TokenRatio: Principled Token-Level Preference Optimization via Ratio Matching

Truong Nguyen [* 1]   Tien-Phat Nguyen [* 1]   Linh Ngo Van [1]   Duy Minh Ho Nguyen [2 3 4]   Khoa D. Doan [5]   Trung Le [6]

Abstract

Direct Preference Optimization (DPO) is a widely used RL-free method for aligning language models from pairwise preferences, but it models preferences over full sequences even though generation is driven by per-token decisions. Existing token-level extensions typically decompose a sequence-level Bradley–Terry objective across timesteps, leaving per-prefix (state-wise) optimality implicit. We study how to recover **token-level** preference optimality using only standard sequence-level pairwise comparisons. We introduce **Token-level Bregman Preference Optimization (TBPO)**, which posits a token-level Bradley–Terry preference model over next-token actions conditioned on the prefix, and derive a Bregman-divergence density-ratio matching objective that generalizes the logistic/DPO loss while preserving the optimal policy induced by the token-level model and maintaining DPO-like simplicity. We introduce two instantiations: TBPO-Q, which explicitly learns a lightweight state baseline, and TBPO-A, which removes the baseline through advantage normalization. Across instruction following, helpfulness/harmlessness, and summarization benchmarks, TBPO improves alignment quality and training stability and increases output diversity relative to strong sequence-level and token-level baselines. Code is available at https://github.com/dinhtruongng/TBPO.

1. Introduction

Aligning large language models with human preferences is a central goal of modern fine-tuning, and recent work has increasingly favored lightweight and stable objectives for this purpose. In particular, Direct Preference Optimization (DPO) (Rafailov et al., 2023) avoids the engineering complexity and variance of RL-style updates while still producing visibly improved behavior. However, a fundamental mismatch remains: language models **generate** text by making a sequence of local decisions, one token at a time conditioned on a prefix, while most preference optimization objectives treat each completion as a single indivisible object. Although this sequence-level view is convenient, it strains the connection between **where** a model is optimized (a whole response) and **how** it actually behaves at inference time (a distribution over next tokens at every state).

A natural response to the mismatch between autoregressive generation and sequence-level preference objectives is to push preference optimization to token granularity. However, existing "token-level" methods (e.g., TDPO (Zeng et al., 2024) and TISDPO (Liu et al., 2025a)) still fundamentally optimize a sequence-level Bradley–Terry preference model, differing mainly in how the sequence-level signal is distributed across timesteps. As a result, token-wise optimality is not explicitly enforced but remains implicit. This limitation becomes increasingly problematic for long generations: preference supervision is only observed at the level of full responses, making credit assignment to individual tokens indirect and allowing small early-token errors to cascade through subsequent decisions, even as the sequence-level objective improves. What we ultimately want, but rarely formalize, is a policy that is **preference-optimal at every state**, i.e., one that makes locally optimal tradeoffs at each state before the remainder of the completion unfolds. Achieving this stronger notion of optimality without sacrificing the simplicity and efficiency that make DPO attractive is therefore both challenging and valuable.

This paper proposes **Token-level Bregman Preference Optimization (TBPO)**, a preference-optimization framework that explicitly models preferences at the token level via density ratio matching while relying solely on sequence-level comparison data. TBPO is built on a token-level Bradley–

---

[*]Equal contribution [1]Hanoi University of Science and Technology [2]University of Stuttgart [3]German Research Center for Artificial Intelligence [4]Max Planck Research School for Intelligent Systems [5]VinUniversity [6]Monash University. Correspondence to: Linh Ngo Van <linhnv@soict.hust.edu.vn>, Trung Le <trunglm@monash.edu>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

Terry preference model that aligns preference learning with the autoregressive decision process of language models, and is designed to recover a token-wise optimal policy instead of implicitly relying on sequence-level credit assignment distributed across timesteps. Depending on the choice of scoring function in the token-level Bradley–Terry model, TBPO naturally yields different variants. We introduce two instantiations: **TBPO-Q**, which defines the token-level score via a state–action value function, and **TBPO-A**, which defines it via an advantage function. By casting TBPO as density-ratio matching under a Bregman-divergence objective, we even obtain a plug-and-play family of pairwise preference losses: small tweaks to the Bregman generator yield different objectives. Crucially, this flexibility does not come at the cost of theory or usability: the resulting objective still supports an optimality characterization in which the minimizer corresponds to the desired optimal token-level policy, while retaining DPO-style simplicity without fitting an explicit reward model or running on-policy RL.

We summarize our contributions as follows:

- **Token-level Bregman Preference Optimization (TBPO).** We propose TBPO, a token-level preference-optimization framework that models preferences directly at the token level and learns from sequence-level comparisons.

- **Theory.** We provide a rigorous derivation and an optimality characterization: under the token-level preference model, minimizing the TBPO objective recovers a token-level optimal policy.

- **Empirical results.** Across benchmarks, model backbones, and tasks, TBPO improves alignment while also enhancing training stability and preserving response diversity relative to strong baselines.

2. Related Work

Preference-based alignment. Aligning large language models (LLMs) with human judgments is commonly approached via preference-based optimization. RLHF learns an explicit reward model from pairwise comparisons and optimizes the policy via reinforcement learning, but suffers from instability and high computational cost (Christiano et al., 2017; Ouyang et al., 2022). DPO provides an RL-free alternative by deriving a closed-form relationship between rewards and optimal policies under a KL-regularized objective, yielding a supervised loss equivalent to fitting a Bradley–Terry model over policy likelihood ratios (Rafailov et al., 2023). Owing to its simplicity and stability, DPO has become a standard alignment baseline.

DPO variants. Numerous extensions modify DPO through alternative feedback forms, regularization schemes, or refer-

ence models (Azar et al., 2024; Ethayarajh et al., 2024; Xu et al., 2024; Hong et al., 2024; Meng et al., 2024; Wu et al., 2024). Recent work also targets more specific failure modes: CDPO applies causal adjustment to reduce confounding in preference data, while logit-space SAM stabilizes DPO by applying sharpness-aware perturbations directly to policy logits (Le et al., 2026; Luo et al., 2026). While improving optimization dynamics or robustness, these methods largely preserve the *sequence-level* preference modeling assumption. Recent surveys summarize this rapidly expanding design space (Xiao et al., 2024; Liu et al., 2025b).

Token-level methods. To better regulate optimization dynamics in autoregressive models, several token-level variants of DPO have been proposed. TDPO decomposes the sequence-level DPO objective into token-wise terms with per-token KL constraints, while TIS-DPO further incorporates token-level importance sampling to improve training stability (Zeng et al., 2024; Liu et al., 2025a). However, both methods still assume a *sequence-level* Bradley–Terry preference model, redistributing sequence-level supervision across timesteps rather than explicitly defining token-level preferences. As a result, per-state optimality remains implicit rather than being directly enforced. Beyond direct preference optimization, SWIFT assigns teacher-derived token importance weights during self-play fine-tuning (Le et al.), while CTPD transfers preference-aligned behavior across heterogeneous tokenizers by adapting token-level importance sampling (Nguyen et al., 2026a). These methods use token-granular signals for weighting or teacher–student transfer, rather than deriving a per-prefix preference-optimal policy from pairwise comparisons. More broadly, granularity-sensitive alignment has also been explored in teacher–student transfer: recent LLM distillation work aligns distributions, layer trajectories, or tokenizer-agnostic spans (Vuong et al., 2026; Chi et al., 2026; Dao et al., 2026), while embedding distillation aligns relations and layers in compact representation spaces (Truong et al., 2025; Vu et al., 2026; Huy et al., 2026). Related topic-modeling work similarly shows that the choice of semantic representation space, including encoder and LLM-derived representations, affects topic quality and consistency (Nguyen et al., 2025; Phat et al., 2026; Xuan et al., 2026). We view this literature as complementary evidence that the unit of representation and alignment matters; TBPO studies this issue in preference learning by moving from sequence-level ratios to prefix- and token-level ratios.

Density ratio matching. Another line of work interprets preference optimization through *density ratio matching*. Bregman Preference Optimization (BPO) formalizes this view by framing preference learning as likelihood-ratio estimation under Bregman divergences, connecting DPO to classical density-ratio estimators (e.g., logistic regression (Gutmann & Hyvärinen, 2012), KLIEP (Sugiyama et al., 2007),

LSIF (Huang et al., 2006)) and showing that DPO is a special case of this framework (Kim et al., 2025). BPO further introduces a scaled Basu's power divergence to interpolate between KLIEP and LSIF while modulating gradient reweighting during optimization. In the safety-constrained setting, BSO derives safety alignment as density-ratio matching and introduces a Bregman family of single-stage safety objectives (Nguyen et al., 2026b). Our focus is complementary: unlike these ratio-matching formulations, TBPO explicitly models token-level preferences and targets per-prefix optimality.

Our contribution. Motivated by these gaps, we study **token-level preference modeling**. By defining preferences via a token-level Bradley–Terry model and adopting a ratio-matching perspective, we learn preference-optimal decisions at each prefix using only sequence-level comparison data.

3. Background

3.1. Preliminary and Notions

When viewing text generation as a Markov decision process (Puterman, 1994), we define the state at step t as the prompt together with the response prefix produced so far, i.e., $s_t = [x, y_{<t}]$. The action is the next token to generate, $a_t = y_t$, and the per-token reward is given by $R_t := R(s_t, a_t) = R([x, y_{<t}], y_t)$. Using these definitions, for a policy π we define the state-action value function Q^π, the state value function V^π and the advantage function A^π as:

$$Q^\pi([x, y_{<t}], y_t) = \mathbb{E}_\pi \left[\sum_{k=0}^\infty \gamma^k R_{t+k} \mid s_t = [x, y_{<t}], a_t = y_t \right],$$
$$V^\pi([x, y_{<t}]) = \mathbb{E}_\pi \left[Q^\pi([x, y_{<t}], y_t) \mid s_t = [x, y_{<t}] \right],$$
$$A^\pi([x, y_{<t}], y_t) = Q^\pi([x, y_{<t}], y_t) - V^\pi([x, y_{<t}]).$$

where γ denotes the discount factor. We set $\gamma = 1$ throughout this work.

Following TDPO (Zeng et al., 2024), we cast next-token generation as a sequential decision problem and optimize a KL-regularized objective at the token level:

$$\max_{\pi_\theta} \mathbb{E}_{x, y_{<t} \sim \mathcal{D}, z \sim \pi_\theta(\cdot|[x, y_{<t}])} \Big[A_{\pi_{\text{ref}}}([x, y_{<t}], z)$$
$$- \beta \, D_{\text{KL}} \Big(\pi_\theta(\cdot \mid [x, y_{<t}]) \| \pi_{\text{ref}}(\cdot \mid [x, y_{<t}]) \Big) \Big]. \quad (1)$$

Intuitively, each prefix state $[x, y_{<t}]$ defines a local decision over the next token. The advantage term encourages increasing the probability of tokens that are estimated (under the reference) to be locally better, while the KL penalty keeps the update conservative by anchoring π_θ to π_{ref}. This is aligned with our goal of **token-level preference optimality**: the optimization objective explicitly targets preference-improving decisions at the state–action level, rather than

only redistributing a sequence-level preference signal across timesteps. We later prove (Theorem 4.3) that, under sufficient model capacity, our token-level density-ratio matching objective recovers the optimal policy for Eq. 1.

Using the objective in Eq. (1), TDPO (Zeng et al., 2024) derives closed-form expressions for the optimal state-action Q and advantage functions under the reference policy.

$$Q_{\pi_{\text{ref}}}([x, y_{<t}], z) =$$
$$\beta \log \frac{\pi_\theta^*(z \mid [x, y_{<t}])}{\pi_{\text{ref}}(z \mid [x, y_{<t}])} + \beta \log Z([x, y_{<t}]; \beta), \quad (2)$$

$$A_{\pi_{\text{ref}}}([x, y_{<t}], z) = \beta \log \frac{\pi_\theta^*(z \mid [x, y_{<t}])}{\pi_{\text{ref}}(z \mid [x, y_{<t}])}$$
$$+ \beta D_{\text{KL}} \Big(\pi_{\text{ref}}(\cdot \mid [x, y_{<t}]) \| \pi_\theta^*(\cdot \mid [x, y_{<t}]) \Big). \quad (3)$$

where $Z([x, y_{<t}]; \beta) = \mathbb{E}_{z \sim \pi_{\text{ref}}(\cdot|[x, y_{<t}])} e^{\frac{1}{\beta} Q_{\pi_{\text{ref}}}([x, y_{<t}], z)}$ is the partition function.

3.2. Density Ratio Matching

Consider two probability distributions $p_{\text{de}}(x)$ and $p_{\text{nu}}(x)$. The objective of likelihood ratio estimation is to construct a parametric model $R_\theta(x)$ that closely approximates the true ratio $R_{\text{data}}(x) := \frac{p_{\text{nu}}(x)}{p_{\text{de}}(x)}$ using independent and identically distributed samples drawn from each distribution. The predominant technique for this task is probabilistic classification through logistic regression (Gutmann & Hyvärinen, 2012). Beyond this, classical approaches including the Kullback-Leibler importance estimation procedure (KLIEP) (Sugiyama et al., 2007) and least-squares importance fitting (LSIF) (Huang et al., 2006) have gained substantial adoption in practice.

A key insight from (Sugiyama et al., 2012) is that these diverse ratio estimation techniques can be understood within a unified framework based on Bregman divergence (Bregman, 1967). This unification leads to the following general objective:

$$D_h(R_{\text{data}}(x) \| R_\theta(x)) = \int p_{\text{de}}(x) \, B_h(R_{\text{data}}(x) \| R_\theta(x)) \, dx.$$

which can be expanded as:

$$D_h(R_{\text{data}}(x) \| R_\theta(x)) = \int p_{de}(x) \Big(h(R_{\text{data}}(x))$$
$$- h(R_\theta(x)) - h'(R_\theta(x)) \left(R_{\text{data}}(x) - R_\theta(x) \right) \Big) \, \mathrm{d}x. \quad (4)$$

where h represents a strictly convex and twice continuously differentiable function whose derivative is denoted by h'. The term B_h corresponds to the pointwise Bregman divergence, which quantifies the approximation error of the first-order Taylor expansion.

Figure 1. Overview comparison between our proposed TBPO and Sequence-level DPO.

4. Methodology

A central modeling choice in preference optimization is the granularity at which preferences are defined. While standard DPO-style derivations posit a **sequence-level** Bradley–Terry model, generation and learning in LMs naturally operate at the **token level**. Our objective is therefore to formulate a token-level preference model, which enables density ratio matching at the token level and culminates in the framework **Token-level Bregman Preference Optimization (TBPO)**. We present in Figure 1 an overview of the proposed method.

4.1. Token-level Bradley–Terry model

For a preference triple (x, y^w, y^l) and token step t, the two responses induce different prefixes

$$s_t^w = [x, y_{<t}^w], \qquad s_t^l = [x, y_{<t}^l], \qquad (5)$$

and the compared actions are $a_t^w = y_t^w$ and $a_t^l = y_t^l$. We posit a token-level Bradley-Terry likelihood

$$p_{\text{data}}(a_t^w \succ a_t^l \mid s_t^w, s_t^l) = \sigma\Big(S(s_t^w, a_t^w) - S(s_t^l, a_t^l)\Big), \ (6)$$

where $S(s, a)$ is a token-level score and the sigmoid function $\sigma(u) = 1/(1 + \exp(-u))$. Note that we call this preference probability as p_{data} because this is computed based on the preference data $y^w \succ y^l$ and we choose an appropriate score function to recover preference samples from data. Note that in DPO, $p_{\text{data}}(y^w \succ p^l) := \sigma(r^*(x, y^w) - r^*(x, y^l))$.

The key idea is to lift a sequence-level preference label $y^w \succ y^l$ to supervision at each generation step. Conditioned on the prompt x and the partial responses generated so far, $y_{<t}^w$ and $y_{<t}^l$, we prefer generating the next token from the winning trajectory, y_t^w, over the next token from the losing trajectory, y_t^l. This provides a preference signal throughout generation and encourages the extended prefixes $y_{\leq t}^w$ and $y_{\leq t}^l$ to separate in quality, as measured by a token-level score S.

Different choices of the score function S yield different instantiations of the token-level Bradley–Terry model. We consider two natural options: the state-action value function Q and the advantage function A, giving two variants: (1) TBPO-Q and (2) TBPO-A.

4.2. TBPO-Q

Here we instantiate the token-level Bradley–Terry model by using the reference-policy state–action value as the score, i.e., $S(s, a) = Q_{\pi_{\text{ref}}}(s, a)$:

$$
\begin{aligned}
& p_{\text{data}}(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l) \\
& = \sigma\big(Q_{\pi_{\text{ref}}}([x, y_{<t}^w], y_t^w) - Q_{\pi_{\text{ref}}}([x, y_{<t}^l], y_t^l)\big)
\end{aligned} \tag{7}
$$

Intuition. Recall that $Q_{\pi_{\text{ref}}}([x, y_{<t}], y_t)$ is the expected cumulative reward obtained by emitting token y_t at prefix $[x, y_{<t}]$ and then rolling out π_{ref} for the remainder of the sequence. Eq. (7) therefore treats each preference label as a noisy comparison of two one-step lookahead utilities: taking y_t^w after the winning prefix versus taking y_t^l after the losing prefix. Equivalently, the Bradley–Terry log-odds are proportional to the difference in these expected returns, so maximizing this likelihood pushes probability mass toward tokens whose downstream continuations are expected to yield higher reward.

We can establish the following equality that links the optimal policy to the reference policy, thereby later encouraging token-level density–ratio matching.

Proposition 4.1. *Let π_{θ^*} denote the optimal policy for the optimization problem in Eq. (1). We have the following equality:*

$$
\begin{aligned}
\frac{\pi_{\theta^*}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\theta^*}(y_t^l \mid [x, y_{<t}^l])} &= \frac{\pi_{\text{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\text{ref}}(y_t^l \mid [x, y_{<t}^l])} \\
&\times \left(\frac{p_{data}(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p_{data}(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}\right)^{\frac{1}{\beta}} w_t,
\end{aligned}
$$
$$(8)$$

with per state weight $w_t = \exp\left(\log \frac{Z([x,y^l_{<t}];\beta)}{Z([x,y^w_{<t}];\beta)}\right)$.

See proof in Appendix A.1.

4.3. TBPO-A

In this variants, we define a token-level BT preference model in terms of the reference-policy advantage.

$$
\begin{aligned}
&p_{\text{data}}\big(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t}\big) \\
&= \sigma\Big(A_{\pi_{\text{ref}}}\big([x,y^w_{<t}], y^w_t\big) - A_{\pi_{\text{ref}}}\big([x, y^l_{<t}], y^l_t\big)\Big)
\end{aligned}
\tag{9}
$$

This likelihood assigns higher probability to the token with larger advantage. Optimizing Eq. (1) therefore encourages the learned policy to increase the advantage gap between the winning and losing extensions $y^w_{\le t}$ and $y^l_{\le t}$, ensure generation quality at every state.

Similarly, we can establish the following equality that links the optimal policy to the reference policy, thereby encouraging token-level density–ratio matching.

Proposition 4.2. *Let $\pi_{\theta*}$ denote the optimal policy for the optimization problem in Eq. (1). We have the following equality:*

$$
\begin{aligned}
\frac{\pi_{\theta*}(y^w_t \mid [x, y^w_{<t}])}{\pi_{\theta*}(y^l_t \mid [x, y^l_{<t}])} &= \frac{\pi_{\text{ref}}(y^w_t \mid [x, y^w_{<t}])}{\pi_{\text{ref}}(y^l_t \mid [x, y^l_{<t}])} \\
&\times \left(\frac{p_{data}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})}{p_{data}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t})}\right)^{\frac{1}{\beta}} w_t,
\end{aligned}
\tag{10}
$$

where we have defined per state weight

$$
\begin{aligned}
w_t = \exp\Big(&D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y^l_{<t}]) \,\|\, \pi^*_\theta(\cdot \mid [x, y^l_{<t}])\big) \\
&- D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y^w_{<t}]) \,\|\, \pi^*_\theta(\cdot \mid [x, y^w_{<t}])\big)\Big).
\end{aligned}
$$

See proof in Appendix A.2.

4.4. Interpretation and estimation of per state weights

In both Proposition 4.1 (TBPO-Q) and Proposition 4.2 (TBPO-A), an extra multiplicative term w_t appears in the policy ratio. This term is *state-dependent but action-independent*: it does not depend on the compared next tokens y^w_t, y^l_t directly, but only on the prefixes $[x, y^w_{<t}]$ and $[x, y^l_{<t}]$. Intuitively, w_t corrects for the fact that we compare actions taken in *two different states* (different prefixes), so state-only baseline terms that would cancel in a same-state comparison no longer cancel.

TBPO-Q: ratio of partition functions. The TBPO-Q weight is

$$
w^{(Q)}_t = \exp\left(\log \frac{Z([x, y^l_{<t}]; \beta)}{Z([x, y^w_{<t}]; \beta)}\right).
$$

therefore measures the *relative continuation value* of the losing prefix versus the winning prefix. If the winning prefix is already a higher-value state (larger $Z([x, y^w_{<t}]; \beta)$, then $w^{(Q)}_t < 1$, which downweights the token ratio update so that we do not over-attribute preference to the single next-token decision when the prefixes already differ substantially in quality.

TBPO-A: difference of KL baselines. The TBPO-A weight is

$$
\begin{aligned}
w^{(A)}_t = \exp\Big(&D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y^l_{<t}]) \,\|\, \pi^*_\theta(\cdot \mid [x, y^l_{<t}])\big) \\
&- D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y^w_{<t}]) \,\|\, \pi^*_\theta(\cdot \mid [x, y^w_{<t}])\big)\Big),
\end{aligned}
$$

which can be interpreted as a *relative policy-improvement cost* between the two prefixes: it compares how far the optimal policy must deviate from the reference at $[x, y^l_{<t}]$ versus $[x, y^w_{<t}]$ under the KL regularizer.

4.4.1. PRACTICAL ESTIMATION.

In practice $Z([x, y_{<t}]; \beta)$ and D_{KL} are either computationally intractable or expensive. We use plug-in estimators that preserve the key property that w_t is state-only.

TBPO-Q (learned baseline). We parameterize a scalar baseline head $b_\phi([x, y_{<t}])$ to predict $\log Z([x, y_{<t}]; \beta)$ and set

$$
\begin{aligned}
\widehat{\log w^{(Q)}_t} &= b_\phi([x, y^l_{<t}]) - b_\phi([x, y^w_{<t}]), \\
\hat{w}^{(Q)}_t &= \exp\left(\widehat{\log w^{(Q)}_t}\right).
\end{aligned}
$$

In practice, we attach a lightweight MLP head on top of the LM's final hidden layer to predict $Z([x, y_{<t}]; \beta)$. We train this head during optimization while stopping gradients into the LM backbone, so the LLM does not allocate capacity to baseline estimation at the expense of preference optimization; see Appendix D.3.

TBPO-A (K3 KL estimator). We approximate π^* by the current policy π_θ and estimate $D_{\text{KL}}(\pi_{\text{ref}}(\cdot \mid s) \,\|\, \pi_\theta(\cdot \mid s))$ using the low-variance unbiased estimator (Schulman, 2020). See more in the Appendix D.3.

TBPO-Q vs. TBPO-A. Both variants require a state-only correction because token comparisons are made across two different prefixes. TBPO-Q's w_t is a *partition-function / soft-value* ratio and is naturally handled by a learned baseline head. TBPO-A removes the partition function analytically but introduces a *KL-baseline* term, which can be estimated directly from π_{ref} and π_θ using an unbiased estimator.

4.5. Token-level Bregman ratio matching

We now use density-ratio matching to derive a tractable optimization objective from the token-level BT preference

Method	HellaSwag	ARC	MMLU	TruthfulQA	Winogrande	GSM8K	Average
Mistral 7B v0.1							
SFT	$80.72_{\pm 0.39}$	$55.54_{\pm 1.45}$	$58.41_{\pm 0.39}$	$43.67_{\pm 1.46}$	$76.71_{\pm 1.18}$	$18.49_{\pm 1.07}$	55.59
DPO	$82.39_{\pm 0.38}$	$59.40_{\pm 1.43}$	$59.15_{\pm 0.38}$	$42.30_{\pm 1.49}$	$77.74_{\pm 1.16}$	$34.87_{\pm 1.21}$	59.30
TDPO	$82.70_{\pm 0.38}$	$\mathbf{61.54}_{\pm 1.41}$	$57.17_{\pm 0.40}$	$43.48_{\pm 1.53}$	$76.95_{\pm 1.18}$	$28.65_{\pm 1.24}$	58.41
TIS-DPO	$\mathbf{83.10}_{\pm 0.37}$	$59.09_{\pm 1.42}$	$57.81_{\pm 0.39}$	$\mathbf{45.91}_{\pm 1.56}$	$77.66_{\pm 1.17}$	$34.42_{\pm 1.31}$	59.66
BPO	$81.03_{\pm 0.39}$	$57.67_{\pm 1.44}$	$59.16_{\pm 0.40}$	$41.71_{\pm 1.46}$	$77.26_{\pm 1.16}$	$32.90_{\pm 1.29}$	58.29
TBPO-Q (Ours)	$82.72_{\pm 0.37}$	$\underline{59.74}_{\pm 1.41}$	$\mathbf{60.82}_{\pm 0.39}$	$\underline{44.47}_{\pm 1.45}$	$\mathbf{78.58}_{\pm 1.11}$	$\mathbf{39.34}_{\pm 1.34}$	$\mathbf{60.95}$
TBPO-A (Ours)	$\underline{82.74}_{\pm 0.38}$	$59.47_{\pm 1.43}$	$60.54_{\pm 0.39}$	$44.40_{\pm 1.46}$	$\underline{78.21}_{\pm 1.21}$	$\underline{39.04}_{\pm 1.35}$	$\underline{60.73}$
Llama3 8B							
SFT	$79.46_{\pm 0.40}$	$51.45_{\pm 1.46}$	$61.19_{\pm 0.39}$	$46.46_{\pm 1.45}$	$75.69_{\pm 1.21}$	$68.67_{\pm 1.37}$	63.82
DPO	$81.23_{\pm 0.38}$	$56.05_{\pm 1.45}$	$61.70_{\pm 0.38}$	$48.47_{\pm 1.53}$	$75.84_{\pm 1.20}$	$74.20_{\pm 1.37}$	66.24
TDPO	$\mathbf{83.29}_{\pm 0.37}$	$59.04_{\pm 1.43}$	$61.86_{\pm 0.39}$	$51.76_{\pm 1.55}$	$76.55_{\pm 1.19}$	$75.10_{\pm 1.36}$	67.91
TIS-DPO	$81.37_{\pm 0.38}$	$58.70_{\pm 1.42}$	$60.77_{\pm 0.39}$	$49.87_{\pm 1.57}$	$75.14_{\pm 1.21}$	$75.12_{\pm 1.37}$	66.82
BPO	$81.68_{\pm 0.37}$	$54.95_{\pm 1.45}$	$61.48_{\pm 0.38}$	$47.88_{\pm 1.52}$	$75.13_{\pm 1.21}$	$73.22_{\pm 1.37}$	65.72
TBPO-Q (Ours)	$81.90_{\pm 0.38}$	$\underline{64.07}_{\pm 1.40}$	$\mathbf{64.92}_{\pm 0.38}$	$\mathbf{53.36}_{\pm 1.52}$	$79.08_{\pm 1.14}$	$78.92_{\pm 1.12}$	$\underline{70.38}$
TBPO-A (Ours)	$\underline{81.97}_{\pm 0.38}$	$\mathbf{64.08}_{\pm 1.40}$	$\underline{64.84}_{\pm 0.39}$	$\underline{53.29}_{\pm 1.53}$	$\mathbf{79.63}_{\pm 1.12}$	$\mathbf{79.22}_{\pm 1.21}$	$\mathbf{70.50}$

Table 1. General capabilities on the Open LLM Leaderboard suite for Mistral 7B v0.1 and Llama 3 8B. Scores are reported by `lm-eval-harness` (higher is better) with $\pm$ standard error. Best is in **bold**, second-best is underlined.

model. Given a time-step t, we could naturally define the BT preference model for the sequences up to the t-th token:

$$p_{\text{data}}\left(y_{\leq t}^w \succ y_{\leq t}^l \mid x\right) = \prod_{i=1}^{t} p_{\text{data}}\left(y_i^w \succ y_i^l \mid x, y_{<i}^w, y_{<i}^l\right). \quad (11)$$

It is evident that we have the recursive formula:

$$p_{\text{data}}\left(y_{\leq t}^w \succ y_{\leq t}^l \mid x\right) = p_{\text{data}}\left(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\right)$$
$$p_{\text{data}}\left(y_{\leq t-1}^w \succ y_{\leq t-1}^l \mid x\right). \quad (12)$$

Preference optimization can be reformulated as a matching problem between the empirical data ratio $R_{\text{data}}^t(y_t^w, y_t^l \mid y_{<t}^w, y_{<t}^l, x)$ and the model ratio $R_\theta^t(y_t^w, y_t^l \mid y_{<t}^w, y_{<t}^l, x)$, where we have defined

$$R_{\text{data}}^t = \frac{p_{\text{data}}(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p_{\text{data}}(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)},$$

$$R_\theta^t = \left(\frac{\pi_\theta\left(y_t^l \mid [x, y_{<t}^l]\right) \pi_{\text{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\text{ref}}\left(y_t^l \mid [x, y_{<t}^l]\right) \pi_\theta(y_t^w \mid [x, y_{<t}^w])} w_t\right)^\beta. \quad (13)$$

We define the t-th time step loss using the Bregman ratio matching framework in Section 3.2 as:

$$D_h^t\left(R_{\text{data}}^t, R_\theta^t\right) =$$
$$\mathbb{E}_{p_{\text{data}}\left(y_{\leq t}^w \succ y_{\leq t}^l \mid x\right)}\left[h\left(R_{\text{data}}^t\right) - h\left(R_\theta^t\right) - h'\left(R_\theta^t\right)\left(R_{\text{data}}^t - R_\theta^t\right)\right].$$

The final loss is the average of the component losses:

$$D_h\left(R_{\text{data}}, R_\theta\right) = \frac{1}{T}\sum_{t=1}^{T} D_h^t\left(R_{\text{data}}^t, R_\theta^t\right), \quad (14)$$

where $T = \min(|y^w|, |y^l|)$ is the length of the shorter response among the chosen and rejected sequences.

In the following theorem, we show that under the sufficient model capacity, by matching the density ratio at the token level, we can recover the optimal policy in the Objective 1.

Theorem 4.3. *Under sufficient model capacity, we have* $\pi_{\theta^*} = \pi_{\hat{\theta}}$ *where* $\pi_{\hat{\theta}} = argmin_{\pi_\theta} D_h\left(R_{data}, R_\theta\right)$.

The loss function $D_h\left(R_{\text{data}}, R_\theta\right)$ is not tractable to train a policy model. We turn it to an equivalent tractable objective function in the following theorem.

Theorem 4.4. *We have* $\mathcal{L}\left(\theta, p_{data}\right) = D_h\left(R_{data}, R_\theta\right) + const$ *where we define*

$$\mathcal{L}\left(\theta, p_{data}\right) = \mathbb{E}_{p_{data}\left(y^w \succ y^l \mid x\right)}\left[\frac{1}{T}\sum_{t=1}^{T}\left(h'\left(R_\theta^t\right) R_\theta^t\right.\right.$$
$$\left.\left. - h\left(R_\theta^t\right) - h'\left(\left(R_\theta^t\right)^{-1}\right)\right)\right] \quad (15)$$

The loss in Eq. (15) depends only on the ratio R_θ^t, hence applies identically to TBPO-Q and TBPO-A. Different ratio estimation methods differ only in their choice of function h. Following BPO (Kim et al., 2025), we employ h function as Scaled Basu's power divergence (SBA), with following formula $h(R) = \frac{R^{1+\lambda} - R}{s\lambda(\lambda+1)}$ for better training flexibility and tuning flexibility. The λ parameter lets you control whether to focus on confident samples (low R_θ) or uncertain ones ($R_\theta \approx 1$), which can be useful for handling noisy preference data.

4.6. Overall principle of our approach

The standard Bradley–Terry (BT) model encourages the entire sequence $y^w \succ y^l$. In contrast, we propose a token-level BT model that encourages progressively better substrings $y^w_{\leq t} \succ y^l_{\leq t}$ as t varies according to the recursive definition in Eq. (12). Our goal is not to treat generation only as a whole-sequence preference problem, but to encourage the model to form good partial reasoning early and maintain it throughout generation. In this way, stronger intermediate prefixes can support further high-quality continuation, which helps address credit assignment in a more fine-grained way and can lead to more concise, shorter, and more precise responses.

From this perspective, sequence-level preference labels are not merely copied to every timestep as hard tokenwise win/loss annotations. Instead, they induce a sequence of prefix-conditioned local comparisons, which is exactly the structure captured by our token-level BT formulation. This is the key difference from prior approaches whose underlying preference model remains sequence-level and mainly redistributes the same global signal across positions. Our formulation instead makes the preference modeling itself token-level, so that better local decisions are encouraged progressively during autoregressive generation.

We also note that a naive decomposition could overemphasize locally plausible but globally suboptimal decisions. TBPO is designed to mitigate exactly this issue. First, supervision is defined over progressively extended prefixes, not isolated tokens, which preserves sequential dependency. Second, the state-dependent correction weight reduces mismatch when comparing next-token decisions made under different prefix contexts, preventing the model from attributing the entire sequence-level preference gap to a single step. In this sense, TBPO should be viewed as a principled subsequence-level surrogate derived from sequence-level preferences, together with an explicit mechanism to reduce the local–global mismatch, rather than as an assumption that improving local token decisions is always sufficient by itself.

5. Experiments

In this section, we present an empirical study of Token-level Bregman Preference Optimization (TBPO). We compare TBPO-Q and TBPO-A to existing preference-optimization methods, with particular attention to alignment quality, training stability, and generation diversity, evaluating whether the theoretical advantages of token-level ratio matching manifest in practice.

5.1. Experiments settings

Backbones and training data. We conduct experiments with two backbone models. Each setup initializes from a publicly available supervised fine-tuning (SFT) checkpoint, followed by preference optimization on the corresponding dataset. Specifically, Mistral 7B v0.1 is initialized from mistral-7b-sft-alpha and trained on UltraFeedback Binarized, while Llama 3 8B is initialized from LLaMA3-SFT-v2 and trained on Llama3-UltraFeedback-ArmoRM.

Baselines. We compare TBPO-Q and TBPO-A against various prior preference-optimization baselines, including SFT, DPO (Rafailov et al., 2023), TDPO (Zeng et al., 2024), TIS-DPO (Liu et al., 2025a), and BPO-SBA (Kim et al., 2025). See more details in Appendix D.2.

LLM judges. Some benchmarks, including Anthropic HH-RLHF, TL;DR, and MT-Bench, require pairwise evaluation with an LLM-as-a-judge framework. Prior work has analyzed the reliability and biases of LLM judges (Zhang et al., 2025; Han et al., 2025); following these findings, we adopt task-specific judge pairs. For Anthropic HH-RLHF and TL;DR, we use Llama 3 70B Instruct and DeepSeek-V3, while for MT-Bench we use Llama 3 70B Instruct and Mixtral-8x22B Instruct v0.1. Additional details are provided in Appendix D.

5.2. General Capabilities

A common concern in preference optimization is over-specialization to preference data, which may harm general capabilities. We evaluate general performance using the **HuggingFace Open LLM Leaderboard** protocol (Hugging Face, 2024) with the Language Model Evaluation Harness (Sutawika et al.) on six benchmarks spanning commonsense reasoning (ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), Winogrande (Sakaguchi et al., 2019)), multi-task understanding (MMLU (Hendrycks et al., 2021)), factuality (TruthfulQA (Lin et al., 2022)), and math reasoning (GSM8k (Cobbe et al., 2021)). Table 1 shows that TBPO improves the overall average across both backbones while remaining competitive on all benchmarks, indicating stronger alignment without sacrificing core capabilities.

On **Mistral 7B v0.1**, sequence-level preference optimization already improves over SFT (DPO: 59.30 vs. SFT: 55.59). TBPO further improves the average: **TBPO-Q is best** (**60.95**) with TBPO-A close behind (60.73), exceeding the strongest non-TBPO baseline (TIS-DPO: 59.66). Gains are broad and strongest on reasoning-heavy tasks: TBPO-Q attains the best GSM8K (39.34 vs. 34.87 for DPO and 18.49 for SFT) and also the best MMLU (60.82) and Winogrande (78.58). While DPO slightly reduces TruthfulQA vs. SFT (42.30 vs. 43.67), TBPO restores and improves it (44.47–44.40). Where other token-level methods peak (e.g., ARC for TDPO, TruthfulQA for TIS-DPO), TBPO remains highly competitive.

On **Llama 3 8B**, **TBPO-A is best on average** (**70.50**) with

Figure 2. MT-Bench pairwise win/tie/lose rates for TBPO-Q (top) and TBPO-A (bottom) against prior preference-optimization baselines, evaluated by two LLM judges. TBPO achieves higher win rates with low loss rates across both judges, and the advantage persists even against the strongest baseline.

TBPO-Q close behind (70.38), improving substantially over TDPO (67.91). Improvements concentrate on ARC (+5 vs. TDPO; 64.08 vs. 59.04), GSM8K (+4; 79.22 vs. 75.10), and Winogrande (+3; 79.63 vs. 76.55). TBPO-A leads on ARC/GSM8K/Winogrande, while TBPO-Q attains the best MMLU (64.92) and TruthfulQA (53.36), showing robustness to the preference-model formulation.

The similar performance of TBPO-Q and TBPO-A is expected from the theory: the two variants differ mainly in how they approximate the same state-only correction term w_t, rather than in the core token-level preference objective. TBPO-Q represents w_t as a partition-function or soft-value ratio and estimates it with a learned baseline head, whereas TBPO-A analytically removes the partition function and replaces it with a KL-based state baseline estimated directly from π_{ref} and π_θ. Since both estimators target the same prefix-level correction, their optimization targets remain nearly identical when the estimators are stable. The near ties in Table 1 therefore suggest that the main gain comes from the shared token-level ratio-matching formulation, while the particular estimator used for w_t is a secondary implementation choice.

Overall, TBPO delivers the strongest mean performance across two model families, with large gains on reasoning-centric benchmarks and no evidence of systematic trade-offs on broader language understanding. To better understand these reasoning gains, we provide qualitative case studies on reasoning questions in Appendix C.

5.3. Generation Quality

MT-Bench. To evaluate generation quality, we run MT-Bench (Zheng et al., 2023) under the standard LLM-as-a-judge protocol and report pairwise win/tie/lose rates against prior preference-optimization baselines. Figure 2 shows that TBPO-Q and TBPO-A consistently improve win rates with low loss rates; against the strongest baseline (TDPO), many comparisons are ties, but TBPO wins the majority of decisive (non-tie) cases. This suggesting the advantage is not driven by a small number of outliers. The overall ordering and effect sizes are consistent across both judges, supporting the conclusion that TBPO improves generation quality in a robust, judge-agnostic manner, and that TBPO-Q and TBPO-A are comparably strong.

Harmless and Helpfulness. To evaluate preference alignment in safety-sensitive dialogue, we test on 200 held-out samples from the Anthropic HH-RLHF dataset. Using an LLM-as-a-judge protocol, we report pairwise and length-controlled (LC) win rates (Li et al., 2023). Figure 3 shows LC win rate versus average response length, highlighting the verbosity–quality trade-off.

Across both judges, TBPO variants achieve the strongest preference alignment while producing shorter responses than competing methods, ruling out verbosity as the primary driver of higher win rates. SFT is the most verbose model yet does not achieve top performance, whereas TBPO attains the best LC win rates with substantially shorter generations. Among TBPO variants, TBPO-Q produces the most concise responses while remaining at the top of the

Anthropic HH - Llama 3 8B

Figure 3. LC win rate vs. average response length against the dataset-preferred completion for Llama 3 8B, evaluated by two LLM judges (Llama 3 70B, DeepSeek-V3); error bars are ±1 s.e. over 200 prompts. TBPO leads with shorter outputs, indicating gains beyond verbosity and consistent across judges.

Figure 4. Generation diversity trade-offs: Distinct-1 vs. predictive entropy (higher is better), colored by self-BLEU (lower is better). TBPO achieves the best three-way trade-off across metrics.

win-rate plot, and TBPO-A achieves the highest (or near-highest) win rates with only a modest increase in length, indicating a controllable conciseness–performance trade-off. In contrast, strong sequence- and token-level baselines improve over SFT but generally require longer responses and still trail TBPO. The consistent ordering across judges and large margins relative to standard errors suggest TBPO's gains reflect genuine improvements in helpful-and-harmless response quality rather than judge bias or verbosity effects.

Generation Diversity. Preference optimization can reduce generation diversity (Kirk et al., 2024). We evaluate diversity on 100 held-out prompts from the test split of UltraFeedback Binarized using predictive entropy (Zeng et al., 2024), Self-BLEU (Zhu et al., 2018), and Distinct-1 (Li et al., 2016). Figure 4 summarizes the trade-off, where higher Distinct-1 and predictive entropy indicate greater

diversity, while lower Self-BLEU indicates less redundancy.

Two patterns emerge. (1) DPO/TDPO/BPO have moderate Distinct-1 but lower predictive entropy and higher self-BLEU, implying samples become more similar and decoding more deterministic (mild mode-seeking). TISDPO amplifies this, achieving the lowest entropy and highest self-BLEU, i.e., the strongest collapse even when Distinct-1 alone does not fully reveal it. (2) TBPO shifts the Pareto frontier toward the desirable upper-right region while keeping self-BLEU favorable, giving the best three-metric trade-off: versus DPO-style methods it raises entropy and Distinct-1 without the usual across-sample diversity loss; versus SFT (high entropy, low self-BLEU, lower Distinct-1) it preserves non-collapsed sampling while increasing lexical variety. Overall, TBPO better balances preference satisfaction with diversity preservation, producing outputs that are less redundant across samples and less over-confident at the token level.

Additional experiments on out-of-domain tasks, ability to handle noisy data, ablation on per-state weight estimation and judge agreement analysis for TBPO are reported in Appendix B.

6. Conclusion

We presented a principled token-level preference optimization framework aligned with the autoregressive nature of language models. We show theoretically that minimizing the proposed objective recovers a token-wise optimal policy, establishing a direct link between preference modeling and optimization behavior. Experiments across benchmarks and model families demonstrate consistent improvements over strong baselines, particularly in challenging reasoning and alignment settings.

Acknowledgements

Trung Le was supported by the Air Force Office of Scientific Research under award number FA2386-25-1-4023 and the ARC Discovery Project grant DP250100262. Linh Ngo Van is funded by Vietnam National Foundation for Science and Technology Development (NAFOSTED) under grant number 102.05-2025.16.

Impact Statement

This work focuses on the development of a principled learning algorithm with theoretical guarantees. The primary impact of this research is methodological, contributing to the design of more reliable and well-understood machine learning systems. While the proposed method is evaluated on standard benchmarks, it is not tied to a specific high-risk application domain. We do not foresee any immediate negative societal consequences arising directly from this work; however, as with most advances in machine learning, downstream applications should be deployed responsibly and with appropriate consideration of ethical and societal implications.

References

Azar, M. G., Guo, Z. D., Piot, B., Munos, R., Rowland, M., Valko, M., and Calandriello, D. A general theoretical paradigm to understand learning from human preferences. In Dasgupta, S., Mandt, S., and Li, Y. (eds.), *International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain*, volume 238 of *Proceedings of Machine Learning Research*, pp. 4447–4455. PMLR, 2024. URL https://proceedings.mlr.press/v238/gheshlaghi-azar24a.html.

Bregman, L. M. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *Ussr Computational Mathematics and Mathematical Physics*, 7: 200–217, 1967. URL https://api.semanticscholar.org/CorpusID:121309410.

Chi, P. K., Dao, Q. P., Nguyen, T., Linh, N. V., Le, T., and Nguyen, H. T. Mta: Multi-granular trajectory alignment for large language model distillation. In *Proceedings of the Association for Computational Linguistics (ACL)*, 2026. URL https://arxiv.org/abs/2605.01374.

Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457.

Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., Hesse, C., and Schulman, J. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

Dao, Q. P., Nguyen, H. S., Chi, P. K., Nguyen, T., Linh, N. V., Diep, N. T. N., and Le, T. Sra: Span representation alignment for large language model distillation. In *Proceedings of the Association for Computational Linguistics (ACL)*, 2026. URL https://arxiv.org/abs/2605.01205.

Ethayarajh, K., Xu, W., Muennighoff, N., Jurafsky, D., and Kiela, D. KTO: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024.

Gutmann, M. U. and Hyvärinen, A. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13(11):307–361, 2012. URL http://jmlr.org/papers/v13/gutmann12a.html.

Han, S., Junior, G. T., Balough, T., and Zhou, W. Judge's verdict: A comprehensive analysis of llm judge capability through human agreement, 2025. URL https://arxiv.org/abs/2510.09738.

Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. Measuring massive multitask language understanding, 2021. URL https://arxiv.org/abs/2009.03300.

Hong, J., Lee, N., and Thorne, J. ORPO: monolithic preference optimization without reference model. In Al-Onaizan, Y., Bansal, M., and Chen, Y. (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pp. 11170–11189. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.EMNLP-MAIN.626. URL https://doi.org/10.18653/v1/2024.emnlp-main.626.

Huang, J., Gretton, A., Borgwardt, K., Schölkopf, B., and Smola, A. Correcting sample selection bias by unlabeled data. In Schölkopf, B., Platt, J., and Hoffman, T. (eds.), *Advances in Neural Information Processing Systems*, volume 19. MIT Press, 2006. URL https://proceedings.neurips.cc/paper_files/paper/2006/file/a2186aa7c086b46ad4e8bf81e2a3a19b-Paper.pdf.

Hugging Face. Open LLM leaderboard v1. https://huggingface.co/docs/leaderboards/open_llm_leaderboard/archive, 2024. Leaderboards documentation; archived June 2024; accessed 2026-01-25.

Huy, P. G., Vu, H. A., Truong, M.-P., Thang, T. D., Linh, N. V., Nguyen, H. T., and Le, T. Mipic: Matryoshka representation learning via self-distilled intra-relational and progressive information chaining. In *Findings of the Association for Computational Linguistics (ACL)*, 2026. URL https://arxiv.org/abs/2604.24374.

Kim, Y., Bae, H., Na, B., and Moon, I.-C. Preference optimization by estimating the ratio of the data distribution, 2025. URL https://arxiv.org/abs/2505.19601.

Kirk, R., Mediratta, I., Nalmpantis, C., Luketina, J., Hambro, E., Grefenstette, E., and Raileanu, R. Understanding the effects of RLHF on LLM generalisation and diversity. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=PXD3FAVHJT.

Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.

Le, T., Vuong, H. T., Tran, Q., Van, L. N., Harandi, M., and Le, T. Token-level self-play with importance-aware guidance for large language models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Le, U., Nguyen, T., Nguyen, T., Doan, T., Le, T., and Le, B. Causal direct preference optimization for language model alignment. In *Findings of the Association for Computational Linguistics: EACL 2026*, 2026.

Li, J., Galley, M., Brockett, C., Gao, J., and Dolan, B. A diversity-promoting objective function for neural conversation models, 2016. URL https://arxiv.org/abs/1510.03055.

Li, X., Zhang, T., Dubois, Y., Taori, R., Gulrajani, I., Guestrin, C., Liang, P., and Hashimoto, T. B. Alpacaeval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval, 2023.

Lin, S., Hilton, J., and Evans, O. Truthfulqa: Measuring how models mimic human falsehoods, 2022. URL https://arxiv.org/abs/2109.07958.

Liu, A., Bai, H., Lu, Z., Sun, Y., Kong, X., Wang, X. S., Shan, J., Jose, A. M., Liu, X., Wen, L., Yu, P. S., and Cao, M. TIS-DPO: token-level importance sampling for direct preference optimization with estimated weights. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025a. URL https://openreview.net/forum?id=oF6e2WwxX0.

Liu, S., Fang, W., Hu, Z., Zhang, J., Zhou, Y., Zhang, K., Tu, R., Lin, T., Huang, F., Song, M., Li, Y., and Tao, D. A survey of direct preference optimization. *CoRR*, abs/2503.11701, 2025b. doi: 10.48550/ARXIV.2503.11701. URL https://doi.org/10.48550/arXiv.2503.11701.

Luo, H., Deng, Z., Do, T.-T., Harandi, M., Phung, D., and Le, T. Sharpness-aware minimization in logit space efficiently enhances direct preference optimization, 2026. URL https://arxiv.org/abs/2603.18258.

Meng, C., Choi, K., Song, J., and Ermon, S. Concrete score matching: Generalized score matching for discrete data, 2023. URL https://arxiv.org/abs/2211.00802.

Meng, Y., Xia, M., and Chen, D. Simpo: Simple preference optimization with a reference-free reward. In Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J. M., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/e099c1c9699814af0be873a175361713-Abstract-Conference.html.

Nguyen, T., Dat, P. V., Nguyen, N., Van, L. N., Le, T., and Nguyen, T. H. CTPD: cross tokenizer preference distillation. In *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI*, pp. 37783–37790. AAAI Press, 2026a.

Nguyen, T. P., Minh, N. V., Nguyen, T., Van, L. N., Nguyen, D. A., Sang, D. V., and Le, T. XTRA: Cross-lingual topic modeling with topic and representation alignments. In Christodoulopoulos, C., Chakraborty, T., Rose, C., and Peng, V. (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 5561–5575, Suzhou, China, November 2025. Association for Computational Linguistics. ISBN 979-8-89176-335-7. doi: 10.18653/v1/2025.findings-emnlp.298. URL https://aclanthology.org/2025.findings-emnlp.298/.

Nguyen, T.-P., Nguyen, T., Nguyen, T., Nguyen, D. M. H., Dinh, N.-T., and Le, T. Bso: Safety alignment is density ratio matching, 2026b. URL https://arxiv.org/abs/2605.12339.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.

Phat, N. T., Minh, N. V., Van, L. N., Diep, N. T. N., and Nguyen, T. H. Gloctm: Cross-lingual topic modeling via a global context space. In Koenig, S., Jenkins, C., and Taylor, M. E. (eds.), *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pp. 32710–32718. AAAI Press, 2026. doi: 10.1609/AAAI.V40I39.40549. URL https://doi.org/10.1609/aaai.v40i39.40549.

Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., 1994. ISBN 9780471619772. doi: 10.1002/9780470316887.

Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C. D., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale, 2019. URL https://arxiv.org/abs/1907.10641.

Schulman, J. Approximating kl divergence. https://joschu.net/blog/kl-approx.html, March 2020. Blog post. Accessed: 2026-01-28.

Sugiyama, M., Nakajima, S., Kashima, H., Buenau, P., and Kawanabe, M. Direct importance estimation with model selection and its application to covariate shift adaptation. In Platt, J., Koller, D., Singer, Y., and Roweis, S. (eds.), *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL https://proceedings.neurips.cc/paper_files/paper/2007/file/be83ab3ecd0db773eb2dc1b0a17836a1-Paper.pdf.

Sugiyama, M., Suzuki, T., and Kanamori, T. Density-ratio matching under the bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64:1009–1044, October 2012. doi: 10.1007/s10463-011-0343-8. URL https://doi.org/10.1007/s10463-011-0343-8.

Sutawika, L., Schoelkopf, H., Gao, L., Abbasi, B., Biderman, S., Tow, J., et al. EleutherAI/lm-evaluation-harness. URL https://doi.org/10.5281/zenodo.5371628.

Truong, M.-P., Vu, H. A., Vu, T., Diep, N. T. N., Van, L. N., Nguyen, T. H., and Le, T. Emo: Embedding model distillation via intra-model relation and optimal transport alignments. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 7605–7617, 2025.

Vu, H. A., Truong, M.-P., Vu, T., and Ngo, L. Mol: Mixture of layers in cross-tokenizer embedding model distillation. *Knowledge-Based Systems*, 343:116001, 2026.

Vuong, H. T., Le, T., Tran, Q., Van, L. N., and Le, T. MCW-KD: multi-cost wasserstein knowledge distillation for large language models. In *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI*, pp. 33332–33340. AAAI Press, 2026.

Wu, J., Xie, Y., Yang, Z., Wu, J., Gao, J., Ding, B., Wang, X., and He, X. β-dpo: Direct preference optimization with dynamic β. In Globersons, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J. M., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/ea888178abdb6fc233226d12321d754f-Abstract-Conference.html.

Xiao, W., Wang, Z., Gan, L., Zhao, S., He, W., Tuan, L. A., Chen, L., Jiang, H., Zhao, Z., and Wu, F. A comprehensive survey of datasets, theories, variants, and applications in direct preference optimization. *CoRR*, abs/2410.15595, 2024. doi: 10.48550/ARXIV.2410.15595. URL https://doi.org/10.48550/arXiv.2410.15595.

Xu, H., Sharaf, A., Chen, Y., Tan, W., Shen, L., Durme, B. V., Murray, K., and Kim, Y. J. Contrastive preference optimization: Pushing the boundaries of LLM performance in machine translation. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=51iwkioZpn.

Xuan, M. C., Nguyen, T.-P., Van, L. N., Sang, D. V., Diep, N. T. N., and Le, T. LLM-XTM: Enhancing cross-lingual topic models with large language models, 2026. URL https://arxiv.org/abs/2605.03299.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence?, 2019. URL https://arxiv.org/abs/1905.07830.

Zeng, Y., Liu, G., Ma, W., Yang, N., Zhang, H., and Wang, J. Token-level direct preference optimization, 2024. URL https://arxiv.org/abs/2404.11999.

Zhang, T., Cao, M., Lam, A., Zhang, S., and Chen, K. Compassjudger-2: Towards generalist judge model via verifiable rewards, 2025. URL https://arxiv.org/abs/2507.09104.

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Brooks, D., Xing, E., Gonzalez, J. E., Jordan, M. I., and Stoica, I. Judging LLM-as-a-judge with MT-Bench and chatbot arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J., and Yu, Y. Texygen: A benchmarking platform for text generation models, 2018. URL https://arxiv.org/abs/1802.01886.

A. Theory Development

A.1. Proof for Prop 4.1

Proof. We prove the proposition by deriving the policy ratio induced by a token-level Bradley–Terry model parameterized by the reference-policy Q-function.

By the token-level Bradley–Terry model in Eq. (7), we have

$$p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l) = \sigma\Big(Q_{\pi_{\mathrm{ref}}}([x, y_{<t}^w], y_t^w) - Q_{\pi_{\mathrm{ref}}}([x, y_{<t}^l], y_t^l)\Big).$$

Using $\sigma(u) = (1 + \exp(-u))^{-1}$ and conditioning on the same pair of prefixes, taking the odds ratio yields

$$\frac{p(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)} = \exp\Big(Q_{\pi_{\mathrm{ref}}}([x, y_{<t}^l], y_t^l) - Q_{\pi_{\mathrm{ref}}}([x, y_{<t}^w], y_t^w)\Big).$$

Substituting the closed-form expression of the reference-policy Q-function from Eq. (2) gives

$$\frac{p(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)} = \exp\left(\beta \log \frac{\pi_\theta^*(y_t^l \mid [x, y_{<t}^l])}{\pi_{\mathrm{ref}}(y_t^l \mid [x, y_{<t}^l])} - \beta \log \frac{\pi_\theta^*(y_t^w \mid [x, y_{<t}^w])}{\pi_{\mathrm{ref}}(y_t^w \mid [x, y_{<t}^w])} \right.$$
$$\left. + \beta \log Z([x, y_{<t}^l]; \beta) - \beta \log Z([x, y_{<t}^w]; \beta) \right).$$

Rearranging terms, we obtain

$$\frac{p(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)} = \left(\frac{\pi_\theta^*(y_t^l \mid [x, y_{<t}^l])\,\pi_{\mathrm{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\mathrm{ref}}(y_t^l \mid [x, y_{<t}^l])\,\pi_\theta^*(y_t^w \mid [x, y_{<t}^w])} \right)^\beta \exp\Big(\beta \log \tfrac{Z([x, y_{<t}^l]; \beta)}{Z([x, y_{<t}^w]; \beta)}\Big).$$

Defining the prefix-dependent weight

$$w_t = \exp\Big(\log \tfrac{Z([x, y_{<t}^l]; \beta)}{Z([x, y_{<t}^w], \beta)}\Big),$$

the above expression can be written as

$$\frac{p(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)} = \left(\frac{\pi_\theta^*(y_t^l \mid [x, y_{<t}^l])\,\pi_{\mathrm{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\mathrm{ref}}(y_t^l \mid [x, y_{<t}^l])\,\pi_\theta^*(y_t^w \mid [x, y_{<t}^w])} \right)^\beta w_t^\beta.$$

Raising both sides to the power $1/\beta$ and rearranging terms yields

$$\frac{\pi_\theta^*(y_t^w \mid [x, y_{<t}^w])}{\pi_\theta^*(y_t^l \mid [x, y_{<t}^l])} = \frac{\pi_{\mathrm{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\mathrm{ref}}(y_t^l \mid [x, y_{<t}^l])} \left(\frac{p(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l)}{p(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l)} \right)^{1/\beta} w_t,$$

which completes the proof. $\qquad\square$

A.2. Proof for Prop 4.2

Proof. We prove Proposition 4.2 by deriving the policy ratio induced by a token-level Bradley–Terry model parameterized by reference-policy advantages.

By definition, the token-level Bradley–Terry likelihood is given by

$$p\big(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\big) = \sigma\Big(A_{\pi_{\mathrm{ref}}}([x, y_{<t}^w], y_t^w) - A_{\pi_{\mathrm{ref}}}([x, y_{<t}^l], y_t^l)\Big).$$

Using $\sigma(u) = (1 + e^{-u})^{-1}$ and taking the odds ratio while conditioning on the same pair of prefixes yields

$$\frac{p\big(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)}{p\big(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)} = \exp\Big(A_{\pi_{\mathrm{ref}}}([x, y_{<t}^l], y_t^l) - A_{\pi_{\mathrm{ref}}}([x, y_{<t}^w], y_t^w)\Big).$$

Substituting the closed-form expression of the reference-policy advantage from Eq. (3) into the above expression gives

$$
\frac{p\big(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)}{p\big(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)} = \exp\left(\beta \log \frac{\pi_\theta^*(y_t^l \mid [x, y_{<t}^l])\,\pi_{\text{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_\theta^*(y_t^w \mid [x, y_{<t}^w])\,\pi_{\text{ref}}(y_t^l \mid [x, y_{<t}^l])} \right.
$$
$$
+ \beta\Big[D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y_{<t}^l]) \,\big\|\, \pi_\theta^*(\cdot \mid [x, y_{<t}^l])\big)
$$
$$
\left. - D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y_{<t}^w]) \,\big\|\, \pi_\theta^*(\cdot \mid [x, y_{<t}^w])\big)\Big] \right).
$$

Raising both sides to the power $1/\beta$ and rearranging terms yields

$$
\frac{\pi_\theta^*\big(y_t^w \mid [x, y_{<t}^w]\big)}{\pi_\theta^*\big(y_t^l \mid [x, y_{<t}^l]\big)} = \frac{\pi_{\text{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\text{ref}}(y_t^l \mid [x, y_{<t}^l])} \left(\frac{p\big(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)}{p\big(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)} \right)^{1/\beta} \cdot w_t,
$$

where

$$
w_t = \exp\Big(D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y_{<t}^l]) \,\big\|\, \pi_\theta^*(\cdot \mid [x, y_{<t}^l])\big) - D_{\text{KL}}\big(\pi_{\text{ref}}(\cdot \mid [x, y_{<t}^w]) \,\big\|\, \pi_\theta^*(\cdot \mid [x, y_{<t}^w])\big)\Big).
$$

This establishes Proposition 4.2. $\qquad\square$

A.3. Proof for Theorem 4.3

Theorem. *Under sufficient model capacity, we have $\pi_{\theta^*} = \pi_{\hat\theta}$ where $\pi_{\hat\theta} = \arg\min_{\pi_\theta} D_h\left(R_{data}, R_\theta\right)$.*

Proof. Let $\arg\min_{\pi_\theta} D_h(R_{\text{data}} \| R_\theta) = \pi_{\hat\theta}$. We have the following decomposition::

$$
D_h\left(R_{\text{data}}, R_{\hat\theta}\right) = \frac{1}{T}\sum_{t=1}^{T} D_h^t\left(R_{\text{data}}^t, R_{\hat\theta}^t\right) = \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}_{p_{\text{data}}\left(y_{\le t}^w \succ y_{\le t}^l \mid x\right)}\left[B_h\left(R_{\text{data}}^t, R_{\hat\theta}^t\right) \right]. \tag{16}
$$

where T denotes the minimum sequence length.

By the non-negativity of the Bregman divergence, $B_h\left(R_{\text{data}}^t, R_{\hat\theta}^t\right)$ attains its minimum value 0 if and only if $R_{\text{data}}^t = R_{\hat\theta}^t$ for a given $(y_{<t}^w, y_{<t}^l, x)$. Assuming that the data distribution has full support, the minimizer $\pi_{\hat\theta}$ therefore satisfies $R_{\text{data}}^t = R_{\hat\theta}^t$ for all $(y_{<t}^w, y_{<t}^l, x)$. Substituting the definitions of R_{data}^t and $R_{\hat\theta}^t$ yields the following condition:

$$
\frac{p_{\text{data}}\big(y_t^w \prec y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)}{p_{\text{data}}\big(y_t^w \succ y_t^l \mid x, y_{<t}^w, y_{<t}^l\big)} = \left[\frac{\pi_{\hat\theta}\big(y_t^l \mid [x, y_{<t}^l]\big)\,\pi_{\text{ref}}(y_t^w \mid [x, y_{<t}^w])}{\pi_{\text{ref}}\big(y_t^l \mid [x, y_{<t}^l]\big)\,\pi_{\hat\theta}(y_t^w \mid [x, y_{<t}^w])}\, w_t \right]^\beta.
$$

Looping over all time steps from $t = 1$ to T and combining with the recursive formula in Eq. 12, we obtain

$$
\frac{p_{\text{data}}\big(y_{\le T}^w \prec y_{\le T}^l \mid x\big)}{p_{\text{data}}\big(y_{\le T}^w \succ y_{\le T}^l \mid x\big)} = \left[\frac{\pi_{\hat\theta}\big(y_{\le T}^l \mid x\big)\,\pi_{\text{ref}}\big(y_{\le T}^w \mid x\big)}{\pi_{\text{ref}}\big(y_{\le T}^l \mid x\big)\,\pi_{\hat\theta}\big(y_{\le T}^w \mid x\big)}\, w_x \right]^\beta,
$$

where $w_x := \prod_{t=1}^{T} w_t$.

$$
\frac{\pi_{\hat\theta}\big(y_{\le T}^w \mid x\big)}{\pi_{\hat\theta}\big(y_{\le T}^l \mid x\big)} = \frac{\pi_{\text{ref}}\big(y_{\le T}^w \mid x\big)}{\pi_{\text{ref}}\big(y_{\le T}^l \mid x\big)} \left(\frac{p_{\text{data}}\big(y_{\le T}^w \succ y_{\le T}^l \mid x\big)}{p_{\text{data}}\big(y_{\le T}^w \prec y_{\le T}^l \mid x\big)} \right)^{\frac{1}{\beta}} w_x.
$$

Similarly, we derive the sequence-level ratio for the optimal policy π_{θ^*} by aggregating the token-level conditions from Propositions 4.1 and 4.2 over $t = 1 \ldots T$ and applying the recursive preference formula (Eq. 12):

$$
\frac{\pi_\theta^*\big(y_{\le T}^w \mid x\big)}{\pi_\theta^*\big(y_{\le T}^l \mid x\big)} = \frac{\pi_{\text{ref}}\big(y_{\le T}^w \mid x\big)}{\pi_{\text{ref}}\big(y_{\le T}^l \mid x\big)} \left(\frac{p_{\text{data}}\big(y_{\le T}^w \succ y_{\le T}^l \mid x\big)}{p_{\text{data}}\big(y_{\le T}^w \prec y_{\le T}^l \mid x\big)} \right)^{\frac{1}{\beta}} w_x.
$$

Comparing this with the expression for $\pi_{\hat{\theta}}$ derived above reveals that both policies satisfy the exact same likelihood ratio for all points $(x, y^w_{\leq T}, y^l_{\leq T})$:

$$\frac{\pi_{\hat{\theta}}\left(y^w_{\leq T} \mid x\right)}{\pi_{\hat{\theta}}\left(y^l_{\leq T} \mid x\right)} = \frac{\pi^*_{\theta}\left(y^w_{\leq T} \mid x\right)}{\pi^*_{\theta}\left(y^l_{\leq T} \mid x\right)}$$

Using the completeness property of concrete score (Meng et al., 2023), we conclude that $\pi_{\hat{\theta}} = \pi_{\theta^*}$. $\qquad\square$

A.4. Proof for Theorem 4.4

Theorem. *We have $\mathcal{L}\left(\theta, p_{data}\right) = D_h\left(R_{data}, R_\theta\right) + C$ with C is a constant with respect to θ where we define*

$$\mathcal{L}(\theta, p_{data}) = \mathbb{E}_{p_{data}(y^w \succ y^l | x)}\left[\frac{1}{T}\sum_{t=1}^{T}\left(h'\left(R^t_\theta\right)R^t_\theta - h\left(R^t_\theta\right) - h'\left(\left(R^t_\theta\right)^{-1}\right)\right)\right].$$

Proof. By definition:

$$
\begin{aligned}
D^t_h\left(R^t_{\text{data}}, R^t_\theta\right) &= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h\left(R^t_{\text{data}}\right) - h\left(R^t_\theta\right) - h'\left(R^t_\theta\right)\left(R^t_{\text{data}} - R^t_\theta\right)\right] \\
&= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h'(R^t_\theta)R^t_\theta - h(R^t_\theta) - h'(R^t_\theta)R^t_{\text{data}}\right] + \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h(R^t_{\text{data}})\right] \\
&= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h'(R^t_\theta)R^t_\theta - h(R^t_\theta) - h'(R^t_\theta)\frac{p_{\text{data}}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t})}{p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})}\right] - C \\
&= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h'(R^t_\theta)R^t_\theta - h(R^t_\theta)\right] - \mathbb{E}_{p_{data}\left(y^w_{\leq t-1} \succ y^l_{\leq t-1} | x\right)}\left[h'(R^t_\theta)p_{\text{data}}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t})\right] - C.
\end{aligned}
$$
(17)

Next, observe that:

$$
\begin{aligned}
h'\left(R_\theta\right)p_{\text{data}}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t}) &= p_{\text{data}}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'\left(R_\theta(x, y^w_t, y^l_t)\right) \\
&= p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'\left(R_\theta(x, y^l_t, y^w_t)\right) \\
&= p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'\left(\left(\frac{\pi_\theta\left(y^w_t \mid [x, y^w_{<t}]\right)\pi_{\text{ref}}\left(y^l_t \mid [x, y^l_{<t}]\right)}{\pi_{\text{ref}}\left(y^w_t \mid [x, y^w_{<t}]\right)\pi_\theta\left(y^l_t \mid [x, y^l_{<t}]\right)}w^{-1}_t\right)^\beta\right) \\
&= p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'\left(\frac{1}{R^t_\theta(x, y^w_t, y^l_t)}\right) \\
&= p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'((R^t_\theta)^{-1})
\end{aligned}
$$

Therefore:

$$
\begin{aligned}
\mathbb{E}_{p_{data}\left(y^w_{\leq t-1} \succ y^l_{\leq t-1} | x\right)}\left[h'(R^t_\theta)\,p_{\text{data}}(y^w_t \prec y^l_t \mid x, y^w_{<t}, y^l_{<t})\right] &= \mathbb{E}_{p_{data}\left(y^w_{\leq t-1} \succ y^l_{\leq t-1} | x\right)}\left[p_{\text{data}}(y^w_t \succ y^l_t \mid x, y^w_{<t}, y^l_{<t})\,h'((R^t_\theta)^{-1})\right] \\
&= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h'((R^t_\theta)^{-1})\right].
\end{aligned}
$$

Substituting the above result into Eq. (17) yields:

$$
\begin{aligned}
D^t_h\left(R^t_{\text{data}}, R^t_\theta\right) &= \mathbb{E}_{p_{data}\left(y^w_{\leq t} \succ y^l_{\leq t} | x\right)}\left[h'(R^t_\theta)R^t_\theta - h(R^t_\theta) - h'((R^t_\theta)^{-1})\right] - C \\
&= \mathbb{E}_{p_{data}(y^w \succ y^l | x)}\left[h'(R^t_\theta)R^t_\theta - h(R^t_\theta) - h'((R^t_\theta)^{-1})\right] - C
\end{aligned}
$$

Figure 5. LC win rate vs. average response length on TLDR for Llama 3 8B, judged by Llama 3 70B and DeepSeek-V3. TBPO achieves highest win rate with shorter outputs, suggesting strong OOD generalization despite training only on UltraFeedback dataset.

Averaging over timesteps, we obtain:

$$
\begin{aligned}
D_h\left(R_{\text{data}}, R_\theta\right) &= \frac{1}{T} \sum_{t=1}^{T} D_h^t\left(R_{\text{data}}^t, R_\theta^t\right) \\
&= \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{p_{\text{data}}(y^w \succ y^l | x)}\left[h'(R_\theta^t)R_\theta^t - h(R_\theta^t) - h'\big((R_\theta^t)^{-1}\big)\right] - C \\
&= \mathbb{E}_{p_{\text{data}}(y^w \succ y^l | x)}\left[\frac{1}{T} \sum_{t=1}^{T} h'(R_\theta^t)R_\theta^t - h(R_\theta^t) - h'\big((R_\theta^t)^{-1}\big)\right] - C \\
&= \mathcal{L}(\theta, p_{\text{data}}) - C
\end{aligned}
$$

$\square$

B. Additional Experimental Result

B.1. Out-of-domain Result on Summarization Task

To stress-test out-of-domain (OOD) transfer, we evaluate on 200 held-out samples from the test split of TLDR. This is a stringent OOD setting: TLDR is a summarization benchmark with long input documents and short, reference-style outputs, and we do not use any TLDR data during training (TBPO is trained only on UltraFeedback).

As with Anthropic HH, we quantify OOD generalization by measuring length-controlled (LC) win rate to reduce verbosity bias, against the dataset-preferred completion and plotting it against average response length (Figure 5). Across both judges (Llama 3 70B and DeepSeek-V3), TBPO (A/Q variants) lies on the Pareto frontier, achieving the strongest win rates while producing substantially shorter outputs than prior preference-optimization baselines. Because TLDR differs in domain, task format, and implicit preference structure, these gains are difficult to attribute to UltraFeedback-specific stylistic overfitting; instead, they suggest TBPO learns preferences that transfer under distribution shift.

The plot also highlights a clear mismatch between length and preference in this OOD task: SFT produces the longest responses yet attains low win rates under both judges, while DPO/BPO/TISDPO/TDPO generally trade longer generations for incremental gains and still fall short of TBPO. Since TLDR implicitly rewards concise, faithful summarization, TBPO's ability to win while staying short strengthens the interpretation that it adapts to task-appropriate brevity rather than relying on generic verbosity. The consistent ordering across two independent judges further supports that the observed advantage reflects genuine OOD generalization.

Method	HellaSwag	ARC	MMLU	TruthfulQA	Winogrande	GSM8K	Average
SFT	79.46	51.45	61.19	46.46	75.69	68.67	63.82
DPO	81.23	56.05	61.70	48.47	75.84	74.20	66.24
TDPO	**83.29**	59.04	61.86	51.76	76.55	75.10	67.91
TIS-DPO	81.37	58.70	60.77	49.87	75.14	75.12	66.82
BPO	81.68	54.95	61.48	47.88	75.13	73.22	65.72
TBPO-Q noisy	81.11	60.24	60.25	50.38	78.24	74.60	67.47
TBPO-Q (Ours)	81.90	**64.07**	**64.92**	**53.36**	79.08	78.92	70.38
TBPO-A noisy	80.92	61.34	60.38	51.08	77.74	74.35	67.63
TBPO-A (Ours)	81.97	**64.08**	64.84	53.29	**79.63**	**79.22**	**70.50**

Table 2. General capability results on the Llama 3 8B backbone when TBPO is trained with 10% noisy preference data, created by swapping the chosen and rejected responses in randomly selected training pairs.

Method	HellaSwag	ARC	MMLU	TruthfulQA	Winogrande	GSM8K	Average
SFT	79.46	51.45	61.19	46.46	75.69	68.67	63.82
DPO	81.23	56.05	61.70	48.47	75.84	74.20	66.24
TDPO	**83.29**	59.04	61.86	51.76	76.55	75.10	67.91
TIS-DPO	81.37	58.70	60.77	49.87	75.14	75.12	66.82
BPO	81.68	54.95	61.48	47.88	75.13	73.22	65.72
TBPO (fixed $w_t = 1$)	80.93	61.35	62.11	52.38	77.24	76.23	68.37
TBPO-Q (Ours)	81.90	64.07	**64.92**	**53.36**	79.08	78.92	70.38
TBPO-A (Ours)	81.97	**64.08**	64.84	53.29	**79.63**	**79.22**	**70.50**

Table 3. General capability results on the Llama 3 8B backbone with the state-dependent correction fixed to $w_t = 1$.

B.2. Training with Noisy Data

We additionally evaluate TBPO under noisy preference labels by randomly swapping the chosen and rejected responses in 10% of the training pairs. Table 2 reports results on the Llama 3 8B backbone.

Performance drops under noisy labels, but TBPO remains competitive with strong token-level baselines and close to the clean TBPO variants. This suggests that TBPO is not brittle to moderate label noise. Broader analyses, such as extreme response lengths and additional distribution shifts, would further clarify the method's practical limits.

B.3. Constant Weight Training

We further conduct an ablation by fixing the state-dependent correction to $w_t = 1$ throughout training. Detailed results on the Llama 3 8B backbone are shown in Table 3.

Relative to both TBPO-Q and TBPO-A, the fixed-w_t variant shows a clear drop, confirming that the per-state correction contributes materially to the final gain. At the same time, it still outperforms all non-TBPO baselines, including TDPO, indicating that the improvement comes from both the token-level ratio-matching formulation and the state-dependent correction.

B.4. Judge Agreement Analysis

To assess judge robustness, we measured agreement across both LLM-as-a-judge setups used in the paper: DeepSeek-V3 vs. Llama 3 70B on HH-RLHF, and Llama 3 70B vs. Mistral-Large on MT-Bench.

Full results are available in Tables 4 and 5.

On HH-RLHF, the two judges show strong agreement (90% raw agreement; PABAK = 0.8 over 1,180 paired judgments).

Method	N	Agree	Agree%	Kappa	PABAK
TBPO-A	197	191	97.0	0.4861	0.9391
TBPO-Q	198	187	94.4	0.3240	0.8889
DPO	196	165	84.2	0.2780	0.6837
SFT	199	174	87.4	0.3516	0.7487
TDPO	195	173	88.7	0.3270	0.7744
TISDPO	195	172	88.2	0.3158	0.7641
Overall	**1180**	**1062**	**90.0**	**0.3471**	**0.8000**

Method	DeepSeek-V3 WR%	Llama 3 70B WR%
TBPO-A	96.00	97.97
TBPO-Q	95.50	95.45
DPO	87.50	86.22
SFT	86.50	91.46
TDPO	90.50	90.26
TISDPO	88.00	92.31

Table 4. HH-RLHF judge agreement results comparing each model response against the reference response. Pearson $r = 0.8056$ ($p = 5.30\mathrm{e}{-}02$), Spearman $\rho = 0.7143$ ($p = 1.11\mathrm{e}{-}01$).

TBPO-Q vs.	N	Agree	Agree%	Kappa	PABAK
DPO	160	117	73.1	0.5058	0.4625
TDPO	160	107	66.9	0.4542	0.3375
TISDPO	160	115	71.9	0.5143	0.4375
Overall	**480**	**339**	**70.6**	**0.4914**	**0.4125**

TBPO-Q vs.	Llama 3 70B WR%	Mistral-Large WR%
DPO	73.44	76.88
TDPO	66.25	67.81
TISDPO	70.00	71.88

Table 5. MT-Bench judge agreement results comparing the TBPO-Q response with the baseline response. Pearson $r = 0.9964$ ($p = 5.40\mathrm{e}{-}02$), Spearman $\rho = 1.0000$ ($p = 0.00\mathrm{e}{+}00$).

The lower Cohen's κ is expected in this highly skewed setting, and aggregate conclusions remain consistent across judges (Pearson $r = 0.81$, Spearman $\rho = 0.71$). On MT-Bench, example-level agreement is moderate (70% raw agreement, $\kappa = 0.49$, PABAK = 0.41), but aggregate conclusions are nearly identical across judges (Pearson $r = 0.99$, Spearman $\rho = 1.0$). Overall, the main conclusions are stable across judges, suggesting that TBPO's gains are not driven by any single judge choice.

C. Reasoning Questions Case Study

We believe the larger gains on reasoning-centric benchmarks are expected because these tasks are especially sensitive to early prefix-conditioned decisions. Once a model commits to an incorrect intermediate step, the rest of the reasoning often follows the wrong trajectory. TBPO is well matched to this setting because it performs token-level ratio matching rather than optimizing only at the sequence level. As a result, it provides more direct learning signal on the local reasoning decisions that determine whether a solution stays on track. In practice, this tends to help the model commit to the correct setup earlier, which is particularly important in multi-step reasoning.

We provide two qualitative examples in Tables 6 and 7. Example 1 shows our model identifying the decisive setup early, while the baseline misses the key relation and propagates the error through the rest of the solution. Example 2 shows a different but related advantage: even when both models reach the correct final answer, our model does so through a shorter

Item	Content
Question	At a certain grade level, three-fourths of students have a desktop computer at home. If 20 students do not have desktop computers, how many students are there at that grade level?
Our model	Let x be the total number of students at that grade level. Three-fourths of the students have desktop computers, so $(3/4)x$ students have desktop computers. The number of students without desktop computers is 20, so $(1/4)x = 20$. To find the total number of students, we can multiply both sides of the equation by 4: $x = 20 \times 4 = 80$.
80
Baseline	If 20 students do not have desktop computers, then $75 - 20 = 55$ students have desktop computers at home. Since $3/4$ of the students have desktop computers, then $3/4 \times 55 = 41.25$ students have desktop computers at home. Since 55 students have desktop computers, then $55 - 41.25 = 13.75$ students do not have desktop computers. Since 13.75 students do not have desktop computers, then $13.75 \times 4/3 = 18.33$ students do not have desktop computers. There are 18.33 students who do not have desktop computers, so there are $18.33 \times 4/3 = 24.44$ students at that grade level.
24

Table 6. A reasoning case where our model identifies the decisive setup immediately, while the baseline commits to an incorrect setup and propagates the error through the rest of the solution.

Item	Content
Question	Helena is a mad scientist testing her latest creation, the Lots-of-Limbs Serum. The serum will make the drinker grow an extra arm every three days and an extra leg every five days. After fifteen days, how many new limbs will Helena's serum cause a person to grow if they drink it?
Our model	At one new arm every 3 days, after 15 days, a person who drank Helena's serum would grow $15/3 = 5$ extra arms. At one new leg every 5 days, after 15 days, the person would grow $15/5 = 3$ extra legs. In all, someone who drank Helena's serum would grow $5 + 3 = 8$ new limbs after 15 days.
8
Baseline	After three days, the drinker will grow one extra arm. After six days, the drinker will grow two extra arms. After nine days, the drinker will grow three extra arms. After twelve days, the drinker will grow four extra arms. After fifteen days, the drinker will grow five extra arms. After five days, the drinker will grow one extra leg. After ten days, the drinker will grow two extra legs. After fifteen days, the drinker will grow three extra legs. After fifteen days, the drinker will grow five extra arms and three extra legs, for a total of 8 new limbs.
8

Table 7. A reasoning case where both models reach the correct final answer, but our model does so through a shorter and more direct reasoning path.

and more direct reasoning path, whereas the baseline relies on a longer enumerative chain.

Together, these examples illustrate why the gains are larger on reasoning-centric benchmarks: such tasks are especially sensitive to early intermediate decisions, and TBPO is designed to train those decisions more directly.

D. Detailed Experimental Setup

D.1. Benchmark and Metrics Details

Open LLM Leaderboard (Hugging Face, 2024). The leaderboard evaluates open-weight LLMs on six core benchmarks using the EleutherAI Language Model Evaluation Harness under a fixed few-shot setup: ARC-Challenge (25-shot), HellaSwag (10-shot), MMLU (5-shot), TruthfulQA (0-shot; implemented as `truthfulqa-mc`), Winogrande (5-shot), and GSM8K (5-shot). Higher scores indicate better performance across these tasks.

MT-Bench (Zheng et al., 2023). MT-Bench-101 is a fine-grained multi-turn dialogue benchmark organized by a three-tier ability taxonomy. It contains 1,388 multi-turn dialogues comprising 4,208 turns across 13 task types. For evaluation, the benchmark uses the *golden context* (the full dialogue history) as input context for each turn, and employs an LLM as an automatic judge to score each turn according to task-specific scoring guidelines. To reflect robustness across the whole

interaction, the final dialogue score is taken as the lowest (minimum) round score among the turns.

Anthropic HH-RLHF. HH-RLHF provides human preference data targeting *helpfulness* and *harmlessness*. The preference-modeling portion is distributed as JSONL records containing paired model responses: a `chosen` response preferred by annotators and a `rejected` response. The dataset is intended for training preference/reward models for RLHF, and the dataset card explicitly warns against using these data for supervised fine-tuning of dialogue agents. The repository also includes red-teaming dialogues and notes that the data may contain potentially offensive or upsetting content.

TLDR. The TL;DR dataset is a processed Reddit summarization corpus built from posts where authors append a "TL;DR" summary. It is provided in a prompt–completion format for summarization training, with a long-form post as input and its short TL;DR as the target summary. The dataset card specifies two columns: `"pompt"` (the full Reddit post text) and `"completion"` (the TL;DR summary).

Length-controlled Win Rate. In a pairwise preference evaluation, each example i consists of an input (e.g., prompt) x_i and two candidate outputs $z_i^{(A)}$ and $z_i^{(B)}$, with a judge label $y_i \in \{0, 1\}$ indicating whether A is preferred ($y_i = 1$) over B. The (raw) win rate is

$$\mathrm{WR}(A, B) \ = \ 100 \cdot \mathbb{E}_i\left[y_i\right]. \tag{18}$$

LLM-based judges (and sometimes humans) can exhibit *length bias*, i.e., a systematic preference for longer outputs, which can inflate win rates for more verbose systems. To control for this confounder, we model pairwise preferences with a logistic regression that includes the output-length difference as a covariate:

$$\Pr(y_i = 1 \mid x_i) \ = \ \sigma\left(\Delta\theta \ + \ \lambda\, g(\Delta\ell_i) \ + \ u_{x_i}\right), \qquad \Delta\ell_i := \ell\left(z_i^{(A)}\right) - \ell\left(z_i^{(B)}\right), \tag{19}$$

where $\sigma(t) = \frac{1}{1+e^{-t}}$, $\ell(\cdot)$ is a length measure (e.g., tokens), $\Delta\theta$ captures the system advantage of A over B, and u_{x_i} optionally absorbs item-specific difficulty via fixed effects (or random effects). We use a bounded transform such as $g(\Delta\ell_i) = \tanh\left(\Delta\ell_i/s\right)$ (with scale s, e.g., the empirical standard deviation) to reduce sensitivity to extreme length gaps. The *length-controlled win rate* is the counterfactual preference probability when both outputs have equal length, i.e., conditioning on $\Delta\ell_i = 0$:

$$\mathrm{LCWR}(A, B) \ = \ 100 \cdot \mathbb{E}_i\left[\Pr(y_i = 1 \mid x_i, \Delta\ell_i = 0)\right] \ = \ 100 \cdot \mathbb{E}_i\left[\sigma\left(\Delta\theta + u_{x_i}\right)\right]. \tag{20}$$

Intuitively, LCWR reports the implied win probability after removing any advantage attributable purely to verbosity.

Diversity metrics. We evaluate the diversity of model responses using three metrics:

- **Predictive Entropy**: Evaluates the diversity of model responses by directly reflecting the probabilities assigned to each generated token. For each prompt, we sample five responses and compute the predictive entropy. Higher entropy indicates more diverse outputs.

- **Self-BLEU**: Measures the diversity of generated sentences. For each prompt, we generate five samples and compute the BLEU score (Zhu et al., 2018) for each pair of responses, then take the average. Lower Self-BLEU indicates higher diversity.

- **Distinct-1**: Assesses the lexical diversity by computing the ratio of unique unigrams to the total number of unigrams (Li et al., 2016). Higher scores indicate greater lexical diversity.

D.2. Baseline details

SFT. Supervised fine-tuning (SFT) adapts a pretrained LM using demonstration data (prompt/instruction $\rightarrow$ reference answer pairs) with standard maximum-likelihood training (token-level cross-entropy / next-token prediction). Its purpose is to directly teach instruction following from supervised targets, without preference comparisons or RL. In RLHF-style pipelines, SFT is commonly used as an initialization step before preference optimization. In our experiments, we perform SFT on the *chosen* responses from the preference dataset.

DPO ([Rafailov et al., 2023](#)).   Direct Preference Optimization (DPO) trains a policy from pairwise preferences (winner vs. loser responses for the same prompt) using a simple logistic objective. It can be derived as a closed-form solution to a KL-regularized RLHF objective, avoiding both an explicit reward model and on-policy RL. A fixed reference policy provides regularization while the learned policy increases the relative likelihood of preferred responses.

TDPO ([Zeng et al., 2024](#)).   Token-level DPO (TDPO) extends DPO from sequence-level preferences to token-level optimization. By introducing token-wise constraints (e.g., per-token forward-KL terms), TDPO better matches autoregressive generation and aims to improve alignment while offering finer control over diversity.

TIS-DPO ([Liu et al., 2025a](#)).   TIS-DPO addresses the limitation of treating an entire response as a single unit: different tokens can contribute unequally to preference. It proposes a token-level importance-sampling objective with per-token weights reflecting reward contribution, yielding an unbiased objective even when the dataset is not ideally balanced. Since token rewards/weights are not observed, TIS-DPO estimates importance weights from probability differences produced by contrastive LLMs (e.g., via contrastive prompting, separate win/lose models, or forward/reverse DPO).

BPO (BPO-SBA) ([Kim et al., 2025](#)).   Bregman Preference Optimization (BPO) generalizes DPO from a likelihood-ratio viewpoint, avoiding an explicit reward model and partition function while retaining theoretical guarantees. It defines a family of tractable objectives induced by Bregman divergences, with DPO as a special case, and introduces gradient scaling (Scaled Basu's power divergence, SBA) to improve optimization behavior for certain instances. Empirically, BPO variants can improve win rate and output entropy relative to DPO, mitigating some fidelity–diversity trade-offs.

Baseline Implementation.   **SFT.** We initialize from the same SFT checkpoints as TBPO and train for 1 epoch on 2 H100 GPUs (batch size 8, gradient accumulation 4, learning rate 2×10^{-4}) using AdamW. **DPO/TDPO/TIS-DPO.** We initialize from the same SFT checkpoints and train on 2 H100 GPUs (batch size 4, gradient accumulation 8, learning rate 5×10^{-7}) using RMSProp with $\beta = 0.1$. Unless otherwise noted, we use the hyperparameters from the official implementations/papers for TDPO and TIS-DPO. **BPO (BPO-SBA).** Since the authors do not release code, we reimplement BPO and match the hyperparameters reported in the paper.

D.3. Implementation Details

D.3.1. Per-state weight approximation

Estimating the state-only correction weight w_t adds negligible overhead relative to standard preference optimization.

TBPO-Q.   We implement the baseline head $b_\phi(s)$ as a 1-hidden-layer MLP on top of the LLM's final hidden state with MLP hidden size is 1024, trained with a separate AdamW optimizer (lr 10^{-3}). We detach the LLM backbone (stop-gradient) when updating b_ϕ to avoid coupling the baseline fitting signal with the main preference-optimization updates.

TBPO-A.   We approximate π^* by the current policy π_θ and estimate $D_{\mathrm{KL}}(\pi_{\mathrm{ref}}(\cdot \mid s) \,\|\, \pi_\theta(\cdot \mid))$ using the K3 Monte-Carlo estimator from [Schulman (2020)](#). For an action $a \sim \pi_{\mathrm{ref}}(\cdot \mid s)$ and ratio $r(a, s) = \pi_\theta(a \mid s)/\pi_{\mathrm{ref}}(a \mid s)$, we use $\widehat{D}_{\mathrm{KL}}^{(k3)}(s; a) = (r(a, s) - 1) - \log r(a, s)$ and plug the result into the KL-baseline difference defining $w_t^{(A)}$.

D.3.2. Hyperparameters.

We train TBPO on 4 H100 GPUs with batch size 32 and max response length 2048, using RMSProp with a cosine learning-rate schedule (warmup ratio 0.05), lr 5×10^{-7}, for 1 epoch. Unless stated otherwise, we use $\beta = 0.1$ and Bregman parameters $\lambda = 0$ and $s = 4$.

D.4. Time Complexity

We report wall-clock runtimes under the same 4×H100 setup and batch size used in our experiments. DPO finishes in 1h52m, while TDPO, TBPO-Q, and TBPO-A take 2h29m, 2h52m, and 2h38m, respectively. TIS-DPO takes 5h40m because it requires training auxiliary positive/negative DPO models before the main preference-learning stage. These results show that TBPO is more expensive than sequence-level DPO, as expected for a token-level objective, but its overhead is modest relative to other token-level baselines: TBPO-A is close to TDPO, TBPO-Q is only moderately slower, and both are

substantially more efficient than TIS-DPO.

Method	Runtime	Relative to DPO
DPO	1h52m	1.00×
TDPO	2h29m	1.33×
TIS-DPO	5h40m	3.04×
TBPO-Q	2h52m	1.54×
TBPO-A	2h38m	1.41×

Table 8. Wall-clock training runtime under the same 4×H100 setup and batch size.