

Beyond Static Best-of-N: Bayesian List-wise Alignment for LLM-based Recommendation

Ruijun Chen
rjchen20@mail.ustc.edu.cn
University of Science and Technology
of China
Hefei, China

Chongming Gao*
chongming.gao@gmail.com
University of Science and Technology
of China
Hefei, China

Jiawei Chen
sleepyhunt@zju.edu.cn
Zhejiang University
Hangzhou, China

Weiqin Yang
futurelover10032@gmail.com
Zhejiang University
Hangzhou, China

Xiangnan He*
xiangnanhe@gmail.com
University of Science and Technology
of China
Hefei, China

Abstract

Large Language Models have revolutionized recommender systems (LLM4Rec) by leveraging their generative capabilities to model complex user preferences. However, existing LLM4Rec methods primarily rely on token-level objectives, making it difficult to optimize list-level and non-differentiable metrics (e.g., NDCG, fairness) that define actual recommendation quality. While Best-of-N (BoN) directly optimizes these metrics during inference, its high computational cost hinders real-world deployment. To address this, BoN Alignment aims to distill the search capability into the model itself, yet current approaches suffer from two critical limitations: (1) **Indiscriminate Supervision**, where the static reference fails to distinguish the relative quality of candidates exceeding its empirical range, leading to a loss of ranking guidance; and (2) **Gradient Decay**, where the effective supervision signal rapidly diminishes as the evolving policy improves, resulting in inefficient optimization.

To overcome these challenges, we propose **BLADE** (Bayesian List-wise Alignment via Dynamic Estimation). Unlike static approaches, BLADE introduces a Bayesian framework that continuously updates the target distribution by fusing historical priors with dynamic evidence from the model's current rollouts. This mechanism constructs a self-evolving target that adapts to the model's growing capabilities, ensuring the training signal remains informative throughout the learning process. Extensive experiments on three real-world datasets demonstrate that BLADE significantly outperforms state-of-the-art baselines. Crucially, it breaks the static performance upper bound, achieving sustained gains in both ranking accuracy (Recall, NDCG) and complex list-wise metrics (Fairness, Diversity). The code is available via <https://github.com/RegionCh/BLADE>.

*Corresponding Author.



CCS Concepts

• Information systems → Recommender systems.

Keywords

Large Language Models for Recommendation; List-wise Alignment; Best-of-N

ACM Reference Format:

Ruijun Chen, Chongming Gao, Jiawei Chen, Weiqin Yang, and Xiangnan He. 2026. Beyond Static Best-of-N: Bayesian List-wise Alignment for LLM-based Recommendation. In *Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26)*, July 20–24, 2026, Melbourne, VIC, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3805712.3809535>

1 Introduction

The integration of Large Language Models (LLMs) into recommender systems, known as LLM4Rec [24, 38], leverages their semantic reasoning and generalization capabilities to enhance personalization [11, 23]. Typically, these methods serialize user historical interactions into natural language prompts, instructing the model to generate the titles of target items sequentially [14]. Since pre-trained LLMs often lack the domain-specific knowledge required to generate valid recommendations [5], current research primarily relies on Supervised Fine-Tuning (SFT) [3] and alignment techniques like Direct Preference Optimization (DPO) [10, 27]. However, these approaches are predominantly grounded in token-level likelihoods or pair-wise preferences, creating a critical misalignment with list-wise objectives. In real-world scenarios, user satisfaction fundamentally relies on holistic list-level metrics such as NDCG and fairness, which capture global item interactions [44]. Consequently, optimization methods restricted to local token objectives fail to directly maximize these intricate, non-differentiable rewards.

Given the challenges of optimizing these non-differentiable rewards directly via gradients, inference-time search has emerged as a robust alternative [36, 37]. Specifically, Best-of-N (BoN) [32] effectively bypasses the differentiability constraint by sampling N candidates and selecting the optimal one based on the list-wise reward (as shown in Fig. 1(a)). This strategy has demonstrated substantial performance gains in aligning LLMs [31, 39]. However, this effectiveness comes at a prohibitive cost. As illustrated in Fig. 1(b), the

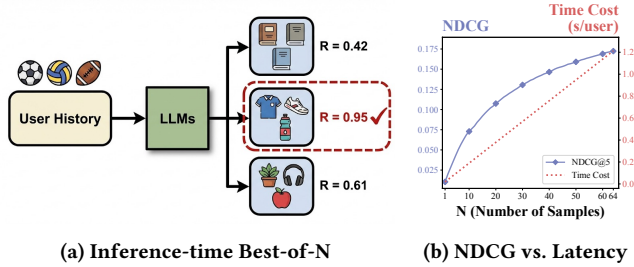


Figure 1: The efficiency bottleneck of Best-of-N in LLM4Rec. (a) BoN selects the optimal candidate from N samples to maximize list-wise rewards. (b) However, increasing N linearly escalates inference latency, making real-time deployment impractical.

inference latency increases linearly with the sample size N , rendering standard BoN impractical for online serving where low latency is mandated. This dilemma motivates the adoption of *BoN Alignment* [2, 29], which aims to internalize the search-and-selection capability into the model’s weights, enabling the direct generation of high-quality responses in a single pass.

Despite the promise of BoN alignment, its application to list-level recommendation faces two critical limitations stemming from the reliance on a static reference distribution. First, it is hindered by **Indiscriminate Supervision**. Since the static reference relies on a fixed support, its reward quantile estimation saturates for candidates outperforming the reference. Consequently, the alignment objective fails to distinguish between these superior candidates, assigning them identical target probabilities and thereby stripping the policy of critical ranking guidance in high-reward regions. Second, we identify the issue of **Gradient Decay**. As the policy evolves and rapidly outpaces the static reference, the proportion of reference samples offering informative signals drastically shrinks. This leads to a vanishing supervision signal, causing the optimization dynamics to stall prematurely.

To overcome these challenges, we propose **BLADE** (Bayesian List-wise Alignment via Dynamic Estimation). Rather than anchoring alignment to a fixed distribution, BLADE introduces a Bayesian framework to dynamically integrate static reference priors with *dynamic evidence* derived from the model’s current rollouts. This enables the online estimation of the evolving BoN distribution, reformulating the learning objective into a *self-evolving target*. Such dynamic alignment ensures that the training signal remains informative as the model’s capability expands, enabling continuous optimization that transcends the limitations of static reference models. To validate our approach, we conduct extensive experiments on three real-world datasets, demonstrating that BLADE achieves sustained improvements across both standard ranking accuracy and complex list-wise metrics.

To summarize, our main contributions are as follows:

- We systematically investigate the potential of BoN strategies in LLM4Rec and identify two critical bottlenecks in static BoN alignment: **Indiscriminate Supervision** stemming from the reference’s inability to distinguish superior candidates, and **Gradient Decay** caused by diminishing supervision signals.

- We propose **BLADE**, a Bayesian framework that dynamically fuses the static reference with dynamic evidence, constructing a *self-evolving target* to ensure sustained, informative gradients throughout training.
- Experiments on three datasets demonstrate that BLADE breaks the static performance upper bound, significantly outperforming baselines in both ranking utility and complex list-wise objectives such as fairness and diversity.

2 Preliminaries

In this section, we introduce the theoretical foundation of our work. We first formulate the recommendation task as a conditional list generation problem. Subsequently, we introduce GRPO, which serves as the optimization backbone of our framework. We conclude by reviewing the Best-of-N strategy, detailing its application in both inference and policy alignment contexts.

2.1 Task Formulation

We formulate the recommendation task as a conditional sequence generation problem. Let \mathcal{U} and \mathcal{I} denote the sets of users and items, respectively. For each user $u \in \mathcal{U}$, we construct a natural language prompt x based on their interaction history. The LLM functions as a stochastic policy π_θ , which takes x as input and autoregressively generates a response $y = (y_1, y_2, \dots, y_T)$ consisting of T tokens. This response y represents an ordered list of K recommended items. The probability of generating the complete response is factorized as:

$$\pi_\theta(y|x) = \prod_{t=1}^T \pi_\theta(y_t|x, y_{<t}). \quad (1)$$

The optimization goal is to maximize a non-differentiable list-wise reward $R(y; u)$, such as NDCG or Diversity, which evaluates the quality of the item sequence embedded within the generated response y .

2.2 GRPO

Standard Reinforcement Learning methods like PPO typically rely on a critic to estimate expected returns, which doubles the memory footprint and introduces training instability due to value approximation errors. To circumvent these issues, we adopt Group Relative Policy Optimization (GRPO) [30].

Formally, in the recommendation scenario, for a given user context x , the policy π_θ samples G independent candidate lists $\{y_1, y_2, \dots, y_G\}$. Rather than relying on an external reward model, GRPO leverages the relative quality of generations within the group. For any specific list-wise reward function $R(y)$ (e.g., NDCG or a customized metric), the advantage A_i for the i -th candidate is computed by normalizing its score against the group statistics:

$$A_i = \frac{R(y_i) - \mu_G}{\sigma_G}, \quad (2)$$

where μ_G and σ_G denote the mean and standard deviation of the rewards $\{R(y_j)\}_{j=1}^G$ within the current group.

The policy is then updated by maximizing the surrogate objective, which incorporates a clipped probability ratio to prevent destructive

updates. The objective function is defined as:

$$\mathcal{J}(\theta) = \mathbb{E}_{\pi_{\theta, \text{old}}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{T_i} \sum_{t=1}^{T_i} \left(\min(\rho_{i,t} A_i, \text{clip}(\rho_{i,t}, 1 - \epsilon, 1 + \epsilon) A_i) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta}(\cdot | h_{<t}) \| \pi_{\text{ref}}(\cdot | h_{<t})) \right) \right] \quad (3)$$

where T_i is the length of response y_i , and $\rho_{i,t} = \frac{\pi_{\theta}(y_{i,t} | x, y_{i,<t})}{\pi_{\theta, \text{old}}(y_{i,t} | x, y_{i,<t})}$ is the token-level importance sampling ratio.

In the context of our work, GRPO serves as an efficient backbone optimization algorithm. By leveraging group-based relative feedback, it enables us to directly optimize non-differentiable list-wise metrics without the computational overhead of a separate critic network. This stability and memory efficiency make it an ideal foundation for implementing our proposed dynamic alignment framework.

2.3 Best-of-N Framework

2.3.1 Best-of-N Inference. To directly optimize non-differentiable list-wise metrics, **BoN sampling** [32] serves as a powerful strategy. Specifically, in the context of recommendation, given a user context x , the policy π_{θ} generates N independent candidate recommendation lists $\{y_1, \dots, y_N\}$. The system then selects the optimal list y^* by maximizing the list-wise reward $R(y)$:

$$y^* = \underset{y \in \{y_1, \dots, y_N\}}{\text{argmax}} R(y). \quad (4)$$

While this approach effectively bypasses gradient limitations, it incurs a substantial computational cost. As demonstrated in Table 1, although increasing N significantly boosts performance, it results in a linear increase in inference latency, rendering it impractical for real-time recommendation.

2.3.2 Best-of-N Alignment. To mitigate the computational bottleneck of sampling, BoN Alignment aims to internalize the search capability into the policy’s parameters. Following the variational perspective of vBoN [2], the exact probability mass that a sample y is selected by BoN sampling is given by the expansion of order statistics:

$$\pi_{\text{BoN}}(y) = \sum_{k=1}^N \binom{N}{k} (F_{\text{ref}}(R(y)))^{N-k} \pi_{\text{ref}}(y)^k, \quad (5)$$

where $F_{\text{ref}}(r) = P_{y' \sim \pi_{\text{ref}}}(R(y') \leq r)$ denotes the Cumulative Distribution Function (CDF) of the reward under the reference policy. Here, the summation accounts for the fact that the same recommendation list y may appear multiple times among the N i.i.d. samples. The term indexed by k corresponds to the event that y is sampled k times, while the remaining $N - k$ samples have rewards no larger than $R(y)$.

Directly optimizing the exact BoN distribution is computationally inconvenient, since the reverse-KL objective involves the logarithm of the above order-statistic summation. This variational formulation therefore introduces a tractable lower-bound surrogate, yielding a quantile-based reward together with KL regularization:

$$\mathcal{J}_{\text{BoN}}(\theta) = \mathbb{E}_{y \sim \pi_{\theta}} \left[(N-1) \log F_{\text{ref}}(R(y)) - \log \frac{\pi_{\theta}(y)}{\pi_{\text{ref}}(y)} \right]. \quad (6)$$

Consequently, aligning the policy π_{θ} with this target is equivalent to minimizing the KL divergence, which translates to maximizing the following objective:

$$\begin{aligned} \mathcal{J}(\theta) &= -\mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{BoN}}) \\ &= \mathbb{E}_{y \sim \pi_{\theta}} \left[\log \frac{\pi_{\text{BoN}}(y)}{\pi_{\theta}(y)} \right] \\ &\simeq \mathbb{E}_{y \sim \pi_{\theta}} \left[\underbrace{(N-1) \log F_{\text{ref}}(R(y))}_{\text{Quantile Reward}} - \underbrace{\log \frac{\pi_{\theta}(y)}{\pi_{\text{ref}}(y)}}_{\text{KL Regularization}} \right]. \end{aligned} \quad (7)$$

Here, the objective implies an RL problem where the model is incentivized to generate responses with high reward quantiles relative to the reference, while remaining anchored to the reference via KL regularization.

In standard implementations, the CDF term $F_{\text{ref}}(\cdot)$ is estimated empirically using a fixed set of samples from a **static reference model**. While computationally convenient, relying on a fixed anchor introduces fundamental theoretical limitations regarding **ranking discrimination** and gradient dynamics. We formally analyze these bottlenecks in Section 3 to motivate the design of our proposed BLADE framework.

Table 1: Comparison of recommendation performance (NDCG@5) and inference efficiency between BLADE, BoN Alignment, and Inference-time BoN.

Method	NDCG@5	Time (s)	Inference Cost
Base	0.0283	44.56	1×
BoN ($N = 2$)	0.0380	89.12	2×
BoN ($N = 4$)	0.0470	178.24	4×
BoN Alignment	0.0333	44.56	1×
BLADE	<u>0.0410</u>	44.56	1×

3 BLADE

In this section, we present **BLADE**. We begin by analyzing the fundamental limitations of static BoN alignment. Subsequently, we outline the overall framework, detail the Bayesian formulation for dynamic target estimations, and finally describe the optimization process using GRPO.

3.1 Limitations of Static Reference

Despite the tractability of the approximation in Eq. (7), existing methods typically estimate the target distribution using a **static reference model**. We formally analyze how this static nature imposes two fundamental limitations.

1. Indiscriminate Supervision. Recall that the BoN target distribution is derived as a re-weighting of the reference. In standard implementations, the CDF $F_{\text{ref}}(\cdot)$ is empirically estimated using a fixed set of samples \mathcal{D}_{ref} . Consequently, the estimation is strictly bounded by the maximum reward observed in the static reference, denoted as $R_{\text{max}} = \max_{y' \in \mathcal{D}_{\text{ref}}} R(y')$.

As the policy π_{θ} evolves, it inevitably explores high-reward regions and generates candidates that outperform the static reference

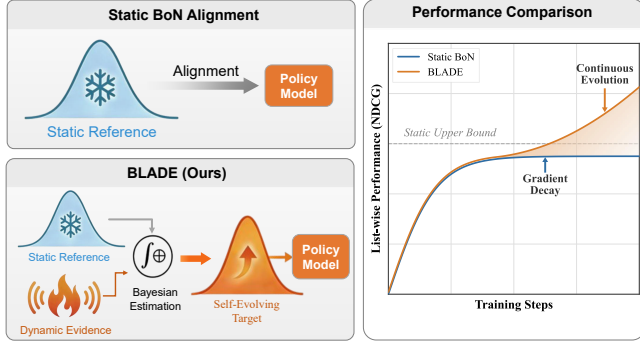


Figure 2: Static BoN Alignment vs. BLADE. (Left) Unlike static methods anchored to a frozen reference, BLADE constructs a Self-Evolving Target by fusing static priors with dynamic evidence, thereby resolving Indiscriminate Supervision in high-reward regions. (Right) This mechanism mitigates Gradient Decay, enabling the model to break the Static Upper Bound for continuous evolution.

(i.e., $R(y) > R_{\max}$). However, for any such superior candidate, the estimated quantile saturates at the upper bound: $F_{\text{ref}}(R(y)) = 1$. This leads to a critical failure in supervision: consider two generated lists y_1 and y_2 s.t. $R(y_1) > R(y_2) > R_{\max}$. Although y_1 is strictly superior, the static estimator assigns them identical quantile scores:

$$(F_{\text{ref}}(R(y_1)))^{N-1} = (F_{\text{ref}}(R(y_2)))^{N-1} = 1. \quad (8)$$

This indicates that the alignment process becomes **indiscriminate** in the high-utility region. By failing to distinguish between varying levels of superior quality, the objective effectively flattens the optimization landscape, preventing the policy from converging towards the globally optimal solution.

2. Gradient Decay. We further analyze the gradient of the primary utility term in Eq. (7), the quantile reward. Let $J_{\text{qt}}(\theta) = \mathbb{E}_{y \sim \pi_\theta} [(N-1) \log F_{\text{ref}}(R(y))]$. According to the policy gradient theorem, the gradient with respect to θ is:

$$\nabla_\theta J_{\text{qt}}(\theta) = \mathbb{E}_{y \sim \pi_\theta} [\nabla_\theta \log \pi_\theta(y) \cdot (N-1) \log F_{\text{ref}}(R(y))]. \quad (9)$$

As the policy π_θ evolves, its probability mass shifts towards high-reward regions, increasingly overlapping with the sparse right tail of the static reference distribution. Consequently, even before the policy reaches the theoretical maximum, the generated responses y fall into high-quantile regions where $F_{\text{ref}}(R(y)) \rightarrow 1$. This leads to a rapid shrinkage of the effective weighting term:

$$\lim_{F_{\text{ref}} \rightarrow 1} \log F_{\text{ref}}(R(y)) = 0. \quad (10)$$

This implies the **gradient decay**, as the policy improves relative to the static reference, the gradient magnitude shrinks rapidly. This vanishing supervision signal leads to inefficient optimization, causing the training to stall prematurely before reaching the global optimum.

3.2 Method Overview

Building upon the theoretical analysis in Section 3.1, we propose **BLADE**. As illustrated in Figure 3, BLADE fundamentally shifts

the alignment paradigm from anchoring to a fixed reference to pursuing a *self-evolving target*.

From a Bayesian perspective, the standard static reference π_{ref} serves as a prior belief about the reward landscape. While providing a stable starting point, this static prior imposes **Indiscriminate Supervision** on high-quality candidates discovered during the policy’s exploration. To rectify this, BLADE reformulates the target construction as a dynamic estimation process. At each training step, we treat the real-time response lists generated by the evolving policy π_θ as **dynamic evidence**. By fusing this up-to-date evidence with the static prior, BLADE constructs a posterior target distribution that adaptively scales with the model’s current capability.

This dynamic mechanism effectively resolves the identified bottlenecks. By incorporating real-time rollouts, BLADE restores the **ranking discrimination** for superior candidates, and by recalibrating the probability scale, it prevents **Gradient Decay**. Fig. 2 illustrates that while the static baseline stalls due to signal saturation, BLADE continuously updates the target to ensure the optimization gradients remain informative, enabling continuous evolution beyond the initial reference.

3.3 Bayesian Dynamic Estimation

To instantiate the self-evolving target distribution, the core challenge lies in robustly estimating the CDF of the rewards, $F(r) = P(R(\mathbf{y}) < r)$. Instead of relying on a static approximation, BLADE casts this estimation as a real-time Bayesian inference problem.

Probabilistic Formulation. For each reward threshold r , we treat the corresponding quantile value $\theta_r = F(r)$ as a random variable. Since θ_r is bounded in $[0, 1]$, we model it using a threshold-specific **Beta distribution**:

$$\theta_r \sim \text{Beta}(\alpha^r, \beta^r). \quad (11)$$

The Beta family is a natural choice because its support matches the range of CDF values and it is conjugate to the binomial likelihood induced by counting how many samples fall below the threshold r , enabling efficient closed-form updates.

Bayesian Update Mechanism. We construct the estimator by fusing statistics from the static reference set \mathcal{D}_{ref} , comprising M samples, and the current training batch $\mathcal{D}_{\text{batch}}$ of size G . Let $N_{\text{ref}}^{<r}$ and $N_{\text{batch}}^{<r}$ denote the count of samples with rewards lower than r in the respective sets.

We initialize the threshold-specific prior belief using the static reference statistics, where $\alpha_0^r = N_{\text{ref}}^{<r}$ and $\beta_0^r = M - N_{\text{ref}}^{<r}$. To balance the stability of the static prior and the adaptivity of online rollouts, we use a **power-scaled likelihood** with a **dynamic coefficient** $\tau \geq 0$. This coefficient controls the effective sample size of the dynamic evidence: $\tau = 0$ recovers the static estimator, $\tau = 1$ corresponds to a standard beta-binomial update using the current batch, and $0 < \tau < 1$ conservatively tempers the batch evidence to reduce estimation variance under limited rollouts. Formally, the posterior

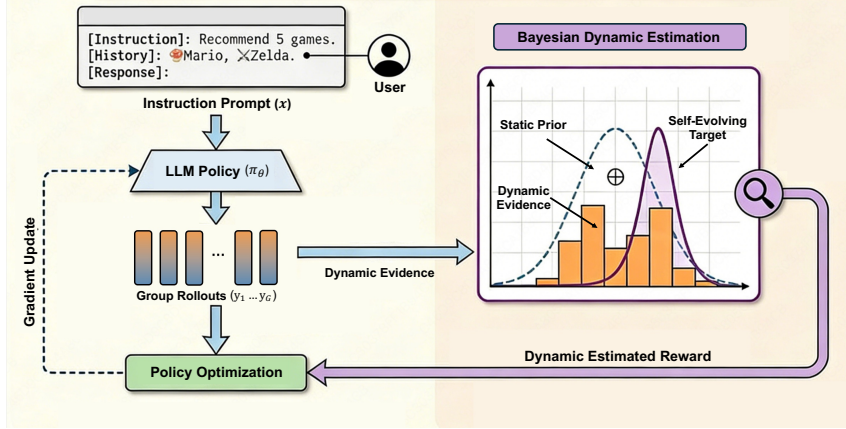


Figure 3: The overall training framework of BLADE. (Left) The pipeline illustrates how BLADE fuses the Static Prior with Dynamic Evidence to construct a self-evolving target distribution. (Right) The detailed training algorithm. For each batch, BLADE computes the Bayesian dynamic target \hat{F}_i via Eq. (14) and converts it into a proxy reward for list-wise policy optimization.

is derived as:

$$\begin{aligned}
 P(\theta_r | \mathcal{D}_{\text{batch}}) &\propto P(\theta_r) \cdot [P(\mathcal{D}_{\text{batch}} | \theta_r)]^\tau \\
 &\propto \underbrace{\theta_r^{\alpha_r^0 - 1} (1 - \theta_r)^{\beta_r^0 - 1}}_{\text{Static Prior}} \cdot \underbrace{\left[\theta_r^{N_{\text{batch}}^{<r}} (1 - \theta_r)^{(G - N_{\text{batch}}^{<r})} \right]^\tau}_{\text{Dynamic Evidence}} \\
 &= \theta_r^{(\alpha_r^0 + \tau N_{\text{batch}}^{<r}) - 1} (1 - \theta_r)^{(\beta_r^0 + \tau(G - N_{\text{batch}}^{<r})) - 1}.
 \end{aligned} \tag{12}$$

Identifying the exponents with the standard Beta distribution, we obtain the updated parameters for the posterior $\text{Beta}(\alpha_{\text{new}}^r, \beta_{\text{new}}^r)$:

$$\alpha_{\text{new}}^r = \alpha_r^0 + \tau \cdot N_{\text{batch}}^{<r}, \quad \beta_{\text{new}}^r = \beta_r^0 + \tau \cdot (G - N_{\text{batch}}^{<r}). \tag{13}$$

Posterior Estimation. Finally, to obtain a robust point estimate for the optimization target at threshold r , we take the expectation of the posterior distribution. This yields the **BLADE Dynamic Estimator**:

$$\hat{F}_{\text{BLADE}}(r) = \mathbb{E}[\theta_r | \mathcal{D}_{\text{batch}}] = \frac{\alpha_{\text{new}}^r}{\alpha_{\text{new}}^r + \beta_{\text{new}}^r} = \frac{N_{\text{ref}}^{<r} + \tau \cdot N_{\text{batch}}^{<r}}{M + \tau \cdot G}. \tag{14}$$

3.4 Efficient Optimization via Shared Sampling

To optimize the policy towards the dynamic target estimated in Eq. (14), we employ Group Relative Policy Optimization (GRPO) as our optimization backbone. We select this algorithm specifically because its group-wise sampling mechanism naturally complements our Bayesian estimation strategy, enabling efficient, critic-free optimization.

Dynamic Reward Integration. Instead of relying on static reward signals, we construct the optimization objective dynamically using the estimator derived in the previous section. For each generated candidate list y_i within the sampled group $\mathcal{D}_{\text{batch}}$, we formulate the proxy reward r_i as:

$$r_i = (N - 1) \log \hat{F}_{\text{BLADE}}(R(y_i)). \tag{15}$$

Algorithm 1 BLADE Training Algorithm

Require: Policy π_θ , Reference D_{ref} , τ , N

- 1: Pre-compute static statistics $\{N_{\text{ref}}^{<r}\}$
- 2: **for** each batch $X \sim \mathcal{D}$ **do**
- 3: Generate candidates $\{y_i\}_{i=1}^G \sim \pi_\theta(\cdot | X)$
- 4: **for** each candidate y_i **do**
- 5: **1. Dynamic Evidence:**
- 6: $N_{\text{batch}}^{<R(y_i)} \leftarrow \sum_{j=1}^G \mathbb{I}(R(y_j) < R(y_i))$
- 7: **2. Estimate Target (Bayesian Fusion):**
- 8: $\hat{F}_i \leftarrow \frac{N_{\text{ref}}^{<R(y_i)} + \tau \cdot N_{\text{batch}}^{<R(y_i)}}{|D_{\text{ref}}| + \tau \cdot G}$
- 9: **3. Proxy Reward:** $r_i \leftarrow (N - 1) \log \hat{F}_i$
- 10: **end for**
- 11: **4. Update:** Optimization step via GRPO
- 12: **end for**

This dynamic reward is then directly substituted into the GRPO objective to compute the group-relative advantage, ensuring that policy updates are consistently driven by the self-evolving target distribution.

Efficiency: The Shared Sampling Mechanism. A critical synergy in our framework is the **Shared Sampling** strategy. Both the Bayesian target estimation and the GRPO baseline calculation require a group of samples to derive distributional statistics. By utilizing the same batch of rollouts $\mathcal{D}_{\text{batch}}$ for a **dual purpose**, BLADE achieves *zero-overhead* target estimation. Specifically, the generated rollouts are first utilized as "dynamic evidence" to update the Bayesian posterior, and immediately reused to compute the policy gradients. Consequently, BLADE maintains the same computational efficiency as static BoN alignment, achieving superior performance **without additional inference or training costs**.

4 Experiments

In this section, we conduct extensive experiments to evaluate BLADE by addressing the following research questions:

- **RQ1 (Overall Performance):** Does BLADE achieve superior performance compared to SFT and state-of-the-art alignment baselines in list-wise recommendation tasks?
- **RQ2 (Mechanism Analysis):** How does the proposed Bayesian dynamic estimation mechanism contribute to the robustness and effectiveness of the alignment process?
- **RQ3 (Generalization):** Can BLADE serve as a general framework to effectively optimize diverse list-wise objectives beyond standard relevance?

4.1 Experimental Setup

4.1.1 Datasets. We evaluate our method on three real-world datasets: *Amazon CDs and Vinyl*¹, *Steam*², and *Goodreads*³, with detailed statistics provided in Table 2. Following [5, 13], we filtered out users with fewer than 10 interactions. For the Steam and Goodreads datasets, we retained the top 15% most popular items to ensure dense interaction history. For list-wise sequence construction, we applied a sliding window approach, utilizing the past 10 items as the user history context and the subsequent 10 items as the target recommendation list. Regarding the training setup, we adhered to a standard two-stage paradigm: the base model was first super-vised fine-tuned on the full dataset to acquire domain knowledge; subsequently, to balance computational costs while leveraging the data efficiency of preference alignment, the alignment phase was conducted on a sampled subset of 4,096 training and 512 validation instances. Crucially, to ensure unbiased evaluation, all results are reported on the complete test set.

Table 2: Summary of dataset statistics.

Dataset	# Users	# Items	# Interactions
Steam	3,050	1,473	23,217
Goodreads	5,762	3,261	101,292
Amazon CDs	21,347	13,078	185,855

4.1.2 Baselines. We compare BLADE against the following baselines:

- **BIGRec:** [3] A standard SFT baseline that fine-tunes LLMs for recommendation tasks, serving as the foundation for subsequent alignment methods.
- **S-DPO:** [10] An adaptation of DPO for recommendation by introducing a softmax loss over multiple negative samples.
- **ReRe:** [34] An RL-based paradigm utilizing constrained beam search for efficient sampling and auxiliary ranking rewards for fine-grained supervision.
- **SPRec:** [13] A self-play framework that treats model generations as negative samples to achieve self-evolution and debiasing.
- **Beam Search:** [33] A representative point-wise decoding strategy that selects the top- K independent hypotheses based solely on sequence probability.
- **D³:** [4] A decoding approach that mitigates amplification bias and homogeneity by calibrating length normalization and penalizing repetitive generations.
- **List DPO:** [27] Follows the standard RLHF protocol to sample candidate list pairs and optimizes them via DPO based on relative rewards.
- **BoN Alignment:** [2] The static alignment strategy described in Section 2.3, which distills Best-of-N capabilities using a fixed reference.

¹http://jmcauley.ucsd.edu/data/amazon/index_2014.html

²https://cseweb.ucsd.edu/~jmcauley/datasets.html#amazon_reviews

³<https://mengtingwan.github.io/data/goodreads>

4.1.3 Evaluation Metrics. We employ two widely adopted metrics to evaluate recommendation accuracy: **NDCG@ k** and **Recall@ k** , reporting performance with $k \in \{3, 5\}$. NDCG assesses the ranking quality by assigning higher scores to relevant items at top positions, while Recall measures the proportion of relevant items successfully retrieved. Additionally, to assess the comprehensive alignment capabilities (RQ3), we utilize:

Mean Group Unfairness (MGU) [19]. MGU measures the disparity between the generated list L and user history H . It calculates the Mean Absolute Error (MAE) of genre distributions over the genre set \mathcal{C} :

$$\text{MGU} = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} |P(c|L) - P(c|H)|, \quad (16)$$

where $P(c|\cdot)$ denotes the normalized frequency of genre c .

Intra-List Diversity (ILD) [45]. ILD evaluates diversity by computing the average pairwise Jaccard distance between items in list L of size N . Let g_i be the genre set of item i :

$$\text{ILD} = \frac{1}{N(N-1)} \sum_{i \neq j} \left(1 - \frac{|g_i \cap g_j|}{|g_i \cup g_j|} \right). \quad (17)$$

A higher ILD indicates lower redundancy among recommended items.

4.1.4 Implementation Details. We utilize **Llama-3.2-1B-Instruct** [16] as the backbone and conduct all experiments on 4 NVIDIA A100 GPUs. The policy is first initialized via standard SFT on the full dataset for 3 epochs with a learning rate of 1×10^{-4} . For BLADE optimization, we train for 3 epochs using the AdamW optimizer with a learning rate of 5×10^{-6} and a linear decay scheduler. We set the group size to $G = 16$ and the reference set size to $M = 128$. The dynamic coefficient τ is tuned within the set $\{0.1, 0.3, 0.5, 1.0\}$. During the training phase, we employ a sampling strategy with a temperature of 0.8 and a repetition penalty of 1.05 to encourage exploration.

Regarding baselines, we generally follow the default configurations of the original papers. Notably, for DPO-based methods (S-DPO, List-DPO), we perform a warm-up SFT on the positive samples of the alignment subset prior to preference optimization. This ensures adequate distribution coverage and training stability. For DPO training, the KL penalty β is set to 0.1, and for S-DPO, we utilize 3 negative samples per positive instance. Crucially, to ensure reproducibility, evaluations for all training-based methods are performed using greedy decoding. For decoding-specific baselines (i.e., Beam Search and D^3), we adhere to their standard protocols.

4.2 Overall Performance

Table 3 presents the overall performance comparison of BLADE against state-of-the-art baselines across three real-world datasets. The results demonstrate that **BLADE consistently outperforms all competitive methods**, achieving superior scores on the majority of metrics across all datasets. Our analysis yields the following key observations:

Breaking the Limits of Static BoN Alignment. The most significant finding is that BLADE substantially surpasses existing list-wise approaches. On the **Amazon CDs** dataset, **BLADE-N** achieves a

Table 3: Overall performance comparison. The background colors distinguish different categories of methods: LLM-based , Inference Strategies , List-wise , and Ours . $R@k$ and $N@k$ denote Recall@ k and NDCG@ k , respectively. Specific to our methods, BLADE-R and BLADE-N utilize Recall and NDCG as the optimization reward signal, respectively. The best results are highlighted in bold, and the second-best results are underlined.

Method	Amazon CDs and Vinyl				Steam				Goodreads			
	R@3	N@3	R@5	N@5	R@3	N@3	R@5	N@5	R@3	N@3	R@5	N@5
BIGRec	0.0079	0.0331	0.0103	0.0283	0.0101	0.0375	0.0158	0.0357	0.0166	0.0602	0.0195	0.0545
S-DPO	0.0099	0.0391	0.0120	0.0319	0.0093	0.0329	0.0147	0.0317	0.0175	0.0642	0.0195	0.0605
ReRe	0.0110	0.0410	0.0135	0.0345	0.0106	0.0382	0.0157	0.0358	0.0176	0.0644	0.0202	0.0584
SPRec	0.0104	0.0409	0.0125	0.0336	0.0109	0.0397	<u>0.0168</u>	0.0374	0.0175	0.0639	0.0195	0.0606
Beam Search	0.0069	0.0251	0.0075	0.0190	0.0070	0.0253	0.0073	0.0187	0.0013	0.0039	0.0013	0.0028
D ³	0.0094	0.0345	0.0119	0.0284	0.0069	0.0232	0.0107	0.0221	0.0118	0.0409	0.0163	0.0358
List DPO	0.0095	0.0380	0.0121	0.0308	0.0096	0.0336	0.0151	0.0325	0.0177	0.0649	0.0196	<u>0.0612</u>
BoN Alignment	0.0098	0.0401	0.0120	0.0333	<u>0.0114</u>	<u>0.0416</u>	0.0164	0.0384	0.0179	0.0653	<u>0.0212</u>	0.0588
BLADE-R	0.0130	<u>0.0451</u>	0.0156	<u>0.0379</u>	0.0116	0.0419	0.0171	0.0395	0.0184	<u>0.0672</u>	0.0219	0.0606
BLADE-N	<u>0.0119</u>	0.0474	<u>0.0144</u>	0.0410	0.0111	0.0407	0.0155	<u>0.0393</u>	<u>0.0183</u>	0.0686	0.0206	0.0618




Figure 4: Study on CDs and Vinyl datasets. (a) Comparison of Recall evolution during training with $\tau = 0$ and $\tau = 0.3$. (b) Performance sensitivity of Test Recall with respect to varying τ .

remarkable **23.1%** improvement in NDCG@5 and an **18.2%** improvement in NDCG@3 compared to the strongest list-wise baseline, BoN Alignment. This empirical evidence strongly validates the effectiveness of our approach in overcoming the performance upper bound of static reference models, unlocking the potential to discover high-quality ranking lists that static methods fail to capture.

Superiority over Standard LLM-based Methods. Compared to standard LLM-based alignment baselines, BLADE demonstrates a clear performance advantage. While methods like S-DPO and BIGRec show improvements over the base SFT model, they still fail to match BLADE’s capability in optimizing holistic list-level utility. This indicates that directly optimizing global list-wise objectives via our Bayesian framework is significantly more effective than optimizing the local preference signals utilized in standard alignment techniques.




Figure 5: Effect of the number of generated items (G) in BLADE

Inadequacy of Inference-time Strategies. While inference-time strategies attempt to optimize recommendation lists without parameter updates, our results reveal their fundamental limitations. The suboptimal performance of Beam Search is directly attributed to its inherent independence assumption. We empirically observe that independently sampling multiple items often leads to severe semantic repetition, as the model greedily collapses into a narrow distribution of high-probability candidates. Although D³ alleviates this homogeneity through explicit decoding penalties, it fundamentally fails to account for the complex relationships between items. Such heuristic constraints merely suppress repetition without capturing the holistic coherence or compatibility required for high-quality lists. In contrast, BLADE internalizes these complex inter-item dependencies directly into the model parameters, naturally generating diverse and coherent recommendations.

Optimization Flexibility and Controllability. A unique advantage of our framework is its ability to align with specific user-defined objectives. As evidenced by the performance trade-off between BLADE-R and BLADE-N, the model effectively adapts to the chosen optimization target. Specifically, on the Goodreads and Amazon CDs datasets, BLADE-R excels in Recall metrics, whereas BLADE-N consistently achieves the highest scores on NDCG metrics. This




Figure 6: Robustness and data efficiency analysis across varying generation sizes (G). The bars (left axis) compare the absolute Recall@5 performance. The overlaid line (right axis) tracks the relative improvement, highlighting substantial gains in low-resource regimes.

distinction highlights the **controllability** of BLADE: it explicitly steers the policy towards the specific reward landscape required by the downstream recommendation task.

4.3 Analysis of Bayesian Dynamic Estimation

To empirically validate our proposed framework, we conduct an in-depth analysis of the Bayesian dynamic estimation mechanism. Specifically, we investigate how the dynamic coefficient τ mitigates the gradient decay problem and how the Bayesian prior **compensates for sampling sparsity** under varying inference budgets.

4.3.1 Impact of Dynamic Coefficient τ . We first verify the critical role of the dynamic coefficient τ in sustaining optimization progress. Fig. 4(a) illustrates the Recall curves during training, while Fig. 4(b) presents the performance sensitivity with respect to τ .

Mitigating Gradient Decay. As shown in Figure 4(a), the training dynamics of the static baseline ($\tau = 0$) and BLADE ($\tau = 0.3$) diverge significantly. The static baseline improves initially but rapidly plateaus. This stagnation empirically validates our hypothesis regarding *Gradient Decay*: as the model improves, the fixed reference distribution fails to provide informative signals, causing the training to stall prematurely. In contrast, BLADE sustains a steady improvement, confirming that updating the target distribution dynamically ensures effective learning throughout the process.

Balancing Prior and Evidence. Figure 4(b) further illustrates the effect of τ . We observe that the performance initially improves as τ increases but subsequently degrades when τ becomes too large. A small τ over-anchors the model to the suboptimal static prior, restricting exploration. Conversely, a large τ introduces high variance due to the limited batch samples. This indicates that a moderate τ strikes an optimal balance, leveraging the stability of the static prior while effectively adapting to the model’s growing capabilities.

4.3.2 Robustness and Estimation Stability.




Figure 7: Training diagnostics under standalone fairness-objective optimization. Recall@5 rapidly decays toward zero, while MGU also decreases during training. This indicates that directly optimizing the auxiliary fairness objective can improve the target list-wise metric but severely degrades recommendation relevance.

Scaling with Compute. First, as shown in Fig. 5, we observe that both NDCG@5 and Recall@5 improve consistently as G increases from 2 to 16. This confirms that our method effectively benefits from a larger search space during training. This scalability implies that the model can achieve further alignment gains when more computational resources are available.

Stabilizing Role of the Prior. More importantly, we highlight the stabilizing role of the Bayesian prior. As shown in Fig. 6, BLADE consistently outperforms the "No Prior" baseline. The relative improvement is most significant when the group size is small. In this scenario, the dynamic evidence from the current batch is unstable due to sampling sparsity, and the static prior acts as a stabilizer, preventing estimation biases. As G increases, the gap between the two methods narrows because the dynamic evidence becomes more reliable. This demonstrates that BLADE is highly robust: it yields substantial improvements over the baseline even with a minimal sampling cost, making it suitable for resource-constrained applications.

4.4 Generalizability to Diverse List-wise Targets

In this section, we investigate the generalizability of BLADE as a metric-agnostic alignment framework. Specifically, we evaluate its performance on two distinct objectives: MGU and ILD.

Composite Reward Formulation. Empirically, we found that optimizing the policy solely with auxiliary list-wise objectives can lead to rapid degradation in recommendation accuracy. These objectives capture distributional properties of the generated list, but do not by themselves provide a relevance-grounding signal for recommendation. As shown in Figure 7, standalone fairness-objective optimization reduces MGU during training, but the Recall@5 signal quickly decays toward zero. This indicates that the policy can improve the auxiliary objective while drifting away from relevant recommendation behavior. Consequently, we formulate a **composite reward** by incorporating the relevance score R_{NDCG} as a

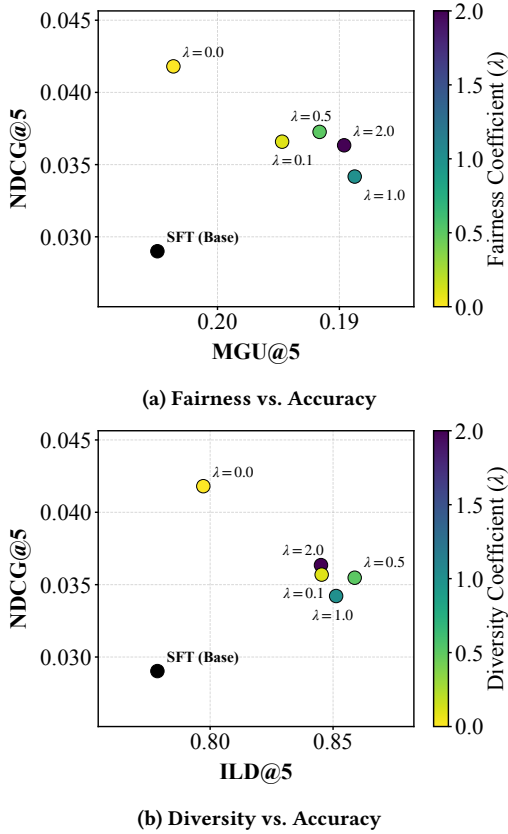


Figure 8: Performance comparison on list-wise metrics. (a) Fairness (MGU, lower is better) vs. Accuracy. (b) Diversity vs. Accuracy. Both figures demonstrate BLADE’s versatility, achieving substantial improvements in the respective list-wise metrics while simultaneously maintaining superior accuracy over the SFT baseline.

grounding signal to maintain utility. The rewards are defined as:

$$R_{\text{fair}}(\mathbf{y}) = R_{\text{NDCG}}(\mathbf{y}) - \lambda \cdot \text{MGU}(\mathbf{y}), \quad (18)$$

$$R_{\text{div}}(\mathbf{y}) = R_{\text{NDCG}}(\mathbf{y}) + \lambda \cdot \text{ILD}(\mathbf{y}), \quad (19)$$

where λ determines the relative weight of the fairness/diversity reward.

Optimization of Fairness. We evaluate the model’s ability to mitigate unfairness using Mean Group Unfairness (MGU). The results in Figure 8(a) reveal two key findings:

- (1) **Pareto Dominance:** The SFT baseline resides in a suboptimal region characterized by low accuracy and high unfairness. In contrast, BLADE achieves strict Pareto dominance: all experimental points, regardless of λ , are located to the top-right of the baseline. This demonstrates that BLADE identifies a superior solution space where accuracy and fairness coexist, rather than strictly trading one for the other.
- (2) **Non-monotonic Sensitivity:** We observe that the improvement in fairness is not strictly monotonic with respect to λ . While introducing λ significantly reduces MGU compared to

the unconstrained case ($\lambda = 0$), increasing λ beyond a certain threshold yields diminishing returns in fairness while incurring a cost in NDCG. This suggests the existence of a “sweet spot” (e.g., $\lambda \approx 0.5$) in the composite reward landscape.

Optimization of Intra-List Diversity. We further assess the model’s capability to optimize Diversity. As shown in Figure 8(b):

- (1) **Significant Diversity Gain:** Both SFT and the unconstrained BLADE ($\lambda = 0$) exhibit low diversity, tending to recommend redundant popular items. However, introducing even a minimal coefficient ($\lambda = 0.1$) triggers a **rapid adaptation**, boosting the ILD score significantly. This confirms that BLADE is highly responsive to sparse high-diversity signals that are typically ignored by static baselines.
- (2) **High-Utility Diversity:** Crucially, even at the peak diversity point ($\lambda = 0.5$), BLADE maintains an NDCG significantly higher than the SFT baseline. This indicates that the diversity gain is not achieved through random exploration, but by successfully retrieving high-quality items that are both relevant and diverse.

In summary, these empirical findings firmly establish BLADE as a robust, metric-agnostic alignment framework capable of transcending the limitations of single-objective optimization. Unlike specialized methods tailored for narrow tasks, BLADE demonstrates exceptional versatility in handling complex, non-differentiable list-wise metrics through a unified Bayesian approach.

5 Related work

In this section, we first briefly review LLM-based recommender systems and the challenges associated with list-wise optimization. Subsequently, we introduce the Best-of-N sampling method and its fine-tuning variants.

5.1 LLMs for Recommendation

The integration of Large Language Models (LLMs) into recommender systems has created a new paradigm, leveraging their extensive world knowledge and semantic reasoning capabilities. Existing research can be broadly categorized into three architectures: (1) **Feature Encoders** [8, 28, 35], which utilize the internal representations of LLMs to enhance the semantic embeddings of traditional collaborative filtering models; (2) **User Modeling** [26, 42], where LLMs serve as user profilers to infer latent interests or generate natural language explanations for recommendations; and (3) **Generative Recommendation** [13, 21, 23, 25], the focus of this work, which formulates recommendation as a sequence-to-sequence generation task. In this paradigm, the model is prompted with user history and directly outputs item titles or identifiers.

Despite the promise of Generative Recommendation, pre-trained LLMs often suffer from domain shift and lack the specific collaborative filtering signals required for accurate ranking [5]. While Supervised Fine-Tuning (SFT) [3, 15] and Direct Preference Optimization (DPO) [10, 22] have been widely adopted to align LLMs with recommendation data, they typically rely on token-level cross-entropy or pair-wise ranking objectives. These local objectives are often

misaligned with the holistic, list-wise metrics (e.g., NDCG, Diversity) that define user satisfaction in real-world scenarios, creating a gap between training objectives and inference-time performance.

5.2 List-wise Recommendation

Standard point-wise recommender systems learn individual preferences [18, 20], misaligning with real-world scenarios where users are exposed to complete item lists [1, 6]. Consequently, user satisfaction relies on holistic list-level metrics (e.g., NDCG) that capture global item interactions. However, optimizing these metrics is fundamentally challenging because they are inherently complex, discrete, and non-differentiable, making them difficult to address with standard gradient-based objectives.

To address this challenge, industrial systems often employ a “Generator-Evaluator” pipeline for re-ranking [40, 43], which is difficult to unify into the end-to-end generative paradigm. Existing approaches utilize differentiable surrogate losses to approximate discrete ranking metrics [7, 41]. However, such approximation inevitably suffers from estimation bias and poor generalization, as the smooth surrogate objectives diverge from the actual, non-differentiable evaluation protocols. Alternatively, Direct Reinforcement Learning theoretically supports non-differentiable rewards but proves notoriously unstable in practice, exhibiting high variance and sensitivity to reward scaling [9]. To overcome these limitations, we propose BLADE, which bypasses gradient constraints by internalizing list-wise search capabilities directly into the model weights.

5.3 Best-of-N Sampling and Alignment

Best-of-N (BoN) is a robust inference-time strategy proposed to optimize language generation quality by selecting the best response from N candidates [32]. Despite its simplicity, BoN has been proven to rival or even surpass complex alignment methods [17]. However, the requirement for N times the inference latency renders it computationally inefficient for practical applications.

To address this efficiency bottleneck, **BoN Alignment** aims to distill the search capability into the model weights. Early approaches employ *offline distillation* by fine-tuning the model on the best samples generated by a teacher policy [12]. More recent methods formulate BoN matching as a constrained optimization problem, solving it via DPO on best-worst pairs or online Reinforcement Learning [2, 17, 29]. While the latter achieves a better balance between maximizing reward and maintaining proximity to the reference model, we argue that it still suffers from Indiscriminate Supervision, and Gradient Decay, limiting their effectiveness in recommendation. Therefore, we propose BLADE, which leverages dynamic estimation of the real-time BoN distribution. This approach surpasses the static performance upper bound without additional training costs, achieving superior optimization for list-wise recommendation.

6 Conclusion

In this work, we addressed the critical limitation of standard LLM4Rec methods, which rely on token-level objectives and consequently struggle to optimize complex, non-differentiable list-wise metrics. While BoN alignment offers a potential solution, we identified two

fundamental bottlenecks in its static implementation: **Indiscriminate Supervision**, where the static reference fails to distinguish superior candidates and **Gradient Decay** due to diminishing supervision signals. To overcome these challenges, we proposed **BLADE**, a Bayesian List-wise Alignment framework via Dynamic Estimation. By reformulating the alignment target as a self-evolving posterior that fuses historical priors with real-time dynamic evidence, BLADE ensures that the optimization process remains informative and effective throughout the training trajectory. Extensive experiments demonstrate that BLADE not only significantly outperforms state-of-the-art baselines in ranking accuracy but also effectively breaks the performance upper bound imposed by static references. Furthermore, our analysis confirms that BLADE serves as a robust, metric-agnostic framework capable of navigating the Pareto frontier between accuracy and complex list-wise objectives such as fairness and diversity.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (62402470, U24B20180, 62525211), the Fundamental Research Funds for the Central Universities of China (WK2100000053), Anhui Provincial Natural Science Foundation (2408085QF189). This research is supported by the advanced computing resources provided by the Supercomputing Center of the USTC.

References

- [1] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 135–144.
- [2] Afra Amini, Tim Vieira, Elliott Ash, and Ryan Cotterell. 2024. Variational best-of-n alignment. *arXiv preprint arXiv:2407.06057* (2024).
- [3] Keqin Bao, Jizhi Zhang, Wenjie Wang, Yang Zhang, Zhengyi Yang, Yanchen Luo, Chong Chen, Fuli Feng, and Qi Tian. 2025. A Bi-Step Grounding Paradigm for Large Language Models in Recommendation Systems. *ACM Transactions on Recommender Systems (TORS)* (2025).
- [4] Keqin Bao, Jizhi Zhang, Yang Zhang, Xinyue Huo, Chong Chen, and Fuli Feng. 2024. Decoding Matters: Addressing Amplification Bias and Homogeneity Issue for LLM-based Recommendation. *EMNLP* (2024).
- [5] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1007–1014.
- [6] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*. 129–136.
- [7] Wen-Shuo Chao, Zhi Zheng, Hengshu Zhu, and Hao Liu. 2024. Make large language model a better ranker. *arXiv preprint arXiv:2403.19181* (2024).
- [8] Junyi Chen, Lu Chi, Bingyue Peng, and Zehuan Yuan. 2024. Hllm: Enhancing sequential recommendations via hierarchical large language models for item and user modeling. *arXiv preprint arXiv:2409.12740* (2024).
- [9] Minmin Chen, Alex Beutel, Paul Covington, Sagar Jain, Francois Belletti, and Ed H Chi. 2019. Top-k off-policy correction for a REINFORCE recommender system. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 456–464.
- [10] Yuxin Chen, Junfei Tan, An Zhang, Zhengyi Yang, Leheng Sheng, Enzhi Zhang, Xiang Wang, and Tat-Seng Chua. 2024. On Softmax Direct Preference Optimization for Recommendation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS '24)*.
- [11] Sunhao Dai, Jiakai Tang, Jiahua Wu, Kun Wang, Yuxuan Zhu, Bingjun Chen, Bangyang Hong, Yu Zhao, Cong Fu, Kangle Wu, et al. 2025. Onepiece: Bringing context engineering and reasoning to industrial cascade ranking system. *arXiv preprint arXiv:2509.18091* (2025).
- [12] Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767* (2023).
- [13] Chongming Gao, Ruijun Chen, Shuai Yuan, Kexin Huang, Yuanqing Yu, and Xiangnan He. 2025. Sprec: Self-play to debias llm-based recommendation. In

- Proceedings of the ACM on Web Conference 2025*. 5075–5084.
- [14] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM conference on recommender systems*. 299–315.
 - [15] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as Language Processing (RLP): A Unified Pretrain, Personalized Prompt & Predict Paradigm (P5). In *Proceedings of the 16th ACM Conference on Recommender Systems (RecSys '22)*. 299–315.
 - [16] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
 - [17] Lin Gui, Cristina Gărbacea, and Victor Veitch. 2024. Bonbon alignment for large language models and the sweetness of best-of-n sampling. *Advances in Neural Information Processing Systems* 37 (2024), 2851–2885.
 - [18] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
 - [19] Meng Jiang, Keqin Bao, Jizhi Zhang, Wenjie Wang, Zhengyi Yang, Fuli Feng, and Xiangnan He. 2024. Item-side Fairness of Large Language Model-based Recommendation System. In *Proceedings of the ACM on Web Conference 2024 (WWW '24)*. 4717–4726.
 - [20] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
 - [21] Xiaoyu Kong, Leheng Sheng, Junfei Tan, Yuxin Chen, Jiancan Wu, An Zhang, Xiang Wang, and Xiangnan He. 2025. MiniOneRec: An Open-Source Framework for Scaling Generative Recommendation. *arXiv preprint arXiv:2510.24431* (2025).
 - [22] Jiayi Liao, Xiangnan He, Ruobing Xie, Jiancan Wu, Yancheng Yuan, Xingwu Sun, Zhanhui Kang, and Xiang Wang. 2024. RosePO: Aligning LLM-based Recommenders with Human Values. *arXiv preprint arXiv:2410.12519* (2024).
 - [23] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. LLaRA: Large Language-Recommendation Assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*. 1785–1795.
 - [24] Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Hao Zhang, Yong Liu, Chuhan Wu, Xiangyang Li, Chenxu Zhu, et al. 2025. How can recommender systems benefit from large language models: A survey. *ACM Transactions on Information Systems* 43, 2 (2025), 1–47.
 - [25] Xinyu Lin, Haihan Shi, Wenjie Wang, Fuli Feng, Qifan Wang, See-Kiong Ng, and Tat-Seng Chua. 2025. Order-agnostic identifier for large language model-based generative recommendation. In *Proceedings of the 48th international ACM SIGIR conference on research and development in information retrieval*. 1923–1933.
 - [26] Zhanyu Liu, Shiyao Wang, Xingmei Wang, Rongzhou Zhang, Jiaxin Deng, Honghui Bao, Jinghao Zhang, Wuchao Li, Pengfei Zheng, Xiangyu Wu, et al. 2025. Onerec-think: In-text reasoning for generative recommendation. *arXiv preprint arXiv:2510.11639* (2025).
 - [27] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* 36 (2024).
 - [28] Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. Representation learning with large language models for recommendation. In *Proceedings of the ACM web conference 2024*. 3464–3475.
 - [29] Pier Giuseppe Sessa, Robert Dadashi, Léonard Hussenot, Johan Ferret, Nino Vieillard, Alexandre Ramé, Bobak Shariari, Sarah Perrin, Abe Friesen, Geoffrey Cideron, et al. 2024. Bond: Aligning llms with best-of-n distillation. *arXiv preprint arXiv:2407.14622* (2024).
 - [30] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and Daya Guo. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *ArXiv abs/2402.03300* (2024). <https://api.semanticscholar.org/CorpusID:267412607>
 - [31] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314* (2024).
 - [32] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in neural information processing systems* 33 (2020), 3008–3021.
 - [33] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems* 27 (2014).
 - [34] Junfei Tan, Yuxin Chen, An Zhang, Junguang Jiang, Bin Liu, Ziru Xu, Han Zhu, Jian Xu, Bo Zheng, and Xiang Wang. 2025. Reinforced Preference Optimization for Recommendation. *arXiv preprint arXiv:2510.12211* (2025).
 - [35] Wenjie Wang, Honghui Bao, Xinyu Lin, Jizhi Zhang, Yongqi Li, Fuli Feng, See-Kiong Ng, and Tat-Seng Chua. 2024. Learnable item tokenization for generative recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 2400–2409.
 - [36] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
 - [37] Sean Welleck, Amanda Bertsch, Matthew Finlayson, Hailey Schoelkopf, Alex Xie, Graham Neubig, Ilya Kulikov, and Zaid Harchaoui. 2024. From decoding to meta-generation: Inference-time algorithms for large language models. *arXiv preprint arXiv:2406.16838* (2024).
 - [38] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, Hui Xiong, and Enhong Chen. 2024. A Survey on Large Language Models for Recommendation. *World Wide Web* 27, 5 (Aug. 2024), 31 pages.
 - [39] Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. 2024. Inference scaling laws: An empirical analysis of compute-optimal inference for problem-solving with language models. *arXiv preprint arXiv:2408.00724* (2024).
 - [40] Hailan Yang, Zhenyu Qi, Shuchang Liu, Xiaoyu Yang, Xiaobei Wang, Xiang Li, Lantao Hu, Han Li, and Kun Gai. 2025. Comprehensive list generation for multi-generator reranking. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2298–2308.
 - [41] Weiqin Yang, Jiawei Chen, Shengjia Zhang, Peng Wu, Yuegang Sun, Yan Feng, Chun Chen, and Can Wang. 2025. Breaking the top-k barrier: Advancing top-k ranking metrics optimization in recommender systems. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*. 3542–3552.
 - [42] Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun, Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2024. Agentcf: Collaborative learning with autonomous language agents for recommender systems. In *Proceedings of the ACM Web Conference 2024*. 3679–3689.
 - [43] Kaike Zhang, Xiaobei Wang, Xiaoyu Yang, Shuchang Liu, Hailan Yang, Xiang Li, Fei Sun, and Qi Cao. 2025. From generation to consumption: Personalized list value estimation for re-ranking. *arXiv preprint arXiv:2508.02242* (2025).
 - [44] Xiangyu Zhao, Liang Zhang, Long Xia, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2017. Deep reinforcement learning for list-wise recommendations. *arXiv preprint arXiv:1801.00209* (2017).
 - [45] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. In *The Web Conference*.