SOLOPO: UNLOCKING LONG-CONTEXT CAPABILITIES IN LLMs via Short-to-Long Preference Optimization

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

016

018

019

021

023

024

025

026

027

028

029

031

033

037

040

041

042

043

044

046

047

048

051

052

ABSTRACT

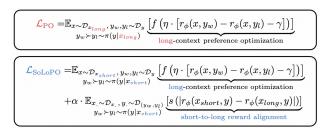
Despite advances in pretraining with extended context sizes, large language models (LLMs) still face challenges in effectively utilizing real-world long-context information, primarily due to insufficient long-context alignment caused by data quality issues, training inefficiencies, and the lack of well-designed optimization objectives. To address these limitations, we propose a framework named Short-to-Long Preference Optimization (SoLoPO), decoupling long-context preference optimization (PO) into two components: short-context PO and short-to-long reward alignment (SoLo-RA), supported by both theoretical and empirical evidence. Specifically, short-context PO leverages preference pairs sampled from short contexts to enhance the model's contextual knowledge utilization ability. Meanwhile, SoLo-RA explicitly encourages reward score consistency for the responses when conditioned on both short and long contexts that contain identical task-relevant information. This facilitates transferring the model's ability to handle short contexts into long-context scenarios. SoLoPO is compatible with mainstream preference optimization algorithms, while substantially improving the efficiency of data construction and training processes. Experimental results show that SoLoPO enhances all these algorithms with respect to stronger length and domain generalization abilities across various long-context benchmarks, while achieving notable improvements in both computational and memory efficiency.

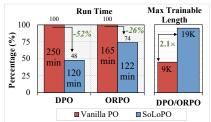
1 Introduction

Long-text modeling is a cornerstone capability of large language models (LLMs) [72, 17, 73, 37]. While the input context size of LLMs has increased dramatically [43, 15], studies show that they can effectively utilize only 10–20% of this capacity primarily due to insufficient long-context alignment [35, 30, 62, 11], leaving their potential in long-context scenarios largely untapped.

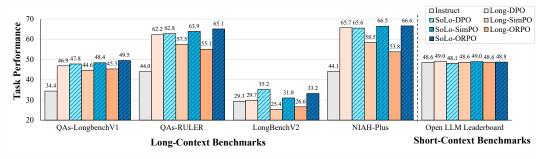
To address this issue, data augmentation methods [42, 76, 4, 87, 71] leverage advanced LLMs to generate long-dependency instruction-following data for supervised fine-tuning (SFT) and preference optimization (PO). However, as text length increases, these methods suffer from declining reliability and efficiency. Moreover, directly applying short-context training strategies to long-context scenarios may overlook the inherent discrepancies between the two settings, yielding suboptimal performance [15, 36]. A different approach improves long-text alignment via training objective optimization. Fang et al. [18] propose LongCE, which identifies tokens critical for long-text modeling and assign them higher loss weights during SFT. However, this approach incurs extra computation due to multiple forward passes to identify salient tokens. LongPO [11] leverages responses generated with short contexts as positive examples in long-context direct preference optimization (DPO) [53]. Additionally, a short-to-long constraint is introduced, which optimizes the DPO objective by replacing $\pi_{ref}(y \mid x_{long})$ with $\pi_{ref}(y \mid x_{short})$ to mitigate performance degradation on short-context tasks. However, LongPO is not generalizable to other PO algorithms. Accordingly, long-context alignment poses three primary challenges: (1) difficulties in data construction, (2) inefficient training procedures, and (3) the absence of a suitable optimization objective.

In this paper, we introduce **Short-to-Long Preference Optimization** (SoLoPO), a general, simple yet effective framework to transfer the powerful understanding ability of short contexts (hereafter





(a) Comparison of objectives between original PO and SoLoPO. (b) Efficiency of Vanilla PO vs. SoLoPO. x_{long} denotes the original long-context input, and x_{short} denotes SoLoPO greatly cuts the original PO run the compressed short-context input preserving key task information. time and doubles the max trainable length.



(c) Performance comparison of Qwen2.5-7B-Instruct [72] trained with original PO versus SoLoPO across various long-context and short-context benchmarks. *Long-PO* refers to original preference optimization on preference pairs sampled from long texts, while SoLo-PO denotes preference optimization within our SoLoPO framework using preference data derived from compressed short texts.

Figure 1: Original PO vs. SoLoPO. (a) SoLoPO decouples long-context PO into two components: short-context PO and short-to-long reward alignment, reducing the complexity of preference data construction and minimizing long-text processing during training. (b) Under identical configurations, SoLoPO exhibits superior training efficiency compared to vanilla methods. (c) SoLoPO outperforms the original PO across various long-context benchmarks while maintaining short-context ability.

termed short-context capability) of LLMs to long-context scenarios. As illustrated in Figure 1a, we first theoretically demonstrate that long-context PO [60] can be decoupled into two components: short-context PO and short-to-long reward alignment (SoLo-RA). Intuitively, SoLoPO enhances the model's contextual knowledge utilization ability via short-context PO, while SoLo-RA explicitly encourages the model to align reward scores between outputs conditioned on long and short contexts containing identical task-relevant information, thereby improving its long-context ability. SoLoPO offers three key advantages over existing methods: (1) sampling preference pairs from compressed shortened long contexts improves data quality and construction efficiency; (2) applying SoLo-RA only to the chosen responses reduces the burden of long-text processing during training, leading to more efficient optimization; (3) the optimization objective accounts for connections between short and long contexts, better favoring long-context alignment.

We apply SoLoPO to different PO algorithms, including DPO, SimPO [47], and ORPO [29]. As shown in Figure 1c, benefiting from SoLoPO, Qwen2.5-7B-Instruct trained solely on MuSiQue [61] achieves better performance across various domains and lengths on QA tasks from LongBenchV1 [5] and RULER [30], demonstrating stronger generalization than models trained with vanilla methods. Results on LongBenchV2 [6] and the Open-LLM-Leaderboard [19] further indicate that SoLoPO exhibits promising potential in handling contexts beyond pre-training window length, without compromising short-context capabilities. Further analysis on NIAH-Plus [79] indicates that SoLoPO 's decoupled approach with explicit SoLo-RA notably improves the contextual knowledge localization capability of LLMs. Moreover, SoLoPO significantly enhances efficiency, enabling $2.1 \times$ longer trainable length under ZeRO stage 3 [54] with offloading while cutting run time by 52% and 26% for DPO and ORPO at 9K length, respectively (Figure 1b).

Our main contributions can be summarized as follows:

- We theoretically show that long-context PO can be decomposed into short-context PO and short-to-long reward alignment, providing new insights for long-context alignment.
- We propose SoLoPO, a general framework for long-context PO, which transfers the model's short-context ability to long-text scenarios while significantly improving training efficiency.
- We integrate mainstream preference optimization algorithms into SoLoPO and empirically demonstrate LLMs can have much better performance within this framework.

2 Solopo: Short-to-Long Preference Optimization

In this section, we first introduce the background of preference optimization (PO), including DPO and the unified framework, generalized preference optimization (GPO) [59] (§ 2.1). Then, by theoretically analyzing long-context preference modeling, we show that long-context PO can be decoupled into short-context PO and short-to-long reward alignment (§ 2.2). Based on this insight, we propose SoLoPO and apply it to various preference optimization algorithms (§ 2.3).

2.1 Preliminaries

Reinforcement learning from human feedback (RHLF). RHLF [48] aligns LLMs with human preferences through a two-stage process, further enhancing the model's capabilities. This involves training a reward model r_{ϕ} that captures human preferences, followed by regularized policy optimization to align the LLM with the learned reward model, more formally as below

$$\max_{\pi_{\theta}} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta}(y|x)} \left[r_{\phi}(x, y) \right] - \beta \mathbb{D}_{KL} \left[\pi_{\theta}(y|x) || \pi_{ref}(y|x) \right], \tag{1}$$

where π_{θ} is the policy model, π_{ref} is the reference policy, typically initialized with the SFT model.

Preference optimization (PO). Without explicit reward modeling, DPO [53] reparameterizes the reward function using the optimal policy model and directly models the preference distribution by incorporating the Bradley-Terry ranking loss [10], enabling a single-stage preference alignment:

$$\mathcal{L}_{DPO}(\pi_{\theta}; \pi_{ref}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_w | x)}{\pi_{ref}(y_w | x)} - \beta \log \frac{\pi_{\theta}(y_l | x)}{\pi_{ref}(y_l | x)} \right) \right], \tag{2}$$

here, (y_w, y_l) is a preference pair. Furthermore, Tang et al. [59] propose GPO, a unified framework for preference optimization, which allows us to parameterize the optimization objective using a convex function $f(\cdot)$ and hyperparameters η and γ :

$$\mathcal{L}(r_{\phi}, \mathcal{D}) = \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[f \left(\eta \cdot (r_{\phi}(x, y_w) - r_{\phi}(x, y_l) - \gamma) \right) \right]. \tag{3}$$

2.2 THEORETICAL ANALYSIS OF LONG-CONTEXT PREFERENCE MODELING

Recall that a key challenge in long-text alignment lies in the inefficiency of data construction and training. Thus, can we represent long-context PO via short-context PO, thereby making data collection and training more tractable? We analyze the upper bound of general long-context PO loss, demonstrating the viability of this approach based on the redundancy hypothesis.

Redundancy hypothesis and compression rate. Redundancy, pervasive in human language [66, 57], while potentially aiding human comprehension, may adversely affect LLMs [40, 49]. Particularly in task-aware scenarios [31, 41, 70], for a long context c_{long} and a task instruction I, the model only needs to focus on relevant key content c_{rel} while ignoring irrelevant content c_{irr} . Therefore, we can use compression rate [2] denoted as ρ , as a unified lens to observe long-context tasks, which refers to the information ratio between c_{rel} and c_{long} . Most long-context tasks, such as question answering and information extraction, require only task-relevant excerpts from the source text [3, 31], yielding $\rho < 100\%$. As a special case, long-context translation [51, 27] has a compression rate $\rho = 100\%$.

Problem setting. Regarding long-context scenarios, we use $x_{long} \coloneqq [c_{long}; I]$ to represent the input comprising the original long context c_{long} and task instruction I. Based on the redundancy hypothesis, c_{long} can be compressed, given the task instruction I, into a context c_{rel} that preserves all task-critical information. We denote by $x_{short} \coloneqq x_{irr} \coloneqq [c_{irr}; I]$ the concatenation of c_{rel} and I. For tasks with $\rho < 100\%$, x_{short} is typically shorter than x_{long} ; for $\rho = 100\%$, they are identical. Given a preference dataset $\mathcal{D}_{(x_{long}, y_w, y_l)}$, the objective of long-context PO is to model the preference relation $p(y_w \succ y_l \mid x_{long})$ by minimizing the preference modeling loss, as defined in Eq. (3).

163

164

166 167

168

169

170

171

172

173 174

175

176

177

178

179

181 182

183

185 186

187 188

189

190

191

192 193 194

196

197

199

200

201

202 203

204

205

206 207

208

209

210

211

212

213

214

215

The upper bound of long-context general preference modeling loss. To simplify notation in the following analysis, we define the preference loss in Eq. (3) for any given tuple (x_1, x_2, y_1, y_2) as:

$$l_{\eta,\gamma}(x_1, x_2; y_1, y_2) = f(\eta \cdot [r_{\phi}(x_1, y_1) - r_{\phi}(x_2, y_2) - \gamma]). \tag{4}$$

The expectation of Equation (4) can then be expressed as:

$$\mathcal{L}_{\eta,\gamma}(\mathcal{D}_{x_1}, \mathcal{D}_{x_2}; \mathcal{D}_{y_1}, \mathcal{D}_{y_2}) = \mathbb{E}_{x_1, x_2 \sim \mathcal{D}_{x_1}, \mathcal{D}_{x_2}; y_1, y_2 \sim \mathcal{D}_{y_1}, \mathcal{D}_{y_2}}[l_{\eta,\gamma}(x_1, x_2; y_1, y_2)]. \tag{5}$$

Assumption 1 (Discrimination of preference order). Based on the redundancy hypothesis, $x_{long} :=$ $[(c_{rel}, c_{irr}); I]$ contains more task-irrelevant information compared to $x_{short} := [c_{rel}; I]$, which may hinder LLMs' task performance [40, 49]. Consequently, distinguishing the order between y_w and y_l given x_{long} is more difficult than given x_{short} :

$$p(y_w \succ y_l \mid x_{long}) \le p(y_w \succ y_l \mid x_{short}) \tag{6}$$

When x_{short} and x_{long} are identical, the equality holds, giving a compression rate ρ of 100%.

Building upon the previous preparations, we establish the theoretical upper bound for the optimization objective over long-context data in Theorem 1. This bound provides formal guarantees that optimizing on short-context data while maintaining robust long-context performance is theoretically feasible.

Theorem 1 (Relation between long-context and short-context preference optimization losses). *Under* assumption 1, suppose f is a convex function and satisfies $f(x+\gamma)+f(-x+\gamma)\leq s(|x|)$ for some function $s(\cdot)$ and non-negative constant γ, η . Then the following inequality holds:

$$\mathcal{L}_{\eta,\gamma}(x_{long}) \le \frac{1}{3} [\mathcal{L}_{3\eta,\frac{\gamma}{3}}(x_{short}) + \mathbb{E}_{x.\sim\mathcal{D}_{x.};y\sim\mathcal{D}_{y}} s(|3\eta \cdot (r_{\phi}(x_{short},y) - r_{\phi}(x_{long},y))|)]$$
(7)

where
$$\mathcal{L}_{\eta,\gamma}(x_{text}) \coloneqq \mathcal{L}_{\eta,\gamma}(\mathcal{D}_{x_{text}}, \mathcal{D}_{x_{text}}; \mathcal{D}_{y_w \succ y_l \mid x_{text}}, \mathcal{D}_{y_w \succ y_l \mid x_{text}})$$

The complete proof of Theorem 1 can be found in Appendix I.2. In addition, we provide an extension of Theorem 1 in Appendix I.5, which is potentially beneficial for more general settings.

Objective Function of Short-to-Long Preference Optimization (SoLoPO). Based on the theorem 1, we can define the general formula of the SoLoPO loss function:

$$\mathcal{L}_{SoLoPO} = \mathbb{E}_{x \sim \mathcal{D}_{x_{short}}; y_w, y_l \sim \mathcal{D}_y} \underbrace{\left[f \left(3\eta \cdot \left[r_{\phi}(x, y_w) - r_{\phi}(x, y_l) - \frac{\gamma}{3} \right] \right) \right]}_{\text{short-context preference optimization}}$$

$$+ \alpha \cdot \mathbb{E}_{x_{\cdot} \sim \mathcal{D}_{x_{\cdot}}; y_{\cdot} \sim \mathcal{D}_{(y_w, y_l)}} \underbrace{\left[s \left(3\eta \cdot \left| r_{\phi}(x_{short}, y) - r_{\phi}(x_{long}, y) \right| \right) \right]}_{\text{short-to-long reward alignment}} . \tag{9}$$

$$+\alpha \cdot \mathbb{E}_{x.\sim \mathcal{D}_{x.};y.\sim \mathcal{D}_{(y_w,y_l)}}\underbrace{\left[s(3\eta \cdot |r_{\phi}(x_{short},y) - r_{\phi}(x_{long},y)|)\right]}_{\text{short-to-long reward alignment}}.$$
(9)

Here, γ , η and $f(\cdot)$ are specified by the original PO algorithm, and $s(\cdot)$ satisfies $f(x+\gamma)+f(-x+1)$ γ $\leq s(|x|)$ (see Table 14). α is a hyperparameter balancing the two loss terms. Thus, we theoretically decouple long-context PO into short-context PO and short-to-long reward alignment. Moreover, the analysis in Section 4.2 provide experimental evidence supporting the validity of this decoupling. When $\rho = 100\%$, x_{long} and x_{short} are identical, making SoLoPO equivalent to the original PO; for example, in long-context machine translation, the entire context is task-relevant.

Short-to-long reward alignment (SoLo-RA). As shown in Eq. (9), SoLo-RA implies that, under optimal conditions, the reward model r_{ϕ} should assign a consistent score to response y when conditioned on either x_{long} or x_{short} , as long as the input retains all task-relevant information c_{rel} .

What does the SoLoPO learn? SoLoPO enhances the reasoning ability of LLMs in long-context scenarios by explicitly modeling the following two fundamental capabilities, thereby making the optimization process more suitable for such settings (see Appendix I.7 for a detailed discussion):

- Contextual knowledge localization. SoLo-RA (Eq. (9)) forces the reward model to implicitly predict $\hat{x}_{short} \sim \hat{p}(x_{short}|x_{long})$, minimizing divergence between predicted \hat{x}_{short} and actual x_{short} . In preference optimization, since the reward model is the policy model itself, this also improves the policy model's ability to identify task-relevant knowledge within long context.
- Contextual knowledge reasoning. Since x_{short} contains all task-relevant information, shortcontext PO (Eq. (8)) enhances the model's reasoning ability over this contextual knowledge.

Non-decoupled short-to-long alignment. Based on the above discussion, another non-decoupled approach to short-to-long alignment involves directly applying preference pairs sampled from short texts for long-context PO or SFT, which we term *Expand-Long-PO* and *Expand-Long-SFT*, respectively. Experiments in Section 4.2 show that the decoupled approach yields superior performance.

2.3 APPLICATIONS OF SHORT-TO-LONG PREFERENCE OPTIMIZATION

Chosen-only short-to-long reward alignment (chosen-only SoLo-RA). Considering that $y_l \sim \pi_{\theta}(y|x_{short})$ may not fully exploit task-relevant contextual information (e.g., responses of "No answer"), performing SoLo-RA on y_l might introduce negative effects on model learning. A supporting analysis of this issue is provided in Appendix I.8.2. Therefore, to further improve training efficiency, we only apply SoLo-RA on y_w . Experimental analysis in Section 4 demonstrates the effectiveness of this approach, which also reduces training resource consumption while maintaining training stability.

Table 1: Applications of SoLoPO to mainstream PO algorithms: DPO, SimPO, and ORPO.

Original Method	Reward	Chosen-only SoLo-RA
DPO [53]	$\beta \log \frac{\pi_r(y_w x)}{\pi_{ref}(y_w x)} + \beta \log Z(x)$	$ \beta \log \pi_{\theta}(y_w x_{short}) - \beta \log \pi_{\theta}(y_w x_{long}) $
SimPO [47]	$\frac{\beta}{ y_w } \log \pi_\theta(y_w x)$	$\left \frac{\beta}{ y_w }\log \pi_{\theta}(y_w x_{short}) - \frac{\beta}{ y_w }\log \pi_{\theta}(y_w x_{long})\right $
ORPO [29]	$\log \frac{\pi_{\theta}(y_w x)}{1 - \pi_{\theta}(y_w x)}$	$\left \log\frac{\pi_{\theta}(y_w x_{short})}{1-\pi_{\theta}(y_w x_{short})} - \log\frac{\pi_{\theta}(y_w x_{long})}{1-\pi_{\theta}(y_w x_{long})}\right $

SoLoPO can be applied to various PO algorithms, provided that the corresponding convergence function $f(\cdot)$ and upper bound function $s(\cdot)$ are specified (see Table 14). We apply SoLoPO to mainstream algorithms, including DPO, SimPO, and ORPO, with their corresponding optimization objectives shown in Table 1. For brevity, we only present the expressions for the chosen-only SoLo-RA, while the objective functions for short-context PO remain consistent with the original methods. For DPO, since $\pi_{ref}(y_w \mid x)$ is constant and not involved in differentiation, we only align $\pi_{\theta}(y_w \mid x)$. See Appendix I.10 for the complete derivation and expressions. Unless otherwise stated, SoLoPO refers to its chosen-only SoLo-RA variant in the remainder of this paper.

Table 2: Composition of different datasets and corresponding trained models. **1.** SoLo denotes short-to-long alignment, where preference pairs derived from short contexts are used for long-context alignment. **2.** "*" means the corresponding PO method used in SoLoPO. **3.** D^{SoLo} is also utilized for training LongPO, which falls under non-decoupled Short-to-Long DPO in our framework.

Method	Dataset	Trained Models
SFT	$D_{ ext{short}}^{ ext{sft}} = \{(q, x_{ ext{short}}, y_w^{ ext{short}})\}$ $D_{ ext{long}}^{ ext{sft}} = \{(q, x_{ ext{long}}, y_w^{ ext{long}})\}$	$M_{ m short}^{ m SFT} \ M_{ m long}^{ m SFT}$
РО	$D_{\text{short}}^{\text{po}} = \{(q, x_{\text{short}}, y_w^{\text{short}}, y_l^{\text{short}})\}$ $D_{\text{long}}^{\text{po}} = \{(q, x_{\text{long}}, y_w^{\text{long}}, y_l^{\text{long}})\}$	$M_{ m short}^{ m PO} \ M_{ m long}^{ m PO}$
SoLo	$D_{\text{expand-long}}^{\text{sft}} = \{(q, x_{\text{long}}, y_w^{\text{short}})\}$ $D_{\text{expand-long}}^{\text{po}} = \{(q, x_{\text{long}}, y_w^{\text{short}}, y_l^{\text{short}})\}$ $D^{\text{SoLo}} = \{(q, x_{\text{short}}, x_{\text{long}}, y_w^{\text{short}}, y_l^{\text{short}})\}$	$M_{ m expand-long}^{ m SFT}$ $M_{ m expand-long}^{ m PO}$ $M_{ m expand-long}^{(*)}$ $M_{ m SoLo}^{(*)}$

3 EXPERIMENTAL SETUP

Dataset Construction We construct $x_{\rm short}$ and $x_{\rm long}$ from the MuSiQue [61] training set using the method in RULER [30]. Specifically, we form a long context by mixing relevant documents with random unrelated ones. On average, the short and long context contain 1.1K and 7.5K tokens, respectively. We use the original QA-pairs (q,a) from Musique as the questions and ground truth answers. To obtain preference pairs, we perform sampling with a temperature of 0.85 using the instruction model. For each input $(x_{\rm short},q,a)$, we sample N=32 Chain-of-Thought [64] outputs and then select the corresponding preference pairs $(y_w^{\rm short},y_l^{\rm short})$ using the sub-em metric. Ultimately, we synthesize 5,000 training samples $D=\{(x_{\rm long},x_{\rm short},q,a,y_w^{\rm short},y_l^{\rm short})\}$. Additionally, we also sample 5,000 real long-context preference pairs $(y_w^{\rm long},y_l^{\rm long})$ based on $x_{\rm long}$. The composition of different datasets is shown in Table 2. Figure 5 shows the pipeline and more details are in Appendix D.

Baselines and Models As shown in Table 2, we compare SoLoPO with other approaches that perform SFT or original PO on different datasets. Additionally, we incorporates results from LongPO [11], which optimizes the reward of DPO based on short-to-long KL constraint, replacing $\pi_{ref}(y \mid x_{long})$ with $\pi_{ref}(y \mid x_{short})$. We also introduce results from Qwen2.5-Instruct-32B/72B and Llama3.1-Instruct-70B for comparative analysis. We use Qwen2.5-7B-Instruct [72] and Llama3.1-8B-Instruct [33] as the backbones for our experiments with per-training context window of 32K and 128K, respectively (hereafter, Qwen2.5-Instruct and Llama3.1-Instruct are referred to as Qwen2.5 and Llama3.1 for brevity). Appendix H.6 presents experiments on Qwen2.5-Instruct-14B to evaluate the scalability of SoLoPO. More training details are provided in the Appendix E.1

Evaluation benchmarks To comprehensively analyze the effectiveness of SoLoPO, we conduct evaluations on both long-context and short-context benchmarks. The long-context benchmarks include: (1) Real-world QA tasks from LongBenchV1 [5], used to evaluate the generalization capability of different methods on multi-document and single-document question answering in real scenarios within a 32K context size. (2) Synthetic QA tasks based on RULER [30], used to evaluate the generalization capability of different methods across various context lengths (4K/8K/16K/32K). (3) We further leverage LongBenchV2 [6] to analyze SoLoPO's potential on longer and more diverse real-world long-context tasks, and employ NIAH-Plus [79] to examine different models' context knowledge utilization ability in Section 4.2. For short-context benchmarks, we use MMLU-Pro [63], MATH [26], GPQA [56], IFEval [82], and BBH [58], following Open LLM Leaderboard [19].

Following previous works [5, 84], we utilize F1-score and multiple-choice accuracy as evaluation metrics, based on task-specific formats. For a fair comparison, we select the best-performing checkpoint on LongBenchV1 within a single training epoch as the final model for cross-benchmark evaluation. See Appendix E.2 for a detailed description of evaluation settings and benchmarks.

Table 3: Performance comparison on QA tasks from LongBenchV1 and RUELR. For LongPO, "*" denotes the publicly released checkpoint, while "reimplemented" indicates our implementation trained on D^{SoLo} . **Bold** and <u>underlined</u> indicate the best and the second-best performance, respectively.

				$\mathbf{Q}A$	s-RUL	ER						
S-Doc QA	M-Doc QA	Avg.	4k	8k	16k	32k	Avg.					
37.8	61.1	49.4	65.7	64.4	61.2	55.0	61.6					
34.1	49.8	41.9	58.4	52.1	46.0	43.9	50.1					
28.5	64.1	46.3	72.1	68.8	52.4	23.2	54.1					
34.8	52.6	43.7	62.4	54.4	48.9	43.1	52.2					
27.5	38.3	32.9	54.7	51.9	40.6	36.3	45.9					
Qwen2.5-7B-Instruct												
29.4	39.4	34.4	53.9	50.1	37.6	34.6	44.0					
28.9	48.4	38.6	63.8	56.7	42.3	31.8	48.7					
34.8	55.8	45.3	65.9	61.4	58.4	52.4	59.5					
34.6	51.8	⁻ 43.2 ⁻	7 0.9 -	63.3	$-45.\overline{3}$	46.9	$^{-}5\overline{6}.\overline{6}$					
35.7	58.2	46.9	$\overline{71.0}$	64.2	60.6	53.2	62.2					
38.0	57.6	47.8	66.4	64.5	62.7	57.7	62.8					
34.7	53.6	44.1	70.8	62.5	$-4\bar{2}.\bar{1}$	48.8	-56.0					
34.2	54.9	44.6	69.8	64.1	49.1	47.2	57.5					
38.1	<u>58.6</u>	<u>48.4</u>	69.2	<u>66.0</u>	<u>62.7</u>	57.8	63.9					
28.9	48.4	38.6	69.1	62.1	50.8	46.6	57.1					
34.8	55.8	45.3	64.9	59.9	50.0	45.6	55.1					
37.6	61.4	49.5	70.8	68.3	64.0	57.3	65.1					
Llama3.1-8B-Instruct												
30.3	49.3	39.8	58.3	49.2	42.9	35.6	46.5					
33.0	<u>56.2</u>	44.6	65.0	61.0	58.5	52.2	59.2					
35.0	57.3	<u>46.1</u>	63.7	59.0	57.5	<u>53.5</u>	58.4					
3 3.1	55.1	$-{44.1}$	- 64.0	59.3	<u>59.7</u>	50.4	$^{-}5\overline{8}.\overline{4}$					
35.4	55.4	45.4	63.2	60.1	58.9	53.2	58.9					
<u>35.2</u>	57.3	46.3	<u>64.4</u>	62.5	60.1	58.2	61.3					
	37.8 34.1 28.5 34.8 27.5 29.4 28.9 34.8 34.6 35.7 38.0 34.7 34.2 38.1 28.9 34.8 37.6 L 30.3 33.0 35.0 33.1 35.4	S-Doc QA M-Doc QA 37.8 61.1 34.1 49.8 28.5 64.1 34.8 52.6 27.5 38.3 Qwen2.5-7B-In 29.4 39.4 28.9 48.4 34.8 55.8 34.6 -51.8 35.7 58.2 38.0 57.6 34.7 53.6 34.2 54.9 38.1 58.6 28.9 48.4 34.8 55.8 37.6 61.4 Llama3.1-8B-Ir 30.3 49.3 33.0 56.2 35.0 55.1 35.4 55.4	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	S-Doc QAM-Doc QAAvg.4k 37.8 61.1 49.4 65.7 34.1 49.8 41.9 58.4 28.5 64.1 46.3 72.1 34.8 52.6 43.7 62.4 27.5 38.3 32.9 54.7 Qwen2.5-7B-Instruct 29.4 39.4 34.4 53.9 28.9 48.4 38.6 63.8 34.8 55.8 45.3 65.9 35.7 58.2 46.9 71.0 38.0 57.6 47.8 66.4 34.7 53.6 44.1 70.8 34.2 54.9 44.6 69.8 38.1 58.6 48.4 69.2 28.9 48.4 38.6 69.1 34.8 55.8 45.3 64.9 37.6 61.4 49.5 70.8 Llama3.1-8B-Instruct 30.3 49.3 39.8 58.3 33.0 56.2 44.6 65.0 33.1 55.1 44.1 -64.0 35.4 55.4 45.4 63.2	S-Doc QA M-Doc QA Avg. 4k 8k 37.8 61.1 49.4 65.7 64.4 34.1 49.8 41.9 58.4 52.1 28.5 64.1 46.3 72.1 68.8 34.8 52.6 43.7 62.4 54.4 27.5 38.3 32.9 54.7 51.9 Qwen2.5-7B-Instruct 29.4 39.4 34.4 53.9 50.1 28.9 48.4 38.6 63.8 56.7 34.8 55.8 45.3 65.9 61.4 28.9 48.4 38.6 63.8 56.7 34.8 55.8 45.3 65.9 61.4 38.0 57.6 47.8 66.4 64.5 34.7 53.6 44.1 70.8 62.5 34.2 54.9 44.6 69.8 64.1 38.1 58.6 48.4 69.2 66.0 28.9	S-Doc QA M-Doc QA Avg. 4k 8k 16k 37.8 61.1 49.4 65.7 64.4 61.2 34.1 49.8 41.9 58.4 52.1 46.0 28.5 64.1 46.3 72.1 68.8 52.4 34.8 52.6 43.7 62.4 54.4 48.9 27.5 38.3 32.9 54.7 51.9 40.6 Qwen2.5-7B-Instruct 29.4 39.4 34.4 53.9 50.1 37.6 28.9 48.4 38.6 63.8 56.7 42.3 34.8 55.8 45.3 65.9 61.4 58.4 34.6 51.8 43.2 70.9 63.3 45.3 35.7 58.2 46.9 71.0 64.2 60.6 38.0 57.6 47.8 66.4 64.5 62.7 34.7 53.6 44.1 70.8 62.5 42.1	S-Doc QA M-Doc QA Avg. 4k 8k 16k 32k 37.8 61.1 49.4 65.7 64.4 61.2 55.0 34.1 49.8 41.9 58.4 52.1 46.0 43.9 28.5 64.1 46.3 72.1 68.8 52.4 23.2 34.8 52.6 43.7 62.4 54.4 48.9 43.1 27.5 38.3 32.9 54.7 51.9 40.6 36.3 Qwen2.5-7B-Instruct 29.4 39.4 34.4 53.9 50.1 37.6 34.6 28.9 48.4 38.6 63.8 56.7 42.3 31.8 34.8 55.8 45.3 65.9 61.4 58.4 52.4 34.6 51.8 43.2 70.9 63.3 45.3 46.9 35.7 58.2 46.9 71.0 64.2 60.6 53.2 38.0 57.6 47.8 </td					

4 EXPERIMENTAL RESULTS

In this section, we present our main experimental results highlighting the effectiveness of the SoLoPO framework across various benchmarks and PO methods (§ 4.1). Through comprehensive comparative analysis, we provide deeper insights into the key components of SoLoPO: decoupling and direct chosen-only short-to-long reward alignment and analyze the impact of the reward alignment coefficient α (§ 4.2). We then experimentally validate the efficiency advantage of SoLoPO (§ 4.3).

4.1 Main Results

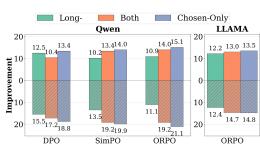
SoLoPO effectively enhances long-context capabilities within pre-training windows. As illustrated in Table 3, SoLoPO achieves substantial performance gains and strong generalization. Qwen2.5-7B, trained solely on the Musique [61] dataset, achieves a score comparable to Qwen2.5-72B on LongBenchV1. Compared to various original algorithms (DPO, SimPO, ORPO), it attains performance improvements of 0.9, 4.3, and 10.9 points, respectively. Furthermore, SoLoPO exhibits consistently superior performance across varying context lengths on RULER, with only a performance drop observed at the length of 4K for DPO and SimPO. Similarly, we conduct experiments with SoLo-ORPO, the best-performing approach, on Llama3.1-8B, and the results further validate our claims. Given that Llama3.1-8B has a context size of 128K, a 32K test length may already lie within its effective range (*i.e.* 25% [30]), and our experimental data has a maximum length of 8K; therefore, the gains are smaller compared to Qwen2.5-7B. Specifically, SoLo-ORPO attains performance gains of 4.2 vs. 0.9 on LongBenchV1 and 10.0 vs. 2.4 on RULER over Long-ORPO, for Qwen2.5-7B and Llama3.1-8B, respectively. See Appendix H.6 for scalability experiments on Qwen2.5-14B.

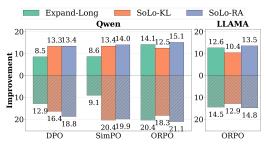
Table 4: Performance comparison of different models on *LongBenchV2* and *Open LLM Leaderboard*. Red values indicate performance degradation on short-context tasks compared to the Instruct model.

Model			LongBo	enchV2				Ope	n LLM Le	aderboard		
	Overall	Easy	Hard	<32k	32k-128k	>128k	MMLU-Pro	IFEval	ввн	MATH	GPQA	Avg.
					Qwe	n2.5-7B-Ins	truct					
Instruct	29.3(±0.7)	30.9	28.3	36.9	24.6	26.1	44.63	74.22	55.25	36.86	31.88	48.56
$M_{\scriptscriptstyle m short}^{\scriptscriptstyle m SFT}$	$30.8(\pm 0.9)$	33.6	29.1	39.2	25.7	27.2	44.81	72.18	55.15	36.71	32.12	48.19
$M_{\cdot \cdot \cdot \cdot \cdot}^{ ext{snort}}$	$30.0(\pm 1.4)$	32.0	28.8	35.7	25.9	28.7	44.74	71.46	55.35	36.55	31.45	47.90
$M_{ ext{long}}^{ ext{man}} \ M_{ ext{C}}^{ ext{ORPO}}$	$29.3(\pm 1.1)$	34.4	26.2	35.0	25.3	27.8	44.78	74.70	55.26	38.21	30.95	48.78
$M_{ ext{long}}^{ ext{short}}$	$26.6(\pm 1.2)$	30.1	24.5	33.8	22.3	23.3	44.64	73.61	55.37	37.16	32.12	48.58
$-\frac{M_{_{\mathrm{long}}}^{\mathrm{ORPO}}}{M_{_{\mathrm{short}}}^{\mathrm{DPO}}}$	$33.2(\pm 1.0)$	36.3	31.2	39.7	28.8	30.9	44.83	75.18	55.23	37.16	31.46	48.77
$\overline{M}_{\cdot}^{ ext{DPO}}$	$-25.5(\pm 1.0)$	27.2	24.5	29.7		22.8	4 4.91	75.06	54.99	39.12	31.21	49.06
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m short}$	$29.7(\pm 0.7)$	34.3	26.9	35.9	25.6	27.6	44.91	75.30	55.06	37.99	31.8	49.01
A / DPO	$35.2(\pm 1.2)$	37.5	33.8	39.3	31.8	35.0	44.66	73.98	54.78	35.57	31.63	48.12
$-\frac{M_{\text{Sol.p}}}{M_{\text{short}}^{\text{SimPO}}}$	$-24.6(\pm 1.5)$	27.1	23.1	29.5	_{21.2}	23.1	4 4. 9 7	73.50	54.94	38.6	31.46	48.69
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m short}$	$25.4(\pm 0.3)$	25.4	25.4	33.0	20.2	23.3	44.74	73.50	55.29	37.61	32.05	48.64
$M_{_{ m Sol.o}}^{^{ m Keig}}$	31.0 (±1.3)	34.1	29.1	37.5	25.7	30.6	44.78	75.90	54.89	37.76	31.54	48.97
					Llam	a3.1-8B-Ins	truct					
Instruct	32.5(±1.0)	35.5	30.7	40.7	27.7	28.5	37.79	62.23	50.98	15.25	31.71	39.59
$M_{\scriptscriptstyle m short}^{\scriptscriptstyle m SFT}$	$31.8(\pm 1.7)$	34.5	30.0	38.3	28.6	27.0	36.37	60.31	50.23	17.90	31.79	39.32
$M_{ ext{long}}^{ ext{SFT}} \ M^{ ext{ORPO}}$	$30.9(\pm 1.3)$	36.8	27.3	36.0	29.6	24.8	37.04	60.67	49.64	16.16	32.38	39.17
$M^{^{\mathrm{orpo}}}$	$33.7(\pm 0.3)$	36.8	31.8	41.2	29.8	28.7	37.26	61.15	50.95	14.88	32.13	39.27
$M_{\scriptscriptstyle m inn}^{\scriptscriptstyle m short}$	$32.1(\pm 1.1)$	34.5	30.6	39.7	27.8	28.0	37.43	60.43	50.86	15.86	31.63	39.24
$M_{ ext{long}}^{ ext{short}}$ $M_{ ext{long}}^{ ext{ORPO}}$ $M_{ ext{SeLo}}^{ ext{ORPO}}$	34.7 (±0.9)	37.5	33.0	40.0	32.1	31.2	37.60	63.43	50.18	15.63	31.38	39.64

SoLoPO shows better length generalization beyond the pre-training window. We test Qwen2.5-7B (w/ YARN [50]) and Llama3.1-8B on LongBenchV2 with results presented in Table 4. For Qwen2.5-7B (w/ YARN), SoLoPO consistently outperforms the original PO algorithms across varying difficulty levels and context lengths. This highlights the superior generalization ability of models trained with SoLoPO, demonstrating its promise in handling longer and more diverse real-world long-context tasks. For Llama3.1-8B, SoLo-ORPO shows improved performance across all evaluation dimensions, except for a slight degradation on tasks with input length <32k words. While short-text training inherently aids length generalization [21, 85], SoLoPO's advantage likely stems from SoLo-RA, which explicitly enhances contextual knowledge localization, as discussed in Appendix I.7. Ablation experiments on NIAH-Plus in Appendix H.3 further support this claim.

SoLoPO maintains short-context performance. Based on the diverse metrics in Table 4, the results demonstrate that SoLoPO maintains performance on short-context tasks. See Appendix I.9 for the supporting analysis of short-context stability in SoLoPO framework.





- (a) Comparison of different SoLo-RA approaches.
- (b) SoLoPO vs. Approximation variants.

Figure 2: Performance improvements of different short-to-long preference optimization frameworks based on various PO algorithms over Qwen2.5-7B on LongBenchV1 (top) and RULER (bottom).

4.2 IN-DEPTH EXPLORATION OF SOLOPO

Empirical evidence for SoLoPO and superior performance of chosen-only SoLo-RA. We investigate the impact of chosen-only SoLo-RA versus applying SoLo-RA jointly to both y_w and y_l (both), as defined in original SoLoPO (Eq.(9)). As shown in Figure 2a, SoLoPO with both SoLo-RA consistently outperforms Long-PO, except for a slight drop on DPO, supporting the validity of its decoupling strategy theoretically established in Section 2.2. Moreover, the chosen-only SoLo-RA surpasses the both approach across diverse algorithms and models. Appendix H.2 reveals that, compared with the both setting, the chosen-only version yields larger reward margins (Figure 10a) and assigns lower prediction probability to y_l (Figure 10b). These results suggest that chosen-only SoLo-RA achieves stronger fitting capacity and more stable convergence, leading to better performance.

Direct reward alignment matters. To validate the effectiveness of SoLo-RA, we compare it with an approximation we termed short-to-long KL divergence (SoLo-KL):

$$\alpha \cdot \mathbb{E}_{x_{\cdot} \sim \mathcal{D}_{x_{\cdot}}; y_{w} \sim \pi_{\theta}(y \mid x_{\text{short}})} |\log \pi_{\theta}(y_{w} \mid x_{\text{short}}) - \log \pi_{\theta}(y_{w} \mid x_{\text{long}})|.$$

$$\tag{10}$$

Here, we employ the absolute function to ensure non-negative training loss. From the reward expressions in Table 1, one can observe that SoLo-KL also promotes the convergence between $r_{\phi}(x_{\text{short}}, y_w)$ and $r_{\phi}(x_{\text{long}}, y_w)$. For DPO and SimPO, SoLo-KL is equivalent to SoLo-RA, as the reward coefficients β can be integrated into the coefficient α in Eq. (9). As shown in Figure 2b, for DPO and SimPO, the performance of SoLo-KL and SoLo-RA is comparable, with minor differences attributed to the slight variations in α . However, for ORPO, SoLo-RA significantly outperforms SoLo-KL on both benchmarks, demonstrating the effectiveness of direct reward alignment.

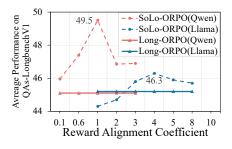
Decoupling long-context alignment yields better results. As shown in Figure 2b, SoLoPO outperforms Expand-Long-PO, a non-decoupled approach across different PO algorithms. We posit

that the decoupled approach, through SoLo-RA, explicitly improves the model's contextual knowledge utilization ability by enabling direct comparison between short and long contexts, while explicitly strengthening the model's perception of rewards and preferences. We further test the contextual knowledge utilization ability of different models on *NIAH-Plus* [79] to validate the above claim. As shown in Table 5, SoLoPO achieves consistently greater improvements over Expand-Long-ORPO, confirming the rationality and effectiveness of our decoupled approach. Similar trends can be observed in the results of DPO and SimPO presented in Table 9.

Table 5: Performance gains of various ORPO over Instruct model on *NIAH-Plus* [79].

Model	S-Doc QA	M-Doc QA	AVG.									
Qwen2.5-7B-Instruct												
M ^{ORPO} _{expand-long}	23.94	17.29	20.62									
$M_{\scriptscriptstyle{ ext{SoLo}}}^{\scriptscriptstyle{ ext{ORPO}}}$	25.98	18.83	22.41									
	LLama3.1-8B-Instruct											
$M_{\scriptscriptstyle m expand-long}^{\scriptscriptstyle m ORPO}$	11.65	9.77	10.71									
$M_{ ext{ iny SoLo}}^{ ext{ iny ORPO}}$	11.82	20.16	15.99									

Impact of reward alignment coefficient α **in SoLoPO** To evaluate the influence of α in Eq. (9), we progressively adjust α in SoLo-ORPO across two distinct foundation models. As shown in Figure 3, all architectures exhibit characteristic response curves with clear peaks in performance metrics, exceeding Long-ORPO in most settings. See Appendix H.4 for analysis of DPO and SimPO.



433

434

435 436

437

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

471 472

473

474

475

476 477 478

479 480

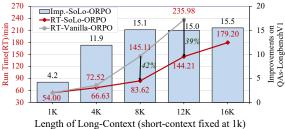
481

482

483

484

485



SoLo-ORPO.

Figure 3: Performance w/ different α in Figure 4: Run time (RT) and performance gains (Imp.) under varying lengths of x_{long} , with x_{short} fixed at 1K.

4.3 EFFICIENCY ADVANTAGE OF SOLOPO

The chosen-only SoLo-RA reduces the processing of long texts during training, thereby improving overall efficiency. As illustrated in Figure 4, we fix the length of x_{short} at 1K and investigate how varying lengths of x_{long} affect the performance and efficiency gains of SoLo-ORPO in the Qwen2.5-7B setting using $4\times A100$ GPUs. Results show that as the length of x_{long} increases, SoLo-ORPO achieves significant efficiency gains over the vanilla ORPO, cutting run time by 42% and 39% for x_{long} 's lengths of 8K and 16K, respectively. Moreover, for Qwen2.5-7B with a 32K pretrained context size, setting the length of x_{short} and x_{long} to 1K and 8K, respectively, yields a favorable trade-off between model performance and computational efficiency. Notably, as shown in Figure 1b, with only ZeRO stage 3 and offloading enabled, SoLoPO supports trainable length up to 19K tokens, while vanilla methods are limited to 9K. See Appendix H.5 for more details and discussions.

RELATED WORK 5

Long-Context Data Augmentation This approach leverages advanced LLMs to synthesize highquality, long-dependency instruction-following data, used with PO or SFT for long-context alignment [42, 76, 86, 71, 4]. However, as text length increases, it becomes less reliable and efficient [60]. Moreover, short-context alignment methods may underperform in long-context settings [15, 18, 36].

Long-Context Alignment Objective Optimization Fang et al. [18] propose LongPPL and LongCE loss to identify key tokens in long-text modeling and increase the loss weights for these critical tokens, thereby improving the effectiveness of long-context SFT. LongPO [11] searches for preference pairs based on short texts and applies them to long-context DPO training to achieve the non-decoupled short-to-long alignment discussed in our paper. Additionally, LongPO introduces a short-to-long constraint, replacing $\pi_{ref}(y \mid x_{long})$ with $\pi_{ref}(y \mid x_{short})$, thereby maintaining the short-context ability. However, LongPO focuses on context size expansion while its optimization is restricted to DPO. Both LongPO and LongCE suffer from limited training efficiency due to their reliance on long text processing, with LongCE incurring additional overhead associated with critical token detection.

SoLoPO introduces a theory-based decoupling strategy for long-context preference optimization, enabling more effective modeling of contextual knowledge localization and reasoning. The *chosen-only* SoLo-RA variant improves performance while facilitating data construction and training efficiency. Moreover, integrating SoLoPO with long-context data augmentation may further improve its alignment performance. A more comprehensive review of related work is provided in Appendix B.

CONCLUSION

In this work, we propose SoLoPO, a general framework for long-context preference optimization (PO). Our method decouples long-context PO into short-context PO and short-to-long reward alignment, supported by both theoretical and empirical evidence. Experimental results demonstrate that the chosen-only variant of SoLoPO consistently outperforms vanilla PO methods and enhances the model's generalization ability in handling long contexts across diverse domains and lengths, while significantly improving both data and training efficiency. SoLoPO highlights the importance of the connection between short and long texts, paving the way for more effective long-context alignment.

ETHICS STATEMENT

This work focuses on optimizing preference learning objectives for large language models (LLMs) in long-context scenarios. It does not introduce unexpected ethical risks beyond those commonly considered in standard NLP research. Although LLMs are trained on large amounts of Internet text that may contain harmful content, our study targets long-context understanding rather than direct deployment, which greatly reduces the risk of propagating biased information. All models and datasets used in our experiments are open-sourced and publicly available, ensuring transparency and minimizing potential ethical concerns.

REPRODUCIBILITY STATEMENT

We have provided comprehensive details of our method (Section 2.2 and Appendix I), data synthesis pipeline (Section D), model training settings (Appendix E.1), and evaluation benchmarks and configurations (Appendix E.2) in the main paper. All datasets and models used in this work are openly available, with direct access links provided in our paper and the supplementary materials. To further facilitate reproduction of our results, we release the complete source code, data examples, and step-by-step usage instructions in the supplementary materials. These resources are intended to enable other researchers to fully replicate and verify our experiments.

REFERENCES

- [1] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pp. 4447–4455. PMLR, 2024.
- [2] Yu Bai, Heyan Huang, Kai Fan, Yang Gao, Yiming Zhu, Jiaao Zhan, Zewen Chi, and Boxing Chen. Unifying cross-lingual summarization and machine translation with compression rate. In SIGIR 2022 Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2022 Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1087–1097. Association for Computing Machinery, Inc, July 2022. doi: 10.1145/3477495.3532071. Publisher Copyright: © 2022 ACM.; 45th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2022; Conference date: 11-07-2022 Through 15-07-2022.
- [3] Yu Bai, Xiyuan Zou, Heyan Huang, Sanxing Chen, Marc-Antoine Rondeau, Yang Gao, and Jackie CK Cheung. CItruS: Chunked instruction-aware state eviction for long sequence modeling. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 5908–5930, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.338. URL https://aclanthology.org/2024.emnlp-main.338/.
- [4] Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. LongAlign: A recipe for long context alignment of large language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), Findings of the Association for Computational Linguistics: EMNLP 2024, pp. 1376–1395, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.74. URL https://aclanthology.org/2024.findings-emnlp.74/.
- [5] Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. LongBench: A bilingual, multitask benchmark for long context understanding. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3119–3137, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.172. URL https://aclanthology.org/2024.acl-long.172/.
- [6] Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longbench v2: Towards deeper understanding and reasoning on realistic long-context multitasks, 2025. URL https://arxiv.org/abs/2412.15204.
- [7] Yushi Bai, Jiajie Zhang, Xin Lv, Linzhi Zheng, Siqi Zhu, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. Longwriter: Unleashing 10,000+ word generation from long context LLMs. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=kQ5s9Yh0WI.

- [8] Masha Belyi, Robert Friel, Shuai Shao, and Atindriyo Sanyal. Luna: A lightweight evaluation model to catch language model hallucinations with high accuracy and low cost. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, Steven Schockaert, Kareem Darwish, and Apoorv Agarwal (eds.), *Proceedings of the 31st International Conference on Computational Linguistics: Industry Track*, pp. 398–409, Abu Dhabi, UAE, January 2025. Association for Computational Linguistics. URL https://aclanthology.org/2025.coling-industry.34/.
 - [9] Baolong Bi, Shaohan Huang, Yiwei Wang, Tianchi Yang, Zihan Zhang, Haizhen Huang, Lingrui Mei, Junfeng Fang, Zehao Li, Furu Wei, Weiwei Deng, Feng Sun, Qi Zhang, and Shenghua Liu. Context-dpo: Aligning language models for context-faithfulness. CoRR, abs/2412.15280, 2024. URL https://doi.org/10.48550/arXiv.2412.15280.
 - [10] Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952. ISSN 00063444, 14643510. URL http://www.jstor.org/stable/2334029.
 - [11] Guanzheng Chen, Xin Li, Michael Shieh, and Lidong Bing. LongPO: Long context self-evolution of large language models through short-to-long preference optimization. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=qTrEq31Shm.
- [12] Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. LongloRA: Efficient fine-tuning of long-context large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=6PmJoRfdaK.
- [13] Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning, 2023. URL https://arxiv.org/abs/2307.08691.
- [14] Pradeep Dasigi, Kyle Lo, Iz Beltagy, Arman Cohan, Noah A. Smith, and Matt Gardner. A dataset of information-seeking questions and answers anchored in research papers. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 4599–4610, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.365. URL https://aclanthology.org/2021.naacl-main.365/.
- [15] Zican Dong, Tianyi Tang, Lunyi Li, and Wayne Xin Zhao. A survey on long text modeling with transformers. ArXiv, abs/2302.14502, 2023. URL https://api.semanticscholar.org/CorpusID: 257232619.
- [16] Zican Dong, Junyi Li, Jinhao Jiang, Mingyu Xu, Wayne Xin Zhao, Bingning Wang, and Weipeng Chen. Longred: Mitigating short-text degradation of long-context large language models via restoration distillation, 2025. URL https://arxiv.org/abs/2502.07365.
- [17] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, and et al. Ahmad Al-Dahle. The llama 3 herd of models. *ArXiv*, abs/2407.21783, 2024. URL https://api.semanticscholar.org/CorpusID:271571434.
- [18] Lizhe Fang, Yifei Wang, Zhaoyang Liu, Chenheng Zhang, Stefanie Jegelka, Jinyang Gao, Bolin Ding, and Yisen Wang. What is wrong with perplexity for long-context language modeling? In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=fL4qWkSmtM.
- [19] Clémentine Fourrier, Nathan Habib, Alina Lozovskaya, Konrad Szafer, and Thomas Wolf. Open llm leaderboard v2. https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard, 2024.
- [20] Yao Fu, Rameswar Panda, Xinyao Niu, Xiang Yue, Hannaneh Hajishirzi, Yoon Kim, and Hao Peng. Data engineering for scaling language models to 128k context. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024.* OpenReview.net, 2024. URL https://openreview.net/forum?id=TaAqeo7lUh.
- [21] Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively). In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7376–7399, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.366. URL https://aclanthology.org/2025.acl-long.366/.

- [22] Tianyu Gao, Alexander Wettig, Howard Yen, and Danqi Chen. How to train long-context language models (effectively), 2025. URL https://openreview.net/forum?id=nwZHFKrYTB.
- [23] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In First Conference on Language Modeling, 2024. URL https://openreview.net/forum?id=tEYskw1VY2.
- [24] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=uYLFoz1vlAC.
- [25] Zhenyu He, Guhao Feng, Shengjie Luo, Kai Yang, Liwei Wang, Jingjing Xu, Zhi Zhang, Hongxia Yang, and Di He. Two stones hit one bird: bilevel positional encoding for better length extrapolation. In Proceedings of the 41st International Conference on Machine Learning, ICML'24. JMLR.org, 2024.
- [26] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021. URL https://arxiv.org/abs/2103.03874.
- [27] Christian Herold and Hermann Ney. Improving long context document-level machine translation. In Michael Strube, Chloe Braud, Christian Hardmeier, Junyi Jessy Li, Sharid Loaiciga, and Amir Zeldes (eds.), Proceedings of the 4th Workshop on Computational Approaches to Discourse (CODI 2023), pp. 112–125, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.codi-1.15. URL https://aclanthology.org/2023.codi-1.15/.
- [28] Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps. In Donia Scott, Nuria Bel, and Chengqing Zong (eds.), *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6609–6625, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.580. URL https://aclanthology.org/2020.coling-main.580/.
- [29] Jiwoo Hong, Noah Lee, and James Thorne. ORPO: Monolithic preference optimization without reference model. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 11170–11189, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.626. URL https://aclanthology.org/2024.emnlp-main.626/.
- [30] Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekesh, Fei Jia, and Boris Ginsburg. RULER: What's the real context size of your long-context language models? In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=kIoBbc76Sy.
- [31] Xijie Huang, Li Lyna Zhang, Kwang-Ting Cheng, Fan Yang, and Mao Yang. Fewer is more: Boosting math reasoning with reinforced context pruning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 13674–13695, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.758. URL https://aclanthology.org/2024.emnlp-main.758/.
- [32] Sam Ade Jacobs, Masahiro Tanaka, Chengming Zhang, Minjia Zhang, Shuaiwen Leon Song, Samyam Rajbhandari, and Yuxiong He. Deepspeed ulysses: System optimizations for enabling training of extreme long sequence transformer models, 2023. URL https://arxiv.org/abs/2309.14509.
- [33] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL https://arxiv.org/abs/2310.06825.
- [34] Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. The NarrativeQA reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018. doi: 10.1162/tacl_a_00023. URL https://aclanthology.org/Q18-1023/.
- [35] Yuri Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), Advances in Neural Information Processing Systems, volume 37, pp. 106519-106554. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/c0d62e70dbc659cc9bd44cbcf1cb652f-Paper-Datasets_and_Benchmarks_Track.pdf.

- [36] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdel rahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Annual Meeting of the Association for Computational Linguistics*, 2019. URL https://api.semanticscholar.org/CorpusID:204960716.
 - [37] Jiawei Li, Yang Gao, Yizhe Yang, Yu Bai, Xiaofeng Zhou, Yinghao Li, Huashan Sun, Yuhang Liu, Xingpeng Si, Yuhao Ye, Yixiao Wu, Yiguan Lin, Bin Xu, Bowen Ren, Chong Feng, and Heyan Huang. Fundamental capabilities and applications of large language models: A survey. ACM Comput. Surv., 58 (2), September 2025. ISSN 0360-0300. doi: 10.1145/3735632. URL https://doi.org/10.1145/3735632.
 - [38] Jiawei Li, Xinyue Liang, Junlong Zhang, Yizhe Yang, Chong Feng, and Yang Gao. Pspo*: An effective process-supervised policy optimization for reasoning alignment, 2025. URL https://arxiv.org/abs/2411.11681.
 - [39] Yanyang Li, Shuo Liang, Michael Lyu, and Liwei Wang. Making long-context language models better multi-hop reasoners. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2462–2475, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024. acl-long.135. URL https://aclanthology.org/2024.acl-long.135/.
 - [40] Yucheng Li, Bo Dong, Frank Guerin, and Chenghua Lin. Compressing context to enhance inference efficiency of large language models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 6342–6353, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.391. URL https://aclanthology.org/2023.emnlp-main.391/.
 - [41] Yuhong Li, Yingbing Huang, Bowen Yang, Bharat Venkitesh, Acyr Locatelli, Hanchen Ye, Tianle Cai, Patrick Lewis, and Deming Chen. SnapKV: LLM knows what you are looking for before generation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=poE54G0q21.
 - [42] Gabrielle Kaili-May Liu, Bowen Shi, Avi Caciularu, Idan Szpektor, and Arman Cohan. Mdcure: A scalable pipeline for multi-document instruction-following. *CoRR*, abs/2410.23463, 2024. doi: 10.48550/ARXIV. 2410.23463. URL https://doi.org/10.48550/arXiv.2410.23463.
 - [43] Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Moore Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, Yuanxing Zhang, Zhuo Chen, Hangyu Guo, Shilong Li, Ziqiang Liu, Yong Shan, Yifan Song, Jiayi Tian, Wenhao Wu, Zhejian Zhou, Ruijie Zhu, Junlan Feng, Yang Gao, Shizhu He, Zhoujun Li, Tianyu Liu, Fanyu Meng, Wenbo Su, Ying Tan, Zili Wang, Jian Yang, Wei Ye, Bo Zheng, Wangchunshu Zhou, Wenhao Huang, Sujian Li, and Zhaoxiang Zhang. A comprehensive survey on long context language modeling. 2025. URL https://api.semanticscholar.org/CorpusID:277271533.
 - [44] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
 - [45] Enzhe Lu, Zhejun Jiang, Jingyuan Liu, Yulun Du, Tao Jiang, Chao Hong, Shaowei Liu, Weiran He, Enming Yuan, Yuzhi Wang, Zhiqi Huang, Huan Yuan, Suting Xu, Xinran Xu, Guokun Lai, Yanru Chen, Huabin Zheng, Junjie Yan, Jianlin Su, Yuxin Wu, Neo Y. Zhang, Zhilin Yang, Xinyu Zhou, Mingxing Zhang, and Jiezhong Qiu. Moba: Mixture of block attention for long-context llms, 2025. URL https://arxiv.org/abs/2502.13189.
 - [46] Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. When less is more: Investigating data pruning for pretraining llms at scale, 2023. URL https://arxiv.org/ abs/2309.04564.
 - [47] Yu Meng, Mengzhou Xia, and Danqi Chen. SimPO: Simple preference optimization with a reference-free reward. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=3Tzcot1LKb.
 - [48] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=TG8KACxEON.

- [49] Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Dongmei Zhang. LLMLingua-2: Data distillation for efficient and faithful task-agnostic prompt compression. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), Findings of the Association for Computational Linguistics: ACL 2024, pp. 963–981, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.57. URL https://aclanthology.org/2024.findings-acl.57/.
 - [50] Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. YaRN: Efficient context window extension of large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=wHBfxhZulu.
 - [51] Ziqian Peng, Rachel Bawden, and François Yvon. Handling Very Long Contexts in Neural Machine Translation: a Survey. Technical Report Livrable D3-2.1, Projet ANR MaTOS, June 2024. URL https://inria.hal.science/hal-04652584.
 - [52] Ziheng Qin, Kai Wang, Zangwei Zheng, Jianyang Gu, Xiangyu Peng, xu Zhao Pan, Daquan Zhou, Lei Shang, Baigui Sun, Xuansong Xie, and Yang You. Infobatch: Lossless training speed up by unbiased dynamic data pruning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=C61sk5LsK6.
 - [53] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *ArXiv*, abs/2305.18290, 2023. URL https://api.semanticscholar.org/CorpusID:258959321.
 - [54] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models, 2020. URL https://arxiv.org/abs/1910.02054.
 - [55] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In Jian Su, Kevin Duh, and Xavier Carreras (eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264. URL https://aclanthology.org/D16-1264/.
 - [56] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. Gpqa: A graduate-level google-proof q&a benchmark, 2023. URL https://arxiv.org/abs/2311.12022.
 - [57] Neil J. A. Sloane and Aaron D. Wyner. Prediction and entropy of printed english. 1951. URL https://api.semanticscholar.org/CorpusID:9101213.
 - [58] Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them, 2022. URL https://arxiv.org/abs/2210. 09261.
 - [59] Yunhao Tang, Zhaohan Daniel Guo, Zeyu Zheng, Daniele Calandriello, Remi Munos, Mark Rowland, Pierre Harvey Richemond, Michal Valko, Bernardo Avila Pires, and Bilal Piot. Generalized preference optimization: A unified approach to offline alignment. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=qu3nacA9AH.
 - [60] Zecheng Tang, Zechen Sun, Juntao Li, Qiaoming Zhu, and Min Zhang. LOGO long context alignment via efficient preference optimization. CoRR, abs/2410.18533, 2024. doi: 10.48550/ARXIV.2410.18533. URL https://doi.org/10.48550/arXiv.2410.18533.
 - [61] Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Musique: Multihop questions via single-hop question composition. *Transactions of the Association for Computational Linguistics*, 10:539–554, 2022. doi: 10.1162/tacl_a_00475. URL https://aclanthology.org/2022.tacl-1.31/.
 - [62] Minzheng Wang, Longze Chen, Fu Cheng, Shengyi Liao, Xinghua Zhang, Bingli Wu, Haiyang Yu, Nan Xu, Lei Zhang, Run Luo, Yunshui Li, Min Yang, Fei Huang, and Yongbin Li. Leave no document behind: Benchmarking long-context LLMs with extended multi-doc QA. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 5627–5646, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.322. URL https://aclanthology.org/2024.emnlp-main.322/.

- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark, 2024. URL https://arxiv.org/abs/2406.01574.
 - [64] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), Advances in Neural Information Processing Systems, 2022. URL https://openreview.net/forum?id=_VjQlMeSB_J.
 - [65] Wikipedia contributors. Kullback-Leibler divergence. https://en.wikipedia.org/wiki/Kullback%E2%80%93Leibler_divergence#cite_note-Csiszar-1, 2024. Accessed: 2024-06-10.
 - [66] Ernst C. Wit and Marie Gillette. What is linguistic redundancy. 2013. URL https://api.semanticscholar.org/CorpusID:1425655.
 - [67] Tong Wu, Yanpeng Zhao, and Zilong Zheng. An efficient recipe for long context extension via middle-focused positional encoding. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=aNHEqFMSON.
 - [68] Wenhan Xiong, Jingyu Liu, Igor Molybog, Hejia Zhang, Prajjwal Bhargava, Rui Hou, Louis Martin, Rashi Rungta, Karthik Abinav Sankararaman, Barlas Oguz, Madian Khabsa, Han Fang, Yashar Mehdad, Sharan Narang, Kshitiz Malik, Angela Fan, Shruti Bhosale, Sergey Edunov, Mike Lewis, Sinong Wang, and Hao Ma. Effective long-context scaling of foundation models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 4643–4663, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024. naacl-long.260. URL https://aclanthology.org/2024.naacl-long.260/.
 - [69] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. Large language models for generative information extraction: A survey. Frontiers of Computer Science, 18(6):186357, 2024.
 - [70] Fangyuan Xu, Weijia Shi, and Eunsol Choi. RECOMP: Improving retrieval-augmented LMs with context compression and selective augmentation. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=mlJLVigNHp.
 - [71] Cehao Yang, Xueyuan Lin, Chengjin Xu, Xuhui Jiang, Shengjie Ma, Aofan Liu, Hui Xiong, and Jian Guo. Longfaith: Enhancing long-context reasoning in llms with faithful synthetic data. CoRR, abs/2502.12583, 2025. doi: 10.48550/ARXIV.2502.12583. URL https://doi.org/10.48550/arXiv.2502. 12583.
 - [72] Qwen An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxin Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yi-Chao Zhang, Yunyang Wan, Yuqi Liu, Zeyu Cui, Zhenru Zhang, Zihan Qiu, Shanghaoran Quan, and Zekun Wang. Qwen2.5 technical report. ArXiv, abs/2412.15115, 2024. URL https://api.semanticscholar.org/CorpusID:274859421.
 - [73] Yizhe Yang, Huashan Sun, Jiawei Li, Runheng Liu, Yinghao Li, Yuhang Liu, Yang Gao, and Heyan Huang. Mindllm: Lightweight large language model pre-training, evaluation and domain application. *AI Open*, 5: 1–26, 2024. URL https://api.semanticscholar.org/CorpusID:271818589.
 - [74] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2369–2380, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL https://aclanthology.org/D18-1259/.
 - [75] Jiajie Zhang, Yushi Bai, Xin Lv, Wanjun Gu, Danqing Liu, Minhao Zou, Shulin Cao, Lei Hou, Yuxiao Dong, Ling Feng, and Juanzi Li. Longcite: Enabling LLMs to generate fine-grained citations in long-context QA, 2024. URL https://openreview.net/forum?id=mMXdHyBcHh.

- [76] Jiajie Zhang, Zhongni Hou, Xin Lv, Shulin Cao, Zhenyu Hou, Yilin Niu, Lei Hou, Yuxiao Dong, Ling Feng, and Juanzi Li. Longreward: Improving long-context large language models with AI feedback. CoRR, abs/2410.21252, 2024. doi: 10.48550/ARXIV.2410.21252. URL https://doi.org/10.48550/arXiv.2410.21252.
 - [77] Peitian Zhang, Ninglu Shao, Zheng Liu, Shitao Xiao, Hongjin Qian, Qiwei Ye, and Zhicheng Dou. Extending llama-3's context ten-fold overnight. *CoRR*, abs/2404.19553, 2024. doi: 10.48550/ARXIV. 2404.19553. URL https://doi.org/10.48550/arXiv.2404.19553.
 - [78] Xinghua Zhang, Haiyang Yu, Cheng Fu, Fei Huang, and Yongbin Li. Iopo: Empowering llms with complex instruction following via input-output preference optimization, 2024. URL https://arxiv.org/ abs/2411.06208.
 - [79] Jun Zhao, Can Zu, Xu Hao, Yi Lu, Wei He, Yiwen Ding, Tao Gui, Qi Zhang, and Xuanjing Huang. LONGA-GENT: Achieving question answering for 128k-token-long documents through multi-agent collaboration. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 16310–16324, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.912. URL https://aclanthology.org/2024.emnlp-main.912/.
 - [80] Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*, 2023.
 - [81] Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyan Luo. LlamaFactory: Unified efficient fine-tuning of 100+ language models. In Yixin Cao, Yang Feng, and Deyi Xiong (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pp. 400–410, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-demos.38. URL https://aclanthology.org/2024.acl-demos.38/.
 - [82] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models, 2023. URL https://arxiv.org/abs/2311.07911.
 - [83] Dawei Zhu, Nan Yang, Liang Wang, Yifan Song, Wenhao Wu, Furu Wei, and Sujian Li. PoSE: Efficient context window extension of LLMs via positional skip-wise training. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=3Z1gxuAQrA.
 - [84] Dawei Zhu, Xiyu Wei, Guangxiang Zhao, Wenhao Wu, Haosheng Zou, Junfeng Ran, Xun Wang, Lin Sun, Xiangzheng Zhang, and Sujian Li. Chain-of-thought matters: Improving long-context language models with reasoning path supervision. CoRR, abs/2502.20790, 2025. doi: 10.48550/ARXIV.2502.20790. URL https://doi.org/10.48550/arXiv.2502.20790.
 - [85] Wenhao Zhu, Pinzhen Chen, Hanxu Hu, Shujian Huang, Fei Yuan, Jiajun Chen, and Alexandra Birch. Generalizing from short to long: Effective data synthesis for long-context instruction tuning. *CoRR*, abs/2502.15592, February 2025. URL https://doi.org/10.48550/arXiv.2502.15592.
 - [86] Wenhao Zhu, Pinzhen Chen, Hanxu Hu, Shujian Huang, Fei Yuan, Jiajun Chen, and Alexandra Birch. Generalizing from short to long: Effective data synthesis for long-context instruction tuning, 2025. URL https://arxiv.org/abs/2502.15592.
 - [87] Wenhao Zhu, Pinzhen Chen, Hanxu Hu, Shujian Huang, Fei Yuan, Jiajun Chen, and Alexandra Birch. Generalizing from short to long: Effective data synthesis for long-context instruction tuning. ArXiv, abs/2502.15592, 2025. URL https://api.semanticscholar.org/CorpusID:276557686.

A THE USE OF LARGE LANGUAGE MODELS

We explicitly disclose that large language models (LLMs) were used solely for the following purposes: (1) Writing refinement – limited to minor grammar correction, wording improvement, and stylistic polishing of the manuscript text; (2) Data generation – specifically, the preference data required for our experiments were sampled from the corresponding open-source Instruct models, following the procedures described in the paper.

LLMs were not used for idea conception, methodological design, result analysis, or any other substantive scientific contribution to this work. All research ideas, methodological innovations, experimental executions, analyses, and conclusions were conceived, implemented, and validated entirely by the authors.

B RELATED WORK

Numerous studies focus on extending the limited pre-training context windows of LLMs to support longer inputs, including post-training on long-context corpora [20, 12, 77], designing novel architectures [24, 45, 23], or modifying positional encodings [50, 25, 67, 83]. However, researches reveal that the capability within the pretraining window of current LLMs has not been fully activated, resulting in suboptimal performance on long-context tasks [30, 35, 8, 75]. To address this challenge, existing approaches primarily focus on two aspects: data augmentation and training objective optimization.

Long-Context Alignment based on Data Augmentation Most research [42, 76, 86, 71] synthesizes high-quality, long-dependency instruction-following data for supervised fine-tuning or offline preference optimization. Instruction synthesis [42, 4] directly leverages one or multiple real long documents to prompt advanced LLMs to generate diverse instructions and responses for long-text alignment. Context synthesis [86], on the contrary, is built on real QA data by having models synthesize background contexts based on questions, then randomly concatenates multiple synthetic contexts to create long-context instruction-following data. Although data augmentation demonstrates some effectiveness, directly applying short-context training methods to the long-context setting may overlook differences between short and long contexts, thus failing to fully activate the model's potential capabilities [15, 36]. While SoLoPO incorporates the connection between short and long contexts into its training objective, it can be combined with data enhancement techniques, which has the potential to further enhance model performance.

Long-Context Alignment based on Training Objective Optimization Some works explore optimizing training objectives to further enhance long-context capabilities in LLMs. Fang et al. [18] propose LongCE loss to identify key tokens in long-text modeling and increase the loss weights for these critical tokens, thereby improving the effectiveness of long-context SFT. LOGO [60] employs multiple negative samples and adapts the SimPO objective to minimize the probability of generating various dis-preference instances. LongPO [11] searches for preference pairs based on short texts and applies them to long-context DPO training to achieve the non-decoupled short-to-long alignment discussed in our paper. Additionally, LongPO introduces a short-to-long constraint, utilizing the output distribution of short texts on the reference model as a reference during long-context DPO training (replacing $\pi_{ref}(y \mid x_{long})$) with $\pi_{ref}(y \mid x_{short})$), thereby maintaining short-context capabilities. Although LOGO and LongPO adopt similar data construction strategies to ours, they fundamentally differ in that they do not decouple the optimization objectives. As a result, these methods fall into the category of non-decoupled short-to-long alignment discussed in our work. Moreover, LOGO, LongPO and LongCE suffer from limited training efficiency due to their reliance on long text processing, with LongCE incurring extra overhead from critical token detection.

SoLoPO introduces a theoretically grounded framework for long-context preference optimization. Specifically, SoLoPO explicitly models the connection between short and long contexts, decoupling long-text preference optimization into short-text preference optimization and short-to-long reward alignment. The *chosen-only* variant of SoLoPO not only improves the model's long-context ability, but also significantly enhances the efficiency of both data construction and training procedure.

C LIMITATIONS & FUTURE WORK

Toward Extended Theoretical Analysis SoLoPO is primarily designed for long-context *input* scenarios, and therefore does not directly address the challenges of long-text *generation* [7]. Extending our theoretical analysis to long-text generation settings represents a natural and important direction for future work, which would further broaden the applicability of SoLoPO. Moreover, as observed from Equation (9), original preference optimization (short-context PO) focuses on discrepancies in the output space, whereas our proposed SoLo-RA emphasizes relationships in the input space. This raises an intriguing question: *Can the decoupled preference modeling approach underlying SoLoPO be generalized to other tasks where modeling input-side connections plays a critical role?* (other context-aware tasks, such as complex instruction following [78] and context-faithful alignment [9].) Investigating this direction may yield new insights into the design of more expressive and flexible preference optimization frameworks.

Training Efficiency Enhancements Although SoLoPO leverages the chosen-only SoLoRA strategy, it remains necessary to process long sequences, which can lead to efficiency bottlenecks when dealing with large-scale datasets. Future work could explore the decoupling strategy of SoLoPO in combination with data pruning techniques [52, 46], aiming to appropriately reduce the processing of long-context inputs and thereby improve training efficiency. Moreover, for tasks where compression rate is 100%, such as long-context machine translation, SoLoPO is equivalent to the original PO algorithm, offering no gain in training efficiency. Given that redundancy also exists in the hidden states of LLMs [31, 49, 40], future research could extend token-level compression to hidden-state-level compression, potentially by combining our approach with KV-cache compression techniques [41, 40]. Such an extension would better support a wider variety of long-text applications.

More experimental analysis Due to resource limitations, our current experiments primarily focus on efficiently activating capabilities within the model's pretrianing context window (32K). Future work should further evaluate the effectiveness of SoLoPO on even longer context and larger model scales to fully understand its capabilities. Additionally, SoLoPO introduces two hyperparameters—compression ratio c and reward alignment coefficient a—which require manual tuning. Future work could explore automated methods for determining optimal values for these parameters. Moreover, while we believe that SoLoPO could support self-evolving mechanisms for progressive context window expansion [11, 60], this remains to be validated via more comprehensive analysis.

Higher-Quality Data Synthesis While our current approach to constructing the short-to-long dataset is simple yet effective, it suffers from limited realism and diversity. Future work could explore integrating SoLoPO with existing data augmentation techniques to synthesize more realistic long-context instruction-following data, such as instruction or context synthesis grounded in authentic data sources [42, 86, 71, 4]. Moreover, extending the dataset to cover a broader range of long-text scenarios—such as long-document summarization, long-in-context learning, and long-form dialogue understanding—could provide a more comprehensive improvement of models' long-context processing capability.

D DATASET CONSTRUCTION

In this section, we present the methodology for the short-to-long preference dataset construction. As noted in our preliminary experiments and related studies [86, 22], training on shorter texts can still yield improvements in performance over longer contexts. Inspired by these findings, we heuristically set the average length of short contexts (x_{short}) to approximately 1K tokens and long contexts (x_{long}) to around 8K tokens, which corresponds to 25% of the pretraining context window (32K) of Qwen2.5-7B-Instruct. The overall data construction pipeline is illustrated in Figure 5.

Context Synthesis We follow the strategy proposed by RULER [30] to generate synthetic contexts based on the MuSiQue dataset. Specifically, given an original sample $(q, a, D_{supporting})$, where q and a denote the question and ground-truth answer respectively, and $D_{supporting}$ is the set of supporting evidence documents, we synthesize a context of target length L as follows: we randomly sample a set of irrelevant documents $D_{irrelevant}$, such that the total token count of $D_{supporting} \cup D_{irrelevant}$ approximates L. The token count is calculated using the tokenizer of the corresponding instruct model. In this setup, both short and long contexts contain the necessary information to answer the question; however, the long contexts include more distractor content, thereby simulating the redundancy in natural language.

Preference Pairs Construction To evaluate the effectiveness of short-to-long alignment compared with original long-context alignment, we construct preference pairs based on both short and long contexts using instruct models. Specifically, for each input context and question, we first generate N=32 Chain-of-Thought responses using a sampling temperature of 0.85 to encourage response diversity. We then apply the sub-em method to identify the chosen and rejected responses within the generated samples. A final preference pair is formed by randomly selecting one from the chosen and the rejected candidates, respectively. Samples where all responses are correct/incorrect are discarded.

¹Alternative evaluation methods, such as LLM-as-Judge, may be employed provided they can differentiate preference pairs.

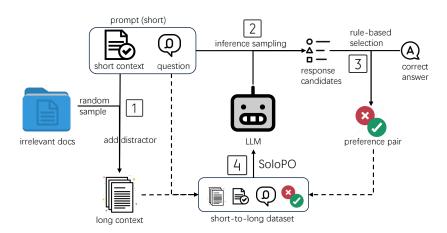


Figure 5: Illustration of the construction pipeline for the short-to-long dataset. (1) Irrelevant documents are randomly sampled and concatenated with the original short input to form long contexts. (2) Multiple candidate responses are generated based on the short context and question via the instruct model. (3) Preference pairs are curated using a sub-em¹ based selection guided by ground-truth answers. (4) The final short-to-long dataset, composed of short contexts, long contexts, questions, and preference pairs, is used for training LLM with SoLoPO.

To enhance the validity of the comparative analysis, we prioritize constructing training examples from the intersection of samples that can be successfully constructed using both short and long contexts. This ensures that differences in model performance stem primarily from alignment strategies rather than data distribution. The final size of each train set is 5000, as detials shown in Table 2 and an example is shown in Table 6. The prompt template used for data construction is illustrated in Figure 6.

Table 6: An example preference pair sampled from Qwen2.5-7B-Instruct using the short context (x_{short}) . Certain reasoning details have been omitted and denoted by "..." for conciseness.

```
When was the institute that owned The Collegian founded?
        1960
answer
        To find when the institute that owned The Collegian was founded,
        let's follow these steps:
chosen
            Identify the owner of The Collegian:...
            Find the founding year of Houston Baptist University:
        The answer is:
                        1960
        To answer the question, we need to identify which institute
        "The Collegian" is associated with and then find its founding date.
        However, the provided passages do not explicitly link The Collegian ...
        no founding date for the university is given in the passages provided.
        The answer is:
                        No answer
```

```
Write an accurate and concise answer for the given question using only the provided search results (some of which might be irrelevant). Start with an accurate, engaging, and concise explanation based only on the provided documents. Must end with 'The answer is: {{your final answer}}'. Use an unbiased and journalistic tone. If the question cannot be answered, end with 'The answer is: No answer'. Input:

Passage "[TITLE-1]":
[DOCUMENT-1]
Passage "[TITLE-2]":
[DOCUMENT-2]
...
Question: [QUERY] Think step-by-step.

Answer:
```

Figure 6: Prompt template used for data construction and training, adapted from Li et al. [39]

E IMPLEMENTATION DETAILS

In this section, we describe the implementation details of our experiments, including the configurations for model training and evaluation.

E.1 MODEL TRAINING CONFIGURATION

General training settings We train our model using LLaMAFactory [81] for data with a maximum length $\leq 8K$. To enhance training efficiency and better utilize GPU memory, we employ Flash-Attention 2 [13] and DeepSpeed ZeRO stage 3 with offloading [54] strategies. All models are fully trained on four NVIDIA A100 80GB GPUs in bf16 precision. We use the AdamW optimizer [44] together with a cosine learning rate scheduler. The warmup_ratio is set to 0.1 and the total batch size is 64. For a fair comparison, for each method we select the checkpoint that achieves the best performance on the QA tasks in LongBench-V1 during a single training epoch² as the final model to be evaluated across different benchmarks.

Supervised fine-tuning (SFT) For SFT, the maximum learning rate is set to 1×10^{-5} .

Original preference optimization We apply DPO, SimPO, and ORPO with a smaller maximum learning rate 1×10^{-6} . Following the original methods [53, 47, 29], for DPO and ORPO, we set the $\beta = 0.1$, and for SimPO, we set $\beta = 2.0$ and $\gamma = 1.4$.

Short-to-Long preference optimization (SoLo-PO) For SoLo-PO, we maintain the same training parameters as in the original method. Specifically, SoLo-DPO, SoLo-SimPO, and SoLo-ORPO achieve optimal performance on LongBenchV1 with the reward alignment coefficient α in Equation (9) set as 3, 1, and 1, respectively, for Qwen2.5-7B, while SoLo-ORPO yields better performance with $\alpha=4$ for LLaMA3.1-8B.

LongPO We train Qwen2.5-7B-Instruct on our constructed short-to-long dataset with the publicly available LongPO codebase³, using the default hyperparameter settings, including the optimizer, learning rate and the learning rate scheduler, except for the batch size, which is set to 64 to match our setup. In addition, we also evaluate the publicly released model checkpoint⁴ for direct comparison.

E.2 EVALUATION DETAILS

E.2.1 EVALUATION BENCHMARKS

To comprehensively evaluate the capabilities of SoLoPO, we conduct experiments across a diverse set of benchmark datasets, as follows:

QA tasks from LongBenchV1 and RULER Given that our training data is derived from the multihop QA dataset MuSiQue [61], with a maximum sequence length of 8K tokens, we primarily assess SoLoPO's performance on long-context QA tasks within a 32K context size. Specifically, we use the QA tasks from LongBenchV1 [5] to evaluate the model's generalization across various real-world domains. These include single-document QA tasks such as Qasper [14], NarrativeQA [34], and MultiFieldQA-En [5], as well as multi-document QA tasks including HotpotQA [74], MuSiQue [61], and 2WikiMQA [28]. Additionally, we incorporate synthetic QA tasks from RULER—SquadQA [55], HotpotQA [74], and MuSiQue [61]—at varying context lengths (4K, 8K, 16K,and 32K tokens)—to further analyze the model's length extrapolation abilities.

LongBenchV2 To explore the potential of SoLoPO in more diverse and longer-context scenarios, we further evaluate it on the full suite of tasks in LongBenchV2 [6]. This benchmark covers a wide range of long-context tasks, including question answering, abstractive summarization, and in-context learning, with input lengths spanning below 32K words, between 32K and 128K words, and beyond 128K words.

²Results show that, on our dataset, all training methods achieve their best performance within a single epoch.

³https://github.com/DAMO-NLP-SG/LongPO

⁴https://huggingface.co/DAMO-NLP-SG/Qwen2.5-7B-LongPO-128K

```
Please read the following text and answer the question below.

<text>
[CONTEXT]

</text>

What is the correct answer to this question: [QUESTION]

Choices:

(A) [CHOICE A]

(B) [CHOICE B]

(C) [CHOICE C]

(D) [CHOICE D]

Let's think step by step. And format your final answer choice as follows: "The correct answer is (insert answer here)".
```

Figure 7: Prompt template used for LongBenchV2 evaluation, adapted from Bai et al. [6]

Open LLM Leaderboard Prior works [68, 16, 11] note that aligning models for long-context tasks may forget their short-context capabilities. To assess short-context performance retention of different methods, we adopt evaluations from the Open LLM Leaderboard⁵ [19]. These include widely used tasks such as MMLU-Pro [63], MATH [26], GPQA [56], IFEval [82], and BBH [58], which valuate general knowledge, mathematical reasoning, scientific (chemistry, biology, physics) knowledge, instruction following, and complex reasoning, respectively.

NIAH-Plus As described in Section 4.2, to better understand how different training strategies affect the contextual knowledge utilization capability, we employ the NIAH-Plus [79] benchmark. This needle-in-a-haystack QA benchmark includes both single-document and multi-document settings, and is designed to directly probe a model's capacity for context-aware retrieval and multi-step reasoning.

E.2.2 EVALUATION SETTINGS

Prompts for Evaluation For QA tasks in LongBenchV1 and RULER, we use the same prompt template as employed during data construction and model training, which is illustrated in Figure 6. For all other benchmarks mentioned in this paper, we adopt their publicly released prompts. Specifically, as shown in Figure 7, for LongBenchV2, we employ a single-stage chain-of-thought prompting strategy to generate answers directly, differing from the official two-stage evaluation protocol. For the Open LLM Leaderboard and NIAH-Plus benchmarks, we follow the default prompts used in the official implementation code (lm-evaluation-harness⁶ and NIAHaystack-PLUS⁷ repositories).

Decoding hyperparameters We use greedy decoding for evaluation on LongBenchV1, RULER, and NIAH-Plus. For other benchmarks, we follow the official decoding settings with temperature values of 0.1 and 0 for LongBenchV2 and the Open LLM Leaderboard, respectively. The error analysis for LongbenchV2 is provided in Appendix F.

F STANDARD DEVIATION OF LONGBENCHV2

As mentioned in Appendix E.2, we use greedy decoding (temperature = 0) for all benchmarks except LongbenchV2, where the temperature is set to 0.1. Here, we report the standard deviation of LongbenchV2 in Figure 8 and Table 4, where all of these metrics are reasonable.

G More Detailed Benchmark Results

We present more detailed results for our experiments presented in Section 4.

 $^{^5 \}rm https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard$

 $^{^6}$ https://github.com/EleutherAI/lm-evaluation-harness

⁷https://github.com/zuucan/NeedleInAHaystack-PLUS

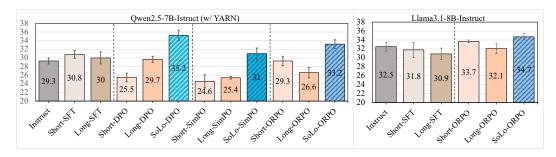


Figure 8: Overall Performance on LongbenchV2. 1. We report the average score with standard deviation across 5 evaluation runs for each model. 2. All of these metrics are reasonable.

Detailed results on QA tasks from LongbenchV1 and RULER Since the training dataset we use, MuSiQue, is designed for multi-hop QA tasks, in Section 4.1, we primarily evaluate various methods on the QA tasks from LongBenchV1 and RULER. This allows us to examine model performance on real-world scenarios and across varying input lengths. Results are presented in Table 7 and Table 8.

Table 7: Detailed Results of QA tasks from LongBenchV1

Model		Single	-Doc QA			Multi-Doc	QA		AVG.	
	NarrativeQA	Qasper	MultiFieldQA-En	Avg.	HotpotQA 2WikiMQA		MuSiQue	Avg.	vg.	
			Qwer	12.5-7B-I	nstruct					
Instruct	15.8	31.3	41.0	29.4	39.6	48.9	29.7	39.4	34.4	
$M_{\scriptscriptstyle{ m short}}^{\scriptscriptstyle{ m SFT}}$	13.4	34.0	39.3	28.9	47.3	63.3	34.5	48.4	38.6	
$M_{\scriptscriptstyle m long}^{\scriptscriptstyle m short}$	21.3	39.4	43.7	34.8	55.4	65.7	46.5	55.8	45.3	
M ^{DPO}	_{17.2}	41.3	45.4	34.6	50.9	68.0	36.4	51.8	43.2	
$M_{\scriptscriptstyle m long}^{\scriptscriptstyle m short}$	22.6	41.1	43.4	35.7	60.2	66.3	48.0	58.2	46.9	
M _{DPO}	26.1	43.0	44.8	38.0	58.8	63.3	50.7	57.6	47.8	
$M^{\scriptscriptstyle ext{SimPO}}$	$ \overline{20.3}$ $ -$	41.4	42.5	34.7	55.5	67.1	38.1	53.6	44.1	
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m Short}$	16.8	41.0	44.8	34.2	56.3	64.9	43.6	54.9	44.6	
M _{SimPO} _	<u>25.5</u>	42.8	46.0	38.1	61.5	68.4	46.0	58.6	48.4	
$M^{ ext{orpo}}$	<u> </u>	34.0	39.3	28.9	47.3	63.3	34.5	48.4	38.6	
$M_{\scriptscriptstyle m long}^{\scriptscriptstyle m MDPO}$	21.3	39.4	43.7	34.8	55.4	65.7	46.5	55.8	45.3	
$M_{_{ m ORPO}}^{_{ m SoLo}}$	24.9	41.4	46.6	37.6	64.3	71.7	<u>48.1</u>	61.4	49.5	
			LLam	a3.1-8B-l	Instruct					
Instruct	12.3	37.3	41.5	30.4	54.8	59.8	33.4	49.3	39.8	
$M^{\scriptscriptstyle ext{SFT}}$	19.1	37.1	42.9	33.0	59.0	65.4	44.4	56.3	44.6	
$M_{\scriptscriptstyle m long}^{\scriptscriptstyle m short}$	20.3	39.6	45.1	35.0	56.6	68.3	47.1	57.3	46.1	
M^{ORPO}	$ {17.1}$	38.8	43.3	33.1	55.4	67.1	42.7	55.1	44.1	
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m short}$	19.6	40.1	46.7	35.4	57.2	68.3	40.8	55.4	45.4	
M SoLo	21.5	40.0	44.2	35.2	54.8	69.0	48.2	57.3	46.3	

Table 8: Detailed Results of QA tasks from RULER

Model		Squ	adQA(O	OD)			Hot	potQA(O	OD)			MuSi(Que (InD	omain)		AVG.
	4k	8k	16k	32k	avg.	4k.	8k	16k	32k	avg.	4k	8k	16k	32k	avg.	
	Qwen2.5-7B-Instruct															
Instruct	61.3	49.5	37.0	35.2	45.8	59.6	61.4	51.3	47.7	55.0	40.9	39.4	24.6	20.8	31.4	44.0
$M_{\scriptscriptstyle m tot}^{\scriptscriptstyle m SFT}$	72.3	57.5	41.4	26.5	49.4	68.2	66.0	52.6	44.8	57.9	50.8	46.5	33.0	24.2	38.6	48.7
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m sint}$	70.9	64.4	61.8	54.1	62.8	71.0	67.7	68.4	66.7	68.5	55.9	52.1	44.9	36.5	47.3	59.5
$\overline{M}_{ ext{short}}^{ ext{DPO}}$	76.2	64.0	41.9	47.8	57.5	74.4	72.2	58.9	59.7	66.3	62.1	53.7	35.0	33.1	46.0	56.6
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m DPO}$	75.6	66.7	62.5	52.7	64.4	74.6	69.8	70.9	67.0	70.6	62.7	56.3	48.5	39.8	51.8	62.2
$M_{ ext{\tiny DPQ}}^{ ext{\tiny SoLo}}$	70.7	64.4	64.8	62.0	65.5	70.5	73.2	71.5	67.3	70.6	58.2	55.8	51.7	44.0	52.4	62.8
$M^{\scriptscriptstyle ext{SimPO}}$	75.8	62.7	39.4	50.0	57.0	74.8	70.1	53.8	62.6	65.3	61.7	54.9	33.2	33.8	45.9	56.0
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m SimPO}$	74.0	64.0	45.0	46.4	57.4	74.1	73.8	62.7	61.1	67.9	61.2	54.7	39.6	34.0	47.4	57.5
M ^{SoLo}	75.2	67.6	66.9	60.6	67.6	74.3	73.9	71.8	69.9	72.5	58.1	56.5	49.3	42.9	51.7	63.9
M ^{ORPO}	75.2	64.3	48.8	45.8	58.5	73.3	69.8	62.8	61.5	66.9	59.0	52.1	40.8	32.5	46.1	- 5 7.1 -
$M_{\scriptscriptstyle m lone}^{\scriptscriptstyle m SBDR}$	70.5	61.2	48.4	45.9	56.5	68.7	69.0	63.4	59.5	65.2	55.4	49.5	38.1	31.4	43.6	55.1
$M_{_{ m ORPO}}^{_{ m SoLo}}$	74.8	67.8	<u>65.6</u>	59.2	66.9	73.3	74.8	70.7	67.2	71.5	64.2	62.4	55.7	45.6	57.0	65.1
							LLama	3.1-8B-Ir	struct							
Instruct	67.4	53.8	44.9	41.7	52.0	68.1	61.4	57.3	49.9	59.2	39.4	32.5	26.3	15.1	28.3	46.5
$M^{\scriptscriptstyle m SFT}$	70.7	66.0	61.7	55.4	63.5	67.6	67.8	66.2	62.5	66.0	56.7	49.1	47.5	38.6	48.0	59.2
$M_{\cdot \cdot \cdot \cdot \cdot}^{ ext{short}}$	69.4	63.6	64.6	58.9	64.1	66.6	64.6	63.7	62.0	64.2	55.0	48.9	44.3	39.5	46.9	58.4
$-\frac{M_{ m long}}{M_{ m short}^{ m ORPO}}-$	70.0	65.5	67.7	52.0	63.8	67.9	- _{64.8} -	65.6	63.0	65.3	- _{54.2} -	-47.7	45.7	36.2	46.0	- 5 8. 4 -
$M^{ m orpo}$	68.5	65.1	64.9	59.2	64.4	68.5	66.7	66.7	61.4	65.8	52.7	48.5	45.1	38.9	46.3	58.9
$M_{_{ m ORPO}}^{^{ m long}}$	67.0	63.5	64.5	59.8	63.7	70.3	<u>67.0</u>	66.0	66.8	67.5	<u>55.8</u>	57.1	49.8	48.1	52.7	61.3

Table 9: Comparative performance of SoLoPO and Expand-Long-PO on NIAH-Plus (within 128K context size). As shown in (\uparrow) , SoLoPO consistently outperforms Expand-Long-PO, validating our decoupled short-to-long preference optimization approach.

Model	Single-document QA	Multi-document QA	AVG.
	Qwen-2	.5-7B-Instruct	
Instruct	35.66	52.63	44.14
$\overline{M}_{\scriptscriptstyle m expand-long}^{\scriptscriptstyle m DPO}$	55.98	68.02	62.00
$M_{\scriptscriptstyle m SoLo}^{\scriptscriptstyle m DPO}$	59.35 († 3.37)	<u>71.76</u> († 3.74)	65.56 († 3.56)
$\overline{M}_{ ext{expand-long}}^{ ext{SimPO}}$	51.81	53.61	52.71
$M_{\scriptscriptstyle m SoLo}^{\scriptscriptstyle m SimPO}$	<u>60.85</u> († 9.04)	72.05 († 18.44)	<u>66.45</u> († 13.74)
$\overline{M}_{ ext{expand-long}}^{ ext{ORPO}}$	59.60	69.92	64.76
$M_{\scriptscriptstyle{\mathrm{SoLo}}}^{\scriptscriptstyle{\mathrm{ORPO}}}$	61.64 († 2.04)	71.46 († 1.54)	66.55 († 1.79)
	LLama3	3.1-8B-Instruct	
Instruct	32.04	32.8	32.42
$\overline{M}_{ ext{expand-long}}^{ ext{ORPO}}$	43.69	42.57	43.13
$M_{\scriptscriptstyle{\mathrm{SoLo}}}^{\scriptscriptstyle{\mathrm{ORPO}}}$	43.86 († 0.17)	52.96 († 10.39)	48.41 († 5.28)

More detailed results on NIAH-Plus In Section 4.2, to validate whether the decoupled approach may more effectively enhance the model's ability to locate contextual knowledge, we evaluate Expand-Long-PO (a non-decoupled approach) and SoLoPO (our decoupled approach) on NIAH-Plus with full results presented in Table 9. SoLoPO consistently outperforms Expand-Long-PO across all evaluation scenarios and preference optimization algorithms, demonstrating the effectiveness of our decoupling-based short-to-long preference optimization approach. Figure 9 presents heatmaps of the performance gains of various ORPO over Qwen2.5-7B.

H MORE EXPERIMENTAL ANALYSIS

H.1 MORE ANALYSIS ABOUT SOLO-DPO/SIMPO ON LONGBENCHV2

Table 4 presents results of Qwen2.5-7B (w/ YARN [50]) on LongBenchV2. For Qwen2.5-7B (w/ YARN), SoLoPO consistently outperforms original PO methods across all evaluation context lengths and difficulty levels. Specifically, (1) SoLo-ORPO also surpasses vanilla PO across all dimensions; (2) SoLo-DPO achieves the best overall performance, particularly on contexts ≥ 32 and hard samples, likely due to the reference model ensuring better generalization; (3) SoLo-SimPO shows relatively

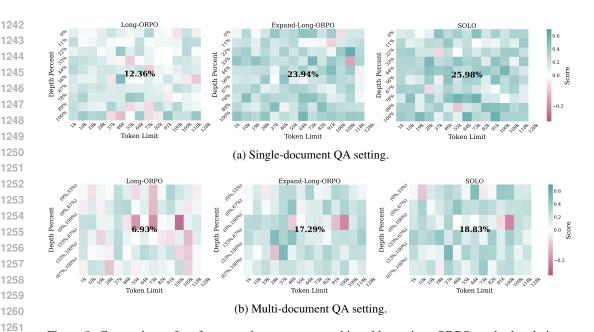


Figure 9: Comparison of performance improvements achieved by various ORPO methods relative to Qwen2.5-7B on the *NIAH-Plus* [79]. Expand-Long-ORPO and SoLoPO demonstrate significantly greater improvements than Long-ORPO, highlighting the effectiveness of short-to-long alignment. Moreover, SoLoPO provides greater gains than Expand-Long-ORPO, validating the effectiveness of our decoupled approach.

weaker performance, possibly due to its reward model relies on normalized prediction probabilities, which can underperform on long-context evaluations like perplexity observed by Fang et al. [18].

H.2 TRAINING DYNAMICS OF DIFFERENT SOLO-RA APPROACHES (chosen-only Vs. both)

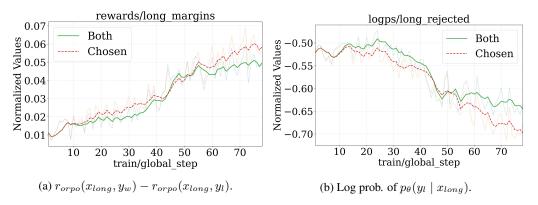


Figure 10: Changes of reward margins and log prob. of rejected response during SoLo-ORPO training given long contexts. Compared to the *both* approach, *chosen-only* short-to-long reward alignment achieves (a) larger reward margins, ensuring higher reward accuracy, while (b) simultaneously applying sufficient penalties to rejected responses, reducing their prediction probability. The chosen-only approach exhibits more precise reward modeling, ultimately achieving better alignment outcomes.

As demonstrated in Figure 10a, the margin curves of the *both* SoLo-RA consistently lie beneath those of the *chosen-only* approach, with its ultimately converged margin values being substantially lower. This observation indicates that the *chosen-only* SoLo-RA exhibits superior fitting capability along the x_{long} dimension. 10b reveals that during the initial optimization phase, the *both* SoLo-RA induces an anomalous transient increase in the probability of $p_{\theta}(y_l \mid x_{long})$, which should theoretically follow a

monotonic decreasing trend. This suboptimal convergence behavior ultimately leads to significantly inferior optimization outcomes compared to the *chosen-only* approach as shown in Figure 2a.

H.3 ABLATION EXPERIMENTS ON SOLO-RA

In Section 2.2, we hypothesize that SoLo-RA enhances the model's contextual knowledge localization ability (see Appendix I.7 for a detailed discussion). In this section, we conduct ablation studies on NIAH-Plus to further validate this hypothesis.

To conduct the ablation analysis of SoLo-RA, we focus on the performance of models trained with Short-PO, Expand-PO, and SoLoPO on NIAH-Plus, for the following reasons:

- Removing SoLo-RA from SoLoPO's optimization objective yields short-PO; however, this also eliminates the x_{long} from the training data, making a direct comparison unable to disentangle the gains from data length on generalization [85].
- Expand-PO and SoLoPO are trained on identical data except for the absence of x_{short} in Expand-PO, allowing their comparison to minimize potential confounding effects from training data length.

Table 10: Ablation studies for SoLo-RA on *NIAH-Plus* (within 128K context size).

Model	S-Doc QA	M-Doc QA	AVG.
	Qwen-2.5-7	B-Instruct	
Instruct	35.66	52.63	44.14
$\overline{M}_{ ext{short}}^{ ext{DPO}}$ —	51.25	64.70	57.98
$M_{\scriptscriptstyle m expand-long}^{\scriptscriptstyle m DPO}$	55.98	68.02	62.00
$M_{ ext{ iny SoLo}}^{ ext{ iny DPO}}$	59.35	71.76	65.56
$\overline{M}_{ ext{short}}^{ ext{SimPO}}$	50.84	63.59	57.22
$M_{\scriptscriptstyle ext{expand-long}}^{\scriptscriptstyle ext{SimPO}}$	51.81	53.61	52.71
$M_{\scriptscriptstyle{\mathrm{SoLo}}}^{\scriptscriptstyle{\mathrm{SimPO}}}$	60.85	72.05	66.45
$\overline{M}_{ ext{short}}^{ ext{ORPO}}$	45.02	59.17	52.10
$M_{\scriptscriptstyle ext{expand-long}}^{\scriptscriptstyle ext{ORPO}}$	59.60	69.92	64.76
$M_{\scriptscriptstyle m SoLo}^{\scriptscriptstyle m ORPO}$	61.64	71.46	66.55
	LLama3.1-8	BB-Instruct	
Instruct	32.04	32.80	32.42
$\overline{M}_{ ext{short}}^{ ext{ORPO}}$	37.20	37.77	37.49
$M_{\scriptscriptstyle ext{expand-long}}^{\scriptscriptstyle ext{ORPO}}$	43.69	42.57	43.13
$M_{\scriptscriptstyle m SoLo}^{\scriptscriptstyle m ORPO}$	43.86	52.96	48.41

As shown in Table 10, SoLo-PO consistently outperforms both Short-PO and Expand-PO in contextual knowledge localization tasks with a 128K context length across different PO algorithms and base models. These results provide further evidence that SoLo-RA more effectively strengthens the capacity of the model to localize contextual knowledge.

H.4 IMPACT OF REWARD ALIGNMENT COEFFICIENT α IN SoLo-DPO and SoLo-SIMPO

Figure 11 shows the performance of SoLo-DPO and SoLo-SimPO under varying α in the Qwen2.5-7B setting. The results indicate that the optimal α values are 3 for SoLo-DPO and 1 for SoLo-SimPO, and in most cases SoLoPO outperforms Long-PO, suggesting its stability across different α values.

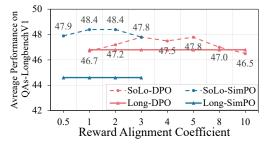


Figure 11: Performance with different α in SoLo-DPO and SoLo-SimPO in the Qwen2.5-7B setting. The optimal values of α for SoLo-DPO and SoLo-SimPO are 3 and 1, respectively.

H.5 EFFICIENCY ANALYSIS OF SOLOPO

Thanks to the reduced overhead in handling long texts, SoLoPO offers notable advantages over the original algorithms in both computational efficiency and memory usage. In this section, we present detailed empirical comparisons of runtime efficiency and provide a theoretical analysis of the computational speedup.

Training Implementation Details The experimental setup for Figure 1b adheres to the same training configuration described in Section E.1. For experiments in Section 4.3, we employ various optimization strategies of ZeRO [54] to maximize GPU memory utilization while enabling training on longer sequences. The specific training configurations are detailed in Table 11.

Table 11: Implementation details and run time of SoLo-ORPO and vanilla ORPO under varying lengths of x_{long} . 1. the default configuration is Zero stage 3 without offloading 2. "2-Stage Forward" indicates sequential forward passes for (y_w, x_{long}) and (y_l, x_{long}) , as opposed to the default concatenated forward strategy. 3. "OOM" denotes CUDA Out-of-Memory errors. 4. SoLoPO significantly improves the training efficiency of the vanilla ORPO.

Length	Method	Offloading	2-Stage Forward for x_{long}	Runtime (min) ↓
1K	vanilla			54.00
4K	vanilla			72.52
4N	SoLo			66.63 (\ 8%)
	vanilla			145.11
8K	SoLo	✓		109.42 (\preceq 25\%)
	SoLo			83.62 (\dagger 42%)
12K	vanilla			235.98
12R	SoLo	✓		144.21 (\psi 39%)
16 <i>K</i>	vanilla			OOM
1011	SoLo	✓		179.20
19K	vanilla			$-\bar{\text{OOM}}$
19 N	SoLo	\checkmark		205.98

- SoLoPO When the lengths of x_{long} ranging from 4K to 16K tokens, we employ a two-stage forward mechanism within LLaMAFactory [81] to perform SoLoPO training. Specifically, for the short-context PO, we apply the <code>concatenated_forward 8</code> function directly on (y_w, x_{short}) and (y_l, x_{short}) to obtain $\log \pi_{\theta}(y_w \mid x_{short})$ and $\log \pi_{\theta}(y_l \mid x_{short})$. Subsequently, we conduct a separate forward pass over (y_w, x_{long}) to compute $\log \pi_{\theta}(y_l \mid x_{long})$. Finally, the SoLoPO loss is calculated based on the corresponding loss function in Table 1.
- Vanilla PO When the lengths of x_{long} is less than or equal to 8K tokens, (y_w, x_{long}) and (y_l, x_{long}) can be efficiently processed using the concatenated_forward function, allowing for straightforward loss computation. However, at x_{long} lengths of 12K tokens, the use of concatenated_forward leads to out-of-memory (OOM) errors. Thus, we adopt a 2-stage forward approach—processing (y_w, x_{long}) and (y_l, x_{long}) sequentially. For sequences as long as 16K tokens, even this serialized method becomes infeasible, necessitating the use of sequence parallelism techniques to enable training. For even longer (x_{long}) exceeding 16K tokens, only a sequence parallelism [32] training strategy becomes feasible. While these alternative approaches mitigate memory constraints, they inevitably increase overall training time.

Computation speedup analysis of SoLoPO Following Bai et al. [2], we define the compression rate $c \in (0,1]$ as the ratio of the length of x_{short} to that of x_{long} . Specifically, if the length of x_{long} is denoted by N, then the length of x_{short} is given by cN. We focus on the computation incurred by the policy model during training and ignore reference models. Since Transformer operations scale quadratically with sequence length (FLOPs $\propto n^2$), the total training computation for vanilla PO and SoLoPO can be expressed as:

$$FLOPs_{PO} = 2N^2$$
, $FLOPs_{SoLoPO} = (2c^2 + 1)N^2$. (11)

Consequently, the computation speedup ratio of SoLoPO over vanilla PO can be expressed as:

$$\operatorname{Speedup}(c) = \frac{\operatorname{FLOPs}_{PO}}{\operatorname{FLOPs}_{SoLoPO}} = \frac{2}{2c^2 + 1}, \quad c \in (0, 1].$$
(12)

This indicates that SoLoPO achieves computational efficiency gains when $c < \frac{1}{\sqrt{2}} \approx 0.707$, i.e., when x_{short} is less than approximately 70.7% of the length of x_{long} .

 $^{^8}$ https://github.com/hiyouga/LLaMA-Factory/blob/main/src/llamafactory/train/dpo/trainer.py#L179

H.6 ON THE SCALABILITY OF SOLOPO TO LARGER MODELS

To examine the scalability of SoLoPO to larger models, we conduct experiments on Qwen2.5-Instruct-14B with different preference optimization methods, constrained by available computational resources. The evaluation is performed on the primary benchmarks used in this paper, LongbenchV1-QAs and RULER-QAs. The training and evaluation settings follow those of Qwen2.5-Instruct-7B, except that the parameter α in SoLoPO is set to 1 and experiments are run on 8× H20 GPUs with LLaMA-Factory (0.9.1). As shown in Table 12, SoLoPO consistently outperforms the original PO algorithms across different benchmarks while significantly improving training efficiency.

Table 12: Performance of Qwen2.5-Instruct-14B trained with different methods on LongbenchV1-QAs and RULER-QAs. Across benchmarks, SoLoPO consistently outperforms vanilla PO algorithms.

Model	QAs-	LongBenchV	1		Q	As-RUL	ER		Run Time
Model	S-Doc QA	M-Doc QA	Avg.	4k	8k	16k	32k	Avg.	/min(↓)
72B-Instruct	37.80	61.10	49.40	65.70	64.4	61.20	55.0	61.60	-
14B-Instruct	34.40	52.07	43.30	56.60	52.27	49.69	43.86	50.60	-
$M_{\scriptscriptstyle{ m short}}^{\scriptscriptstyle{ m SFT}}$	36.20	59.13	47.70	64.47	59.46	55.67	46.31	56.48	-
$M_{\scriptscriptstyle m long}^{\scriptscriptstyle m SFT}$	34.53	61.07	47.80	65.72	61.44	52.56	43.83	55.89	-
$M_{\scriptscriptstyle{ m short}}^{\scriptscriptstyle{ m DPO}}$	37.83	65.33	51.60	71.32	65.79	63.12	55.97	64.05	-
$M_{ ext{long}}^{ ext{DPO}}$	36.63	66.63	51.60	73.37	69.49	67.69	59.83	67.60	249
$M_{\scriptscriptstyle{ ext{SoLo}}}^{\scriptscriptstyle{ ext{DPO}}}$	<u>37.80</u>	67.20	53.40	<u>72.95</u>	70.53	68.52	62.03	68.51	172
$M_{ ext{short}}^{ ext{SimPO}}$	37.47	64.30	50.90	71.43	65.87	63.21	55.69	64.05	-
$M_{ ext{long}}^{ ext{SimPO}}$	36.77	66.10	51.40	70.33	66.24	63.27	56.91	64.19	248
$M_{ ext{SoLo}}^{ ext{SimPO}}$	38.40	<u>65.67</u>	52.00	71.79	67.85	66.52	60.47	66.66	169
$M_{ ext{short}}^{ ext{ORPO}}$	39.37	63.80	51.60	68.63	63.85	61.47	52.96	61.73	-
$M_{ ext{long}}^{ ext{ORPO}}$	38.80	62.73	50.80	68.69	64.19	62.16	52.83	61.97	248
$M_{ ext{SoLo}}^{ ext{ORPO}}$	39.80	65.70	52.80	72.45	67.85	65.10	60.78	66.54	170

I THEORETICAL DERIVATION AND SUPPORTING ANALYSIS OF SOLOPO

In this section, we formulate the theoretical foundation for the short-to-long preference optimization (SoLoPO) method through a novel reward loss function, and develop a distance metric condition for applying the framework for generalized distance metrics. Furthermore, we systematically extend this framework to mainstream preference optimization paradigms, including Direct Preference Optimization (DPO), Simple Preference Optimization (SimPO) and so on, demonstrating its methodological generality.

I.1 Proof of Lemma 1

Lemma 1. If the function $f(\cdot)$ of equation (4) is convex, then the following inequality holds true:

```
l_{\eta,\gamma}(x_{long}, x_{long}; y_w, y_l) \leq \frac{1}{3} [l_{3\eta, \frac{\gamma}{3}}(x_{long}, x_{short}; y_w, y_w) + l_{3\eta, \frac{\gamma}{3}}(x_{short}, x_{short}; y_w, y_l) + l_{3\eta, \frac{\gamma}{3}}(x_{short}, x_{long}; y_l, y_l)] 
(13)
```

Proof.

$$\begin{array}{ll} 1460 \\ 1461 & l_{\eta,\gamma}(x_{long},x_{long};y_w,y_l) \\ 1462 & = f(\eta\cdot [r_{\phi}(x_{long},y_w)-r_{\phi}(x_{long},y_l)-\gamma]) \\ 1463 & = f(\eta\cdot [r_{\phi}(x_{long},y_w)-r_{\phi}(x_{short},y_w) \\ & + r_{\phi}(x_{short},y_w)-r_{\phi}(x_{short},y_w) \\ & + r_{\phi}(x_{short},y_w)-r_{\phi}(x_{long},y_l)-\gamma]) \\ 1466 & + r_{\phi}(x_{short},y_w)-r_{\phi}(x_{long},y_l)-\gamma]) \\ 1467 & = f(\eta\cdot [\Delta_1+\Delta_2+\Delta_3-\gamma]) \\ 1468 & = f(\frac{1}{3}[\eta\cdot (3\Delta_1-\gamma+3\Delta_2-\gamma+3\Delta_3-\gamma)]) \\ 1470 & \leq \frac{1}{3}[f(\eta\cdot (3\Delta_1-\gamma))+f(\eta\cdot (3\Delta_2-\gamma))+f(\eta\cdot (3\Delta_3-\gamma))] \quad \text{by Jensen's Inequality} \\ 1471 & = \frac{1}{3}[l_{3\eta,\frac{\gamma}{3}}(x_{long},x_{short};y_w,y_w)+l_{3\eta,\frac{\gamma}{3}}(x_{short},x_{short};y_w,y_l)+l_{3\eta,\frac{\gamma}{3}}(x_{short},x_{long};y_l,y_l)] \end{array}$$

where

$$\begin{split} &\Delta_1 = r_\phi(x_{long}, y_w) - r_\phi(x_{short}, y_w) \\ &\Delta_2 = r_\phi(x_{short}, y_w) - r_\phi(x_{short}, y_l) \\ &\Delta_3 = r_\phi(x_{short}, y_l) - r_\phi(x_{long}, y_l) \end{split}$$

Proof of Theorem 1

Proof. Directly applying expectation $\mathbb{E}_{x.\sim \mathcal{D}_{x.};y_w,y_l\sim \mathcal{D}_y}$ to inequality 13 and using assumption 1, we can obtain the following inequality:

$$\mathcal{L}_{\eta,\gamma}(\mathcal{D}_{x_{long}}, \mathcal{D}_{x_{long}}; \mathcal{D}_{y_w \succ y_l \mid x_{long}}, \mathcal{D}_{y_w \succ y_l \mid x_{long}}) \\
\leq \frac{1}{3} \begin{bmatrix}
\mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{long}}, \mathcal{D}_{x_{short}}; \mathcal{D}_{y_w \mid x_{short}}, \mathcal{D}_{y_w \mid x_{short}}) \\
+ \mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{short}}, \mathcal{D}_{x_{short}}; \mathcal{D}_{y_w \succ y_l \mid x_{short}}, \mathcal{D}_{y_w \succ y_l \mid x_{short}}) \\
+ \mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{short}}, \mathcal{D}_{x_{long}}; \mathcal{D}_{y_l \mid x_{short}}, \mathcal{D}_{y_l \mid x_{short}})]
\end{bmatrix}.$$
(14)

We prove only the second term of inequality 14, as the proofs for the remaining terms follow in the same manner.

$$\mathbb{E}_{x.\sim\mathcal{D}_{x.};y_{w},y_{l}\sim\mathcal{D}_{y},\ y_{w}\succ y_{l}|x_{long}}[l_{3\eta,\frac{\gamma}{3}}(x_{short},x_{short};y_{w},y_{l})]$$

$$=\mathbb{E}_{x.\sim\mathcal{D}_{x.};y_{w},y_{l}\sim\mathcal{D}_{y}}[P(y_{w}\succ y_{l}|x_{long})l_{3\eta,\frac{\gamma}{3}}(x_{short},x_{short};y_{w},y_{l})]$$

$$P(y_{w}\succ y_{l}|x_{long}) P(y_{w}\succ y_{l}|x_{long}) P(y_{w}\succ y_{l}|x_{long})$$

$$(15)$$

$$= \mathbb{E}_{x_{\cdot} \sim \mathcal{D}_{x_{\cdot}}; y_{w}, y_{l} \sim \mathcal{D}_{y}} \left[\frac{P(y_{w} \succ y_{l} | x_{long})}{P(y_{w} \succ y_{l} | x_{short})} P(y_{w} \succ y_{l} | x_{short}) l_{3\eta, \frac{\gamma}{3}}(x_{short}; y_{w}, y_{l}) \right]$$
(17)

$$\leq \mathbb{E}_{x. \sim \mathcal{D}_{x.}; y_w, y_l \sim \mathcal{D}_y} [P(y_w \succ y_l | x_{short}) l_{3\eta, \frac{\gamma}{3}}(x_{short}; y_w, y_l)]$$
(18)

$$= \mathcal{L}_{3\eta, \frac{\gamma}{3}}(\mathcal{D}_{x_{short}}, \mathcal{D}_{x_{short}}; \mathcal{D}_{y_w \succ y_l \mid x_{short}}, \mathcal{D}_{y_w \succ y_l \mid x_{short}})$$

$$\tag{19}$$

Now, we prove the inequality 7. We only need to consider the sum of the first term and the third term:

1514
$$\mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{long}},\mathcal{D}_{x_{short}};\mathcal{D}_{y_w|x_{short}},\mathcal{D}_{y_w|x_{short}}) + \mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{short}},\mathcal{D}_{x_{long}};\mathcal{D}_{y_l|x_{short}},\mathcal{D}_{y_l|x_{short}})$$
1516 (20)

1517
$$= \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; \\ y_w, y_l \sim \mathcal{D}_y}} \left[P(y_w \succ y_l | x_{short}) (l_{3\eta, \frac{\gamma}{3}}(x_{long}, x_{short}; y_w, y_w) + l_{3\eta, \frac{\gamma}{3}}(x_{short}, x_{long}; y_l, y_l)) \right]$$
(21)

1519
$$\leq \mathbb{E}_{\substack{x_{\cdot} \sim \mathcal{D}_{x, \cdot}; \\ y_{w}, y_{l} \sim \mathcal{D}_{y}}} \left[l_{3\eta, \frac{\gamma}{3}}(x_{long}, x_{short}; y_{w}, y_{w}) + l_{3\eta, \frac{\gamma}{3}}(x_{short}, x_{long}; y_{l}, y_{l}) \right]$$
(22)

1520
$$y_{w}, y_{l} \sim \mathcal{D}_{y}$$
1521
$$= \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.;} \\ y_{w}, y_{l} \sim \mathcal{D}_{y}}} [f(\eta \cdot (3\Delta_{1}(y_{w}) - \gamma)) + f(\eta \cdot (3\Delta_{3}(y_{l}) - \gamma))]$$
1522
$$(23)$$

$$= \mathbb{E}_{\substack{x. \sim \mathcal{D}_x, \\ y_w, y_l \sim \mathcal{D}_y}} \left[f(\eta \cdot (3\Delta_1(y_w) - \gamma)) + f(\eta \cdot (3\Delta_3(y_l) - \gamma)) \right]$$

$$(24)$$

$$= \mathbb{E}_{x.\sim \mathcal{D}_{x.}; y\sim \mathcal{D}_{y}} [f(\eta \cdot (3\Delta_{1}(y) - \gamma)) + f(\eta \cdot (3\Delta_{3}(y) - \gamma))]$$
(25)

$$\begin{array}{ll}
1525 & -\mathbb{E}_{x.\sim \mathcal{D}_{x.};y\sim \mathcal{D}_{y}}[f(\eta \cdot (3\Delta_{1}(y)-\gamma))+f(\eta \cdot (3\Delta_{3}(y)-\gamma))] \\
& = \mathbb{E}_{x.\sim \mathcal{D}_{x.};y\sim \mathcal{D}_{y}}[f(\eta \cdot (3\Delta_{1}(y)-\gamma))+f(\eta \cdot (-3\Delta_{1}(y)-\gamma))]
\end{array} \tag{25}$$

$$= \mathbb{E}_{x, \sim \mathcal{D}_x : y \sim \mathcal{D}_y} [f(3\eta \Delta_1(y) - \eta \gamma)) + f(-3\eta \Delta_1(y) - \eta \gamma))] \tag{27}$$

$$\leq \mathbb{E}_{x.\sim \mathcal{D}_x: y\sim \mathcal{D}_y} s(|3\eta \cdot (r_\phi(x_{short}, y) - r_\phi(x_{long}, y))|)$$
(28)

where $s(\cdot)$ satisfies $f(z+\gamma)+f(-z+\gamma) \leq s(|z|)$ and

$$\Delta_1(y) = r_\phi(x_{long}, y) - r_\phi(x_{short}, y) \tag{29}$$

$$\Delta_3(y) = r_\phi(x_{short}, y) - r_\phi(x_{long}, y). \tag{30}$$

I.3 Solopo for $f(x) = x^2$

Proposition 1. Following the notation of Theorem 1, if we take $f(x) = x^2$, the inequality becomes:

$$\mathcal{L}_{\eta,\gamma}(x_{long}) \le \frac{1}{3} \mathcal{L}_{3\eta,\frac{\gamma}{3}}(x_{short}) + 3\eta^2 \cdot \mathbb{E}_{x.\sim \mathcal{D}_{x.};y\sim \mathcal{D}_y} |r_{\phi}(x_{short},y) - r_{\phi}(x_{long},y)|^2 + \frac{2}{3}\gamma^2$$
(31)

Proof.

$$f(x+\gamma) + f(-x+\gamma) = (x+\gamma)^2 + (-x+\gamma)^2 = 2x^2 + 2\gamma^2$$

Therefore, by using the Theorem 1, we prove this proposition.

I.4 Solopo for $f(x) = -\log \sigma(x)$

Proposition 2. Following the notation of Theorem 1, if we take $f(x) = -\log \sigma(x)$, then the inequality can be:

$$\mathcal{L}_{\eta,\gamma}(x_{long}) \leq \frac{1}{3} \mathcal{L}_{3\eta,\frac{\gamma}{3}}(x_{short}) + \eta \cdot \mathbb{E}_{x.\sim\mathcal{D}_{x.};y\sim\mathcal{D}_{y}} |r_{\phi}(x_{short},y) - r_{\phi}(x_{long},y)| + \frac{2}{3} \log(1 + e^{\eta \cdot \gamma})$$
(32)

Proof.

$$\mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{long}},\mathcal{D}_{x_{short}};\mathcal{D}_{y_w|x_{short}},\mathcal{D}_{y_w|x_{short}}) + \mathcal{L}_{3\eta,\frac{\gamma}{3}}(\mathcal{D}_{x_{short}},\mathcal{D}_{x_{long}};\mathcal{D}_{y_l|x_{short}},\mathcal{D}_{y_l|x_{short}})$$
(33)

(34)

$$\leq \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.;} \\ y_w, y_l \sim \mathcal{D}_y}} \left[l_{3\eta, \frac{\gamma}{3}}(x_{long}, x_{short}; y, y) + l_{3\eta, \frac{\gamma}{3}}(x_{short}, x_{long}; y, y) \right]$$
(34)

(35)

For brevity, we omit the expectation in the following derivation.

$$l_{3\eta,\frac{\gamma}{2}}(x_{long}, x_{short};, y_w) + l_{3\eta,\frac{\gamma}{2}}(x_{short}, x_{long}; y, y)$$

$$\tag{36}$$

$$= f(3\eta \cdot (r_{\phi}(x_{long}, y) - r_{\phi}(x_{short}, y)) - \eta \cdot \gamma) + f(3\eta \cdot (r_{\phi}(x_{short}, y) - r_{\phi}(x_{long}, y)) - \eta \cdot \gamma)$$

$$(37)$$

we denote
$$r_{\phi}(x_{long}, y), r_{\phi}(x_{short}, y)$$
 as r_l, r_s (38)

$$= f(3\eta \cdot (r_l - r_s) - \eta \cdot \gamma) + f(3\eta \cdot (r_s - r_l) - \eta \cdot \gamma) \tag{39}$$

$$= -\log \sigma(3\eta \cdot r_l - 3\eta \cdot r_s - \eta \cdot \gamma) - \log \sigma(3\eta \cdot r_s - 3\eta \cdot r_l - \eta \cdot \gamma) \tag{40}$$

$$= -\log \frac{1}{1 + \exp\left\{-(3\eta \cdot r_l - 3\eta \cdot r_s - \eta \cdot \gamma)\right\}} - \log \frac{1}{1 + \exp\left\{-(3\eta \cdot r_s - 3\eta \cdot r_l - \eta \cdot \gamma)\right\}}$$

$$\tag{41}$$

$$= -\log \frac{e^{3\eta \cdot r_l}}{e^{3\eta \cdot r_s + \eta \cdot \gamma} + e^{3\eta \cdot r_l}} - \log \frac{e^{3\eta \cdot r_s}}{e^{3\eta \cdot r_s} + e^{3\eta \cdot r_l + \eta \cdot \gamma}}$$

$$\tag{42}$$

$$= -3\eta \cdot r_l - 3\eta \cdot r_s + \log(e^{3\eta \cdot r_s + \eta \cdot \gamma} + e^{3\eta \cdot r_l}) + \log(e^{3\eta \cdot r_s} + e^{3\eta \cdot r_l + \eta \cdot \gamma})$$

$$\tag{43}$$

$$\stackrel{(a)}{\leq} 3\mathbb{E}_{x.\sim \mathcal{D}_{x.};} |r_l - r_s| + 2\log(1 + e^{3\gamma}) \tag{44}$$

In the following, we prove the inequality (a).

If $r_l \leq r_s$, then

$$-3\eta \cdot r_l - 3\eta \cdot r_s + \log(e^{3\eta \cdot r_s + \eta \cdot \gamma} + e^{3\eta \cdot r_l}) + \log(e^{3\eta \cdot r_s} + e^{3\eta \cdot r_l + \eta \cdot \gamma}) \tag{45}$$

$$\leq -3\eta \cdot r_l - 3\eta \cdot r_s + \log(e^{3\eta \cdot r_s + \eta \cdot \gamma} + e^{3\eta \cdot r_s}) + \log(e^{3\eta \cdot r_s} + e^{3\eta \cdot r_s + \eta \cdot \gamma})] \tag{46}$$

$$= -3\eta \cdot r_l - 3\eta \cdot r_s + 6\eta \cdot r_s + 2\log(1 + e^{\eta \cdot \gamma})] \tag{47}$$

$$=3\eta \cdot r_s - 3\eta \cdot r_l + 2\log(1 + e^{\eta \cdot \gamma}) \tag{48}$$

By symmetry, we can easily obtain that if $r_s \leq r_l$, then

$$-3\eta \cdot r_l - 3\eta \cdot r_s + \log(e^{3\eta \cdot r_s + \eta \cdot \gamma} + e^{3\eta \cdot r_l}) + \log(e^{3\eta \cdot r_s} + e^{3\eta \cdot r_l + \eta \cdot \gamma}) \le 3\eta \cdot r_l - 3\eta \cdot r_s + 2\log(1 + e^{\eta \cdot \gamma})$$

$$\tag{49}$$

Therefore,

$$-3\eta \cdot r_{l} - 3\eta \cdot r_{s} + \log(e^{3\eta \cdot r_{s} + \eta \cdot \gamma} + e^{3\eta \cdot r_{l}}) + \log(e^{3\eta \cdot r_{s}} + e^{3\eta \cdot r_{l} + \eta \cdot \gamma}) \le 3\eta \cdot |r_{l} - r_{s}| + 2\log(1 + e^{\eta \cdot \gamma})$$

$$(50)$$

I.5 GENERALIZATION OF THEOREM 1

Theorem 2. Let $D_p(x_1, x_2) = (\mathbb{E}_{y \sim \mathcal{D}_y} | r_{\phi}(x_1, y) - r_{\phi}(x_2, y)|^p)^{\frac{1}{p}}$. If any divergence $D(x_1, x_2)$ satisfies

$$D_1(x_1, x_2) \le C_1 \cdot D(x_1, x_2) \tag{51}$$

where C_1 are positive constants, then the following inequality holds true for the convex function $f(x) = -\log \sigma(x)$, as in the settings of DPO, SimPO, and ORPO:

$$\mathcal{L}_{\eta,\gamma}(\mathcal{D}_{long}, \mathcal{D}_{long}; \mathcal{D}_{y|x_{long}}, \mathcal{D}_{y|x_{long}}) \tag{52}$$

$$\frac{1612}{1613} \leq \frac{1}{3} \mathcal{L}_{3\eta, \frac{\gamma}{3}}(\mathcal{D}_{short}, \mathcal{D}_{short}; \mathcal{D}_{y|x_{short}}, \mathcal{D}_{y|x_{short}}) + \eta \cdot C_1 \mathbb{E}_{x. \sim \mathcal{D}_{x.}}[D(x_{short}, x_{long})] + C_2 \quad (53)$$

where $C_2 = \frac{2}{2}(\log(1 + e^{\eta \cdot \gamma})).$

Theorem 2 guarantees that any new distance satisfying the inequality 51 can substitute the absolute distance.

Proof. This theorem can be proved by directly using the proposition 2.

1621 1622

1623 1624

1629

1630

1633 1634

1635

1637

1638

1639

1640 1641

1642

1643

1644

1645 1646

1647 1648 1649

1650 1651

1652

1654

1655

1656

1657

1658

1659

1661

1662

1663 1664 1665

1668

1669

1671

1672

1673

THE GENERAL FORMULA OF SOLOPO LOSS FUNCTION

The general formula of Short-to-Long Preference Optimization (SoLoPO) loss function:

$$\mathcal{L}_{SoLoPO} = \mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} [f(\eta \cdot [r_{\phi}(x, y_w) - r_{\phi}(x, y_l) - \gamma])] + \alpha \cdot \mathbb{E}_{x_{\cdot} \sim \mathcal{D}_{x_{\cdot}}; y \sim \mathcal{D}_y} s(|r_{\phi}(x_{short}, y) - r_{\phi}(x_{long}, y)|)$$

where f is a convex function, and $f(x+\gamma)+f(-x+\gamma)\leq s(|x|)$ for some function s. γ,α,η are hyperparameters.

DISCUSSION ON THE MODELING OF SOLOPO

Two key abilities in long-context scenarios Unlike short-context tasks such as mathematics [38], which can directly leverage the model's inherent (contextual knowledge) reasoning ability, longcontext tasks—such as question answering [61] or information extraction [69]—also require the model to possess contextual knowledge localization skills, i.e., the ability to identify task-relevant information c_{rel} from a long context c_{long} while ignoring irrelevant content c_{irr} . In other words, the model needs to first identify the key task-relevant information c_{rel} from the context c_{long} and subsequently perform reasoning upon it [37].

Explicit modeling of two key abilities in SoLoPO Recall that the optimization objective of SoLoPO is defined as follows:

$$\mathcal{L}_{SoLoPO} = \mathbb{E}_{x \sim \mathcal{D}_{x_{short}}; y_w, y_l \sim \mathcal{D}_y} \underbrace{\left[f \left(3\eta \cdot \left[r_{\phi}(x, y_w) - r_{\phi}(x, y_l) - \frac{\gamma}{3} \right] \right) \right]}_{\text{short-context preference optimization}} + \alpha \cdot \mathbb{E}_{x \sim \mathcal{D}_{x}; y \sim \mathcal{D}_{(y_w, y_l)}} \underbrace{\left[s \left(3\eta \cdot \left| r_{\phi}(x_{short}, y) - r_{\phi}(x_{long}, y) \right| \right) \right]}_{\text{short-to-long reward alignment}}.$$

$$(54)$$

$$+\alpha \cdot \mathbb{E}_{x.\sim \mathcal{D}_{x.};y.\sim \mathcal{D}_{(y_w,y_l)}} \underbrace{\left[s(3\eta \cdot | r_{\phi}(x_{short}, y) - r_{\phi}(x_{long}, y)|) \right]}_{y_w \succ y_l \sim \pi_{\theta}(y|x_{short})}$$
(55)

SoLoPO explicitly models the two abilities in a decoupled manner:

- Contextual knowledge reasoning: Since x_{short} consists of the task instruction I and the taskrelevant content c_{rel} , i.e., $x_{short} := [c_{\text{rel}}; I]$, SoLoPO directly enhances the model's inherent reasoning ability via short-context preference optimization (Eq. (54)).
- Contextual knowledge localization: SoLo-RA (Eq. 55) encourages the reward model r_{ϕ} to implicitly predict $\hat{x}_{short} \sim \hat{p}(x_{short} \mid x_{long})$ by minimizing the divergence between $\hat{x}_{short} :=$ $[\hat{c}_{rel}; I]$ and the ground-truth $x_{short} \coloneqq [c_{rel}; I]$. In preference optimization, where the reward model r_{ϕ} and the policy model π_{θ} coincide, this process also strengthens the policy model's ability to locate relevant knowledge c_{rel} for task I within a long context c_{long} .

Taking SimPO as an example, where s(|x|) = |x| + C and the reward is defined as $r_{\theta}(x,y) =$ $\frac{\beta}{|y|} \log \pi_{\theta}(y|x)$. For brevity, we set $\eta = \frac{1}{3}$ and omit constant C, the SoLo-RA loss becomes:

SoLo-RA_{SimPO} =
$$\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y \mid x_{short})}} \left[\frac{\beta}{|y|} |\log \pi_{\theta}(y \mid x_{short}) - \log \pi_{\theta}(y \mid x_{long})| \right],$$
 (56)

which encourages π_{θ} to produce an output y with the same likelihood whether it is conditioned on x_{short} or on the full input x_{long} . Under this objective, SoLo-RA **implicitly** guides the model to extract from x_{long} only the minimal sufficient information needed to behave as if conditioned on x_{short} . That is, it enforces:

$$\pi_{\theta}(y \mid x_{long}) \approx \pi_{\theta}(y \mid x_{short}) \stackrel{implicitly}{\Longrightarrow} g_{\theta'}(x_{long}) = x_{short}$$
 (57)

where $q_{\theta'}$ can be interpreted as an internal attention mechanism or latent projection that compresses x_{long} into a representation functionally equivalent to x_{short} .

This learned behavior aligns with the principle of contextual knowledge localization.

SUPPORTING ANALYSIS FOR CHOSEN-ONLY SOLO-RA

THE ORIGINAL OBJECTIVE OF SOLOPO IS THEORETICAL AND EMPIRICAL SUPPORTED

From an theoretical perspective, SoLoPO decomposes long-context preference optimization (PO) into short-context PO and short-to-long reward alignment (SoLo-RA) (§ 2.2). The experimental analysis in Section 4.2 and Figure 2a shows that directly using the SoLoPO objective—applying SoLo-RA jointly to both the chosen and rejected responses (both SoLo-RA)—achieves superior performance to Long-PO across most settings. These results offer both theoretical proof and empirical validation for the original optimization objective of **SoLoPO** with both **SoLo-RA**, as defined in Eq. (55).

I.8.2 SUPPORTING ANALYSIS FOR CHOSEN-ONLY SoLo-RA

In practical scenarios, we further consider a variant, SoLoPO with chosen-only SoLo-RA, which is motivated by two factors:

- 1. y_l may not always exploit the key information in the context. For example, y_l might simply respond, "No task-relevant content can be found in the context, so the question cannot be answered." In such cases, enforcing SoLo-RA may not improve, and could even degrade, the model's ability to locate or reason over relevant long-context information.
- 2. We aim to reduce the amount of long-text processing during training, thereby improving training efficiency.

In addition, we further examine the theoretical validity of the first motivation from a **data-sampling** perspective.

a. Relation $\pi(y \mid x_{short}) \geq \pi(y \mid x_{long})$ typically holds owing to the practical data sampling strategy. Recall that preference pairs in SoLoPO are obtained based on short contexts, as defined in Eq. (55):

$$y_w \succ y_l \sim \pi_\theta(y|x_{short}).$$
 (58)

For arbitrary x_{short} and x_{long} , the Kullback–Leibler divergence [65] satisfies:

$$D_{KL}(\pi(y|x_{short}) \parallel \pi(y|x_{long})) = \int \left[\log \pi(y|x_{short}) - \log \pi(y|x_{long})\right] \pi(y|x_{short}) \, dy \ge 0. \tag{59}$$

This implies that, when sampling $y \sim \pi_{\theta}(y \mid x_{short})$, the quantity $\log \pi(y \mid x_{short}) - \log \pi(y \mid x_{long})$ is more likely than not to be non-negative. Since the logarithm is strictly monotonic, the following inequality tends to hold (more accurately, holds in expectation; see Table 13 for empirical evidence):

$$\pi_{\theta}(y \mid x_{short}) \ge \pi_{\theta}(y \mid x_{long}), \quad \text{for} \quad y \sim \pi_{\theta}(y \mid x_{short}).$$
 (60)

b. Applying SoLo-RA to y_l may harm long-context capability. Referring to Table 1, the SoLo-RA loss becomes:

SoLo-RA_{SimPO} =
$$\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}: y \sim \mathcal{D}_y \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\frac{\beta}{|y|} \left| \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \right| \right]$$
 (61)

$$\begin{aligned}
& solar RA_{DPO} = \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\beta \left| \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} + \log \frac{\pi_{ref}(y|x_{long})}{\pi_{ref}(y|x_{short})} + \log \frac{Z(x_{short})}{Z(x_{long})} \right| \right] \\
& solar RA_{ORPO} = \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\left| \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} + \log \frac{1 - \pi_{\theta}(y|x_{long})}{1 - \pi_{\theta}(y|x_{short})} \right| \right] \end{aligned} (62)$$

SoLo-RA_{ORPO} =
$$\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\left| \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} + \log \frac{1 - \pi_{\theta}(y|x_{long})}{1 - \pi_{\theta}(y|x_{short})} \right| \right]$$
 (63)

Since the latter two terms in SoLo-RA_{DPO} are independent of the learnable parameters θ , and the reward coefficient does not affect the subsequent analysis, we treat the SoLo-RA in DPO as equivalent to that in SimPO. Expanding the expectation, we separate **cases** based on the likelihood ratio:

SoLo-RA_{DPO/SimPO} =
$$\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\mathbf{1} \left[\frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \ge 1 \right] \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \right]$$
 (64)
- $\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\mathbf{1} \left[\frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} < 1 \right] \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \right]$ (65)

$$-\mathbb{E}_{\substack{x.\sim\mathcal{D}_{x.};y\sim\mathcal{D}_{y}\\y\sim\pi_{\theta}(y|x_{short})}} \left[\mathbf{1} \left[\frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} < 1 \right] \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \right]$$
(65)

SoLo-RA_{ORPO} =
$$\mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\mathbf{1} \left[\frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} \ge 1 \right] \left(\log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} + \log \frac{1 - \pi_{\theta}(y|x_{long})}{1 - \pi_{\theta}(y|x_{short})} \right) \right]$$

$$- \mathbb{E}_{\substack{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_{y} \\ y \sim \pi_{\theta}(y|x_{short})}} \left[\mathbf{1} \left[\frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} < 1 \right] \left(\log \frac{\pi_{\theta}(y|x_{short})}{\pi_{\theta}(y|x_{long})} + \log \frac{1 - \pi_{\theta}(y|x_{long})}{1 - \pi_{\theta}(y|x_{short})} \right) \right]$$
(67)

We restrict our analysis to the cases in SoLo-RA $_{\rm DPO/SimPO}$, as SoLo-RA $_{\rm ORPO}$ can be examined in the same manner. We next consider two cases:

- Case 1: Winning Responses $(y=y_w)$ Based on Eq. (60), we have $\pi_{\theta}(y_w|x_{long}) \leq \pi_{\theta}(y_w|x_{short})$. Here, SoLo-RA minimizes $\log \frac{\pi_{\theta}(y_w|x_{short})}{\pi_{\theta}(y_w|x_{long})}$, which:
 - **Decrease** $\pi_{\theta}(y_w|x_{short})$ (nominator) but is counterbalanced by first loss term (short-context preference optimization);
 - Increases $\pi_{\theta}(y_w|x_{long})$ (denominator), aligning the objective of long-context alignment.

Since SoLoPO also includes short-context PO (Eq. (54)), this term dominates and can prevent $\pi_{\theta}(y_w|x_{short})$ from decreasing, thereby mitigating the negative impact on the model's short-context performance.

- Case 2: Losing Responses $(y = y_l)$ Based on Eq. (60), we have $\pi_{\theta}(y_l|x_{long}) \leq \pi_{\theta}(y_l|x_{short})$. SoLo-RA again minimizes $\log \frac{\pi_{\theta}(y_l|x_{short})}{\pi_{\theta}(y_l|x_{long})}$, leading to:
 - **Decrease** in $\pi_{\theta}(y_l|x_{short})$ (desirable for reducing poor generation)
 - Increase in $\pi_{\theta}(y_l|x_{long})$ (undesirable, as it promotes y_l in long contexts).

However, there is no other loss here that can reduce the prediction probability of $\pi_{\theta}(y_l|x_{long})$, therefore, using **Case 2** would bring certain negative impacts to the model's long-text capabilities.

Based on the above analysis and our goal of improving training efficiency, we adopt the *chosen-only* SoLo-RA in practical applications of SoLoPO. The experimental results in Section 4.2 and Figure 2a further demonstrate the effectiveness of the *chosen-only* SoLo-RA.

- c. No conflict with the original objective of SoLoPO It is important to note that the above analysis does not conflict with the original optimization objective of SoLoPO (with both SoLo-RA). As long as the objective is perfectly optimized—e.g., the SoLo-RA loss (Eq. (55)) reaches zero—the short-context PO (Eq. (54)) will simultaneously reduce $\pi_{\theta}(y_l|x_{short})$ and $\pi_{\theta}(y_l|x_{long})$, thereby resolving the issue in Case 2 and aligning with the objective of preference optimization.
- **d. Empirical validation of relationship** $\pi_{\theta}(y \mid x_{\text{short}}) \geq \pi_{\theta}(y \mid x_{\text{long}})$ We further empirically validate the relationship $\pi_{\theta}(y \mid x_{\text{short}}) \geq \pi_{\theta}(y \mid x_{\text{long}})$ for $y \sim \pi_{\theta}(y \mid x_{\text{short}})$ (Eq. (60)). Specifically, 100 preference pairs are sampled from x_{short} (length 1K) using Qwen-2.5-7B-Instruct, and the corresponding contexts are then extended to lengths from 4K to 32K in 4K increments to form x_{long} . For each length, we measure the proportion of instances in which the relationship holds. As shown in Table 13, the relationship is preserved in 100% of the cases across all context lengths tested.

Length of x_{long}	4K	8K	12K	16K	20K	24K	28K	32K
Ratio of $\pi_{\theta}(y_w \mid x_{\text{short}}) \geq \pi_{\theta}(y_w \mid x_{\text{long}})$	100%	100%	100%	100%	100%		100%	100%
Ratio of $\pi_{\theta}(y_l \mid x_{\text{short}}) \geq \pi_{\theta}(y_l \mid x_{\text{long}})$	100%	100%	100%	100%	100%		100%	100%

Table 13: Proportion of cases where relationship $\pi(y \mid x_{\text{short}}) \geq \pi(y \mid x_{\text{long}})$ holds for preference pairs $(y_w \succ y_l)$ sampled from x_{short} (length 1K) evaluated on x_{long} with varying lengths.

1.9 SUPPORTING ANALYSIS OF SOLOPO'S STABILITY IN SHORT-CONTEXT CAPABILITY

Appendix I.8.2 analyzes the rationale of the *chosen-only* SoLo-RA. In **Case 1**, it is noted that during optimization, $\pi_{\theta}(y_w|x_{short})$ may decrease. However, since SoLoPO also incorporates short-context PO (Eq. (54)), this term dominates and prevents $\pi_{\theta}(y_w|x_{short})$ from dropping, thereby

alleviating distributional shift in short contexts [16] and mitigating its adverse impact on short-context performance.

I.10 APPLICATIONS OF SHORT-TO-LONG PREFERENCE OPTIMIZATION LOSS

Methods	f(x)	s(x)	$r_{\phi}(x,y)$
DPO[53]	$\log(1 + e^{-x})$	$ x + 2\log(1 + e^{3\gamma})$	$\beta \log \frac{\pi_{\theta}(y x)}{\pi_{ref}(y x)} + \beta \log Z(x)$
SimPO[47]	$\log(1 + e^{-x})$	$ x + 2\log(1 + e^{3\gamma})$	$\frac{\beta}{ y } \log \pi_{\theta}(y x)$
ORPO[29]	$\log(1 + e^{-x})$	$ x + 2\log(1 + e^{3\gamma})$	$\log \frac{\pi_{\theta}(y x)}{1-\pi_{\theta}(y x)}$
IPO [1]	x^2	$2x^2 + 2\gamma^2$	$\log \frac{\pi_{\theta}(y x)}{\pi_{ref}(y x)}$
SLiC [80]	max(0, -x)	$ x -\gamma$	$\log \pi_{ heta}(y x)$
-	e^{-x}	$2e^{-\gamma}\cosh(x)$ $x^2 - \gamma^2$	-
-	$(\max(0, -x))^2$	$x^2 - \gamma^2$	-

Table 14: Variants of Convex function f(x) and upper bound function s(x).

In Section 2.3, we apply SoLoPO to several mainstream preference optimization (PO) algorithms, including DPO, SimPO, and ORPO. More generally, SoLoPO is compatible with any PO algorithm for which there exists an upper-bound function s(x) of its convergence function f(x) such that

$$f(x+\gamma) + f(-x+\gamma) \le s(x). \tag{68}$$

Table 14 lists each PO algorithm with its corresponding f(x) and s(x). It is worth noting that if the inequality is tight, the performance would be further enhanced theoretically. If s(x) is perfectly tight, i.e. $f(x+\gamma) + f(-x+\gamma) = s(x)$, then s(x) should satisfy the following properties:

- s(x) is an even function
- $\forall x, \ s(x) \geq 2 \cdot f(\gamma) = s(0)$

We subsequently present the complete theoretical objectives for all considered PO algorithms.

DPO Setting:

$$\mathcal{L}_{SoLoPO}^{DPO} = -\mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} [\log \sigma (3 \cdot [\beta \log \frac{\pi_{\theta}(y_w|x)}{\pi_{ref}(y_w|x)} - \beta \log \frac{\pi_{\theta}(y_l|x)}{\pi_{ref}(y_l|x)} - \gamma])] + 3 \cdot \mathbb{E}_{x \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_y} |\beta \log \frac{\pi_{\theta}(y|x_{short})}{\pi_{ref}(y|x_{short})} - \beta \log \frac{\pi_{\theta}(y|x_{long})}{\pi_{ref}(y|x_{long})}|$$

SimPO Setting:

$$\mathcal{L}_{SoLoPO}^{SimPO} = -\mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} [\log \sigma(3 \cdot [\frac{\beta}{|y_w|} \log \pi_{\theta}(y_w|x) - \frac{\beta}{|y_l|} \log \pi_{\theta}(y_l|x) - \gamma])]$$

$$+ 3 \cdot \mathbb{E}_{x \sim \mathcal{D}_{x_*}; y \sim \mathcal{D}_y} |\frac{\beta}{|y|} \log \pi_{\theta}(y|x_{short}) - \frac{\beta}{|y|} \log \pi_{\theta}(y|x_{long})|$$

ORPO Setting:

To maintain consistency with the vanilla ORPO, we add the conventional causal language modeling negative log-likelihood (NLL) loss.

$$\begin{split} \mathcal{L}_{SoLoPO}^{ORPO} &= -\mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} [\log \sigma(3 \cdot [\log \frac{\pi_{\theta}(y_w|x)}{1 - \pi_{\theta}(y_w|x)} - \log \frac{\pi_{\theta}(y_l|x)}{1 - \pi_{\theta}(y_l|x)} - \gamma])] \\ &+ 3 \cdot \mathbb{E}_{x \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_y} |\log \frac{\pi_{\theta}(y|x_{short})}{1 - \pi_{\theta}(y|x_{short})} - \log \frac{\pi_{\theta}(y|x_{long})}{1 - \pi_{\theta}(y|x_{long})}| \\ &+ \mathbb{E}_{x \sim \mathcal{D}_{x_{short}}; \mathcal{L}_{NLL}(\pi_{\theta}; x; y_w)} \end{split}$$

IPO Setting:

$$\mathcal{L}_{SoLoPO}^{IPO} = \mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} [3 \cdot [\log \frac{\pi_{\theta}(y_w|x)}{\pi_{ref}(y_w|x)} - \log \frac{\pi_{\theta}(y_l|x)}{\pi_{ref}(y_l|x)} - \gamma]]^2 + 18 \cdot \mathbb{E}_{x_{\cdot \cdot \sim \mathcal{D}_{x_{\cdot \cdot}}}; y \sim \mathcal{D}_y} [\log \frac{\pi_{\theta}(y|x_{short})}{\pi_{ref}(y|x_{short})} - \log \frac{\pi_{\theta}(y|x_{long})}{\pi_{ref}(y|x_{long})}]^2$$

SLiC Setting:

$$\begin{split} \mathcal{L}_{SoLoPO}^{SLiC} &= \mathbb{E}_{\substack{x \sim \mathcal{D}_{x_{short}}; \\ y_w, y_l \sim \mathcal{D}_y; y_w \succ y_l}} \left[\max(0, -\log \pi_{\theta}(y_w|x) + \log \pi_{\theta}(y_l|x) + \gamma) \right] \\ &+ \mathbb{E}_{x. \sim \mathcal{D}_{x.}; y \sim \mathcal{D}_y} |\log \pi_{\theta}(y|x_{short}) - \log \pi_{\theta}(y|x_{long})| \end{split}$$