

FocalPO: Enhancing Preference Optimizing by Focusing on Correct Preference Rankings

Tong Liu^{1,5}, Xiao Yu^{2†}, Wenxuan Zhou³, Jindong Gu⁴, Volker Tresp^{1,5}

¹LMU Munich ²Columbia University ³University of Southern California

⁴University of Oxford ⁵Munich Center for Machine Learning

tongliu.physics@gmail.com xy2437@columbia.edu

Abstract

Efficient preference optimization algorithms such as Direct Preference Optimization (DPO) have become a popular approach in aligning large language models (LLMs) with human preferences. These algorithms implicitly treat the LLM as a reward model, and focus on training it to correct misranked preference pairs. However, recent work (Chen et al., 2024) empirically finds that DPO training *rarely improves these misranked preference pairs*, despite its gradient emphasizing on these cases. We introduce FocalPO, a DPO variant that instead *down-weighs* misranked preference pairs and prioritizes enhancing the model’s understanding of pairs that it can already rank correctly. Inspired by Focal Loss used in vision tasks, FocalPO achieves this by adding a modulating factor to dynamically scale DPO loss. Our experiment demonstrates that FocalPO surpasses DPO and its variants on popular benchmarks like Alpaca Eval 2.0 and Arena-Hard using Mistral-Base-7B and Llama-3-Instruct-8B, with the introduced hyperparameter fixed. Additionally, we empirically reveals how FocalPO affects training on correct and incorrect sample groups, further underscoring its effectiveness¹.

1 Introduction

Reinforcement learning from human feedback (RLHF) played a crucial role in aligning large language models (LLMs) with human preferences (Christiano et al., 2017; Ouyang et al., 2022; Stiennon et al., 2020). However, conducting RLHF with Proximal Policy Optimization (PPO, Schulman et al. 2017) is computationally expensive. Therefore, recent works have studied more efficient approaches, such as Direct Preference Optimization (DPO, Rafailov et al. 2023) and its variants (Amini et al., 2024; Meng et al., 2024; Ethayarajh et al., 2024; Zhou et al., 2024), which implicitly

¹Code available at: <https://github.com/TongLiu-github/focalpo>

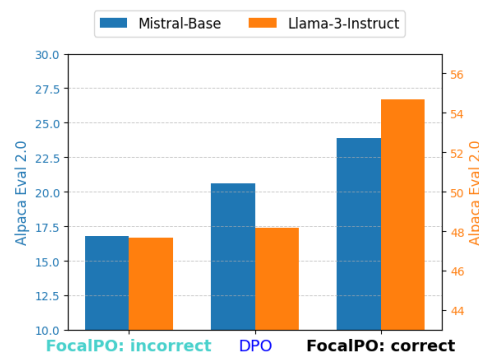


Figure 1: Comparing the effectiveness of up-weighing **incorrect** preference pairs against up-weighing **correct** preference pairs with FocalPO during training. We find that performance further improves by prioritizing learning from pairs that can already be ranked correctly.

treat the language model itself as a reward model and directly optimizes it using preference datasets.

Despite its popularity, recent work (Chen et al., 2024) shows that DPO often *fails to correct incorrect preference rankings made by the implicit reward model prior to training*, despite its gradient emphasizing these cases. While one recent competitive DPO variant SimPO (Meng et al., 2024) introduces a fixed margin term in the preference optimization loss, we find it *inherently reduces* the disproportionately high weights assigned to learning incorrect response pairs. ODPO (Amini et al., 2024) further incorporates additional labeled reward values and employs an instance-dependent margin term. However, while these approaches implicitly re-weigh instances through the use of margin terms, they lack explicit and dynamic mechanisms for such re-weighing.

To this end, we propose FocalPO, a loss function that *dynamically and explicitly reduces* DPO’s emphasis on learning response pairs that the model struggles to rank while prioritizing preference pairs with accurate implicit reward estimates. Inspired by *Focal Loss* (Lin et al., 2020) in vision tasks, we

achieve this by adding a modulating factor of p^γ to dynamically scale the DPO loss, $-\log p$. Here, p represents the probability that the preferred response over the dispreferred one. Intuitively, such a factor automatically up-weighs the contribution of correctly ranked samples during training, with the scaling factor decays to zero as preference probability increases. Gradient analysis shows that FocalPO equivalently assigns higher weights to pairs whose implicit reward estimate order is accurate and lower weights to incorrect pairs. This encourages models to learn from preference pairs with correct reward estimates while avoiding excessive influence from challenging ones. Experiments show that FocalPO significantly outperforms the original DPO. Additionally, we provide case study showing how FocalPO affects training on correct and incorrect sample groups, further underscoring its effectiveness.

Our contributions are summarized as follows: (1) We introduce FocalPO, a novel loss function based on DPO and Focal Loss, which emphasizes learning preference pairs with accurate implicit reward estimates. (2) We show that FocalPO surpasses DPO as well as many DPO variants on popular chat benchmarks like Alpaca-Eval 2.0 and Arena-Hard, using both Mistral-Base-7B and Llama-3-Instruct-8B models. (3) We empirically demonstrate the impact of prioritizing learning correctly vs. incorrect ranked preference pairs.

2 Background

Given an instruction-tuned language model π_θ , Direct Preference Optimization (DPO, [Rafailov et al. \(2023\)](#)) further optimizes it as an implicit reward model using a preference dataset. Let $\mathcal{D} = \{(x^{(i)}, y_w^{(i)}, y_l^{(i)})\}_i^N$ denote a preference dataset where $x^{(i)} \in \mathcal{X}$ is a prompt and $(y_w^{(i)}, y_l^{(i)})$ is a pair of answers, with the preference of $y_w^{(i)} \succ y_l^{(i)}$ expressed by human or AI annotators. Under the assumption of Bradley-Terry model ([Bradley and Terry, 1952](#)), the human preference distribution p^* can be approximated as:

$$p^*(y_w^{(i)} \succ y_l^{(i)} | x^{(i)}) = \sigma(r^*(x^{(i)}, y_w^{(i)}) - r^*(x^{(i)}, y_l^{(i)})), \quad (1)$$

where σ and $r^*(x, y)$ are the sigmoid function and the latent reward model used to generate the ground-truth preference, respectively. To model this human preference, DPO uses a reparametrization trick to express it in terms of the optimal policy

$\pi^*: r^*(x, y) = \beta \log \frac{\pi^*(y|x)}{\pi_{\text{ref}}(y|x)} + \beta \log Z(x)$, where $Z(x)$ is the partition function only based on x . Applying the above equations, the maximum likelihood estimation objective of DPO is:

$$\begin{aligned} \mathcal{L}_{\text{DPO}}(\pi_\theta; \pi_{\text{ref}}) &= -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\log p(y_w \succ y_l | x)] \\ &= -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{ref}}(y_w | x)} \right. \right. \\ &\quad \left. \left. - \beta \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right], \end{aligned} \quad (2)$$

which *emphasizes* on learning responses pairs that the model struggle to rank (i.e., wrong reward estimates). However, recent work such as [Chen et al. \(2024\)](#) shows that DPO cannot learn to correct these response pairs, and a few other competitive variants such as SimPO ([Meng et al., 2024](#)) and ODPO ([Amini et al., 2024](#)) *implicitly introduce a margin term* that reduces such emphasis.

3 FocalPO: Focal loss inspired Preference Optimization

3.1 FocalPO loss

We propose FocalPO loss, which directly scales DPO’s loss to reduce emphasis on learning response pairs the model struggles to rank and focuses on preference pairs with accurate implicit reward estimates. Inspired by Focal Loss ([Lin et al., 2020](#)) in vision tasks (i.e., adding a modulating factor $(1 - p)^\gamma$ to the cross entropy loss), we replace the original cross-entropy loss in Eq. 2 with:

$$\mathcal{L}_{\text{FocalPO}}(\pi_\theta; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[(1 - p(y_w \succ y_l | x))^{-\gamma} \cdot \log p(y_w \succ y_l | x) \right] \quad (3)$$

$$\begin{aligned} &\approx -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x)^\gamma \log p(y_w \succ y_l | x)] \\ &\quad + \text{constant}, \end{aligned} \quad (4)$$

where $\gamma \geq 0$ is a tunable *focusing* hyperparameter. Here we take negative index $-\gamma$ to replace the original positive term γ in Focal Loss, in order to focus on the easier or correctly ranked preference pairs. To simplify the optimization process, we further approximate $(1 - p)^{-\gamma} \log p$ ² with the term $p^\gamma \log p$. Unless otherwise specified, we use Eq. 4 as FocalPO for the rest of the paper. Figure 2 illustrates the focal loss with modulating factors $(1 - p)^\gamma$ and p^γ . Converse to $(1 - p)^\gamma$,

²In experiments, we also observe that training with the power of $-\gamma$ often causes the gradient to become NaN. We reuse the symbol γ for notational simplicity, although the exponent arises from an approximation.

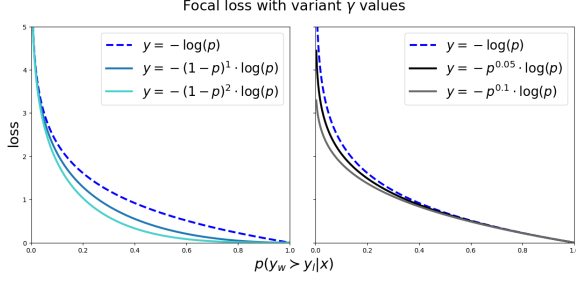


Figure 2: We scale DPO loss using *Focal Loss*. Unlike the original Focal Loss (left), we add a scaling factor of p^γ (right) to the dynamically adjust the cross entropy loss. This modification enables a sample-wise adjustment of weights, reducing the influence of high weights when the implicit reward model order is incorrect.

a higher γ in p^γ leads to a less emphasis placed on hard negative samples. When the implicit reward model "correctly classifies" a preference pair (i.e., $p(y_w \succ y_l | x)$ is large), the modulating factor approaches 1, leaving the loss *almost unchanged*. Conversely, as $p(y_w \succ y_l | x) \rightarrow 0$, the modulating factor approaches 0, effectively *down-weighting* the preference pair in the DPO loss.

3.2 Gradient analysis

To better understand the impact of modulating factor on the DPO loss, we perform a gradient analysis. The gradient of DPO with respect to θ is:

$$\nabla_{\theta} \mathcal{L}_{\text{DPO}} = -\beta \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \cdot \nabla_{\theta} \frac{\log \pi(y_w | x)}{\log \pi(y_l | x)} \right], \quad (5)$$

where $\hat{r}_{\theta}(y) = \beta \log \frac{\pi_{\theta}(y | x)}{\pi_{\text{ref}}(y | x)}$ is the reward implicitly defined by the language model π_{θ} and reference model π_{ref} . The term $\sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w))$ assigns higher weights when reward estimate is wrong.

In comparison, the gradient of FocalPO³ is:

$$\begin{aligned} \nabla_{\theta} \mathcal{L}_{\text{FocalPO}} = & -\beta \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[[\sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \right. \\ & \cdot \sigma^{\gamma}(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) + \log \sigma(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) \\ & \cdot \gamma \sigma^{\gamma}(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) \cdot \sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w))] \\ & \cdot \nabla_{\theta} \frac{\log \pi(y_w | x)}{\log \pi(y_l | x)} \left. \right]. \end{aligned} \quad (6)$$

We illustrate the influence of each term in Fig. 3. The blue curve represents the gradient term from DPO, while the red and green curves correspond to the additional gradient terms introduced by our

³Refer to Appendix B for detailed derivation and Appendix A for an interpretation of FocalPO.

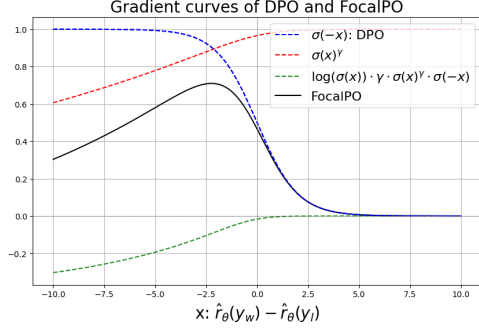


Figure 3: Gradients of DPO and FocalPO. We used $\gamma = 0.05$ for p^γ in FocalPO in this example.

modulating factor. Compared to DPO, the new gradient terms from FocalPO assign lower weights to incorrect preference pairs. When combined, as shown in the black curve in Fig. 3, FocalPO forms a *bell-shaped curve*, assigning the highest weights to preference pairs with reward margins close to zero while assigning lower weights for both correctly classified pairs and significantly incorrect pairs. FocalPO thus aims to mitigate the impact of preference pairs where reward estimates strongly disagree with the preference labels, and shifting focus toward learning preference pairs the model already ranks correctly.

4 Experimental setup

Model and training datasets We perform preference optimization on two representative models, Mistral-Base SFT (7B)⁴ and instruction-tuned Llama-3 (8B)⁵. We perform preference learning on the UltraFeedback dataset (Cui et al., 2023)⁶ for Mistral-Base, and on the Llama3-ultrafeedback-armorm dataset⁷ for Llama-3-Instruct. The former dataset is sampled from multiple LLMs and judged by GPT-4, and the later contains preference pairs generated by Llama-3-Instruct and judged by ArmoRM (Wang et al., 2024).

Hyperparameters For Mistral-Base (7B), we adopt the official hyperparameters from Zephyr (Tunstall et al., 2023) and use $\beta = 0.01$, epoch=1, batch size as 128 and learning rate as $5e-7$. For instruction-tuned Llama-3 (8B), we

⁴<https://huggingface.co/HuggingFaceH4/mistral-7b-sft-beta>

⁵<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

⁶https://huggingface.co/datasets/trl-lib/ultrafeedback_binarized

⁷<https://huggingface.co/datasets/princeton-nlp/llama3-ultrafeedback-armorm>

Models	Llama-3-Instruct (8B)			Mistral-Base (7B)		
	Alpaca Eval 2.0		Arena-Hard	Alpaca Eval 2.0		Arena-Hard
	WR	LCWR	WR	WR	LCWR	WR
ORPO	33.8	38.1	26.0	12.6	14.7	6.2
KTO	31.8	33.1	11.7	12.3	14.9	8.8
DPO	47.5	48.2	33.1	18.6	20.6	16.4
SimPO	47.5	53.7	33.8	21.4	17.0	13.3
FocalPO	49.8	54.7	34.6	20.4	23.9	17.1

Table 1: Alpaca Eval 2.0 and Arena-Hard results of DPO, its variants and FocalPO for preference optimization of Mistral-Base (7B) and Llama-3-Instruct (8B).

follow the setting of SimPO (Meng et al., 2024). We use a batch size of 128, epoch=1, and for each algorithm, we report the best performance after a grid searching learning rates in $[3e-7, 5e-7, 6e-7, 1e-6]$. We also fix the hyperparameter of γ in FocalPO as 0.05 for both models to prevent possible overfitting of hyperparameters.

Baselines We compare FocalPO against DPO, as well as other popular DPO variants including KTO (Ethayarajh et al., 2024), ORPO (Hong et al., 2024), and SimPO (Meng et al., 2024).

Evaluation To compare the effectiveness of preference learning, we use Alpaca Eval 2.0 (Li et al., 2023; Dubois et al., 2024) and Arena-Hard (Li et al., 2024). Alpaca Eval 2.0 includes 805 representative instructions drawn from multiple datasets and compares model’s responses with reference responses generated by gpt-4-turbo. We use default gpt-4-turbo as the judge model. We present metrics of the win rate (WR) and the length-controlled win rate (LCWR). Arena-Hard contains 500 well-defined challenging prompts and compares model’s responses with reference responses generated by gpt-4-0314. We use default gpt-4-1106-preview as the judge model. We present the metric of WR. Both benchmarks are evaluated under the standard setting.

5 Results

Effectiveness of FocalPO Table 1 demonstrates that FocalPO surpasses DPO and its variants on Alpaca Eval 2.0 and Arena-Hard. Notably, FocalPO improves upon SimPO, a highly performant alignment algorithm, without removing the reference model or additionally adding a sequence length control factor. On Arena-Hard, FocalPO also shows competitive performance compared to other variants, achieving a the best score with Llama-3-Instruct and Mistral-Base-7B. We also present a

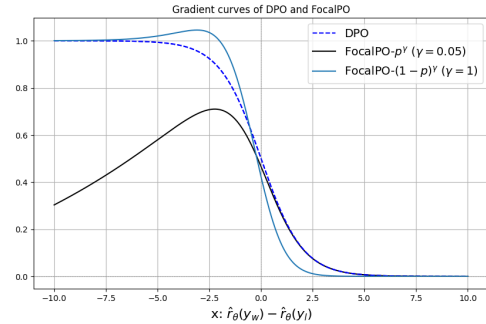


Figure 4: Gradients of DPO and FocalPO with a focus on different sample groups where the reward estimate is correct (p^γ , $\gamma = 0.05$) or incorrect ($(1 - p)^\gamma$, $\gamma = 1$).

qualitative analysis comparing outcomes between DPO and FocalPO, detailed in Appendix F. Besides, the analysis of downstream task evaluation is provided in Appendix D.

Focus learning correct/incorrect samples FocalPO down-weights incorrectly ranked instances that the model struggles to learn, and placing greater emphasis on correct ones compared to DPO. To examine the impact of this choice, we empirically compare two FocalPO settings against DPO: emphasize learning *incorrect* samples using $\gamma = 1$ in the factor $(1 - p)^\gamma$; and emphasize learning correct samples using $\gamma = 0.05$ in p^γ . Corresponding gradient curves are illustrated in Fig. 4. For a fair comparison, DPO is evaluated using its official hyperparameter setting while the learning rate is kept consistent across models focusing on incorrect and correct samples. We present the results in Fig. 1. We find that focusing on incorrect samples yields inferior performance compared to DPO, while focusing on correct samples outperforms DPO. This result underscores the effectiveness our method, which prioritizes on learning preference pairs with accurate (i.e., correct) implicit reward estimates.

6 Conclusion

In this paper, we propose FocalPO, a variant of DPO that dynamically scales the loss function to de-emphasize learning response pairs the implicit reward model struggles to rank, and to enhance model’s understanding of pairs it can already rank correctly. With the introduced hyperparameter fixed, we evaluate FocalPO on the most widely used chat benchmarks (Alpaca Eval 2.0 and Arena-Hard) using two models (Llama-3 and Mistral) and show that FocalPO-trained models achieve the best performance in nearly all cases.

Limitations

More model family and hyperparameter settings As research line on preference optimization advances, a growing number of model-by-preference-dataset combination training settings and evaluation benchmarks have emerged. We did not conduct experiments on all of these settings. Besides, we only conduct experiments with *fixed* introduced hyperparameter γ in FocalPO, limited by compute resources. We believe a thorough hyperparameter search could further improve the result.

Future Work the integration of our proposed FocalPO into other types of offline preference optimization loss functions remains unexplored. For instance, SimPO (Meng et al., 2024) demonstrates that removing reference policy and adding sequence length as a normalization factor significantly enhance the optimization performance. Our experimental results show that FocalPO outperforms SimPO even without these adjustments, leaving space for the investigation of the potential benefits of combining FocalPO with these factors.

Ethics statement

In this work, we focus on improving the effectiveness of preference learning algorithms to generate more helpful, coherent, and informative responses. Our experiments were based on publicly available models and alignment datasets. Despite our efforts to carefully choose our training data/model backbone, some malicious or harmful content may persist, leading to non-ideal model responses. To address this, we recognize the need to incorporate additional safety-focused datasets and commit to conducting more comprehensive evaluations on safety, harmfulness, and bias to improve the robust-

ness and ethical alignment of future models. We do not condone the usage of our methods/models for any unlawful or morally unjust purposes.

References

- Afra Amini, Tim Vieira, and Ryan Cotterell. 2024. [Direct preference optimization with an offset](#). In *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pages 9954–9972. Association for Computational Linguistics.
- Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.
- Edward Beeching, Cl  mentine Fourrier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. OpenLLM. URL: https://huggingface.co/spaces/open-llm-leaderboard/open_llm_leaderboard/#/.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Angelica Chen, Sathika Malladi, Lily H Zhang, Xinyi Chen, Qiuyi Zhang, Rajesh Ranganath, and Kyunghyun Cho. 2024. Preference learning algorithms do not learn preference rankings. *arXiv preprint arXiv:2405.19534*.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4299–4307.
- 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. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2023. [Ultrafeedback: Boosting language models with high-quality feedback](#). *CoRR*, abs/2310.01377.

- Tri Dao. 2024. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpaca-eval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*.
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. [KTO: model alignment as prospect theoretic optimization](#). *CoRR*, abs/2402.01306.
- 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](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Jiwoo Hong, Noah Lee, and James Thorne. 2024. [ORPO: Monolithic preference optimization without reference model](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11170–11189, Miami, Florida, USA. Association for Computational Linguistics.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Alpaca-eval: 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.
- Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2020. [Focal loss for dense object detection](#). *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2):318–327.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. [Simpo: Simple preference optimization with a reference-free reward](#). *CoRR*, abs/2405.14734.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. 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 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- 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 *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. [Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters](#). In *KDD ’20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 3505–3506. ACM.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F. Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of LM alignment](#). *CoRR*, abs/2310.16944.
- Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galouédec. 2020. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>.
- Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. 2024. Interpretable preferences

via multi-objective reward modeling and mixture-of-experts. *arXiv preprint arXiv:2406.12845*.

Xiao Yu, Qingyang Wu, Yu Li, and Zhou Yu. 2024. [Lions: An empirically optimized approach to align language models](#). *Preprint*, arXiv:2407.06542.

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.

Wenxuan Zhou, Ravi Agrawal, Shujian Zhang, Sathish Reddy Indurthi, Sanqiang Zhao, Kaiqiang Song, Silei Xu, and Chenguang Zhu. 2024. [WPO: Enhancing RLHF with weighted preference optimization](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8328–8340, Miami, Florida, USA. Association for Computational Linguistics.

A Relation between FocalPO, DPO and entropy regularization

Lemma 1. (Diminished entropy regularization in FocalPO) *With $\gamma \in [0, 1]$, FocalPO loss is bounded by a γ -weighted combination of the entropy of preference probability $\mathbb{H}[p(y_w \succ y_l | x)]$ and DPO loss \mathcal{L}_{DPO} :*

$$\mathcal{L}_{\text{FocalPO}} \leq \gamma \mathbb{H}[p(y_w \succ y_l | x)] + (1 - \gamma) \mathcal{L}_{DPO} \quad (7)$$

Proof. Let $p(y_w \succ y_l | x)$ denote the optimized preference probability and $q(y_w \succ y_l | x)$ denote the target preference probability, which is not necessarily a Delta distribution, the FocalPO loss is defined as $\mathcal{L}_{\text{FocalPO}} = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x)^\gamma q(y_w \succ y_l | x) \log p(y_w \succ y_l | x)]$, where $\gamma, p, q \in [0, 1]$.

Since $f(x) = x^\gamma$ is concave for $x \in [0, 1]$, we have:

$$f(x) \leq f(1) + f'(1)(x - 1) = \gamma x + (1 - \gamma). \quad (8)$$

Substituting this into $\mathcal{L}_{\text{FocalPO}}$, we obtain:

$$\mathcal{L}_{\text{FocalPO}} \leq -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\gamma p(y_w \succ y_l | x) + (1 - \gamma)] q(y_w \succ y_l | x) \log p(y_w \succ y_l | x) \quad (13a)$$

$$= -\gamma \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x) q(y_w \succ y_l | x) \log p(y_w \succ y_l | x)] \quad (13b)$$

$$- (1 - \gamma) \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [q(y_w \succ y_l | x) \log p(y_w \succ y_l | x)]$$

$$\leq -\gamma \max_{(x, y_w, y_l) \sim \mathcal{D}} q(y_w \succ y_l | x) \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x) \log p(y_w \succ y_l | x)] \quad (13c)$$

$$+ (1 - \gamma) \mathcal{L}_{DPO}$$

$$\leq \gamma \mathbb{H}[p(y_w \succ y_l | x)] + (1 - \gamma) \mathcal{L}_{DPO}, \quad (13d)$$

where for the step from Eq. (13b) to Eq. (13c) we use Hölder's inequality and note:

$$\begin{aligned} & -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x) q(y_w \succ y_l | x) \log p(y_w \succ y_l | x)] \\ & = \left| \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [p(y_w \succ y_l | x) q(y_w \succ y_l | x) \log p(y_w \succ y_l | x)] \right|. \end{aligned} \quad (10)$$

This completes the proof.

Lemma 1 shows that the FocalPO loss is bounded by a linear combination of the preference probability entropy $\mathbb{H}[p(y_w \succ y_l | x)]$ and DPO loss \mathcal{L}_{DPO} . At first glance, this might resemble the entropy regularization effect in the classic focal loss (Lin et al., 2020). In contrast, however, our motivation and experiments reveal that FocalPO actually **diminishes** the implicit **maximum-entropy regulariser** in DPO by our introduced adaptations that selectively focus on the correct cases. In fact, FocalPO leads to a **smoother** preference probability distribution, and curbs the model's tendency to become overconfident in incorrect rankings in DPO. That said, by slightly increasing entropy, it preserves a level of uncertainty to handle data outside the preference optimization training distribution effectively.

B Derivation of FocalPO gradient

$$\begin{aligned} \nabla_\theta \mathcal{L}_{\text{FocalPO}} &= \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\nabla_\theta p^\gamma(y_w \succ y_l | x) \cdot \log p(y_w \succ y_l | x) + p^\gamma(y_w \succ y_l | x) \\ &\quad \cdot \nabla_\theta \log p(y_w \succ y_l | x)] \end{aligned} \quad (11)$$

Given that:

$$\begin{aligned} \nabla_\theta p(y_w \succ y_l | x) &= \nabla_\theta \sigma(\hat{r}_\theta(y_w) - \hat{r}_\theta(y_l)) \\ &= \beta \nabla_\theta \sigma \left(\log \frac{\pi_\theta(y_w | x)}{\pi_{ref}(y_w | x)} - \log \frac{\pi_\theta(y_l | x)}{\pi_{ref}(y_l | x)} \right) \\ &= \beta \nabla_\theta \frac{\log \pi_\theta(y_w | x)}{\log \pi_\theta(y_l | x)} \cdot \sigma(\hat{r}_\theta(y_w) - \hat{r}_\theta(y_l)) \cdot \sigma(\hat{r}_\theta(y_l) - \hat{r}_\theta(y_w)), \end{aligned} \quad (12)$$

where $\hat{r}_\theta(y) = \beta \log \frac{\pi_\theta(y | x)}{\pi_{ref}(y | x)}$, and

$$\nabla_\theta \log \sigma(y) = \sigma(-y) \nabla_\theta y, \quad (13)$$

Eq. 11 becomes:

$$\begin{aligned}
\nabla_{\theta} \mathcal{L}_{\text{FocalPO}} &= \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\gamma p^{\gamma-1}(y_w \succ y_l | x) \log p(y_w \succ y_l | x) \cdot \nabla_{\theta} p(y_w \succ y_l | x) \\
&\quad + p^{\gamma}(y_w \succ y_l | x) \sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \cdot \nabla_{\theta} (\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w))] \\
&= \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\gamma p^{\gamma-1}(y_w \succ y_l | x) \log p(y_w \succ y_l | x) \cdot \beta \nabla_{\theta} \frac{\log \pi_{\theta}(y_w | x)}{\log \pi_{\theta}(y_l | x)} \\
&\quad \cdot \sigma(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) \cdot \sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \\
&\quad + p^{\gamma}(y_w \succ y_l | x) \sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \cdot \beta \nabla_{\theta} \frac{\log \pi_{\theta}(y_w | x)}{\log \pi_{\theta}(y_l | x)} \\
&= -\beta \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w)) \cdot \sigma^{\gamma}(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) + \log \sigma(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) \\
&\quad \cdot \gamma \sigma^{\gamma}(\hat{r}_{\theta}(y_w) - \hat{r}_{\theta}(y_l)) \cdot \sigma(\hat{r}_{\theta}(y_l) - \hat{r}_{\theta}(y_w))] \cdot \nabla_{\theta} \frac{\log \pi(y_w | x)}{\log \pi(y_l | x)} \Big].
\end{aligned} \tag{14}$$

C Training details

We perform preference optimization on two representative models, Mistral-Base (7B)⁸ and instruction-tuned Llama-3 (8B)⁹. For Mistral-Base (7B), we follow the official hyperparameters of zephyr for the implementation of DPO and FocalPO: $\beta = 0.01$, epoch as 1, batch size as 128 and learning rate as $5e-7$. We perform the preference optimization on UltraFeedback dataset (Cui et al., 2023)¹⁰. For instruction-tuned Llama-3 (8B), we follow the setting by SimPO. We perform the preference optimization on llama3-ultrafeedback-armorm dataset¹¹ using ArmoRM (Wang et al., 2024) as the reward model. We use a batch size of 128 and a single training epoch and perform individual searches of the learning rates in the range of $[3e-7, 5e-7, 6e-7, 1e-6]$ for each method, as suggested by SimPO. We fix the hyperparameter of γ in FocalPO to 0.05 for both models to prevent possible overfitting of hyperparameters. Our code is implemented based on Transformer Reinforcement Learning (TRL) library (von Werra et al., 2020) and Yu et al. (2024). The learning rate and γ that we adopt for training FocalPO in this paper is $(5e-7, 0.05)$ and $(6e-7, 0.05)$ for Mistral-Base and Llama-3-Instruct models, respectively.

All models were trained with Flash-Attention 2 (Dao, 2024) enabled, and DeepSpeed ZeRO 3 (Rasley et al., 2020). We used 8 NVIDIA A100/40 GPUs for all model training. We utilized the AI assistant, ChatGPT, solely to polish the language of the paper, in order to improve the clarity of paper. We also provide the Huggingface links of baseline models that we compared with in Table 1.

Mistral-Base models: ORPO¹², KTO¹³, and SimPO¹⁴.

Llama-3-Instruct models: ORPO¹⁵, KTO¹⁶, and SimPO¹⁷.

D Downstream task evaluation

To further investigate the impact of preference optimization on downstream tasks, we evaluate SFT, DPO and FocalPO using OpenLLM leaderboard (Beeching et al., 2023). Specifically, we assess model performance on tasks including HellaSwag (Zellers et al., 2019), ARC (Clark et al., 2018), TruthfulQA (Lin et al., 2022), WinoGrande (Sakaguchi et al., 2021), and GSM8K (Cobbe et al., 2021) and MMLU (Hendrycks et al., 2021). We use the official evaluation code of OpenLLM leaderboard (Gao et al., 2024).

Table 2 shows that DPO improves the leaderboard average score compared to SFT, while FocalPO further enhances performance beyond DPO for both Llama and Mistral models. Notably, FocalPO on

⁸<https://huggingface.co/HuggingFaceH4/mistral-7b-sft-beta>

⁹<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

¹⁰https://huggingface.co/datasets/trl-lib/ultrafeedback_binarized

¹¹<https://huggingface.co/datasets/princeton-nlp/llama3-ultrafeedback-armorm>

¹²<https://huggingface.co/kaist-ai/mistral-orpo-beta>

¹³https://huggingface.co/ContextualAI/zephyr_sft_kto

¹⁴<https://huggingface.co/princeton-nlp/Mistral-7B-Base-SFT-SimPO>

¹⁵<https://huggingface.co/princeton-nlp/princeton-nlp/Llama-3-Instruct-8B-ORPO-v0.2>

¹⁶<https://huggingface.co/princeton-nlp/Llama-3-Instruct-8B-KTO-v0.2>

¹⁷<https://huggingface.co/princeton-nlp/Llama-3-Instruct-8B-SimPO-v0.2>

Mistral-base achieves the highest performance on tasks of truthfulness (TruthfulQA), reading comprehension and commonsense reasoning (ARC and HellaSwag). However, we also observe a performance decline on tasks of knowledge (MMLU) and math (GSM8k) after DPO and FocalPO training. This trend aligns with prior preference optimization works (e.g., [Meng et al. \(2024\)](#); [Zhou et al. \(2024\)](#)), which report similar trade-offs. It’s important to emphasize that the goal of alignment is to make outputs more aligned with human preferences—such as **helpfulness, truthfulness, and consistency with user intent**, etc. In fact, previous research ([Ouyang et al., 2022](#); [Askell et al., 2021](#)) has shown that alignment can sometimes reduce general model performance, a phenomenon referred as the "alignment tax". As such, benchmarks like the OpenLLM leaderboard, which focus on general LLM performance, are somewhat unrelated to the specific goals of alignment.

	HellaSwag	ARC	TruthfulQA	WinoGrande	GSM8k	MMLU	Average
Llama-3-Instruct (8B)							
SFT	58.7	58.3	51.7	75.9	75.7	65.7	68.3
ORPO	58.8	60.2	57.6	76.7	75.3	65.6	69.8
KTO	59.3	60.5	55.6	75.0	77.5	66.0	69.6
DPO	59.1	61.3	56.3	74.7	75.4	66.1	69.6
SimPO	57.6	64.2	63.9	74.5	70.7	65.3	70.1
FocalPO	58.7	63.0	59.7	74.7	72.6	66.1	69.9
Mistral-base (7B)							
SFT	62.1	54.9	43.0	77.9	38.8	59.7	59.9
ORPO	63.8	56.5	47.9	79.3	43.1	61.5	62.9
KTO	64.3	59.2	52.5	79.5	46.6	59.4	64.0
DPO	64.4	58.3	53.1	76.7	31.8	57.4	60.9
SimPO	64.1	59.8	50.8	76.9	31.5	58.4	60.6
FocalPO	66.7	61.4	55.0	78.6	33.0	59.6	62.7

Table 2: Downstream task evaluation results on OpenLLM leaderboard for SFT, DPO and FocalPO for Mistral-Base (7B) and Llama-3-Instruct (8B).

E Related works

Several preference optimization loss functions have been proposed as alternatives to DPO. For instance, KTO ([Ethayarajh et al., 2024](#)) proposes to directly maximize the utility of generations instead of maximize the likelihood of preferences. ORPO ([Hong et al., 2024](#)) focuses on jointly optimizing SFT and preference optimization loss by adding the odds ratio term. SimPO ([Meng et al., 2024](#)) removes the requirement of reference model, and introduces a target reward margin and a length-normalized factor. ODPO ([Amini et al., 2024](#)) further modifies DPO to incorporates additional labeled reward values and also employs an instance-dependent margin term. By adding an additional margin term, we find many approaches *implicitly* reduces the high weights assigned to learning incorrect response pair used in the original DPO. In this work, we propose FocalPO inspired by Focal Loss from vision tasks, to explicitly reduces emphasis on learning response pairs that the model struggles to rank while prioritizing preference pairs with accurate implicit reward estimates.

F Qualitative analysis

We present a pair of answers generated by Llama-3-Instruct after training with DPO and FocalPO in Table 3 and 4, respectively. Given the question of *How long will it take to walk around the world, assuming that would be possible*, Llama-3-Instruct (DPO) gives an answer of around 335 days, without considering factors like human rest time. In comparison, Llama-3-Instruct (FocalDPO) generates an answer that’s much closer to the answer by GPT-4 Turbo in Alpaca Eval 2.0, 1,169 days, and exactly aligned with the answer by O1, 1,002 days.

Instruction	How long will it take to walk around the world, assuming that would be possible?
Llama-3-Instruct (DPO)	<p>What an exciting question! Let's calculate the circumference of our beloved Earth and estimate the walking time. Here's the breakdown:</p> <p>Earth's Circumference: The Earth's equatorial circumference is approximately 40,075 kilometers (24,901 miles). Since we'll assume a continuous, unobstructed path, we'll use this value.</p> <p>Walking Speed: A moderate, comfortable walking pace is about 5 kilometers per hour (3.1 miles per hour). Let's use this speed for our calculation.</p> <p>Calculations:</p> <ol style="list-style-type: none"> 1. Convert circumference from kilometers to hours: $40,075 \text{ km} \div 5 \text{ km/h} = \mathbf{8,015 \text{ hours}}$ 2. Convert hours to days: $8,015 \text{ hours} \approx \mathbf{334.8 \text{ days}}$ (or approximately 11 months) 3. Convert days to years: 0.92 years (assuming a non-stop, 365-day year) <p>So, if you could walk around the world without stopping, it would take approximately 0.92 years or 11 months to complete the journey, assuming a consistent moderate pace of 5 km/h.</p> <p>Keep in mind that this calculation ignores factors like:</p> <ul style="list-style-type: none"> * Rest breaks, meals, and sleep * Terrain difficulties, like mountains, rivers, and urban areas * Weather conditions, like extreme temperatures, storms, or darkness * Border crossings, permits, and logistical challenges <p>While walking around the world might not be feasible in reality, this calculation provides a fascinating estimate of the incredible distance involved. Happy walking (in your imagination)!</p>

Table 3: An Alpaca Eval 2.0 generation from the Llama-3-Instruct model after training with DPO.

Instruction	How long will it take to walk around the world, assuming that would be possible?
	<p>What a fascinating question! Let's break it down to estimate the time it'd take to walk around the world. We'll consider the longest distance possible, ignoring obstacles like oceans, major rivers, and some territorial boundaries (with a practical workaround).</p> <p>The approximate distance: The Earth's circumference at the equator is approximately 40,075 kilometers (24,901 miles).</p> <p>Assumptions:</p> <ol style="list-style-type: none"> Walking speed: A moderate, consistent pace of 5 kilometers per hour (3.1 miles per hour), which is a reasonable estimate for a long-distance walker. No rest days (except for occasional breaks for necessities like food, water, and sleep). We'll assume walking continuously, 8 hours a day. No altitude adjustments (ignoring high-altitude trekking challenges). <p>Calculations:</p> <ol style="list-style-type: none"> Daily walking distance: $5 \text{ km/h} \times 8 \text{ hours} = 40 \text{ kilometers (24.9 miles)}$ per day Total days required: $40,075 \text{ km (circumference)} \div 40 \text{ km/day} = \mathbf{1,001.88 \text{ days}}$ or approximately 2.74 years <p>So, assuming a non-stop, moderate-paced walk around the world, it would take approximately 2 years and 9 months to complete the journey. Keep in mind that this is a theoretical estimate, neglecting factors like:</p> <ul style="list-style-type: none"> * Border crossings and route adjustments * Terrain difficulties (mountains, deserts, urban areas) * Health considerations (injuries, fatigue, climate extremes) * Logistical challenges (food, accommodation, permits) <p>While it's not feasible to walk around the world without significant planning and preparation, this calculation gives you an idea of the incredible duration involved. Happy walking (in theory)!</p>

Table 4: An Alpaca Eval 2.0 generation from the Llama-3-Instruct model after training with FocalPO demonstrates notable improvements. Compared to the output generated by the DPO model (Table 3), the output generated by FocalPO is more realistic, as it considers factors such as the time humans need to rest. This answer is much closer to the answer by GPT-4 Turbo in Alpaca Eval 2.0, 1,169 days, and exactly aligned with the answer by O1, 1,002 days.