

Finding the Sweet Spot: Preference Data Construction for Scaling Preference Optimization

Yao Xiao^{1,4,*} Hai Ye^{2,4} Linyao Chen³
Hwee Tou Ng² Lidong Bing⁴ Xiaoli Li⁵ Roy Ka-Wei Lee¹

¹Singapore University of Technology and Design

²National University of Singapore

³The University of Tokyo ⁴MiroMind

⁵Institute for Infocomm Research, A*Star, Singapore

Abstract

Iterative data generation and model retraining are widely used to align large language models (LLMs). It typically involves a policy model to generate on-policy responses and a reward model to guide training data selection. Direct Preference Optimization (DPO) further enhances this process by constructing preference pairs of chosen and rejected responses. In this work, we aim to *scale up* the number of on-policy samples via repeated random sampling to improve alignment performance. Conventional practice selects the sample with the highest reward as chosen and the lowest as rejected for DPO. However, our experiments reveal that this strategy leads to a *decline* in performance as the sample size increases. To address this, we investigate preference data construction through the lens of the underlying normal distribution of sample rewards. We categorize the reward space into seven representative points and systematically explore all 21 (C_7^2) pairwise combinations. Through evaluations on four models using AlpacaEval 2, we find that selecting the rejected response at reward position $\mu - 2\sigma$, rather than the minimum reward, is crucial for optimal performance. We finally introduce a scalable preference data construction strategy that consistently enhances model performance as the number of samples scales up.¹

1 Introduction

Large language models (LLMs) have significantly advanced natural language processing, demonstrating remarkable capabilities across various tasks (Brown et al., 2020; Wei et al., 2022; Bubeck et al., 2023). However, these models still generate unintended outputs due to their unsupervised nature (Bai et al., 2022; Weidinger et al., 2022). To

*This work was partially done during the internship of YX at MiroMind.

¹Our code and data can be found at https://github.com/XYao0000/DPO_Pair.

mitigate these issues, recent efforts have focused on improving LLM alignment with human preferences (Ouyang et al., 2022; Rafailov et al., 2023; Gheshlaghi Azar et al., 2024; Ethayarajh et al., 2024; Meng et al., 2024).

Reinforcement learning from human feedback (RLHF) has become a widely adopted framework to align LLMs with human preferences. RLHF involves first training a reward model, which then provides feedback signals to optimize a policy model through reinforcement learning, typically using Proximal Policy Optimization (PPO) (Schulman et al., 2017; Ahmadian et al., 2024). However, this approach is complex and often unstable as it requires substantial memory and computational resources to accommodate three separate models simultaneously (Xu et al., 2024; Ivison et al., 2024). To address these challenges, Rafailov et al. (2023) introduced Direct Preference Optimization (DPO), which bypasses the need for reward models. Subsequent research has focused on improving preference optimization efficiency (Liu et al., 2024b; Gheshlaghi Azar et al., 2024; Ethayarajh et al., 2024; Han et al., 2024).

Currently, with the increasing availability of powerful reward models (Gao et al., 2023; Wang et al., 2024b; Liu et al., 2024a), an effective pipeline (Figure 1) to further enhance the alignment capabilities of LLMs without human annotations has gained popularity. To start with, n on-policy responses (Tajwar et al., 2024; Guo et al., 2024; Song et al., 2024; Tang et al., 2025) are first sampled from LLMs and subsequently scored by a reward model. The response with the highest reward is selected as the chosen response, while the one with the lowest reward is selected as the rejected response to construct a preference dataset. The constructed preference dataset can be used to train the policy model through DPO in return. In practice, five samples per prompt can achieve significant performance gains (Meng et al., 2024).

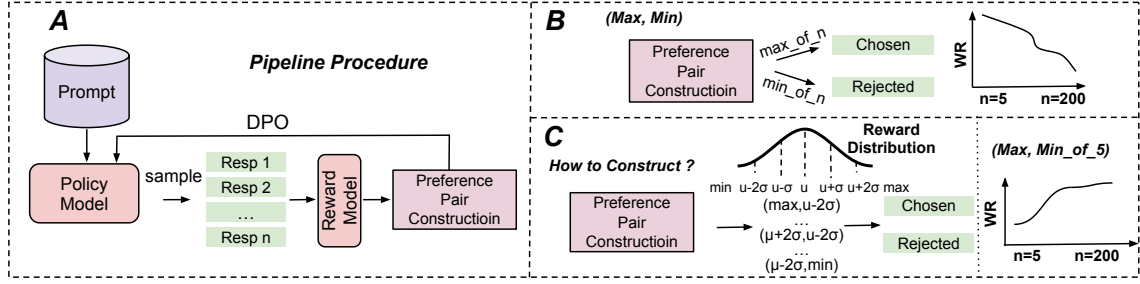


Figure 1: We orderly show (A) the procedure of preference data construction; (B) the conventional preference data construction strategy; (C) our exploration and proposed method.

In this paper, we focus on the construction of preference pairs (Kim et al., 2025) given sufficient sample budgets by scaling up the number of samples per prompt. We first construct preference pairs following the pipeline described above and train policy models. However, increasing the number of samples does not lead to any significant performance improvement and can even result in a decline, as shown in Figure 2. As preference pairs play a pivotal role in DPO training, we then investigate the construction of preference data based on the underlying normal distribution of rewards. Specifically, we classify the samples per prompt into seven categories and construct 21 preference datasets. We systematically explore the performance of trained models through DPO with each preference dataset and observe consistent empirical results across different settings. Finally, we propose a simple preference data construction strategy that can steadily improve the performance of trained models through DPO if we increase the number of samples.

Contributions. This work makes the following key contributions:

- We point out that the conventional preference pair construction strategy fails to improve the performance of models when increasing the number of samples per prompt.
- We, for the first time, construct preference pairs based on the distribution of rewards and explore their effects on policy models. We find that selecting the rejected response at reward position $\mu - 2\sigma$ is a key factor for optimal results.
- We propose a scalable preference pair construction method to improve the performance of models with an increasing number of samples per prompt. Its effectiveness can be fur-

ther demonstrated by comparing with previous work.

2 Background

2.1 Direct Preference Optimization

Different from conventional RLHF which first compresses human preferences into reward models, direct preference optimization is a RL-free algorithm for training language models to align with human preferences. DPO is recognized as one of the most widely used methods for preference optimization. It reformulates the reward function r into a closed-form expression aligned with the optimal policy:

$$r(x, y) = \beta \log \frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \log Z(x)$$

where π_{θ} denotes the policy model, π_{ref} represents the reference model (usually the supervised fine-tuned model) and $Z(x)$ is the partition function. By embedding this reward formulation into the Bradley-Terry (BT) ranking framework, the probability of preference $p(y_w > y_l|x)$ is computed as $\sigma(r(x, y_w) - r(x, y_l))$, where σ is the sigmoid function. Therefore, DPO replaces the reliance on a reward model with the policy model, resulting in the following objective:

$$\begin{aligned} \mathcal{L}_{\text{DPO}}(\pi_{\theta}; \pi_{\text{ref}}) = & \\ & - \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma(r(x, y_w) - r(x, y_l)) \right] \end{aligned}$$

$$\text{where } r(x, y) = \beta \log \frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)}.$$

2.2 Preference Data Construction

Recently, a method for constructing preference pairs without relying on human annotation has been gaining popularity (Dong et al., 2023; Meng et al., 2024). As shown in Figure 1, given a language model policy π_{θ} , a reward function r and k prompts $\{x_i\}_{i=1}^k$, we sample n generations $\{y_{ij}\}_{j=1}^n$ for the




Figure 2: Alpaca evaluation results. The conventional approach which selects the response with the highest reward as the chosen response and the response with the lowest reward as the rejected response for DPO fails to improve the performance of models when we increase the number of samples. The x-axis represents the number of samples, while the y-axis shows the score (%).

i -th prompt from π_θ . The given reward model will be used to score the sampled generations. The reward of n candidate samples of the i -th prompt is $\{r_{ij}\}_{j=1}^n$. Afterwards, the completion of the highest reward score $\max_{j=1}^n \{r_{ij}\}$ is selected as the chosen response, while the completion of the lowest reward score $\min_{j=1}^n \{r_{ij}\}$ is selected as the rejected response to construct a preference pair $(y_w^{(i)}, y_l^{(i)})$ for x_i . In practice, $n = 5$ can achieve significant performance gains (Meng et al., 2024). In this work, we explore the effects of increasing the number of samples n .

3 Case Study on Conventional Preference Data Construction

In this part, we follow the preference data strategy described in Section 2.2 to construct preference data and train policy models.

3.1 Experimental Setup

We follow the setup in Meng et al. (2024) to conduct experiments using two setups: **Base** and **Instruct**. For Base setting, we first fine-tune Llama-3-8B and Mistral-7B-v0.1 on the UltraChat-200k dataset (Ding et al., 2023) to obtain the SFT model. For Instruct setting, we use Llama-3-8B-Instruct and Mistral-7B-Instruct-v0.2 as the SFT models. With these SFT models, we sample responses for prompts from UltraFeedback (Cui et al., 2024), constructing preference datasets as described in Section 2.2. We then use the preference dataset scored with the Absolute-Rating Multi-Objective Reward Model (Armorm) (Wang et al., 2024a) to train the SFT model via DPO (Rafailov et al., 2023). For sampling, we use a temperature of 0.8 and scale the number of samples from 5 to 200. We employ vLLM (Kwon et al., 2023) for efficient inference. Following Meng et al. (2024), we set the evaluation temperature to 0.8 for Llama-3-8B and

Llama-3-8B-Instruct, while using a temperature of 0.7 for Mistral-7B-v0.1 and 0.5 for Mistral-7B-Instruct-v0.2. Additional training hyperparameters are provided in Appendix A.1.

We mainly evaluate models on **AlpacaEval 2** (Li et al., 2023; Dubois et al., 2024), which is the most widely used benchmark for instruction following. AlpacaEval 2 consists of 805 questions from multiple domains and tasks, which enables a comprehensive assessment of LLMs. We report both *win rate* and *length-controlled win rate* results (*LC win rate*) for it.

3.2 Experimental Results

The experimental results are presented in Figure 2. As the number of samples increases, the performance of the trained models, measured by both win rate and LC win rate, exhibits fluctuations rather than consistent improvements in the base setting of Llama. In the instruct setting of Llama, the degradation is even more pronounced, with performance declining as the sample size increases. A similar trend is observed for Mistral-7B-v0.1. Mistral-7B-Instruct-v0.2 shows a slight improvement before eventually declining, reinforcing the instability of this conventional preference pair construction method.

These results indicate that the max-min construction strategy cannot improve alignment of LLMs as the sample size increases. This finding highlights a limitation in existing preference construction strategies and motivates the need for alternate approaches to constructing preference pairs, particularly when ample samples are available.

4 Preference Data Construction via Reward Distribution

In this section, we explore alternative ways to categorize sampled responses based on their reward

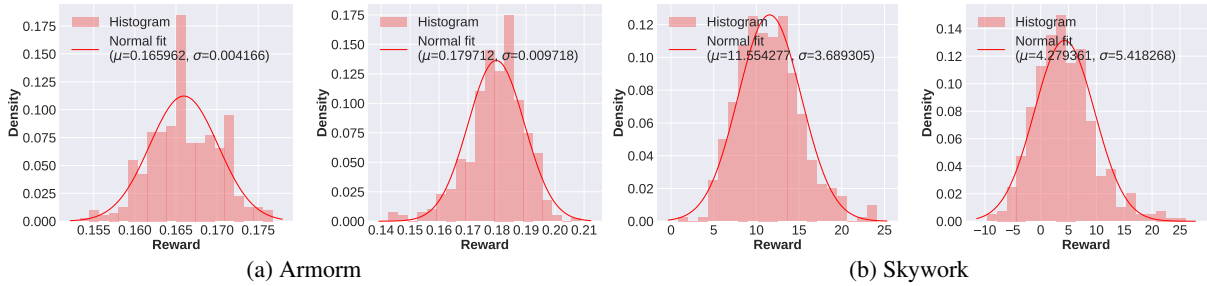


Figure 3: We plot the reward distribution of 400 completions for 2 random prompts and their approximate Gaussian distribution, calculating with Absolute-Rating Multi-Objective reward model (Armorm) (Wang et al., 2024a) and Skywork reward models (Skywork) (Liu et al., 2024a).

scores, focusing on a distribution-based approach. We first discuss the limitations of ranking-based categorization and introduce a reward distribution-based strategy. We then describe the preference pair construction process, followed by experimental validation and the key insights derived.

4.1 Reward Distribution

In reality, reward scores often exhibit a skewed or clustered distribution, making it challenging to establish meaningful distinctions using fixed ranking intervals. Instead of dividing samples into equal-sized bins, we define preference categories based on the mean (μ) and standard deviation (σ) of the underlying normal distribution, as illustrated in Figure 3. This method ensures that preference pairs are drawn from statistically meaningful intervals.

By sampling responses at key points in the distribution, such as $\mu \pm 2\sigma$, $\mu \pm \sigma$, and μ , we can capture variations in reward scores that reflect quality distinctions of responses. Another advantage of this approach is that it allows for precise control over the reward margin between chosen and rejected responses. By leveraging distribution-aware categorization, we aim to construct preference pairs systematically, enabling a more comprehensive understanding of trained models.

4.2 Preference Data Construction

We propose a structured approach to constructing preference pairs and training policy models through DPO. Our method leverages the statistical properties of reward distributions to systematically select responses for preference pair construction.

For each prompt, we first generate n responses from an SFT model and compute their reward scores using a given reward model. Given the reward scores of responses for the i -th prompt, we approximate the distribution as $N(\mu_i, \sigma_i^2)$, where μ_i and σ_i denote the mean and standard deviation

of the rewards, respectively. To ensure a representative selection of responses, we extract samples at key points in the reward distribution. Specifically, we select responses closest to the values $\mu_i - 2\sigma_i$, $\mu_i - \sigma_i$, μ_i , $\mu_i + \sigma_i$, $\mu_i + 2\sigma_i$, along with responses with minimum and maximum reward scores. This process results in a set of seven different sample points: $\{min, \mu \pm 2\sigma, \mu \pm \sigma, \mu, max\}$. Although μ and σ are prompt specific, we drop i for brevity in the rest of this paper.

The preference pairs are then constructed by considering all possible pairwise combinations of these seven points, following the principle that the chosen response should have a higher reward than the rejected response. This results in $C_7^2 = 21$ distinct preference pairs per prompt, also 21 preference datasets as a whole. We subsequently train 21 different policy models through DPO, each optimized on a unique preference dataset. Figure 1 illustrates the overall preference construction process.

4.3 Experimental Setup

We largely follow the same experimental and implementation setup described in Section 3.1. We generate 200 samples per prompt and apply the proposed preference data construction strategy. For comparison, we also evaluate models trained with conventional preference pair selection, as detailed in Section 2.2. The results of these baseline models are reported in Appendix A.2.

4.4 Experimental Results

We evaluate the performance of 84 policy models trained with the constructed preference datasets, with results presented in Figure 4. To mitigate biases introduced by response length, we primarily focus on the LC win rate as our evaluation metric (Dubois et al., 2024). In the following, we summarize our key findings and their implications for preference pair construction in DPO.

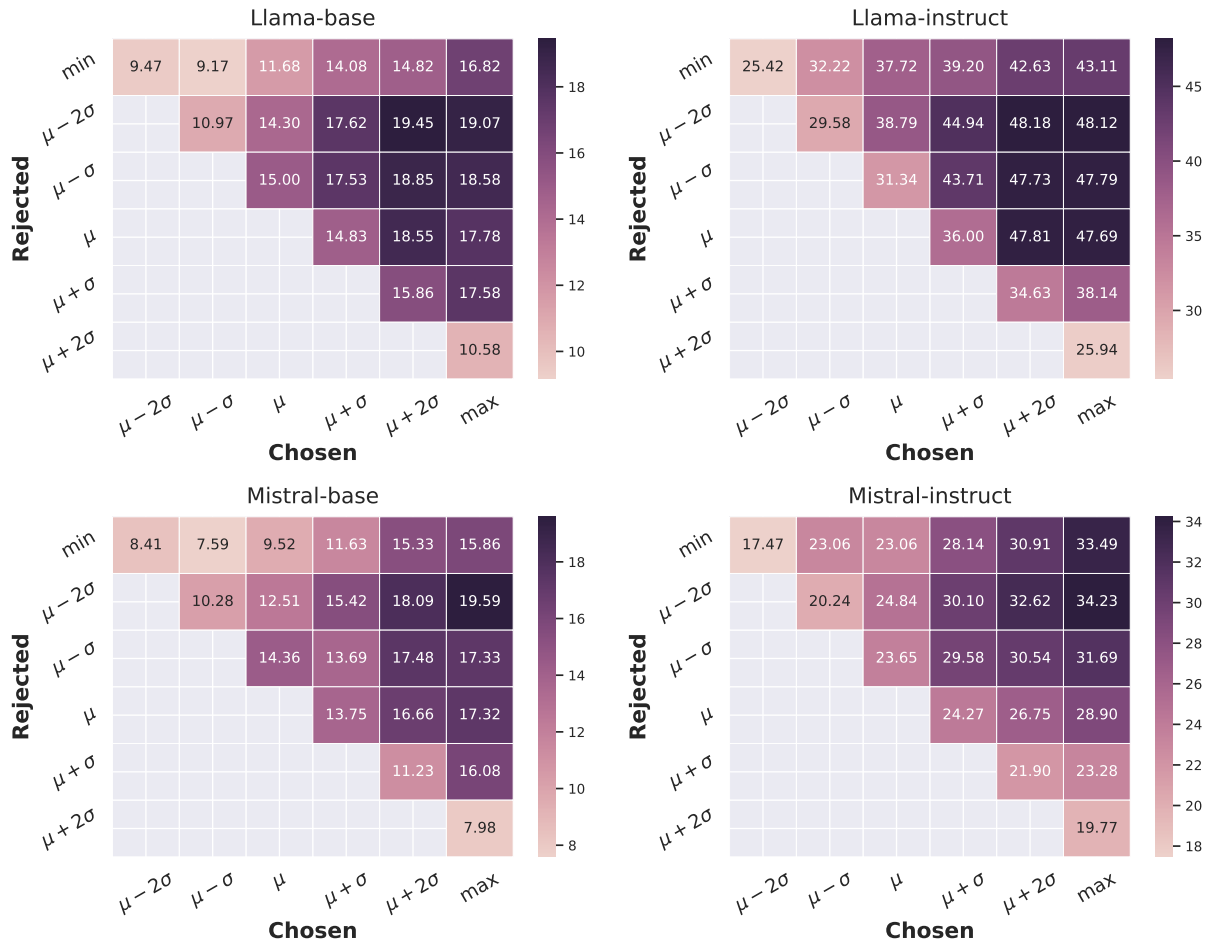


Figure 4: Alpaca evaluation results. We report the length-controlled win rate for each preference dataset here. y-axis is the reward point where the rejected response is selected, while x-axis is the reward point where the chosen response is selected.

Impact of Preference Pair Construction on Performance. Our results indicate that the chosen response should be selected from $\{max, \mu + 2\sigma\}$. In addition, the rejected response should be selected at reward position $\mu - 2\sigma$ instead of the minimum reward to produce the optimal performance. Among all preference pairs, the pair $(\mu + 2\sigma, \mu - 2\sigma)$ consistently outperforms others in most cases. For example, Llama-3-8B-Instruct trained with this preference pair achieves a length-controlled win rate of 48.18%, surpassing the conventional preference data construction strategy by about 3 percentage points. These findings suggest that preference pairs constructed from well-separated reward intervals improve preference optimization of policy models more effectively than the naive max-min strategy.

Effect of Reward Margins on Performance. A key observation from our experiments is that the performance of trained models improves as the reward of the chosen response increases, provided

that the rejected response is appropriately selected. When the rejected response is at reward position $\mu - 2\sigma$, the length-controlled win rate increases as the chosen response moves toward higher reward values. This trend is witnessed across multiple models and settings, as illustrated in Figure 4. These results reinforce the importance of ensuring a sufficiently large reward margin between the chosen and rejected responses, which contributes to more effective preference optimization.

Limitation of Small Reward Margins. We further observe that preference pairs with small reward margins perform poorly. When the reward of the chosen response is only slightly higher than that of the rejected response, the model struggles to learn meaningful distinctions. For example, training Llama-3-8B-Instruct with the pair $(\mu + 2\sigma, \mu + \sigma)$ results in a length-controlled win rate of 34.63%, significantly lower than pairs with larger reward differences.




Figure 5: We record the training loss for six datasets (max, min) , $(max, \mu - 2\sigma)$, $(max, \mu - \sigma)$, (max, μ) , $(max, \mu + \sigma)$, and $(max, \mu + 2\sigma)$ for Llama-3-8B-Instruct and Mistral-7B-Instruct-v0.2 every five steps. x-axis is the step and y-axis is the loss.

Robustness of DPO Training. Notably, none of the preference pairs degrades the performance of the SFT checkpoint. This confirms that DPO training remains robust to different preference pairs. Even for suboptimal preference pairs, model performance does not regress below the baseline established by the SFT checkpoint, highlighting the stability of DPO.

4.5 Analysis

Extending Reward Positions. To further explore the impact of preference data construction, we extend our data construction to include additional reward points at $\mu \pm 4\sigma$ and $\mu \pm 3\sigma$. Experiments conducted on Llama-3-8B-Instruct reveal that sample points at $\mu + 4\sigma$ and $\mu + 3\sigma$ show no significant difference from selecting max-reward responses. Similarly, $\mu - 4\sigma$ and $\mu - 3\sigma$ exhibit no substantial difference from selecting min-reward responses. These findings suggest that expanding the reward range beyond $\mu \pm 2\sigma$ does not provide additional benefits for preference optimization, reinforcing the sufficiency of our selected reward points. The experimental results can be found in Appendix A.4.

Scaling to 400 Samples Per Prompt. While our main experiments use 200 samples per prompt due to computational constraints, we also evaluate the scalability of our findings by conducting experiments with 400 samples per prompt. Based on the experiment with Llama-3-8B-Instruct as the SFT model, we observe that our conclusions remain consistent across both sample sizes. More details on these results are provided in Appendix A.5.

Training Dynamics and Loss Analysis. To better understand how different preference pairs influence DPO training, we record the training loss of six datasets, corresponding to the pairs (max, min) , $(max, \mu - 2\sigma)$, $(max, \mu - \sigma)$,

(max, μ) , $(max, \mu + \sigma)$, and $(max, \mu + 2\sigma)$. The loss curves, presented in Figure 5, reveal several important trends. First, larger reward margins facilitate training by enabling the model to converge more effectively. Models trained with larger reward gaps achieve lower loss values, which correlate with improved performance. By contrast, training loss for the pair $(max, \mu + 2\sigma)$ stagnates, indicating underfitting. We assume the reason is that it is difficult for the model to distinguish the chosen and rejected responses in this pair, leading to ineffective optimization. Interestingly, the preference dataset (max, min) exhibits the lowest training loss. While this may suggest faster convergence, it also raises concerns about overfitting, as models trained on this dataset fail to perform as well as those trained with intermediate reward pairs. These findings highlight the trade-off between reward margins, optimization efficiency, and generalization performance. A more detailed empirical and theoretical analysis is provided in Appendix A.6.

5 Scaling Samples to Improve Alignment

We established that selecting the rejected response at reward position $\mu - 2\sigma$ is a key factor and model performance increases as the quality of the chosen response improves. Building on these insights, we propose a simple and effective preference pair construction strategy for DPO. We further validate the effectiveness of this strategy across multiple reward models to ensure its robustness.

5.1 Data Construction Strategy

Given a language model policy, a reward function, and k prompts $\{x_i\}_{i=1}^k$, we sample n responses per prompt, denoted as $\{y_{ij}\}_{j=1}^n$, from the policy model π_θ . Each response is scored using the reward function. For the rejected response, we se-




Figure 6: Alpaca evaluation results. The rejected response is selected as the one of the minimal rewards in 5 samples, while the chosen response is selected as the one of the maximal rewards in n samples. We can improve the performance of models when increasing n within a range. x-axis is the number of samples (n), and y-axis is the performance.




Figure 7: Alpaca evaluation results. We demonstrate the effectiveness of preference data construction strategy on the Skywork reward model. x-axis is the number of samples (n), and y-axis is the performance.

lect the response with the lowest reward from 5 random samples. We find this approach to be an effective proxy for $\mu - 2\sigma$ if the sample size is insufficient to approximate the true normal distribution of rewards. For the chosen response, we select the response with the highest reward among all n samples. This ensures that as n increases, the quality of the chosen responses improves naturally, leading to better preference optimization. An illustration of the data construction process is provided in Figure 1. We further analyze how the quality of the chosen responses evolves with increasing sample size in Appendix A.7.

5.2 Experimental Setup

We evaluate our proposed preference data construction method by comparing it with the conventional approach, where the chosen response is selected as the one with the highest reward and the rejected response is the one with the lowest reward among five samples.

For our method, we begin by sampling 5 responses per prompt. The response with the lowest reward is designated as the rejected response. As we progressively increase the number of sam-

pled responses, we continue to select the chosen response as the one with the highest reward among all available candidates. This approach ensures that as the sample size grows, the quality of the chosen response improves, allowing us to examine the impact of a larger sample pool on model alignment. All experiments are conducted by following the implementation details outlined in Section 3.1, unless specified otherwise.

5.3 Experimental Results and Analysis

Scaling the Number of Samples. The results of our proposed preference data construction strategy are presented in Figure 6. For reference, the first point in each line represents the performance of the conventional approach. Since our method is only identical to the baseline when using five samples per prompt, performance differences emerge as n increases. We observe a steady improvement in performance across all models as we increase the number of samples from 5 to 200, even though the rate of improvement may diminish in some cases. The only exception occurs in Llama-3-8B-Instruct, where performance experiences a slight drop when increasing the number of samples from 100 to 200.

Comparison with Prior Work. To further validate the effectiveness of our method, we compare it with the results of Meng et al. (2024) (first 2 rows) in Table 1, which employs the conventional data construction strategy. Our method can outperform baseline results of DPO in both benchmarks, AlpacaEval 2 and Arena-hard. Furthermore, it can also surpass the baseline results of SimPO in terms of Alpaca win rate and Arena-hard win rate.

Evaluation on Skywork Reward Model. While our previous experiments used Armorm as the reward model, we also evaluate our preference data

Data (Method)	#Sample	AE LC	AE WR	AH WR
Baseline* (SimPO)	5	53.7	47.5	36.5
Baseline* (DPO)	5	48.2	47.5	35.2
Baseline† (SimPO)	400	44.6	43.9	34.8
Baseline† (DPO)	400	42.0	42.0	34.5
Ours (DPO)	400	49.1	50.2	37.3

Table 1: We compare our method with reported baseline scores from Meng et al. (2024) on Llama-3-8B-Instruct. AE denotes alpaca evaluation. AH represents arena-hard evaluation (Li et al., 2024a). LC denotes length controlled win rate, while WR denotes win rate. * means the original results from Meng et al. (2024). † means our own implementation.

construction strategy using the Skywork reward model to ensure its general applicability. We adopt Llama-3-8B-Instruct as the SFT model and record the results of AlpacaEval 2 in Figure 7. We can see that model performance improves as the number of samples increases before reaching a plateau, which confirms that our strategy is robust across different reward models.

Evaluation on Academic Benchmarks. To assess whether our preference data construction method negatively affects performance on established NLP benchmarks, we evaluate our trained model based on Llama-3-8B-Instruct on a set of widely used academic tasks, including ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), TruthfulQA (Lin et al., 2022), and GSM8K (Cobbe et al., 2021). We use the Language Model Evaluation Harness (Gao et al., 2024) for evaluation. More details of our results are presented in Appendix A.8. We observe that our policy models do not show performance drops on academic benchmarks.

6 Related Work

Reinforcement Learning from Human Feedback. RLHF is a dominant approach to align large language models with human preferences in the generation of natural language (Ouyang et al., 2022; Touvron et al., 2023). The RLHF process generally involves three stages: initial supervised fine-tuning, reward modeling (Lambert et al., 2024), and policy optimization. This approach has been extended to address various challenges such as reducing toxicity, improving safety and reasoning capabilities (Qi et al., 2024; Wu et al., 2023; Dai et al., 2024; Yu et al., 2024). However, RLHF

has issues such as training instability and complexities due to the nature of reinforcement learning and multi-stage pipeline, potentially leading to biases and verbose model outputs.

DPO (Rafailov et al., 2023) and its variants (Meng et al., 2024; Ethayarajh et al., 2024; Han et al., 2024) were developed to overcome these limitations by directly integrating preferences into the policy model using pairwise preference data. This approach simplifies the policy optimization process by eliminating the need for a surrogate reward phase. As an increasing number of powerful reward models become publicly available (Jiang et al., 2023; Wang et al., 2024b,a; Liu et al., 2024a), a popular practice (Dong et al., 2023; Liu et al., 2024b; Meng et al., 2024; Ye and Ng, 2024) has gained focus for models to enhance capability training through DPO, which employ reward models to select self-generated samples.

Synthetic Data for LLMs. Human-curated data has consistently been a highly effective resource to enhance model performance in natural language processing (Bai et al., 2022; Köpf et al., 2023). Although human-curated data are typically of high quality, obtaining sufficient amounts can be prohibitively expensive. To address this challenge, the use of synthetic data has gained attention as a cost-effective alternative to human data (West et al., 2022; Hsieh et al., 2023; Wang et al., 2023; Dong et al., 2024; Li et al., 2024b). This approach often involves the use of advanced LLMs to generate high-quality synthetic datasets (Tajwar et al., 2024; Dong et al., 2023; Agarwal et al., 2024; Chen et al., 2025). In particular, on-policy data has emerged as a highly effective and efficient approach, drawing considerable interest in recent work. Previous work mainly aims to generate more high-quality data to improve the capabilities of LLMs. In this paper, we focus on how to use self-generated data to construct optimal preference pairs for DPO.

Scaling Inference. Recently, many studies have explored how scaling inference (sample budgets) impacts the performance of large language models (Wu et al., 2024; Brown et al., 2024; Snell et al., 2024; Zhang et al., 2024). They have shown that increasing the number of samples improves the solve rate with a simple best-of-N strategy (Amini et al., 2024) and a powerful reward. Specifically, improvements in the performance of these models in mathematical problems have been achieved by repeatedly sampling potential solutions by ma-

nipulating temperatures. Research has also been conducted to investigate the scaling effects related to various inference algorithms, reward functions, and model sizes (Chi et al., 2024; Wu et al., 2024).

7 Conclusion

We first point out the failure case of a conventional preference data construction strategy for DPO. To address this, we then classify the samples into seven categories based on the underlying normal distribution of rewards per prompt and construct 21 preference datasets as a whole to systematically evaluate their impact on model performance. Our findings demonstrate that selecting the rejected response at reward position $\mu - 2\sigma$ is critical for effective optimization of DPO. We finally propose a simple preference data construction strategy that can steadily improve the performance of trained models through DPO as the sample size increases.

Limitations

Our method assumes the existence of a strong reward model, which may not always be readily available or easily trainable for all tasks. The quality of the reward model directly impacts the effectiveness of our approach, and inaccuracies in the reward model can lead to suboptimal performance or biased outcomes. Another limitation is the computational cost associated with generating a large number of responses for each prompt to construct the preference dataset. For tasks with a large input space or high complexity, this can become resource-intensive and time-consuming. We mainly focus on DPO in this paper and will explore other preference optimization methods in the future.

Acknowledgements

This research/project is supported by the National Research Foundation, Singapore under its National Large Language Models Funding Initiative (AISG Award No: AISG-NMLP-2024-004). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. This research/project is supported by the Ministry of Education, Singapore, under its SUTD-SMU Joint Grant Call, if applicable. We also thank Professor Lu Wei for providing GPUs so that we can carry out preliminary experiments.

References

- Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. 2024. [On-policy distillation of language models: Learning from self-generated mistakes](#). In *The Twelfth International Conference on Learning Representations*.
- Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. 2024. [Back to basics: Revisiting REINFORCE-style optimization for learning from human feedback in LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12248–12267, Bangkok, Thailand. Association for Computational Linguistics.
- Afra Amini, Tim Vieira, and Ryan Cotterell. 2024. [Variational best-of-n alignment](#).
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#).
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. 2024. [Large language monkeys: Scaling inference compute with repeated sampling](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#).
- Guanzheng Chen, Xin Li, Michael Shieh, and Lidong Bing. 2025. [LongPO: Long context self-evolution of](#)

- large language models through short-to-long preference optimization. In *The Thirteenth International Conference on Learning Representations*.
- Yizhou Chi, Kevin Yang, and Dan Klein. 2024. [Thoughtsculpt: Reasoning with intermediate revision and search](#).
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#).
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#).
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Bingxiang He, Wei Zhu, Yuan Ni, Guotong Xie, Ruobing Xie, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2024. [ULTRAFEEDBACK: Boosting language models with scaled AI feedback](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 9722–9744. PMLR.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. [Safe RLHF: Safe reinforcement learning from human feedback](#). In *The Twelfth International Conference on Learning Representations*.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. [Enhancing chat language models by scaling high-quality instructional conversations](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, KaShun SHUM, and Tong Zhang. 2023. [RAFT: Reward ranked finetuning for generative foundation model alignment](#). *Transactions on Machine Learning Research*.
- Qingxiu Dong, Li Dong, Xingxing Zhang, Zhifang Sui, and Furu Wei. 2024. [Self-boosting large language models with synthetic preference data](#).
- Yann Dubois, Percy Liang, and Tatsunori Hashimoto. 2024. [Length-controlled alpacaeval: A simple debiasing of automatic evaluators](#). In *First Conference on Language Modeling*.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. [Model alignment as prospect theoretic optimization](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 12634–12651. PMLR.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. [Scaling laws for reward model overoptimization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10835–10866. PMLR.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2024. [A framework for few-shot language model evaluation](#).
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. [A general theoretical paradigm to understand learning from human preferences](#). In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4447–4455. PMLR.
- Shangmin Guo, Biao Zhang, Tianlin Liu, Tianqi Liu, Misha Khalman, Felipe Llinares, Alexandre Rame, Thomas Mesnard, Yao Zhao, Bilal Piot, Johan Ferret, and Mathieu Blondel. 2024. [Direct language model alignment from online ai feedback](#).
- Jiaqi Han, Mingjian Jiang, Yuxuan Song, Jure Leskovec, Stefano Ermon, and Minkai Xu. 2024. [f-po: Generalizing preference optimization with f-divergence minimization](#).
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. [Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8003–8017, Toronto, Canada. Association for Computational Linguistics.
- Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A. Smith, Yejin Choi, and Hannaneh Hajishirzi. 2024. [Unpacking DPO and PPO: Disentangling best practices for learning from preference feedback](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. [LLM-blender: Ensembling large language models with pairwise ranking and generative fusion](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14165–14178, Toronto, Canada. Association for Computational Linguistics.
- Joongwon Kim, Anirudh Goyal, Aston Zhang, Bo Xiong, Rui Hou, Melanie Kambadur, Dhruv Mahajan, Hannaneh Hajishirzi, and Liang Tan. 2025.

- A systematic examination of preference learning through the lens of instruction-following. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Albuquerque, New Mexico. Association for Computational Linguistics.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Minh Nguyen, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Alexandrovich Glushkov, Arnav Varma Dantururi, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Julian Mattick. 2023. [Openassistant conversations - democratizing large language model alignment](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Rewardbench: Evaluating reward models for language modeling](#).
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024a. [From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline](#).
- Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason E Weston, and Mike Lewis. 2024b. [Self-alignment with instruction back-translation](#). In *The Twelfth International Conference on Learning Representations*.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024a. [Skywork-reward: Bag of tricks for reward modeling in llms](#).
- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and Jialu Liu. 2024b. [Statistical rejection sampling improves preference optimization](#). In *The Twelfth International Conference on Learning Representations*.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. [Simpo: Simple preference optimization with a reference-free reward](#). In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal, and Peter Henderson. 2024. [Safety alignment should be made more than just a few tokens deep](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#).
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. [Scaling llm test-time compute optimally can be more effective than scaling model parameters](#).
- Yuda Song, Gokul Swamy, Aarti Singh, Drew Bagnell, and Wen Sun. 2024. [The importance of online data: Understanding preference fine-tuning via coverage](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, and Aviral Kumar. 2024. [Preference fine-tuning of LLMs should leverage suboptimal, on-policy data](#). In *Forty-first International Conference on Machine Learning*.
- Yunhao Tang, Taco Cohen, David W. Zhang, Michal Valko, and Rémi Munos. 2025. [RL-finetuning llms from on- and off-policy data with a single algorithm](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhoale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller,

- Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and finetuned chat models](#).
- Haoxiang Wang, Yong Lin, Wei Xiong, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao, and Tong Zhang. 2024a. [Arithmetic control of llms for diverse user preferences: Directional preference alignment with multi-objective rewards](#). In *ACL*.
- Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. 2024b. [Interpretable preferences via multi-objective reward modeling and mixture-of-experts](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10582–10592, Miami, Florida, USA. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. [Taxonomy of risks posed by language models](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’22*, page 214–229, New York, NY, USA. Association for Computing Machinery.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. [Symbolic knowledge distillation: from general language models to commonsense models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4602–4625, Seattle, United States. Association for Computational Linguistics.
- Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. 2024. [Scaling inference computation: Compute-optimal inference for problem-solving with language models](#). In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS’24*.
- Zequ Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A. Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. [Fine-grained human feedback gives better rewards for language model training](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Shusheng Xu, Wei Fu, Jiaxuan Gao, Wenjie Ye, Weilin Liu, Zhiyu Mei, Guangju Wang, Chao Yu, and Yi Wu. 2024. [Is DPO superior to PPO for LLM alignment? a comprehensive study](#). In *Forty-first International Conference on Machine Learning*.
- Hai Ye and Hwee Tou Ng. 2024. [Preference-guided reflective sampling for aligning language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 21646–21668, Miami, Florida, USA. Association for Computational Linguistics.
- Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2024. [Metamath: Bootstrap your own mathematical questions for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [HellaSwag: Can a machine really finish your sentence?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Kexun Zhang, Shang Zhou, Danqing Wang, William Yang Wang, and Lei Li. 2024. [Scaling llm inference with optimized sample compute allocation](#).

A Appendix

A.1 Hyperparameters for Training

Here, we report the details of our training. The learning rate and batch size are listed in Table 2. Our experiments are running on 8 H100 and 8 A100 GPUs.

A.2 Baseline Results

For comparison, we report the results of SFT models and models trained with the conventional preference data construction strategy in Table 5. In base

Hyperparameters	Llama-Base		Mistral-Base		Llama-Inst	Mistral-Inst
	SFT	DPO	SFT	DPO	DPO	DPO
Batch size	128	128	128	128	128	128
Epochs	1	1	1	1	1	1
Learning rate	2e-5	5e-7	2e-5	3e-7	3e-7	3e-7
Beta	-	0.01	-	0.01	0.01	0.1
Warm-up ratio	0.1	0.01	0.1	0.01	0.01	0.01

Table 2: Hyperparameters for SFT and DPO training.

setting, the SFT models are trained on UltraChat, while it is the original instruct model in the instruct setting. For the conventional preference data construction strategy, we sample five responses for each prompt.

A.3 Win Rate Results

We report the win rate results of our policy models, which correspond to Section 4.2 in Figure 10. We find that our findings in Section 4.2 also hold in most cases.

A.4 Extend Reward Points

We have tried data construction based on the values $\{min, \mu \pm 2\sigma, \mu \pm \sigma, \mu, max\}$. Here, we extend this set to further include the values $\{\mu \pm 4\sigma, \mu \pm 3\sigma\}$. We experiment on Llama-3-8B-Instruct. We find that sample points on $\{\mu + 4\sigma, \mu + 3\sigma\}$ have no significant differences from $\{max\}$ when used to train models through DPO. In addition, sample points on $\{\mu - 4\sigma, \mu - 3\sigma\}$ have no significant differences from $\{min\}$. We report some of our results in Table 3.

A.5 400 Samples per Prompt

In this part, we experiment with 400 samples per prompt to explore whether our results of 200 samples can still hold. As we have emphasized, we are given sufficient sample budgets. The SFT model is Llama-3-8B-Instruct. We report the results in Table 4. To save some computation and evaluation costs, we only train and evaluate 11 models. Our conclusions of constructed datasets are consistent for both scenarios, 200 samples and 400 samples per prompt.

A.6 Overfitting

We find that the data construction strategy, which selects the generation of maximal reward as the chosen response and selects the one with minimal reward as the rejected response among n responses, may lead to overfitting of policy models when n reaches a point.

Chosen	Rejected	Llama-Inst	
		LC	WR
max	min	43.11	43.97
max	$\mu - 4\sigma$	43.61	44.30
max	$\mu - 3\sigma$	44.10	43.78
max	$\mu - 2\sigma$	48.12	49.66
$\mu + 4\sigma$	$\mu - 2\sigma$	48.11	49.46
$\mu + 3\sigma$	$\mu - 2\sigma$	48.26	49.20

Table 3: Evaluation results of extended data points.

Chosen	Rejected	Llama-Inst	
		LC	WR
$none^*$	$none^*$	16.58	11.29
max_of_5	min_of_5	45.23	47.86
max	min	42.01	42.04
max	$\mu - 2\sigma$	49.09	50.23
max	$\mu - \sigma$	48.27	49.76
max	μ	48.21	48.76
max	$\mu + \sigma$	42.32	38.35
max	$\mu + 2\sigma$	28.42	27.81
$\mu + 2\sigma$	min	39.35	40.08
$\mu + 2\sigma$	$\mu - 2\sigma$	50.71	51.37
$\mu + 2\sigma$	$\mu - \sigma$	48.41	49.56
$\mu + 2\sigma$	μ	46.79	45.39
$\mu + 2\sigma$	$\mu + \sigma$	35.15	32.57

Table 4: 400 Samples per prompt data construction results on Alpaca evaluation.




Figure 8: Loss of Llama-3-8B-Instruct when training with the data by selecting the response of maximal reward as the chosen response and selecting the response of minimal reward as the rejected response among n responses.

Empirical Results. As shown in Figure 8, although the training loss can reach a lower bound when we increase n from 5 to 400, the length-controlled win rate does not improve accordingly, 45.23, 46.88, 42.98, 42.01 for 5, 20, 60, 400 samples, respectively.

Theoretical Support. The reward model produces scores that are approximately normally dis-

Chosen	Rejected	Llama-Base		Mistral-Base		Llama-Inst		Mistral-Inst	
		LC	WR	LC	WR	LC	WR	LC	WR
<i>none*</i>	<i>none*</i>	7.50	4.34	6.90	4.32	23.75	24.23	17.63	14.68
<i>max_of_2</i>	<i>min_of_2</i>	10.69	6.83	15.12	8.36	31.86	32.61	21.97	17.55
<i>max_of_5</i>	<i>min_of_5</i>	16.58	11.29	16.03	10.08	45.23	47.86	30.96	25.32

Table 5: We report the results of baselines, SFT models, and DPO models trained with the conventional data construction strategy. The method with * denotes the results of SFT model in the base setting and instruct model in the instruct setting.




Figure 9: The average reward of 3 top ranking responses.

tributed:

$$r(y) \sim \mathcal{N}(\mu, \sigma^2).$$

When evaluating a large number n , the extreme values—namely, the maximum $\max_{j=1}^n r(y_j)$ and the minimum $\min_{j=1}^n r(y_j)$ —tend to be statistical outliers rather than typical examples from the distribution. According to extreme value theory, for large n , the expected maximum and minimum rewards are approximately given by

$$\mathbb{E} \left[\max_{j=1}^n r(y_j) \right] \approx \mu + \sigma \sqrt{2 \log n},$$

$$\mathbb{E} \left[\min_{j=1}^n r(y_j) \right] \approx \mu - \sigma \sqrt{2 \log n}.$$

In the DPO framework, the log-likelihood loss is sensitive to the differences in reward scores. Specifically, consider the term

$$\log \sigma(r(y_w) - r(y_l)),$$

where $r(y_w)$ and $r(y_l)$ represent the rewards for the winning and losing samples, respectively. Using the approximations above, the difference between

Tasks	ARC_C(5)	HS(10)	TQA(0)	GSM(5)
Llama-Inst	57.25	58.71	35.99	75.06
Ours				
5	61.43	59.19	40.64	76.88
20	61.52	58.90	39.78	75.15
60	61.52	58.79	39.66	76.19
100	61.26	59.02	39.41	76.19
200	61.43	58.84	39.66	77.26

Table 6: Performance of trained models with 5, 20, 60, 100, 200 samples per prompt on academic benchmarks. We observe no performance drops. HS denotes HellaSwag, while TQA denotes TruthfulQA.

an extreme high and an extreme low reward scales roughly as

$$r(y_w) - r(y_l) \approx 2\sigma \sqrt{2 \log n},$$

so that

$$\log \sigma(r(y_w) - r(y_l)) \approx \log \sigma(2\sigma \sqrt{2 \log n}).$$

As n increases, this term becomes saturated, which diminishes the effective learning signal. In other words, the presence of extreme outliers can lead the optimization process to overfit to these statistical artifacts rather than capturing improvements that generalize well.

A.7 Reward with More Samples

To support our view that the quality of top-ranking responses is getting better, we record the average reward score of the three highest-ranking responses for 1,000 prompts in Ultrafeedback. As shown in Figure 9, the reward scores of both Armorm and Skywork reward models confirm our hypothesis that the quality of top-ranking responses improves as the number of samples increases.

A.8 Evaluation on Academic Benchmarks

In Table 6, we evaluate our trained model based on Llama-3-8B-Instruct on a set of widely used academic tasks, including ARC (Clark et al., 2018),

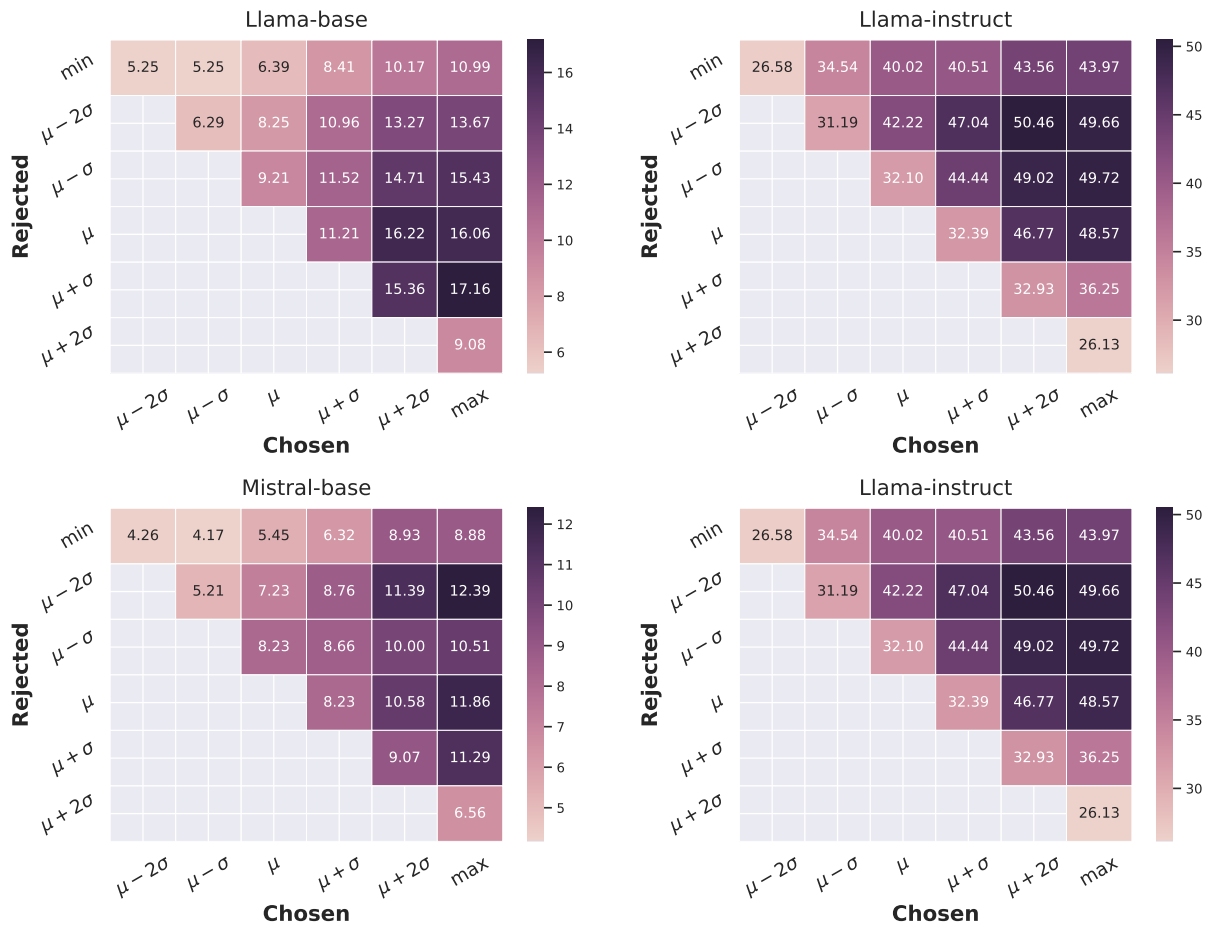


Figure 10: Win rate results of Alpaca evaluation.

HellaSwag (Zellers et al., 2019), TruthfulQA (Lin et al., 2022), and GSM8K (Cobbe et al., 2021).