

Stackelberg Learning from Human Feedback: Preference Optimization as a Sequential Game

Barna Pásztor

ETH Zürich

Thomas Kleine Buening

The Alan Turing Institute, ETH Zürich

Andreas Krause

ETH Zürich

BARNA.PASZTOR@AI.ETHZ.CH

Abstract

We propose Stackelberg Learning from Human Feedback (SLHF), a new framework for preference optimization. SLHF frames the alignment problem as a sequential-move game between two policies: a Leader, which commits to an action, and a Follower, which responds conditionally on the Leader’s action. This formulation departs from prior approaches such as Reinforcement Learning from Human Feedback (RLHF), which rely real-valued reward models, and Nash Learning from Human Feedback (NLHF), which seek to compute a Nash equilibrium. The sequential structure of SLHF naturally enables *test-time improvement*, as the Follower learns to refine the Leader’s actions. We compare the solution concepts of SLHF, RLHF and NLHF, and lay out key advantages in consistency, data sensitivity, and robustness to intransitive preferences. Our experiments demonstrate that SLHF effectively aligns large language models with diverse, potentially intransitive, human preferences, and its test-time improvement generalizes across models without further training.

1. Introduction

Reinforcement Learning from Human Feedback (RLHF) has become the dominant paradigm for aligning Large Language Models (LLMs) with human preferences [10, 24]. The standard pipeline involves two stages: first, a reward model is trained on a dataset of pairwise human comparisons, and second, a policy is optimized via reinforcement learning to maximize this reward [11]. While this approach has achieved impressive results, it hinges on a critical, and often flawed, assumption: that diverse human preferences can be faithfully represented by a single, real-valued reward function. In practice, this assumption often fails. Scalar reward models cannot capture intransitive structures, and even when preferences are transitive, widely used formulations such as the Bradley-Terry model [8] can cause the learned rewards to diverge from the actual preferences [7].

A common alternative to relying on reward models and the Bradley-Terry assumption are preference models [23], which directly predict the probability of one action being preferred over another in a given context. However, such models might exhibit preference cycles so that the notion of policy optimality becomes ambiguous. Nash Learning from Human Feedback (NLHF) proposes the Nash Equilibrium (NE) as a solution to this problem by framing preference optimization as a two-player simultaneous-move game, where the Nash equilibrium (NE) corresponds to a usually stochastic policy whose actions are preferred to any other policy’s actions on average [31]. We extend this game-theoretic framing by introducing Stackelberg Learning from Human Feedback (SLHF). In contrast to NLHF, SLHF formulates preference optimization as a *sequential-move* game

inspired by Stackelberg dynamics [45]: a Leader first commits to an action, and a Follower then responds conditional on the Leader’s action. This asymmetric structure offers a distinct advantage as the Follower learns to refine a given action, an easier objective than optimizing directly or against a non-stationary opponent as in NLHF. Consequently, the Leader’s objective becomes anticipating the outcome of this refinement process and choosing the action least exploitable by the Follower.

This sequential game structure naturally enables *test-time improvement*: the ability to refine model outputs at inference time via repeated sampling. This is particularly valuable when a model, trained on preferences aggregated across diverse annotators, must ultimately align with an individual’s taste, especially since such aggregation can induce intransitive preference cycles (Section 3). SLHF realizes this refinement through its two components: the Leader policy produces an initial response, and the Follower policy generates refined responses conditional on the previous output. Unlike sampling from a static distribution, this produces a sequence of outputs that can efficiently explore the preference space. Crucially, this allows for performance gains through test-time computation alone, without any need for additional training or external feedback. Our main contributions are:

- We introduce Stackelberg Learning from Human Feedback (SLHF), a preference optimization framework that models alignment as a two-player sequential game. We formalize this game over a learned pairwise preference model and show that SLHF admits a unique Stackelberg equilibrium under standard regularity assumptions (Section 3).
- We propose STACKELBERGGDA, an algorithm that approximates the Stackelberg equilibrium via two-timescale gradient descent ascent. Our algorithm benefits from online RL optimization without the need of an explicit reward model or expensive inference with a mixture policy (Section 4).
- Our experimental results show that the Follower, conditioned on the Leader’s output, consistently outperforms both RLHF and NLHF baselines, whereas the Leader performs similarly to the approximated Nash policy. Furthermore, we show that the Follower generalizes across models, improving outputs from independently trained policies without additional fine-tuning (Section 5).

We defer a discussion of related work to Appendix B.

2. Problem Statement

We consider a contextual bandit setting with a finite set of contexts \mathcal{X} and actions \mathcal{Y} . The contexts x are sampled from a fixed and known distribution $\rho \in \Delta_{\mathcal{X}}$, where $\Delta_{\mathcal{X}}$ denotes the probability simplex over \mathcal{X} . A policy $\pi : \mathcal{X} \rightarrow \Delta_{\mathcal{Y}}$ maps each context $x \in \mathcal{X}$ to a discrete probability distribution $\pi(\cdot | x) \in \Delta_{\mathcal{Y}}$, where $\Delta_{\mathcal{Y}}$ is the probability simplex over \mathcal{Y} . We let $\Pi := \{\pi : \mathcal{X} \rightarrow \Delta_{\mathcal{Y}}\}$ denote the set of all policies. In the case of LLM fine-tuning, \mathcal{X} corresponds the set of prompts, \mathcal{Y} to the set of responses, and π to the LLM.

Preferences are observed as pairwise comparisons between two actions $y \in \mathcal{Y}$ and $y' \in \mathcal{Y}$ given a context $x \in \mathcal{X}$ and provided by a human or AI annotator a . The annotator is drawn at random from a distribution $\nu \in \Delta_{\mathcal{A}}$, where $\mathcal{A} := \{1, \dots, A\}$ denotes the set of all annotators. The feedback of annotator a is then given by $y^w \succ_a y^l$, where y^w and y^l are the preferred and non-preferred actions, respectively. Let the *preference function* $p(y \succ y' | x)$ define the probability that y is preferred over y' given x , where the randomness is due to the choice of annotator, that is, the preference function is given by $p(y \succ y' | x) := \mathbb{E}_{a \sim \nu} [\mathbb{1}\{y \succ_a y' | x\}]$.¹ Overloading notation, we define the preference

1. This can be straightforwardly generalized to the case where each annotator is associated with a probabilistic preference function $p_a(y \succ y' | x)$ that describes the probability that annotator a prefers y' over y given x .

over two policies π and π' given context x as $p(\pi \succ \pi' | x) := \mathbb{E}_{y \sim \pi(\cdot | x), y' \sim \pi'(\cdot | x)} [p(y \succ y' | x)]$. Averaging over the context distribution ρ , the expected probability that π is preferred over π' is then given by $p(\pi \succ \pi') := \mathbb{E}_{x \sim \rho} [p(\pi \succ \pi' | x)]$.

Reinforcement Learning from Human Feedback (RLHF). RLHF as proposed by Christiano et al. [11] and adapted to language modeling by Ziegler et al. [63] splits preference optimization into two steps. First, it assumes that the preference function p follows the Bradley-Terry model [8] so that

$$p(y \succ y' | x) = \sigma(r(x, y) - r(x, y')), \quad (1)$$

where $\sigma(x) = \frac{1}{1 + \exp(-x)}$ is the sigmoid function and $r : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ is a real-valued reward function. The reward function r is unknown so that an estimator \hat{r} is used that maximizes the log-likelihood of a given dataset $\mathcal{D} = \{(x_i, y_i^w, y_i^l)\}_{i=1}^N$.

In a second step, the policy π^* is chosen to maximize the expected reward with respect to \hat{r} regularized by the Kullback-Leibler (KL) divergence against a fixed reference policy $\pi^{\text{ref}} \in \Pi$:

$$\pi^* \in \arg \max_{\pi \in \Pi} \mathbb{E}_{x \sim \rho} \left[\mathbb{E}_{y \sim \pi(\cdot | x)} [\hat{r}(x, y)] - \tau \text{KL}_x(\pi \| \pi^{\text{ref}}) \right], \quad (2)$$

where $\tau \geq 0$ and $\text{KL}_x(\pi \| \pi^{\text{ref}})$ is computed between $\pi(\cdot | x)$ and $\pi^{\text{ref}}(\cdot | x)$. Under the Bradley-Terry assumption, Equation (2) admits a unique closed-form solution [39]. However, additive score models like Bradley-Terry are provably limited in expressing cyclic or intransitive preference structures, which have been empirically observed in both strategic games [7] and human preference data [3, 10]. Consequently, the optimal policy π^* depends critically on the data distribution in the training set \mathcal{D} , especially its sampling biases [31], which we elaborate more on in Section C.

Nash Learning from Human Feedback (NLHF). NLHF provides an alternative approach to optimizing preferences by formulating the problem as a simultaneous-move game between two policies $\pi \in \Pi$ and $\pi' \in \Pi$ [31]. The optimization problem is given by:

$$\max_{\pi \in \Pi} \min_{\pi' \in \Pi} \mathbb{E}_{x \sim \rho} \left[p(\pi \succ \pi' | x) - \tau \text{KL}_x(\pi \| \pi^{\text{ref}}) + \tau \text{KL}_x(\pi' \| \pi^{\text{ref}}) \right]. \quad (3)$$

The solution to Equation (3) is a *Nash equilibrium* (π^*, π'^*) , where neither side can be improved unilaterally. The existence and uniqueness of this equilibrium follows from the concave-convex nature of the objective [31]. Unlike RLHF, which relies on a fixed dataset \mathcal{D} of labeled comparisons and is sensitive to its sampling distribution, NLHF collects feedback during training and does not assume a specific preference structure. This makes NLHF better suited to situations where online feedback is available and preferences are diverse. Notably, if there is no action that is majority-preferred, the NE must be mixed even in the absence of KL regularization [28]. This inherent stochasticity can be undesirable in applications where consistency and reliability are critical.

3. Stackelberg Learning from Human Feedback (SLHF)

We now present Stackelberg Learning from Human Feedback (SLHF), a novel perspective on the preference optimization problem. Inspired by Stackelberg games [45], we adopt a game-theoretic approach and model the optimization as a *sequential-move* game between two players: the *Leader* and the *Follower*. In our formulation, the Leader first observes the context x and chooses its action $y \sim \pi(\cdot | x)$. Then, the Follower observes both the context x and the Leader's action y and chooses

its own action $y' \sim \omega(\cdot | x, y)$. The Follower’s policy ω is chosen from the set of policies that are conditioned on both inputs $\Omega = \{\omega : \mathcal{X} \times \mathcal{Y} \rightarrow \Delta_{\mathcal{Y}}\}$. Formally, the optimization problem given two reference policies $\pi^{\text{ref}} \in \Pi$ and $\omega^{\text{ref}} \in \Omega$ is defined as:

$$\max_{\pi \in \Pi} \min_{\omega \in \Omega} \mathbb{E}_{x \sim \rho} \left[\mathbb{E}_{y \sim \pi(\cdot | x)} \left[\mathbb{E}_{y' \sim \omega(\cdot | x, y)} [p(y \succ y' | x)] + \tau^F \text{KL}_{x,y}(\omega \| \omega^{\text{ref}}) \right] - \tau^L \text{KL}_x(\pi \| \pi^{\text{ref}}) \right] \quad (4)$$

where $\tau^L, \tau^F \geq 0$ are player-specific regularization parameters, and the Follower’s regularization term, $\text{KL}_{x,y}(\omega \| \omega^{\text{ref}})$, is computed between $\omega(\cdot | x, y)$ and $\omega^{\text{ref}}(\cdot | x, y)$. Note that in the absence of regularization, i.e., $\tau^L = \tau^F = 0$, Equation (4) defines a sequential-move constant-sum game.

It is worth highlighting that Equation (4) differs from many common Stackelberg formulations studied in the algorithmic game theory literature [12], where the Leader commits to a stochastic policy π , and the Follower selects a response ω conditional on π but *without* access to the realization $y \sim \pi(\cdot | x)$. In contrast, the Follower gets to condition on y in SLHF, which provides strictly more information when π is stochastic and yields an easier stationary problem. In line with previous results that the RLHF optimization problem (2) admits a closed-form solution, we show that there exists a unique solution to the SLHF problem (4).

Proposition 1 *Let $\tau^L, \tau^F > 0$ and suppose that $\pi^{\text{ref}}(y | x) > 0$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. For any preference function $p(y \succ y' | x)$, there exists a unique solution (π^*, ω^*) to the preference optimization problem in Equation (4).*

The solution (π^*, ω^*) is called a *Stackelberg equilibrium*. It is folklore in the algorithmic game theory literature that there exists a deterministic Stackelberg equilibrium when the Leader’s realized action is observed by a best responding Follower, as there always exists a deterministic best response for the Follower and thus no point in randomizing for the Leader. In the absence of additional regularization (and only then), the analogue is also true for SLHF: the optimization problem (4) admits a deterministic solution whenever $\tau^L = \tau^F = 0$ (see Appendix A.2 for a short proof). This stands in contrast to the NE, which is generally stochastic.

Test-time Improvement. Motivated by applications such as text summarization, open-ended generation, and media creation, where users can reject outputs and ask for a new sample, we introduce the concept of *test-time improvement* for preference optimization. While previous work address this problem in verifiable domains and for reward models (see Appendix B), to the best of our knowledge, we are the first to consider it in the context of general preference models. We assume that at test-time, a single user (annotator) $a \sim \nu$ and a context x are sampled. The user then has the option to resample the action until they receives one that suits their preference, analogously to the *pass@k* metric for verifiable domains. This is a non-trivial task as the models are trained on the average preferences of the population ν while at test-time the task is to optimize for a single user. The Stackelberg solution (π^*, ω^*) naturally extends to this setting: it returns an initial action $y_1 \sim \pi^*(\cdot | x)$, and subsequent actions are sampled as $y_i \sim \omega^*(\cdot | x, y_{i-1})$ for $i = 2, 3, \dots$, thereby progressively improving the output.

We provide a thorough example to demonstrate the differences between solution concepts in Appendix C.

4. Stackelberg Gradient Descent Ascent (STACKELBERGGDA)

We now introduce STACKELBERGGDA, a two-timescale Gradient Descent-Ascent (GDA) algorithm tailored to the sequential preference optimization problem in Equation (4). STACKELBERGGDA

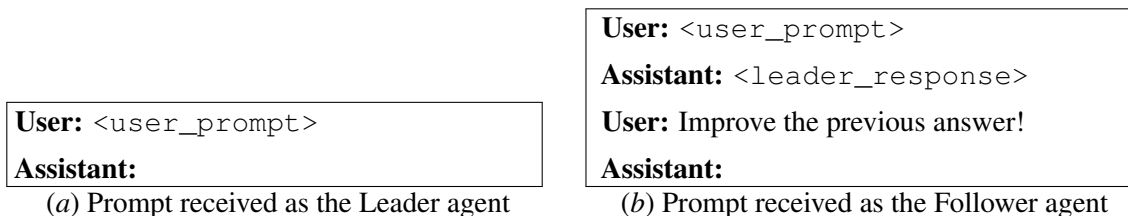


Figure 1: Prompt templates used to train a single model for both Leader (left) and Follower (right). performs simultaneous ascent on the Leader’s objective and descent on the Follower’s, but with a larger step size for the Follower. Let

$$f(\pi, \omega) := \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x), y' \sim \omega(\cdot|x,y)} [p(y \succ y' | x) + \tau^F \text{KL}_{x,y}(\omega \| \omega^{\text{ref}}) - \tau^L \text{KL}_x(\pi \| \pi^{\text{ref}})] \quad (5)$$

denote the objective of the optimization in Equation (4). STACKELBERGGDA performs simultaneously gradient ascent and descent update steps on π and ω with step size α^L and α^F , respectively, to find the max min solution to f defined in Equation (5). It is a two-timescale algorithm as we choose $\alpha^F > \alpha^L$. After each gradient step, we restore feasibility by projecting both π and ω back onto their respective probability simplexes. We denote the two-timescale coefficient as $\kappa = \frac{\alpha^F}{\alpha^L}$.

Observe that $f(\pi, \omega)$ is concave in π and convex in ω .² In this setting, standard gradient descent-ascent with equal step sizes guarantees convergence of the averaged iterates under mild assumptions, but the last iterates can oscillate even in simple bilinear games [33]. Last-iterate convergence can be recovered if the objective is strongly concave–convex (as ensured here by our KL regularizers), or by using algorithmic variants such as extragradient or optimistic GDA [13, 25]. We adopt a two-timescale variant of GDA, which under strong concave–convexity yields last-iterate convergence, and which more generally has stronger guarantees in nonconvex-concave regimes [27] and is empirically successful in actor-critic methods [36] and GAN training [20].

Scalable Implementation of STACKELBERGGDA for LLM Fine-Tuning. When working with a large context and action spaces \mathcal{X} and \mathcal{Y} , e.g., when fine-tuning LLMs, optimizing over Π and Ω becomes intractable. To address this, π and ω can be parametrized and the gradients estimated from batches. Most importantly for LLM fine-tuning, the Leader and the Follower can share the same parametrization by using the prompt template in Figure 1, thereby reducing the memory requirements. Details and the pseudocode for this practical implementation can be found in Appendix D.

5. Experiments

We experimentally evaluate and compare RLHF, NLHF, and SLHF for fine-tuning LLMs to align with human preferences. Our results show that the Leader model trained with STACKELBERGGDA achieves comparable preference scores to NASH-MD-PG [31] and higher scores than RLOO [1], which serve as representative algorithms of the NLHF and RLHF frameworks, respectively. Moreover, the Follower model significantly improves upon the Leader’s responses, and this improvement generalizes to outputs from other models without requiring additional training. Throughout this section, we refer to the context x as a *prompt* and the action y as a *response*. We provide the key results on fine-tuning in Section 5.1 and test-time improvements in Section 5.2. We defer additional

2. This follows from results in Munos et al. [31]. A formal proof is provided in Appendix A.3.

Table 1: Pairwise preference comparisons between the responses of QWEN2.5-0.5B, RLOO, NASH-MD-PG, and STACKELBERGGDA algorithms. Each cell represents the preference model’s average score for the row algorithm over the column algorithm.

	QWEN2.5-0.5B	RLOO	NASH-MD-PG	STACKELBERGGDA	
				LEADER	FOLLOWER
QWEN2.5-0.5B	0.000	0.407	0.279	0.266	0.200
RLOO	0.593	0.000	0.393	0.387	0.344
NASH-MD-PG	0.721	0.607	0.000	0.497	0.406
STACKELBERGGDA -LEADER	0.734	0.613	0.503	0.000	0.395
STACKELBERGGDA -FOLLOWER	0.800	0.656	0.594	0.605	0.000

results on iterative improvements to Appendix F.2, ablations w.r.t. κ to Appendix F.3, and model scaling to Appendix F.4.

Experimental Setup. We use the HELPSTEER2 dataset [53], which contains 11,826 human-annotated single-turn dialogues, each rated by human annotators along five attributes: helpfulness, correctness, coherence, complexity, and verbosity. We treat these attributes as distinct annotators, denoted by the set \mathcal{A} , and define ν as a uniform distribution over \mathcal{A} . For each attribute $a \in \mathcal{A}$, we estimate a reward function \hat{r}_a using the Bradley-Terry model. The overall preference function p is then defined as $p(y \succ y' | x) = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} \mathbb{1}\{\hat{r}_a(x, y) \geq \hat{r}_a(x, y')\}$. Each reward function is trained independently, initialized from the QWEN2.5-1.5B³ model with a single linear head. Although each \hat{r}_a is transitive, their aggregation into p can produce cycles, analogous to the Condorcet paradox discussed in Appendix C. Further details on the preference model specification and the resulting intransitivity are provided in Appendix F.1.

Compared Methods. We use RLOO [1] and NASH-MD-PG [31] as representative algorithms for the RLHF and NLHF approach, respectively. Since RLHF requires a scalar feedback signal, we use the mean of the attribute reward models \hat{r}_a as the reward function for RLOO. All models are fine-tuned from the QWEN2.5-0.5B⁴ model and run for 1,000 gradient steps with a batch size of $B = 32$. We sweep over learning rates $\eta \in \{1e-6, 5e-6, 1e-5\}$ and KL penalties $\tau \in \{0.001, 0.01, 0.1\}$ for all algorithms. For NASH-MD-PG, we additionally sweep over the mixture parameter $\beta \in \{0, 0.25, 0.5, 0.75, 1\}$, and for STACKELBERGGDA, the two-timescale coefficient $\kappa \in \{1, 5, 10\}$.

5.1. Round-Robin Tournament

Table 1 reports pairwise preference scores between the initial QWEN2.5-0.5B and the three fine-tuned models. Both the first responses of STACKELBERGGDA (Leader) and NASH-MD-PG achieve approximately 73% preference over QWEN2.5-0.5B and 61% over RLOO, while they tie at 50% when compared to each other. This outcome is consistent with scenarios in which multiple high-quality responses exist and the Stackelberg and Nash solutions coincide, as discussed in Appendix C.

Importantly, applying the Follower model of STACKELBERGGDA to improve its own initial responses leads to a substantial performance gain. It achieves 80% preference over QWEN2.5-0.5B,

3. <https://huggingface.co/unsloth/Qwen2.5-1.5B-Instruct>

4. <https://huggingface.co/unsloth/Qwen2.5-0.5B-Instruct>

Table 2: Test-time improvement using different models for the initial response (Leader) and the improvement (Follower). Each cell represents the preference model’s average score for the Follower’s responses over the Leader’s responses.

		Leader			
		QWEN2.5-0.5B	RLOO	NASH-MD-PG	STACKELBERGGDA
Follower	QWEN2.5-0.5B	0.549	0.443	0.363	0.362
	RLOO	0.534	0.403	0.369	0.360
	NASH-MD-PG	0.708	0.600	0.493	0.476
	STACKELBERGGDA	0.803	0.665	0.600	0.606

66% over RLOO, 60% over NASH-MD-PG, and even outperforms the responses it was conditioned on in 60.5% of comparisons. Thus, the two-turn inference yields significant gains at the cost of an additional inference step.

5.2. Test-time Improvements

We next evaluate the ability of each model to improve responses at test time, specifically by acting as a Follower that refines an initial response produced by another model. Although only STACKELBERGGDA is explicitly trained for this task (though only to best respond to itself), we apply the same refinement procedure to all models to assess whether they can generalize to this setting. To do so, we evaluate every pair of Leader and Follower models selected from {QWEN2.5-0.5B, RLOO, NASH-MD-PG, STACKELBERGGDA} as follows. For every validation prompt, we first generate a response with the chosen Leader model, and then apply the Follower prompting template from Figure 1 to generate a potentially improved response using the Follower model. Table 2 reports the resulting preference scores, showing how often the Follower output is preferred over the Leader’s generation.

STACKELBERGGDA consistently improves across all Leader models; most notably over QWEN2.5-0.5B and RLOO, but also achieving 60% gains over both NASH-MD-PG and itself. By contrast, QWEN2.5-0.5B and RLOO only improve on the initial responses from QWEN2.5-0.5B and even produce worse outputs when applied to other Leader models. NASH-MD-PG is able to improve on responses from QWEN2.5-0.5B and RLOO when used as a Follower; however, even its 70% preference score over QWEN2.5-0.5B falls short of its own 73% score reported in Table 1. These results extend previous works on verifiable domains [22, 37, 50, 62] by underscoring the importance of explicitly learning to improve given outputs and that mere instruction prompting is not sufficient to improve with respect to preferences either.

6. Conclusion

In this paper, we introduced Stackelberg Learning from Human Feedback (SLHF), a two-player sequential-move framework that directly optimizes pairwise preference signals without relying on real-valued reward models. Our proposed algorithm, STACKELBERGGDA, efficiently approximates the unique Stackelberg equilibrium and scales to complex tasks such as aligning large language models with human preferences. Empirically, STACKELBERGGDA’s Leader matches or exceeds standard baselines, while its Follower consistently refines outputs at test-time, even when paired with models it was not trained with.

References

- [1] Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting REINFORCE-style optimization for learning from human feedback in LLMs. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2024.
- [2] Afra Feyza Akyurek, Ekin Akyurek, Ashwin Kalyan, Peter Clark, Derry Tanti Wijaya, and Niket Tandon. RL4F: Generating natural language feedback with reinforcement learning for repairing model outputs. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2023.
- [3] Carlos Alós-Ferrer, Ernst Fehr, and Michele Garagnani. Identifying nontransitive preferences. Technical report, Working Paper, 2022.
- [4] Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. A General Theoretical Paradigm to Understand Learning from Human Preferences. *Proceedings of Machine Learning Research*, 238:4447–4455, 10 2023.
- [5] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- [6] Ahmad Beirami, Alekh Agarwal, Jonathan Berant, Jacob Eisenstein, Chirag Nagpal, Ananda Theertha Suresh, Google Research, and Google DeepMind. Theoretical guarantees on the best-of-n alignment policy. *arXiv preprint arXiv:2401.01879*, 2024.
- [7] Quentin Bertrand, Wojciech Marian Czarnecki, and Gauthier Gidel. On the limitations of the elo, real-world games are transitive, not additive. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, 2023.
- [8] Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [9] Daniele Calandriello, Daniel Guo, Rémi Munos, Mark Rowland, Yunhao Tang, Bernardo Avila Pires, Pierre Harvey Richemond, Charline Le Lan, Michal Valko, Tianqi Liu, Rishabh Joshi, Zeyu Zheng, and Bilal Piot. Human Alignment of Large Language Models through Online Preference Optimisation. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [10] Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Wang, Samuel Marks, Charbel-Raphaël Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen, Mehul Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J. Michaud, Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem Bıyık, Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. Open Problems and

Fundamental Limitations of Reinforcement Learning from Human Feedback. *arXiv preprint arXiv:2307.15217*, 7 2023.

- [11] Paul F Christiano, Jan Leike, Tom B Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep Reinforcement Learning from Human Preferences. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, 2017.
- [12] Vincent Conitzer and Tuomas Sandholm. Computing the optimal strategy to commit to. *Proceedings of the ACM Conference on Electronic Commerce*, 2006:82–90, 2006.
- [13] Constantinos Daskalakis and Ioannis Panageas. Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization. 7 2018. URL <http://arxiv.org/abs/1807.04252>.
- [14] Jean Antoine Nicolas de Caritat Mis et al. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie royale, 1785.
- [15] Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori Hashimoto. AlpacaFarm: A simulation framework for methods that learn from human feedback. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [16] Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled alpacaEval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475*, 2024.
- [17] Jacob Eisenstein, Chirag Nagpal, Alekh Agarwal, Ahmad Beirami, Alex D’Amour, DJ Dvijotham, Adam Fisch, Katherine Heller, Stephen Pfohl, Deepak Ramachandran, Peter Shaw, and Jonathan Berant. Helping or herding? reward model ensembles mitigate but do not eliminate reward hacking, 2024.
- [18] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A Theory of Regularized Markov Decision Processes. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- [19] Amelia Glaese, Nat McAleese, Maja Trębacz, John Aslanides, Vlad Firoiu, Timo Ewalds, Mari-beth Rauh, Laura Weidinger, Martin Chadwick, Phoebe Thacker, et al. Improving alignment of dialogue agents via targeted human judgements. *arXiv preprint arXiv:2209.14375*, 2022.
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the Thirty-first International Conference on Neural Information Processing Systems*, 2017.
- [21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

- [22] Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
- [23] Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 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, 2023.
- [24] Timo Kaufmann, Paul Weng, Viktor Bengs, and Eyke Hüllermeier. A Survey of Reinforcement Learning from Human Feedback. *arXiv preprint arXiv:2312.14925*, 2023.
- [25] Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- [26] Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, Lei M Zhang, Kay McKinney, Disha Shrivastava, Cosmin Paduraru, George Tucker, Doina Precup, Feryal Behbahani, and Aleksandra Faust. Training language models to self-correct via reinforcement learning. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- [27] Tianyi Lin, Chi Jin, and Michael I Jordan. Two-timescale gradient descent ascent algorithms for nonconvex minimax optimization. *Journal of Machine Learning Research*, 26(11):1–45, 2025.
- [28] Kaizhao Liu, Qi Long, Zhekun Shi, Weijie J Su, and Jiancong Xiao. Statistical impossibility and possibility of aligning llms with human preferences: From condorcet paradox to nash equilibrium. *arXiv preprint arXiv:2503.10990*, 2025.
- [29] Jacob Makar-Limanov, Arjun Prakash, Denizalp Goktas, Nora Ayanian, and Amy Greenwald. Sta-rlhf: Stackelberg aligned reinforcement learning with human feedback. In *Coordination and Cooperation for Multi-Agent Reinforcement Learning Methods Workshop*, 2024.
- [30] Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, et al. Teaching language models to support answers with verified quotes. *arXiv preprint arXiv:2203.11147*, 2022.
- [31] Rémi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J Mankowitz, Doina Precup, Bilal Piot, and Google Deepmind. Nash Learning from Human Feedback. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [32] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- [33] Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

- [34] Leo Gao Openai, John Schulman, Openai Jacob, and Hilton Openai. Scaling Laws for Reward Model Overoptimization. In *International Conference on Machine Learning, 2023*. ISBN 2210.10760v1.
- [35] Long Ouyang, Jeff 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 Askill, Peter Welinder Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Proceedings of the 36th Conference on Neural Information Processing Systems, 2022*.
- [36] HL Prasad, Prashanth LA, and Shalabh Bhatnagar. Two-timescale algorithms for learning nash equilibria in general-sum stochastic games. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, 2015*.
- [37] Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching language model agents how to self-improve. *Advances in Neural Information Processing Systems, 2024*.
- [38] Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115v2*, 12 2024.
- [39] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In *Advances in Neural Information Processing Systems, 2023*. ISBN 2305.18290v2.
- [40] Corby Rosset, Ching-An Cheng, Arindam Mitra, Michael Santacrose, Ahmed Awadallah, and Tengyang Xie. Direct Nash Optimization: Teaching Language Models to Self-Improve with General Preferences. In *arXiv preprint arXiv:2404.03715*, 2024.
- [41] William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. Self-critiquing models for assisting human evaluators. *arXiv preprint arXiv:2206.05802*, 2022.
- [42] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [43] Pier Giuseppe Sessa, Robert Dadashi, Léonard Hussenot, Johan Ferret, Nino Vieillard, Alexandre Ramé, Bobak Shariari, Sarah Perrin, Abe Friesen, Geoffrey Cideron, Sertan Girgin, Piotr Stanczyk, Andrea Michi, Danila Sinopalnikov, Sabela Ramos, Amélie Héliou, Aliaksei Severyn, Matt Hoffman, Nikola Momchev, Olivier Bachem, and Google Deepmind. BOND: Aligning LLMs with Best-of-N Distillation. *arXiv preprint arXiv:2407.14622*, 2024.
- [44] Han Shen, Zhuoran Yang, and Tianyi Chen. Principled penalty-based methods for bilevel reinforcement learning and rlhf. In *Proceedings of the 41st International Conference on Machine Learning, 2024*.

- [45] Heinrich von Stackelberg. *Theory of the market economy*. Oxford University Press, 1952.
- [46] Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M Ziegler, Ryan Lowe, Chelsea Voss, and Alec Radford Dario Amodei Paul Christiano. Learning to summarize from human feedback. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, 2020.
- [47] Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, and Alekh Agarwal. A Minimaximalist Approach to Reinforcement Learning from Human Feedback RLHF / PbRL. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- [48] Xiaohang Tang, Sangwoong Yoon, Seongho Son, Huizhuo Yuan, Quanquan Gu, and Ilija Bogunovic. Game-theoretic regularized self-play alignment of large language models. *arXiv preprint arXiv:2503.00030*, 2025.
- [49] Vinzenz Thoma, Barna Pásztor, Andreas Krause, Giorgia Ramponi, and Yifan Hu. Contextual bilevel reinforcement learning for incentive alignment. In *Proceedings of the Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [50] Gladys Tyen, Hassan Mansoor, Victor Cărbune, Yuanzhu Peter Chen, and Tony Mak. Llms cannot find reasoning errors, but can correct them given the error location. In *Findings of the Association for Computational Linguistics ACL 2024*, 2024.
- [51] Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Gallouédec. Trl: Transformer reinforcement learning, 2020.
- [52] Mingzhi Wang, Chengdong Ma, Qizhi Chen, Linjian Meng, Yang Han, Jiancong Xiao, Zhaowei Zhang, Jing Huo, Weijie J Su, and Yaodong Yang. Magnetic preference optimization: Achieving last-iterate convergence for language model alignment. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [53] Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. Helpsteer 2: Open-source dataset for training top-performing reward models. In *Proceedings of the Thirty-eight Conference on Neural Information Processing Systems*, 2024.
- [54] Sean Welleck, Ximing Lu, Peter West, Faeze Brahman, Tianxiao Shen, Daniel Khashabi, and Yejin Choi. Generating Sequences by Learning to Self-Correct. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.
- [55] Sean Welleck, Amanda Bertsch, Matthew Finlayson, Hailey Schoelkopf, Alex Xie, Graham Neubig, Ilija Kulikov, and Zaid Harchaoui. From decoding to meta-generation: Inference-time algorithms for large language models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.
- [56] Christian Wirth, Riad Akrou, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *Journal of Machine Learning Research*, 18 (136):1–46, 2017.

- [57] Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. Self-Play Preference Optimization for Language Model Alignment. *arXiv preprint arXiv:2405.00675*, 2024.
- [58] Jiancong Xiao, Ziniu Li, Xingyu Xie, Emily Getzen, Cong Fang, Qi Long, and Weijie J Su. On the algorithmic bias of aligning large language models with rlhf: Preference collapse and matching regularization. *arXiv preprint arXiv:2405.16455*, 2024.
- [59] Yi Xu, Laura Ruis, Tim Rocktäschel, and Robert Kirk. Investigating non-transitivity in llm-as-a-judge. *arXiv preprint arXiv:2502.14074*, 2025.
- [60] Yuheng Zhang, Dian Yu, Baolin Peng, Linfeng Song, Ye Tian, Mingyue Huo, Nan Jiang, Haitao Mi, and Dong Yu. Iterative Nash Policy Optimization: Aligning LLMs with General Preferences via No-Regret Learning. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- [61] Yunxiang Zhang, Muhammad Khalifa, Lajanugen Logeswaran, Jaekyeom Kim, Moontae Lee, Honglak Lee, and Lu Wang. Small language models need strong verifiers to self-correct reasoning. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Findings of the Association for Computational Linguistics: ACL 2024*, 2024.
- [62] Huaixiu Steven Zheng, Swaroop Mishra, Hugh Zhang, Xinyun Chen, Minmin Chen, Azade Nova, Le Hou, Heng-Tze Cheng, Quoc V Le, Ed H Chi, et al. Natural plan: Benchmarking llms on natural language planning. *arXiv preprint arXiv:2406.04520*, 2024.
- [63] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-Tuning Language Models from Human Preferences. *arXiv preprint arXiv:1909.08593*, 2019.

Contents of Appendix

A Proofs	14
A.1 Proof of Theorem 1	14
A.2 SLHF Admits a Deterministic Solution	15
A.3 Concave-Convex Property of f	15
B Related Work	15
C Comparison of Solution Concepts	16
D Scalable Implementation of STACKELBERGGDA	18
E Implementation Details	19
F Additional Experimental Results	19
F.1 Preference Model	19
F.2 Iterative Improvements at Test-time	20
F.3 Ablation on the two-timescale coefficient	23
F.4 Model Scaling	24

Appendix A. Proofs

A.1. Proof of Theorem 1

Proof First, assume that the Leader’s policy π is fixed and consider the Follower’s optimization problem

$$\min_{\omega} \mathbb{E}_{x \sim \rho, y \sim \pi(\cdot|x)} [\mathbb{E}_{y' \sim \omega(\cdot|x,y)} [p(y \succ y' | x)] + \tau^F \text{KL}_{x,y}(\omega \| \omega^{\text{ref}})]. \quad (6)$$

The optimization problem in (6) is equivalent to Equation (2) for the reward function $r(\tilde{x}, y') := p(y \succ y' | x)$ with contexts $\tilde{x} = (x, y)$ and context distribution $\tilde{x} \sim \rho \otimes \pi$. As a result, Equation (6) has a unique closed-form solution [4, 18, 39] given by

$$\omega^*(y' | x, y) = \frac{1}{Z(x, y)} \omega^{\text{ref}}(y' | x, y) \exp\left(\frac{1}{\tau^F} p(y' \succ y | x)\right)$$

where $Z(x, y) = \sum_{y' \in \mathcal{Y}} \omega^{\text{ref}}(y' | x, y) \exp\left(\frac{1}{\tau^F} p(y' \succ y | x)\right)$ is a partition factor that depends only on (x, y) and π^{ref} . Hence, ω^* can be expressed as a function of (x, y) and ω^{ref} without explicit dependence on π .

Now, define the following reward function for the Leader’s optimization problem

$$r(x, y) := \mathbb{E}_{y' \sim \omega^*(\cdot|x,y)} [p(y \succ y' | x)]. \quad (7)$$

Note that ω^* is unique so that $r(x, y)$ is a scalar. We can now restate Equation (4) for the Leader’s optimization problem as

$$\max_{\pi} \mathbb{E}_{x \sim \rho} [\mathbb{E}_{y \sim \pi(\cdot|x)} [r(x, y)] - \tau^L \text{KL}_x(\pi \| \pi^{\text{ref}})]$$

which is again a KL-regularized optimization problem that admits a closed-form solution

$$\pi^*(y | x) = \frac{1}{Z(x)} \pi^{\text{ref}}(y | x) \exp\left(\frac{1}{\tau^L} r(x, y)\right).$$

■

A.2. SLHF Admits a Deterministic Solution

Proof This lemma is folklore in the algorithmic game theory community and can be quickly verified.

Let $x \in \mathcal{X}$. Given any action $y \in \mathcal{Y}$, there exists a not necessarily unique $y' \in \mathcal{Y}$ minimizing $p(y \succ y' | x)$. Hence, irrespective of the Leader’s policy $\pi(\cdot | x)$, there always exists a Follower’s deterministic best response policy $\omega_{\text{br}}(\cdot | x, y)$ with $\omega_{\text{br}}(y' | x, y) = 1$ for some y' . In other words, the Follower always has a deterministic best response policy.

Similarly, the optimization problem for the Leader given some context x reduces to finding y that maximizes $\mathbb{E}_{y' \sim \omega_{\text{br}}(\cdot | x, y)}[p(y \succ y' | x)]$ so that the SLHF optimization problem admits a deterministic solution. ■

A.3. Concave-Convex Property of f

We show here that the objective function f in Equation (5) of the Stackelberg optimization problem is concave-convex. Similar results were established in the context of NLHF by Munos et al. [31].

Throughout this section, we assume $|X| = 1$ and omit x from the notation for clarity. All results extend directly to the general case with a finite context space \mathcal{X} .

Then, the objective function of Equation (4) is given by

$$f(\pi, \omega) = \mathbb{E}_{y \sim \pi(\cdot), y' \sim \omega(\cdot | y)}[p(y \succ y')] - \tau^L \text{KL}(\pi \| \pi^{\text{ref}}) + \tau^F \mathbb{E}_{y \sim \pi(\cdot)}[\text{KL}_y(\omega \| \omega^{\text{ref}})]. \quad (8)$$

The first term is bilinear in π and ω , as shown by expanding the expectation:

$$\mathbb{E}_{y \sim \pi(\cdot), y' \sim \omega(\cdot | y)}[p(y \succ y')] = \sum_{y \in \mathcal{Y}} \pi(y) \sum_{y' \in \mathcal{Y}} p(y \succ y') \omega(y' | y).$$

The KL terms are convex in their respective arguments. Hence, f is bilinear when $\tau^L = \tau^F = 0$, and strongly concave-convex when $\tau^L, \tau^F > 0$.

Appendix B. Related Work

Reinforcement Learning from Human Feedback (RLHF). RLHF aims to optimize a policy based on human preferences by learning from comparisons or rankings rather than explicit numeric rewards [24, 56]. Most RLHF methods follow a two-step pipeline introduced by Christiano et al. [11]: first, a reward model is trained from pairwise comparisons; then, this model is used as a proxy reward function for policy optimization, typically using PPO [42]. This framework has been widely adopted for tasks such as text summarization [46], question answering [30, 32], and language model fine-tuning [5, 19, 35, 63]. Recent work has also combined these two steps into a bilevel optimization problem to jointly optimize reward model and policy [29, 44, 49].

Limitations of Reward Modeling. Most RLHF methods reduce preference learning to scalar reward estimation, typically using the Bradley-Terry model [8]. While effective for transitive, single-objective preferences, such models fail to represent intransitive structures and may even misrank transitive ones under model misspecification [7]. As a result, they struggle to capture the diversity and ambiguity of real human preferences [10]. These issues are compounded during optimization: the final policy can be highly sensitive to the distribution of training comparisons [31] and may collapse to a single preference mode under continued training [58]. Interestingly, even when preferences are elicited from LLMs, as in the AlpacaEval framework [16], the feedback can exhibit intransitive preference structures [59]. In contrast, we work directly with pairwise preference signals without imposing an underlying reward model. This enables optimization with respect to diverse and potentially intransitive feedback.

Preference Optimization. To address the limitations of reward modeling in RLHF, IPO [4] extends Direct Preference Optimization (DPO) [39] by optimizing for the win rate against a reference policy. Nash Learning from Human Feedback (NLHF) casts the learning problem as a two-player simultaneous-move game and introduces NASH-MD-PG and NASH-EMA-PG to approximate the Nash Equilibrium (NE) of a learned preference model via mirror descent [31]. Subsequent work has extended this perspective, proposing various algorithms to optimize for (approximate) NE, including ONLINE-IPO [9], SPPO [57], SPO [47], INPO [60], DNO [40], RSPO [48], NASH-RS [28], and MPO [52]. These methods typically converge to mixed strategies unless one option is majority-preferred, even in non-regularized settings [28], due to the symmetry of simultaneous-move games. In contrast, we formulate preference optimization as a sequential-move game, inspired by Stackelberg (Leader-Follower) dynamics. Here, the Leader commits first, and the Follower responds conditionally on the Leader’s action. This asymmetric structure leads to a different solution concept than the NE, admitting deterministic policies in the non-regularized limit.

Test-time Preference Improvement. Improving the capabilities of LLMs through additional computation at test-time has received significant attention recently, especially in verifiable domains such as coding or mathematics [55]. Closest to our work are self-correction algorithms that aim to improve their responses without external feedback at test-time. A natural approach to self-correction is to provide instructions only without further training, which, however, can lead to performance degradation [22, 37, 50, 62]. Other work on training models for self-correction either assumes human or AI revisions [37, 41] or a reward function scoring responses [2, 26, 54, 61]. Similarly to SLHF, Kumar et al. [26] also propose to train an LLM in a sequential manner, however, assume a reward model and train in two-stages instead of a single loop. On the other hand, STACKELBERGGDA can self-improve on general preferences and train in a single loop.

Appendix C. Comparison of Solution Concepts

Before we describe how to approximate the Stackelberg equilibrium, we present the differences between RLHF, NLHF, and SLHF at the example of the so-called Condorcet paradox [14].

Consider a problem with $|\mathcal{X}| = 1$ and $\mathcal{Y} = \{A, B, C\}$, and let the population of annotators \mathcal{A} consist of three distinct types. Table 3 defines the annotators’ preferences over \mathcal{Y} . Each type of annotator has a strict preference ranking and the proportion of the types according to the distribution over annotators ν is such that $\alpha + \beta + \gamma = 1$. Table 4 shows the aggregated preferences of the whole population. For example, row A column B states that $1 - \beta$ of the population prefers A over B ,

Table 3: Transitive individual annotator preferences over three options: $\{A, B, C\}$.

Preference Relationship	Proportion
$A \succ B \succ C$	α
$B \succ C \succ A$	β
$C \succ A \succ B$	γ

Table 4: The preference function derived from the individual rankings in Table 3.

	A	B	C
A	0.5	$1 - \beta$	α
B	β	0.5	$1 - \gamma$
C	$1 - \alpha$	γ	0.5

i.e., $p(A \succ B) = 1 - \beta$. A common example is to choose $\alpha = \beta = \gamma = 1/3$, which leads to a cyclic relationship between three actions in which $A \succ B \succ C$ but $C \succ A$. Hence, there exists no Condorcet winner in this case.⁵ This example is often referred to as the *Condorcet paradox*, because the annotators individually have transitive preferences (Table 3), but their aggregated preferences form a cycle (Table 4). For ease of presentation, we consider a non-regularized problem in the rest of this section so that $\tau = \tau^L = \tau^F = 0$.

RLHF Solution. Our first observation is that the estimated reward function $\hat{r} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ depends heavily on the sampling distribution of the dataset \mathcal{D} used to estimate \hat{r} . Consider a case where \mathcal{D} consists only of observations between pairs (A, B) and (B, C) and the pair (A, C) is not included. Since $A \succ B$ and $B \succ C$, the reward estimate that maximizes the log-likelihood satisfies $\hat{r}(A) > \hat{r}(B) > \hat{r}(C)$. Hence, the RLHF solution to Equation (2) is $\pi^*(A) = 1$. However, similar outcomes can occur for actions B and C when only pairs (A, C) , (B, C) or (A, C) , (A, B) are sampled, respectively. This illustrates a key limitation of RLHF, namely that its solutions are sensitive to the specific comparisons present in \mathcal{D} .

Nash Equilibrium. Under the assumption that $\alpha, \beta, \gamma \leq 0.5$, the NE of the matrix game defined in Table 4 is given by $\pi^*(A) = 1 - 2\gamma$, $\pi^*(B) = 1 - 2\alpha$, $\pi^*(C) = 1 - 2\beta$. In the special case of $\alpha = \beta = \gamma = 1/3$, the NE is uniform over \mathcal{Y} , i.e., it has the highest possible entropy. Even though this solution does not depend on the distribution of any dataset, it leads to a stochastic policy that might be undesirable depending on the application area.

Stackelberg Equilibrium. Since the Follower gets to observe the Leader’s action, it can play a best response $\omega(\cdot | y)$, which we can in turn use to derive the Leader’s optimal policy π^* :

$$\omega^*(\cdot | y) = \begin{cases} C & \text{if } y = A \text{ w.p. } 1 \\ A & \text{if } y = B \text{ w.p. } 1 \\ B & \text{if } y = C \text{ w.p. } 1 \end{cases} \quad \pi^*(\cdot) = \begin{cases} A & \text{if } \alpha > \max\{\beta, \gamma\} \text{ w.p. } 1 \\ B & \text{if } \beta > \max\{\alpha, \gamma\} \text{ w.p. } 1 \\ C & \text{if } \gamma > \max\{\alpha, \beta\} \text{ w.p. } 1 \end{cases}$$

Unlike RLHF, this solution does not rely on any offline dataset \mathcal{D} . Furthermore, in contrast to NLHF, it admits a deterministic solution in the absence of KL regularization. However, when $\alpha = \beta = \gamma = 1/3$, the Leader is indifferent among the three actions, meaning that any distribution over A, B, C constitutes a Stackelberg equilibrium. This includes the Nash equilibrium $\pi^*(A) = \pi^*(B) = \pi^*(C) = 1/3$.

5. Other common optimality conditions in social choice theory are also undefined for this case, e.g., the Borda or plurality winner.

Test-time Improvements. Finally, we consider how each solution can be improved at test-time with respect to the preferences of a single annotator $a \in \mathcal{A}$. Suppose $\alpha = \beta = \gamma = 1/3$, and without loss of generality, let a belong to the first annotator type with preference $A \succ B \succ C$. The RLHF solution may return A , which aligns with a 's most preferred action. However, as previously discussed, the outcome of RLHF depends on the dataset \mathcal{D} ; it could just as easily produce B or C . Since the RLHF policy is deterministic, repeated sampling does not change this outcome. In contrast, the NLHF solution is uniform over A, B, C , so the probability of sampling A in a single draw is $1/3$. By sampling N times, the probability of observing at least one A is $1 - (2/3)^N$, which equals 56% for $N = 2$ and 70% for $N = 3$. The Stackelberg solution starts similarly: the first action is sampled from the Leader's potentially uniform policy. However, subsequent actions are drawn from the Follower's policy, i.e., $y_i \sim \omega^*(\cdot | x, y_{i-1})$ for $i \geq 2$. In this case, the probability of sampling A within $N = 2$ steps increases to 67%, and for $N = 3$, the entire preference cycle is traversed regardless of the Leader's initial choice.

Appendix D. Scalable Implementation of STACKELBERGGDA

When fine-tuning large language models, the context and action spaces \mathcal{X} and \mathcal{Y} are far too large to optimize over Π and Ω directly. To address this, we introduce a practical variant of STACKELBERGGDA in Algorithm 1.

Policy Parameterization. We replace the tabular policies π and ω with neural parameterizations π_θ and ω_ϕ (e.g., transformer networks). This renders the policy spaces tractable via their parameter vectors θ and ϕ , however, the concave-convex property does not necessarily carry over to the parameters θ and ϕ .

Batched, Variance-reduced Gradient Estimates. Exact evaluation of the expectations in ∇f is infeasible due to the expectation over the context and action spaces. Instead, at each iteration we sample a batch of size B :

$$\{(x_i, y_i, y'_i, p_i)\}_{i=1}^B, \quad x_i \sim \rho, \quad y_i \sim \pi_\theta(\cdot | x_i), \quad y'_i \sim \omega_\phi(\cdot | x_i, y_i), \quad p_i = p(y_i \succ y'_i | x_i).$$

We then form unbiased estimates as

$$\widehat{\nabla}_\theta f = \frac{1}{B} \sum_{i=1}^B (p_i - \tau^L k_i^L) \nabla_\theta \log \pi_\theta(y_i | x_i), \quad \widehat{\nabla}_\phi f = \frac{1}{B} \sum_{i=1}^B (p_i - \tau^F k_i^F) \nabla_\phi \log \omega_\phi(y'_i | x_i, y_i),$$

with likelihood ratios $k_i^L = \frac{\pi_\theta(y_i | x_i)}{\pi^{\text{ref}}(y_i | x_i)}$ and $k_i^F = \frac{\omega_\phi(y'_i | x_i, y_i)}{\omega^{\text{ref}}(y'_i | x_i, y_i)}$. The gradient estimators are naturally compatible with additional variance reduction techniques such as subtracting a constant baseline.

Single-Model Instantiation. Simultaneously training two billion-parameter transformer models is memory-prohibitive. Similarly to SCORE [26], we collapse both Leader and Follower into one model π_θ by using distinct chat templates (Figure 1). When the model is only given the context x , we use the template in Figure 1(a) that only includes x as the prompt. When the model is given both the context x and an action y , we use the template in Figure 1(b) that includes both the context x and the action y , as well as a predefined instruction to improve the action y .

Then, letting $\kappa = \frac{\alpha^F}{\alpha^L}$ denote the two-timescale weight coefficient, we optimize the model to minimize the following loss function

$$\mathcal{L}(\theta) = -\frac{1}{B} \sum_{i=1}^B (p_i - \tau^L k_i^L) \log \pi_\theta(y_i | x_i) + \frac{\kappa}{B} \sum_{i=1}^B (p_i - \tau^F k_i^F) \log \pi_\theta(y'_i | x_i, y_i). \quad (9)$$

Gradient steps on $\mathcal{L}(\theta)$ realize the two-time-scale gradient descent-ascent updates via a single network, thereby substantially reducing memory usage.

Algorithm 1 STACKELBERGGDA (Practical)

Input: Datasets \mathcal{X}, \mathcal{Y} , distribution ρ , learning rate η

Output: Trained policy π with parameters θ

Initialize the policy π and weights ω

for $i = 1, 2, \dots$ **do**

for $b = 1, \dots, B$ **do**

 Sample prompt $x_b \sim \rho$

 Sample Leader response using the prompt in Figure 1(a): $y_b \sim \pi_\theta(\cdot | x_b)$

 Sample Follower response using the prompt in Figure 1(b): $y'_b \sim \pi_\theta(\cdot | x_b, y_b)$

 Observe preference feedback $p_b = p(y_b \succ y'_b | x_b)$

end

 Update the weights θ according to the loss in Equation (9):

$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}(\theta)$

end

Appendix E. Implementation Details

Implementation We trained RLOO⁶ and NASH-MD-PG⁷ using their implementations in the TRL Python package [51]. For all training runs, including reward modeling, we used Low-Rank Adaptation (LoRA) [21] with rank $r = 32$, scaling factor $\alpha = 64$, and dropout rate set to 0.1.

Compute Resources All experiments were conducted on a cluster with 8 NVIDIA GeForce RTX 4090 GPUs, 16 CPU cores, and 64 GB of RAM. The total compute time, including hyperparameter sweeps, was approximately 4,000 GPU-hours.

Appendix F. Additional Experimental Results

F.1. Preference Model

We estimate the preference model used to fine-tune the LLMs by treating the five attributes in the HELPSTEER2 datasets [53] as distinct annotators, denoted by the set \mathcal{A} , and define ν as a uniform distribution over \mathcal{A} . For each attribute $a \in \mathcal{A}$, we estimate a reward function \hat{r}_a using the Bradley-Terry model and maximize the log-likelihood on the training dataset $\mathcal{D} = \{(x_i, y_i^w, y_i^l)\}_{i=1}^N$

$$\min_r \sum_{i=1}^N \sigma(r(x_i, y_i^w) - r(x_i, y_i^l)) + \lambda(r(x_i, y_i^w) + r(x_i, y_i^l))^2.$$

We here decided which response is the winning and losing one in the dataset by comparing the attribute scores provided by the annotators. The additional regularization ensures that the rewards are centralized around zero [17]. For the attributes correctness, helpfulness, and coherence, we consider higher scores to be better while for verbosity and complexity lower values are more preferable. This

6. https://huggingface.co/docs/trl/main/en/rloo_trainer

7. https://huggingface.co/docs/trl/main/en/nash_md_trainer

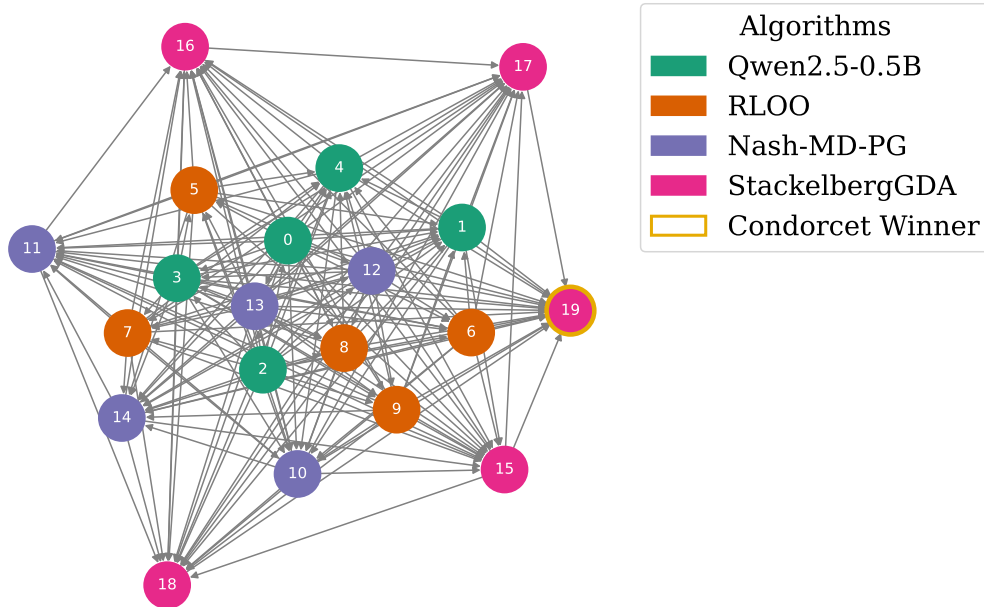


Figure 2: Directed graph based with completions generated by the fine-tuned models and edge directions representing the preference between them.

is in accordance with the scoring criteria described in Wang et al. [53]. We set the regularization parameter to $\lambda = 0.01$. Centralizing the reward values is crucial for our RLOO implementation, which uses the average reward across the five attributes as its training signal. It ensures that no attribute heavily dominates due to scale differences and avoids bias toward any particular attribute.

We train each model for 5 epochs on the training prompts and completions with batch size 32 and learning rate $1e-4$. The final accuracies of the models on the validation dataset are 78%, 65%, 61%, 60%, and 59% for verbosity, complexity, correctness, helpfulness, and coherence, respectively.

We evaluate the non-transitivity of the preference model p on the validation prompts and five responses from each of the four models used for comparison in Section 5. For each prompt, we construct a complete directed graph between the 20 completions as nodes and edges directed from the non-preferred completion towards the preferred one. Figure 2 illustrates this directed graph on the first prompt of the validation dataset. 57% of the directed graphs include cycles, which illustrate the intransitivity of the preference function p on the completion space \mathcal{Y} .

F.2. Iterative Improvements at Test-time

We extend the experimental results from Section 5 by analyzing iterative improvements and performance scaling with increased test-time computation. Our results suggests that with increasing test-time computation the benefit of fine-tuning using both RLOO or NASH-MD-PG is negligible compared to using the base model QWEN2.5-0.5B. In contrast, STACKELBERGGDA yields strict improvements. Building on the example in Appendix C, we assume that at test-time, a single annotator $a \sim \nu$ and a context $x \sim \rho$ are sampled.

Table 5: Preference scores for RLOO versus QWEN2.5-0.5B across all attributes as a function of the number of test-time samples N .

Attribute	Number of Samples N				
	1	2	3	4	5
Coherence	0.521	0.338	0.262	0.207	0.164
Complexity	0.892	0.842	0.801	0.771	0.754
Correctness	0.388	0.212	0.139	0.098	0.070
Helpfulness	0.322	0.150	0.087	0.057	0.036
Verbosity	0.846	0.777	0.728	0.695	0.669
Average	0.594	0.464	0.403	0.366	0.338

For the base model QWEN2.5-0.5B and the models with fine-tuned with RLOO and NASH-MD-PG, we independently sample N responses y_1, \dots, y_N . For STACKELBERGGDA, which inherently supports iterative refinement, we generate the first sample from the Leader policy $y_1 \sim \pi^*(\cdot | x)$, and subsequent responses from the Follower policy $y_i \sim \omega^*(\cdot | x, y_{i-1})$ for $i \geq 2$. We define $y_{1:N} := (y_1, \dots, y_N)$.

In line with prior work on Best-of- N sampling [6, 16, 34, 43], we evaluate the quality of the N samples by computing the maximum reward obtained for each attribute under the sampled annotator’s reward model, that is,

$$\hat{r}_a^N(x, y_{1:N}) := \max_{y_1, \dots, y_N} \hat{r}_a(x, y_i). \quad (10)$$

Analogous to the preference function p defined in Section 2, in this section, we compare two models π and π' w.r.t. the preference functions derived from the annotator-specific reward functions under Best-of- N sampling:

$$p_a^N(\pi \succ \pi' | x) := \mathbb{E}_{\substack{y_{1:N} \sim \pi(\cdot | x) \\ y'_{1:N} \sim \pi'(\cdot | x)}} \left[\mathbb{1} \{ \hat{r}_a^N(x, y_{1:N}) \geq \hat{r}_a^N(x, y'_{1:N}) \} \right]. \quad (11)$$

Notational Note. We here adopt a slight abuse of notation. Specifically, we write $y_{1:N} \sim \pi(\cdot | x)$ to denote the sampling of N responses from a model π , even though this notation does not faithfully represent the sampling procedure used by STACKELBERGGDA. For QWEN2.5-0.5B, RLOO, and NASH-MD-PG, the samples y_1, \dots, y_N are drawn i.i.d. from a single model $\pi(\cdot | x)$. In contrast, for STACKELBERGGDA, the sampling process is inherently autoregressive: we first draw $y_1 \sim \pi^*(\cdot | x)$ from the Leader policy, and then generate $y_i \sim \omega^*(\cdot | x, y_{i-1})$ for $i \geq 2$ using the Follower policy. Despite this difference, we overload the notation $y_{1:N} \sim \pi(\cdot | x)$ to unify the presentation in Equations (10) and (11). In the case of STACKELBERGGDA, this notation should be interpreted as shorthand for the autoregressive sampling process described above.

Previous work has shown that Best-of- N sampling can rival the performance of RLHF-based fine-tuning [6, 15, 43]. Motivated by this, we compare the preference scores defined in Equation (11) of RLOO, NASH-MD-PG, and STACKELBERGGDA with respect to the base model QWEN2.5-0.5B. Throughout this section, we consider the maximum number of samples to be $N = 5$.

Results. Table 5 reports the preference scores for RLOO. While the model initially (i.e. $N = 1$) performs competitively on complexity and verbosity attributes, iterative sampling reveals a collapse

Table 6: Preference scores for NASH-MD-PG versus QWEN2.5-0.5B across all attributes as a function of the number of test-time samples N .

Attribute	Number of Samples N				
	1	2	3	4	5
Coherence	0.731	0.743	0.757	0.763	0.760
Complexity	0.761	0.743	0.732	0.726	0.727
Correctness	0.633	0.615	0.620	0.617	0.615
Helpfulness	0.645	0.633	0.640	0.641	0.638
Verbosity	0.858	0.855	0.850	0.849	0.852
Average	0.726	0.718	0.720	0.719	0.718

into a single preference mode. In particular, we observed deterministic outputs for RLOO the generic response: *"I apologize, but I'm unable to engage in conversations about political topics. If you have any other questions or need further assistance with a different subject, feel free to ask."* As a result, the model's diversity and coverage deteriorate, and its overall preference scores (relative to QWEN2.5-0.5B) decline sharply as N increases.

In contrast, NASH-MD-PG demonstrates some benefit from additional sampling, as shown in Table 6. Its preference score for verbosity remains stable and it shows moderate improvement in coherence. However, for the remaining attributes (correctness, helpfulness, and complexity) its gains are slower than those achieved by QWEN2.5-0.5B with Best-of- N sampling. Consequently, the overall preference score of NASH-MD-PG declines with increasing N , suggesting that the model fine-tuned with NASH-MD-PG does not improve notably (compared to the base model) when the number of samples drawn at test-time increases.

On the other hand, STACKELBERGGDA exhibits more favorable behavior. As shown in Table 7, while the preference score on verbosity and complexity taper off with more samples, STACKELBERGGDA achieves notably faster gains on coherence, correctness, and helpfulness. For these attributes, the preference rate improves by 10 percentage points or more, making STACKELBERGGDA the only method among the three to demonstrate consistent improvement over QWEN2.5-0.5B as N increases. This means that the performance of STACKELBERGGDA effectively scales with test-time compute.

The strong emphasis on complexity and verbosity by RLOO is expected, as it optimizes the average reward across all five attributes, and these two dimensions yield the highest values. However, for NASH-MD-PG, this outcome is less expected. We hypothesize that it stems from its training objective, which pits the policy against a mixture of the reference policy QWEN2.5-0.5B and the most recent iteration. Once NASH-MD-PG outperforms QWEN2.5-0.5B on all attributes, it begins focusing on attributes where further improvement over itself is possible, namely, complexity and verbosity. Nevertheless, this skewed emphasis is suboptimal: annotators prefer models that perform well on all attributes. In fact, a policy that focuses on coherence, correctness, and helpfulness is preferred by 60% of the annotators. STACKELBERGGDA's asymmetric formulation that trains a Leader and a Follower policy separately (though potentially unified in a single model) helps mitigate this imbalance across attributes. This leads to a more balanced policy that is preferred by a wider range of annotators.

Table 7: Preference scores for STACKELBERGGDA versus QWEN2.5-0.5B across all attributes as a function of the number of test-time samples N .

Attribute	Number of Samples N				
	1	2	3	4	5
Coherence	0.778	0.850	0.865	0.875	0.873
Complexity	0.714	0.666	0.628	0.600	0.592
Correctness	0.670	0.762	0.791	0.795	0.803
Helpfulness	0.692	0.767	0.783	0.786	0.791
Verbosity	0.833	0.820	0.798	0.777	0.765
Average	0.738	0.773	0.773	0.767	0.765

 Table 8: Ablation on the follower weight parameter κ in STACKELBERGGDA’s loss function (9) fine-tuning the QWEN2.5-0.5B model. Scores show the average preference over the base model.

Follower Weight κ	Train		Validation	
	Leader	Follower	Leader	Follower
1	0.768	0.804	0.761	0.784
5	0.743	0.814	0.723	0.806
10	0.725	0.800	0.710	0.783

F.3. Ablation on the two-timescale coefficient

Table 8 and Table 9 present ablations on the follower weight parameter κ in STACKELBERGGDA’s loss function (9) when fine-tuning the QWEN2.5-0.5B and QWEN2.5-1.5B models, respectively. Each row reports the average preference scores over the corresponding initial policy, for both the Leader and Follower policies, on the training and validation splits. These results highlight the importance of balancing the two components of STACKELBERGGDA’s asymmetric training objective. In general, moderate values of κ can help the Follower improve without compromising the Leader too severely, but excessively large weights may impair both players.

In Table 8, we observe that increasing κ leads to a gradual decline in the Leader’s performance. While the Follower benefits from increasing κ from 1 to 5, performance worsens at $\kappa = 10$ for both the Leader and Follower, indicating an overemphasis on the Follower’s loss can destabilize the overall optimization.

Table 9 shows a similar trend for the larger QWEN2.5-1.5B model. Due to the decrease of performance above $\kappa = 5$ in Table 8, we carry out the ablation on a finer grid $\kappa \in \{1, 2, 3, 4, 5\}$. Moreover, we evaluate each model after 2000 training steps as a larger base model requires more gradient updates to converge. While the performance of $\kappa = 1$ stands out in Table 9, we observe that it is overfitting to verbosity and complexity by responding to every prompt with short, non-informative answers asking for further information such as *"Certainly! If you need detailed insights on technical topics like that, feel free to ask—I'm here to assist with informatively aligned queries!"*. On the contrary to the collapse observed for RLOO in Appendix F.2, the model remains stochastic

Table 9: Ablation on the follower weight parameter κ in STACKELBERGGDA’s loss function (9) fine-tuning the QWEN2.5-1.5B model. Scores show the average preference over the base model.

Follower Weight κ	Train		Validation	
	Leader	Follower	Leader	Follower
1	0.848	0.852	0.850	0.851
2	0.718	0.737	0.719	0.730
3	0.767	0.806	0.771	0.803
4	0.733	0.811	0.736	0.807
5	0.720	0.819	0.720	0.818

with the responses having similar information content. This outcome demonstrates the effectiveness of STACKELBERGGDA in optimizing its objective despite the qualitatively undesirable responses.

F.4. Model Scaling

We extend our round-robin comparison from Section 5.1 to larger models within the Qwen2.5 family, specifically, QWEN2.5-1.5B and QWEN2.5-3B [38]. These evaluations demonstrate that STACKELBERGGDA continues to be on par or outperform baselines even as model size increases. Since larger models require more training updates to reach convergence in our setup, we train NASH-MD-PG for 1,500 steps and STACKELBERGGDA for 2,000 steps. The RLOO method converges earlier and requires only 1,000 steps even for these larger models. We fix the follower-weight parameter at $\kappa = 5$ for both scales, based on our ablation results in Appendix F.3.

Table 10 summarizes results for models fine-tuned from QWEN2.5-1.5B. Both NASH-MD-PG and STACKELBERGGDA clearly outperform the base model and the RLOO baseline. While the Leader policy of STACKELBERGGDA underperforms compared to NASH-MD-PG, the Follower policy conditioned on the Leader’s responses matches or exceeds NASH-MD-PG’s performance, mirroring the improvements observed when starting from the QWEN2.5-0.5B in Table 1. As noted in Appendix F.3, this performance gap between the Leader and NASH-MD-PG could likely be reduced by tuning κ , albeit at the potential cost of Follower quality.

Table 11 shows analogous comparisons for models initialized from QWEN2.5-3B. Here, STACKELBERGGDA again performs strongly, with its Follower policy matching or surpassing NASH-MD-PG across most pairwise matchups, and both algorithms outperforming the base model. NASH-MD-PG and STACKELBERGGDA are closely matched when compared directly. Due to compute limitations, we capped training at 2,000 steps for these larger models. Nonetheless, the Leader policy continued to improve near the end of training, suggesting further gains in preference score may be possible with additional updates.

Table 10: Pairwise preference comparisons between the responses of QWEN2.5-0.5B, QWEN2.5-1.5B, RLOO, NASH-MD-PG, and STACKELBERGGDA algorithms. Fine-tuned models are trained from the QWEN2.5-1.5B. Each cell shows the average preference score of the row model over the column model.

	QWEN2.5-0.5B	QWEN2.5-1.5B	RLOO	NASH-MD-PG	STACKELBERGGDA	
					LEADER	FOLLOWER
QWEN2.5-0.5B	0.000	0.479	0.379	0.188	0.271	0.166
QWEN2.5-1.5B	0.521	0.000	0.401	0.209	0.293	0.187
RLOO	0.621	0.599	0.000	0.197	0.310	0.175
NASH-MD-PG	0.812	0.791	0.803	0.000	0.623	0.489
STACKELBERGGDA LEADER	0.729	0.707	0.690	0.377	0.000	0.313
STACKELBERGGDA FOLLOWER	0.834	0.813	0.825	0.511	0.687	0.000

Table 11: Pairwise preference comparisons between the responses of QWEN2.5-0.5B, QWEN2.5-3B, RLOO, NASH-MD-PG, and STACKELBERGGDA algorithms. Fine-tuned models are trained from the QWEN2.5-3B. Each cell shows the average preference score of the row model over the column model.

	QWEN2.5-0.5B	QWEN2.5-3B	RLOO	NASH-MD-PG	STACKELBERGGDA	
					LEADER	FOLLOWER
QWEN2.5-0.5B	0.000	0.504	0.399	0.187	0.304	0.187
QWEN2.5-3B	0.496	0.000	0.412	0.199	0.319	0.179
RLOO	0.601	0.588	0.000	0.173	0.338	0.201
NASH-MD-PG	0.813	0.801	0.827	0.000	0.638	0.507
STACKELBERGGDA LEADER	0.696	0.681	0.662	0.362	0.000	0.312
STACKELBERGGDA FOLLOWER	0.813	0.821	0.799	0.493	0.688	0.000