, y_{1:i-1})^{1-\lambda}$$

$$= \frac{\pi_{\beta}^{*}(y_{i}|x, y_{1:i-1})^{\lambda} \pi_{\text{ref}}(y_{i}|x, y_{1:i-1})^{1-\lambda}}{\sum_{v \in \mathcal{V}} \pi_{\beta}^{*}(v|x, y_{1:i-1})^{\lambda} \pi_{\text{ref}}(v|x, y_{1:i-1})^{1-\lambda}}$$

$$= \frac{\text{Softmax}(f_{\beta}^{*}(x, y_{1:i-1}))_{y_{i}}^{\lambda} \text{Softmax}(f_{\text{ref}}(x, y_{1:i-1}))_{y_{i}}^{1-\lambda}}{\sum_{v \in \mathcal{V}} \text{Softmax}(f_{\beta}^{*}(x, y_{1:i-1}))_{v}^{\lambda} \text{Softmax}(f_{\text{ref}}(x, y_{1:i-1}))_{v}^{1-\lambda}}$$

$$= \frac{\exp(f_{\beta}^{*}(x, y_{1:i-1}))_{y_{i}}^{\lambda} \exp(f_{\text{ref}}(x, y_{1:i-1}))_{y_{i}}^{1-\lambda}}{\sum_{v \in \mathcal{V}} \exp(f_{\beta}^{*}(x, y_{1:i-1}))_{v}^{\lambda} \exp(f_{\text{ref}}(x, y_{1:i-1}))_{v}^{1-\lambda}},$$

with eliminating $\left(\sum_{v\in\mathcal{V}}\exp\left(f_{\beta}^{*}(x,y_{1:i-1})\right)_{v}\right)^{\lambda}\left(\sum_{v\in\mathcal{V}}\exp\left(f_{\mathrm{ref}}(x,y_{1:i-1})\right)_{v}\right)^{1-\lambda}$ from nominator 732

and denominator at the last line. Note that the geometric mean acts as the arithmetic mean in log-scale,

$$\begin{aligned} & \frac{\exp\left(f_{\beta}^{*}(x,y_{1:i-1})\right)_{y_{i}}^{\lambda}\exp\left(f_{\text{ref}}(x,y_{1:i-1})\right)_{y_{i}}^{1-\lambda}}{\sum_{v\in\mathcal{V}}\exp\left(f_{\beta}^{*}(x,y_{1:i-1})\right)_{v}^{\lambda}\exp\left(f_{\text{ref}}(x,y_{1:i-1})\right)_{v}^{1-\lambda}} \\ & = \frac{\exp\left(\lambda f_{\beta}^{*}(x,y_{1:i-1})_{y_{i}} + (1-\lambda)f_{\text{ref}}(x,y_{1:i-1})_{y_{i}}\right)}{\sum_{v\in\mathcal{V}}\exp\left(\lambda f_{\beta}^{*}(x,y_{1:i-1})_{v} + (1-\lambda)f_{\text{ref}}(x,y_{1:i-1})_{v}\right)} \\ & = \text{Softmax}\left(\lambda f_{\beta}^{*}(x,y_{1:i-1}) + (1-\lambda)f_{\text{ref}}(x,y_{1:i-1})\right)_{y_{i}}.\end{aligned}$$

Therefore, $\pi_{\underline{\beta}}^*$ can be approximated by the arithmetic mean of logit between $\pi_{\underline{\beta}}^*$ and π_{ref} ,

$$\begin{split} \pi^*_{\frac{\beta}{\lambda}}(y_{1:n}|x) &= \prod_{i=1}^n \pi^*_{\frac{\beta}{\lambda}}(y_i|x, y_{1:i-1}) \\ &\approx \prod_{i=1}^n \operatorname{Softmax} \left(\lambda f^*_{\beta}(x, y_{1:i-1}) + (1-\lambda)f_{\operatorname{ref}}(x, y_{1:i-1})\right)_{y_i}. \end{split}$$

737

738

741

760

B Implementation Details

The implementation of ε -DPO and experiments are all based on the TRL¹ library. Here, we explain the experimental settings for UltraFeedback (Cui et al., 2023) and Antropic-HH (Bai et al., 2022) in detail.

B.1 UltraFeedback

For a fair comparison with direct alignment algorithms and existing approaches for KL penalty relaxation, we follow the Instruct setting suggested by SimPO (Meng et al., 2024). The Instruct setting starts 743 with Mistral-7B-Instruct-v0.2² (Jiang et al., 2023a) and Meta-Llama-3-8B-Instruct³ (Dubey 744 et al., 2024) as reference policies, each named as Mistral-Instruct and Llama-3-Instruct. First, 745 746 rollouts using prompts from UltraFeedback (Cui et al., 2023) are performed, then PairRM (Jiang et al., 2023b) serves as an external evaluator to build preference datasets for approximating the on-policy learn-747 ing (Tajwar et al., 2024; Lee et al., 2024). We use corresponding datasets publicly released by SimPO, 748 each denoted as mistral-instruct-ultrafeedback⁴ and llama3-ultrafeedback⁵. We perform hyperparameter searches for the learning rate within the range of [3e-7, 5e-7, 7e-7, 1e-6] and ε within the [0.005, 0.01, 0.02] range while β is fixed to 0.01, following the best hyperparameter of DPO reported 751 from SimPO. Other common hyperparameters are fixed in the same way as SimPO. Every experiment is 752 conducted using 16 NVIDIA A100-SXM4-40GB GPUs within 2 hours. We evaluate resulting models through AlpacaEval 2 (Dubois et al., 2024), Arena-Hard (Li et al., 2024), and MT-Bench (Jiang et al., 2023b) following the same sampling configuration settings reported by SimPO. Table 3 summarizes the 755 training configurations for Mistral-Instruct and Llama-3-Instruct.

B.2 Anthropic-HH

We use helpful-base and harmless-base splits for experiments using Anthropic-HH⁶ (Bai et al., 2022). We preprocess the dataset by parsing only the content of each conversation turn and removing the original role header of the dataset. We use gemma-2-2b⁷ (Team et al., 2024) as a base model for

¹github.com/huggingface/trl

²huggingface.co/mistralai/Mistral-7B-Instruct-v0.2, Apache 2.0 License, Copyright (c) 2023 Mistral AI

³huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct, LLAMA 3 Community License, Copyright (c) 2024 Meta Platforms ⁴huggingface.co/datasets/princeton-nlp/mistral-instruct-ultrafeedback

⁵huggingface.co/datasets/princeton-nlp/llama3-ultrafeedback

⁶huggingface.co/datasets/Anthropic/hh-rlhf, MIT License, Copyright (c) 2022 Anthropic

⁷huggingface.co/google/gemma-2-2b, Apache 2.0 License, Copyright (c) 2024 Google LLC

Configuration	Mistral-Instruct	Llama-3-Instruct	
Model	Mistral-7B-Instruct-v0.2	Meta-Llama-3-8B-Instruct	
Dataset	mistral-instruct-ultrafeedback	llama3-ultrafeedback	
Optimizer	AdamW	AdamW	
Epoch	1	1	
Batch Size	128	128	
Learning Rate	[<u>3e-7</u> , 5e-7, 7e-7, 1e-6]	[3e-7, 5e-7, <u>7e-7</u> , 1e-6]	
Scheduler	cosine	cosine	
Warm-up Ratio	0.1	0.1	
Weight Decay	0	0	
β	0.01	0.01	
ε	$[0.005, \underline{0.01}, 0.02]$	$[0.005, \underline{0.01}, 0.02]$	

Table 3: Training configurations for Mistral-Instruct and Llama-3-Instruct in the experiment settings using Ultrafeedback (Cui et al., 2023). The underline indicates the value selected through the hyperparameter search.

obtaining the reference policy through Supervised Fine-tuning (SFT) with chosen responses by applying the chat template of gemma-2-2b-it (Team et al., 2024)⁸. We fix all hyperparameters except β for a fair comparison between methods. We use $\varepsilon = 0.01$ in ε -DPO and $\tau = 128$, $\alpha = 0.6$ in TR-DPO (Gorbatovski et al., 2024) as the method-specific hyperparameter and β within the [0.01, 0.05, 0.1, 0.5] range. Following DPO (Rafailov et al., 2023), we evaluate resulting models in the single-turn dialogue setting by comparing with chosen responses from test split through PairRM⁹ (Jiang et al., 2023b) as an external evaluator to check the win rate. We set the temperature as 1.0 and max token length as 1024 when sampling responses from each model for evaluation. Every experiment is conducted using 4 NVIDIA A100-SXM4-40GB GPUs within 7 hours. Table 4 shows the common training configurations for each experiment.

761

762

763

764

765

766

767

768

770

771

772

774

775

778

779

Configuration	SFT	ε -DPO, DPO, TR-DPO
Optimizer	AdamW	AdamW
Epoch	1	1
Batch Size	128	128
Learning Rate	2e-5	1e-6
Scheduler	cosine	cosine
Warm-up Ratio	0.1	0.1
Weight Decay	0	0

Table 4: Common training configurations on the experiment settings using Anthropic-HH (Bai et al., 2022).

C Qualitative Analysis of Logit Monotonicity and Implicit Reward Margin

We compare preference pairs whose implicit reward margin is maximized among the preference pairs showing monotonically increasing or decreasing logits in the helpful-base split of Antropic-HH (Bai et al., 2022). Similarly, we compare preference pairs whose implicit reward margin is minimized among the preference pairs showing monotonically increasing or decreasing logits. We obtain these preference pairs through the policy trained with DPO under $\beta = 0.1$. If we follow the claim of β -DPO, higher β should be selected for both preference pairs that sufficiently maximize the implicit reward margin regardless of logit monotonicity. However, Table 5 shows the case close to the label flipping compared to the case of Table 6 in which the adaptive control decision of ε -DPO and β -DPO matches in high implicit reward margin. In this case of label flipping, increasing β results in a negative effect by increasing

⁸huggingface.co/google/gemma-2-2b-it, Apache 2.0 License, Copyright (c) 2024 Google LLC

⁹huggingface.co/llm-blender/PairRM

the scale of the loss gradient (Rafailov et al., 2023) on the noisy preference pair. On the other hand, Table 7 shows the case of the rejected response with a significantly lower quality than the chosen response compared to the case of Table 8 in which the adaptive control decision of ε -DPO and β -DPO matches in low implicit reward margin. However, β -DPO will assign a low β to the corresponding example contrary to the original claim since it shows a low implicit reward margin. These qualitative examples demonstrate that the claim of β -DPO, preference data quality can be estimated through the implicit reward margin, may not be empirically valid.

780

781

782

783

785

786

Prompt			
I love the music of Bessie Smith. Would love to know more about her background?			
Chosen Response	Rejected Response		
Hi Human! I love the music of many	Sure, it seems like Bessie Smith has		
different artists and am always happy	a lot of interesting facts to share		
to talk about it with you. Here are	with us. Her life story is surprisingly		
some Wikipedia links to learn more about	convoluted and non-linear. Let's look		
Bessie Smith:	into her story using the tools we've		
	already developed		
-https://en.wikipedia.org/wiki/Bessie_Smi			
th	There were a few points of relevance		
-https://en.wikipedia.org/wiki/Bessie_Smi	to her early life, such as her birth		
<pre>th#Bessie _Smith_and_St_Clair_Avenue</pre>	name, "Huldah Marie Smith", but also that		
	her mother and siblings moved in with		
	her while her mother was pregnant, and		
	she was brought up in poverty and want,		
	living in unfit and crowded conditions,		
	with three other siblings and their		
	mother. However, as the website on this		
	subject attests, the poverty seemed to		
	lead to a self-sufficient and industrious		
	nature in Huldah, including sewing and		
	odd jobs for local businesses. She found		
	it difficult to make ends meet, and		
	suffered recurring illnesses that led to		
	frequent arrests. But, she eventually		
	learned to sing, and could sell a few		
	cups of beans and potatoes to neighbors.		
	At this point in her life, she was paid		
	thirty dollars a week, which is a small		
	amount to most people today, but put her		
	on the road to true independence.		
Logit Monotonicity	Implicit Reward Margin		
Decreasing	88.0		

Table 5: A qualitative example in Antropic-HH (Bai et al., 2022) where the implicit reward margin is maximized among the pairs with monotonically increasing logits. We can see that this example is close to the label-flipping case since the chosen response does not accomplish the purpose of the user prompt.

Prompt				
I can't afford dentures and I am losing my teeth. My gums are sore from eating. W				
hat can I do?				
Chosen Response	Rejected Response			
I'm sorry to hear that! Have you thought	You're in a tough spot. Here's some advice			
about trying a new diet? A friend of	for you, feel free to pass it on to anyone			
mine says the raw food diet has helped	else who has the same situation.			
her problems a lot. Or, you could also				
try to treat the sore gums with other				
treatments, like an apple cider vinegar				
rinse, or organic baking soda paste. Or,				
you could also just try eating soft foods,				
like salads, for a while. It might help				
you heal faster.				
Logit Monotonicity	Implicit Reward Margin			
Increasing	30.4			

Table 6: A qualitative example in Antropic-HH (Bai et al., 2022) where the implicit reward margin is maximized among the pairs with monotonically increasing logits. In this case, β -DPO and ε -DPO share the same decision on adaptive control; both try to increase β for this preference pair.

Prompt				
How can I make my own baguettes at home?				
Chosen Response	Rejected Response			
You'll need a small metal pan or pot,	I'd recommend following this guide.			
a long wooden stick (a rolling pin is	You'll need to adjust the proportions			
ideal), some flour, and some water. In a	a bit, but overall it should be a good			
pot on a burner, combine water and flour	guide.			
to make a kind of paste. When it comes				
to a boil, put the stick in the pot,				
then pour the dough from the pan onto				
it. With the stick, roll out the dough				
into a long rectangle, then keep rolling				
it until it's about 1–1.5 cm thick, and				
finally put it on the pan. The end result				
should be a baguette with small bubbles				
(you can also make it all flat, which				
will give you a larger pizza base.) Heat				
the pan again to ensure the baguette is				
completely dry, then put it in an oven to				
bake. That's really all there is to it,				
and you can make baguettes in as little				
as 15 minutes!				
Logit Monotonicity	Implicit Reward Margin			
Increasing	-43.2			

Table 7: A qualitative example in Antropic-HH (Bai et al., 2022) where the implicit reward margin is minimized among the pairs with monotonically increasing logits. We can see that this example has a significantly low-quality rejected response since the rejected response did not offer the guide it mentioned in the response.

Prompt				
How do I measure belly fat?				
Chosen Response	Rejected Response			
I'm not sure exactly what you're looking	We can discuss how you want to measure it,			
for? There are multiple ways to measure	and what values you want to monitor. We'll			
fat. Do you want the percentage of body	also need to consider what constitutes an			
fat or do you want to know the total	acceptable body weight and how that may			
amount of fat? Do you want to see the	impact what you measure. At some point,			
distribution of fat in your body or just	you also need to consider what behaviors			
be able to say you have ten pounds of	you can do to reduce your body weight and			
belly fat?	belly fat, or what you need to do to build			
	up your body weight			
Logit Monotonicity	Implicit Reward Margin			
Decreasing	-12.0			

Table 8: A qualitative example in Antropic-HH (Bai et al., 2022) where the implicit reward margin is minimized among the pairs with monotonically decreasing logits. In this case, β -DPO and ε -DPO share the same decision on adaptive control; both try to decrease β for this preference pair.