

Optimizing User Profiles via Contextual Bandits for Retrieval-Augmented LLM Personalization

Linfeng Du^{1,6} Ye Yuan^{1,6} Zichen Zhao²

Fuyuan Lyu^{1,6} Emiliano Penaloza^{3,6} Xiuying Chen² Zipeng Sun^{1,6}

Jikun Kang⁴ Laurent Charlin^{5,6} Xue Liu^{1,2,6} Haolun Wu^{1,6§}

¹McGill University ²Mohamed bin Zayed University of Artificial Intelligence

³Université de Montréal ⁴Salesforce ⁵HEC Montréal ⁶Mila - Quebec AI Institute

{linfeng.du, haolun.wu}@mail.mcgill.ca

Abstract

Large language models (LLMs) excel at general-purpose tasks, yet adapting their responses to individual users remains challenging. Retrieval augmentation provides a lightweight alternative to fine-tuning by conditioning LLMs on user history records, and existing approaches typically select these records based on semantic relevance. We argue that relevance serves as an unreliable proxy for utility: a record may be semantically similar to a query yet fail to improve generation quality or even degrade it due to redundancy or conflicting information. To bridge this gap, we propose **PURPLE**, a contextual bandit framework that optimizes User Profiles for LLM personalization. In contrast to a greedy selection of the most relevant records, PURPLE treats profile construction as an order-sensitive generation process and utilizes a Plackett-Luce ranking model to capture complex inter-record dependencies. By training with semantically rich feedback provided by the likelihood of the reference response, our method aligns retrieval directly with generation quality. Extensive experiments on nine personalization tasks demonstrate that PURPLE consistently outperforms strong heuristic and retrieval-augmented baselines in both effectiveness and efficiency, establishing a principled and scalable solution for optimizing user profiles¹.

1 Introduction

Large language models (LLMs) have demonstrated remarkable success in various natural language processing tasks. As these models are increasingly applied to personalized applications, tailoring responses to individuals based on their own preferences has become a crucial challenge. Existing approaches for personalizing LLMs, such

[§]Corresponding author.

¹Our code is available [here](#).

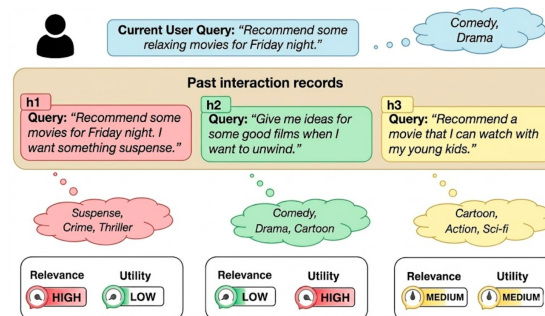


Figure 1: An illustration of the discrepancy between relevance and utility for a user’s past interaction records given a query. The thought bubbles illustrate the potential movie genres that satisfy the underlying user intent.

as reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), generally require modifying model parameters. These approaches incur high computational costs, demand frequent updates, and are impractical for real-time personalization at scale, especially when the LLM is not fully open-sourced. Moreover, continually fine-tuning models for different individuals would complicate safety evaluation since each personalized variant would require separate testing.

In this paper, we study a lightweight approach to LLM personalization via retrieval augmentation (Wu et al., 2025), in which user profiles are constructed by retrieving and injecting past user records into the prompt. Prior work shows that such profiles can effectively steer LLM outputs toward individual preferences (Salemi et al., 2024; Jiang et al., 2025). Compared to parameter-updating methods, retrieval-augmented personalization is lightweight, transparent, and readily deployable, as users can directly inspect and modify the records guiding generation. However, naively including the full user history introduces redundancy and noise, and may exceed the model’s context window, while aggressive pruning risks discarding critical personalization signals. This tension raises a central challenge: **which user records should be selected to**

construct a user profile that maximizes personalization utility, i.e., downstream task performance when conditioned on by the LLM?

Existing strategies for building user profiles often rely on heuristics, such as selecting records with the highest *relevance*, defined as lexical or semantic similarity to the query (Karpukhin et al., 2020). However, we argue that *relevance* is not a reliable proxy for *utility*, which measures the extent to which a record helps the LLM generate an appropriate personalized response. Relevance-based heuristics conflate surface-level contextual similarity with the underlying preference signals that drive personalization, prioritizing records that look similar to the query rather than those that support the user’s current intent. To illustrate this dichotomy, consider Figure 1, where the user seeks “*relaxing movies for Friday night*”. Relevance-based heuristics prioritize record h_1 due to the high lexical overlap of “*Friday night*”, ignoring the latent preference for “*suspense*” which contradicts the current need for relaxation. Conversely, record h_2 offers high utility by capturing the semantic intent to “*unwind*”, despite lacking surface-level keyword matches. This highlights that relevance scores prioritize how similar a record looks to the query, whereas utility depends on whether it supports the underlying user intent.

Beyond evaluating individual records, constructing an effective user profile requires list-aware reasoning rather than independent selection (Ai et al., 2018; Pang et al., 2020). Even if the utility of each record could be estimated accurately in isolation, greedy top- k selection remains suboptimal because record utilities do not compose additively. Instead, the value of a record is contingent on the context provided by others, as irrelevant or conflicting information within the retrieved profile can significantly degrade downstream generation (Liu et al., 2024; Yoran et al., 2024). Revisiting Figure 1, while h_2 and h_3 individually appear useful by reflecting a shared preference for lighthearted content, their combination is suboptimal. Jointly prompting with h_2 and h_3 introduces additional constraints by combining “*comedy*” preferences with “*young kids*” requirements, which could dilute the conditioning signal, thereby degrading generation quality relative to using h_2 alone. This confirms that personalization utility is non-monotonic. Simply accumulating *useful* records may introduce interference and actively degrade performance.

The limitations highlighted above call for a

reranking mechanism that is directly optimized for downstream generation and sensitive to interactions among records. Existing approaches fall short of these requirements: heuristic retrieval relies on static proxies like similarity, while recent listwise rerankers, though capable of modeling dependencies, remain constrained by relevance-oriented supervision. To bridge this gap, we propose **PURPLE**, a framework that optimizes User Profiles for LLM personalization via contextual bandits (Langford and Zhang, 2007). In this formulation, the context consists of both the current query and the user’s past records. The selection policy is guided by a reward function reflecting downstream personalized generation performance. PURPLE outputs a *propensity score* for each user record, which parameterizes a Plackett-Luce ranking model to produce the final user profile. This formulation enables the model to capture interactions between records and adaptively select those that are most beneficial as a whole. We train PURPLE end-to-end using the policy gradient method (Sutton et al., 1999).

Our main contributions are as follows:

- We introduce PURPLE, a framework that casts retrieval-augmented LLM personalization as a contextual bandit problem, adaptively optimizing user profiles beyond static heuristics.
- We show through extensive experiments on nine personalization tasks, covering classification, regression, and short- and long-text generation, that PURPLE consistently outperforms strong baselines in both effectiveness and efficiency.
- We perform comprehensive ablation studies on PURPLE’s key design choices and further validate its effectiveness through extended analysis and human evaluation.

2 Related work

LLMs for Personalization. LLMs demonstrate strong performance across domains (OpenAI et al., 2024), yet their outputs often diverge from user expectations because pre-training captures general rather than individual needs. Reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022) is widely employed to align pre-trained models with user preferences; however, while parameter-efficient fine-tuning (PEFT) (e.g., LoRA (Hu et al., 2022)) drastically lowers its training costs, the necessity of model fine-tuning remains an impractical barrier for individual end

users. A complementary direction personalizes LLMs through user profiles (Salemi et al., 2024) constructed from prior user interactions. Integrating these profiles has demonstrated substantial benefits across various tasks requiring personalization, including recommendation (Penaloza et al., 2025), summarization (Zhang et al., 2024), question answering (Wu et al., 2024), content generation (Shen et al., 2024), and chatbot interaction (Jiang et al., 2025). Yet, it remains unclear which specific user records within a profile truly drive performance improvements, particularly in retrieval-augmented generation (RAG), where efficacy hinges on retrieving semantically relevant context. Moreover, little analysis has been conducted on how to optimally select and assemble user records into profiles with high personalization utility. Our work addresses this gap by investigating how user profiles shape personalization in retrieval-augmented LLMs, and by proposing strategies for selecting user records to maximize downstream performance.

Retrieval-Augmented Generation. Retrieval-augmented generation (RAG) augments LLMs with external knowledge sources to enhance factuality and coverage. Early architectures, such as REALM (Guu et al., 2020) and RAG (Lewis et al., 2020), jointly optimize the retriever and the LM, while Re2G (Glass et al., 2022) further incorporates a reranking module trained end-to-end with the other components. To mitigate computational costs, subsequent methods freeze the LM and perform in-context retrieval. For instance, IC-RALM (Ram et al., 2023) leverages LLMs for reranking, whereas REPLUG (Shi et al., 2024) distills retrievers from LLMs. More recently, instruction-tuned variants like SelfRAG (Asai et al., 2024) and RankRAG (Yu et al., 2024) jointly model retrieval and generation, but their reliance on large-scale fine-tuning renders them impractical for user-level personalization.

The methods most relevant to our work are IC-RALM and REPLUG; however, both methods process retrieved records in isolation, which is a limitation our method directly addresses. Specifically, REPLUG approximates joint reasoning by weighting individual generation outputs with retrieval probabilities, while IC-RALM periodically triggers retrieval during decoding and replaces previously used context. These designs arise because jointly reasoning over multiple records leads to a combinatorial explosion in the number of possible profiles. In contrast, our approach overcomes this limitation

by modeling inter-record dependencies and directly optimizing over multi-record profiles without resorting to such approximations.

LLMs for Reranking. LLM-based rerankers are broadly categorized into pointwise, pairwise, and listwise paradigms. Pointwise models, such as MonoBERT (Nogueira et al., 2019) and MonoT5 (Nogueira et al., 2020), score each query–document pair independently, whereas pairwise models like DuoT5 (Pradeep et al., 2021) evaluate candidates relative to one another. In contrast, listwise approaches such as ListT5 (Yoon et al., 2024) and EBCAR (Yuan et al., 2026) jointly model the full candidate set. As LLMs continue to expand their context lengths, the listwise paradigm has seen rapid advancement through prompt-only ranking (RankGPT (Sun et al., 2023)), model distillation (e.g., RankVicuna, RankZephyr, Lit5Distill, FIRST (Pradeep et al., 2023a,b; Tamber et al., 2023; Gangi Reddy et al., 2024)), and inference-time relevance extraction (ICR (Chen et al., 2025)). However, these methods inherently conflate relevance with utility, which is not ideal for personalization. In this work, we instead train rerankers using downstream generation quality as a feedback signal, prioritizing task utility over semantic similarity.

3 Methodology

We formulate retrieval-augmented LLM personalization as a contextual bandit problem (Langford and Zhang, 2007), where the goal is to learn a policy that selects informative user records. Unlike classic multi-armed bandits, contextual bandits incorporate auxiliary information (e.g., the current query and user history) before making a selection, enabling direct optimization of retrieval strategies through policy gradient reinforcement learning. This aligns the selection of user records with downstream personalization objectives.

3.1 Problem Formulation

We consider a dataset $\mathcal{D} = \{(\mathcal{H}^u, x^u, y^u)\}_{u=1}^{|\mathcal{D}|}$, where each example consists of a user’s collection of history records \mathcal{H}^u , a query x^u , and a reference personalized response y^u . Personalization is achieved by retrieving informative records from \mathcal{H}^u and supplying them as context to a frozen LLM, which then generates the final response. In practice, we apply PURPLE as a reranking module on top of a candidate pool selected by lightweight heuristics, ensuring low-latency inference compatible with

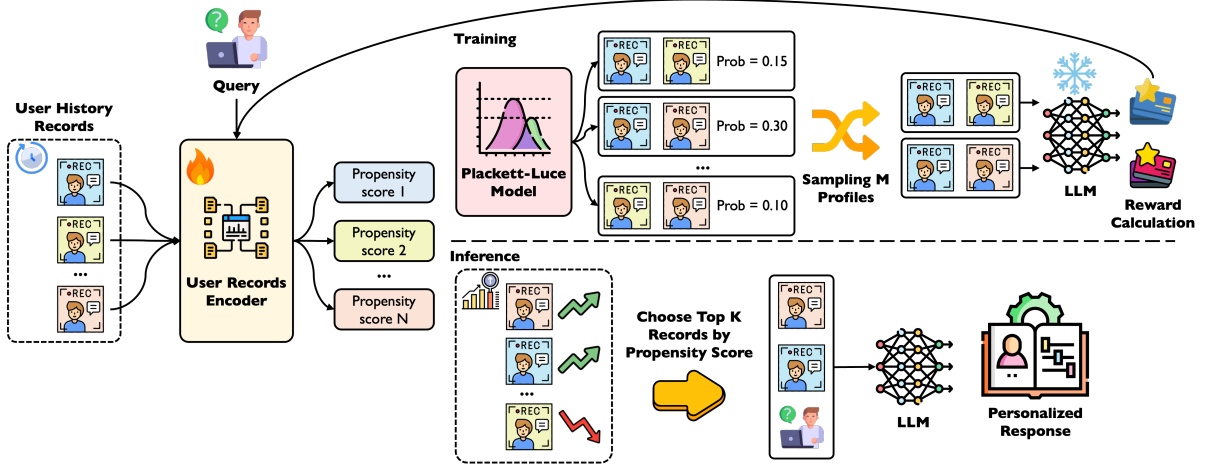


Figure 2: Workflow of the proposed PURPLE framework. User records encoder takes a user query and a list of user history records as input, outputting the propensity scores of all records. During **training**, a Plackett-Luce model is employed to convert the propensity scores into a probability distribution over all possible profiles, followed by sampling M profiles for policy gradient estimation. During **inference**, records with the top- K propensity scores are provided to the LLM along with the user query to generate a personalized response.

large-scale systems. In the following development, we focus on a single user and omit the superscript index for brevity.

Let $\mathcal{H} = \{h_1, \dots, h_N\}$ denote the set of N history records for a user, where each record $h_i = (x_i, y_i)$ is an input–output pair (e.g., a query and its answer from the user). Given a new query x , our goal is to construct a user profile from \mathcal{H} to condition the LLM for generating a personalized response. Formally, a profile is an ordered tuple $\mathcal{P} = \langle p_1, \dots, p_K \rangle \in \text{Perm}_K(\mathcal{H})$, which is a K -permutation of \mathcal{H} . We stress that the profile is order-sensitive: different permutations of the same K records correspond to distinct profiles and thus provide different inputs to the downstream LLM.

We formulate the selection of \mathcal{P} as a *contextual bandit problem*, where the context is given by the user’s history \mathcal{H} and the query x , and the action corresponds to selecting K records from \mathcal{H} to construct a profile. Formally, this formulation consists of the following key components:

- **Context:** $\mathcal{C} = (\mathcal{H}, x)$, where \mathcal{H} is the user’s collection of history records and x is the query. This representation captures both past user preferences and the immediate intent.
- **Actions:** $\mathcal{P} = \langle p_1, \dots, p_K \rangle \in \text{Perm}_K(\mathcal{H})$, which corresponds to selecting K distinct records from \mathcal{H} in a particular order. The action thus determines not only which records to use but also how they are arranged. The size of the action space is $N!/(N - K)!$.

- **Reward:** $R(\text{LLM}(\mathcal{P} \parallel x), y)$, a function that measures the quality of the LLM-generated response conditioned on the concatenation of the user profile \mathcal{P} and the query x , relative to the reference personalized response y .

We model the policy with a neural distribution $\pi_\theta(\cdot \mid \mathcal{C})$, parameterized by θ , which assigns probabilities to candidate user profiles given the context \mathcal{C} . The objective is to learn parameters θ such that the policy assigns higher probabilities to more informative profiles, which ultimately enhance personalized text generation. To this end, we maximize the expected reward over sampled user profiles, optimizing the following objective on a dataset \mathcal{D} spanning multiple users, each associated with a set of history records, a query, and the corresponding reference response:

$$\mathcal{J}(\theta) = \mathbb{E}_{(\mathcal{H}, x, y) \sim \mathcal{D}, \mathcal{P} \sim \pi_\theta(\cdot \mid \mathcal{C})} [R(\text{LLM}(\mathcal{P} \parallel x), y)]. \quad (1)$$

It is challenging to directly optimize Equation 1 since the reward is not differentiable. To address this, we employ the likelihood ratio gradient estimator from reinforcement learning and stochastic optimization (Williams, 1992; Sutton et al., 1999), which allows us to compute the gradient as:

$$\nabla_\theta \mathcal{J}(\theta) = \mathbb{E}_{(\mathcal{H}, x, y) \sim \mathcal{D}, \mathcal{P} \sim \pi_\theta(\cdot \mid \mathcal{C})} [\nabla_\theta \log \pi_\theta(\mathcal{P} \mid \mathcal{C}) R(\text{LLM}(\mathcal{P} \parallel x), y)]. \quad (2)$$

Since it is intractable to enumerate all profiles $\mathcal{P} \in \text{Perm}_K(\mathcal{H})$ during the optimization process,

we estimate Equation 2 by randomly sampling $M = 32$ profiles. To stabilize training and reduce variance in gradient estimation, we apply z-score normalization across the M sampled rewards for each example. The precise formulation of this estimator is detailed in Equation 5 of Appendix A.

3.2 Model and Function Design

Design of $\pi_\theta(\cdot | \mathcal{C})$ Since different permutations of the selected records may lead to different responses, we adopt the Plackett–Luce (PL) ranking model, which assigns probabilities to profiles based on the scores of individual user records. Therefore, $\pi_\theta(\cdot | \mathcal{C})$ defines a distribution over all $N!/(N - K)!$ permutations of length K drawn from the user’s N history records. The probability assigned to a specific profile \mathcal{P} is given by:

$$\pi_\theta(\mathcal{P} | \mathcal{C}) = \prod_{k=1}^K \frac{f_\theta(p_k; \mathcal{C})}{S - \sum_{j=1}^{k-1} f_\theta(p_j; \mathcal{C})}, \quad (3)$$

where $S = \sum_{i=1}^N f_\theta(h_i; \mathcal{C})$, and $f_\theta(\cdot)$ is the user record encoder that outputs a propensity score in $[0, 1]$ for each record, indicating the model’s tendency to include that record in the user profile. During training, profiles are generated by sampling K records without replacement based on Equation 3. At inference time, the top- K records ranked by propensity scores are selected. Because the PL model is order-sensitive and rewards are assigned to ordered sets, the learned propensity score can be interpreted as each record’s contribution to the selected profile. We empirically demonstrate in Section 5.3 that this ordering achieves higher final utility compared with other baselines.

Design of f_θ For the user record encoder f_θ , we aim to capture the interdependencies among user records. A key design consideration is the trade-off between modeling dependencies at the token level versus the record level. While the former could, in principle, capture finer-grained interactions, it would quickly exceed the encoder’s context length. To address this, we adopt a late interaction strategy (Khattab and Zaharia, 2020), where we first obtain record embeddings with a pre-trained encoder, and then apply a Transformer encoder to model dependencies across records. Figure 2 illustrates the overall workflow of our method. We utilize a pre-trained Contriever (Izacard et al., 2022) to obtain token embeddings for both the query and the records. Each record first cross-attends to the

query at the token level, producing query-fused token representations that incorporate query information. A subsequent pooling operation is applied to produce fixed-size record embeddings, which are then processed by a Transformer encoder to model inter-record dependencies. We omit positional encodings to avoid ordering bias among records.

Design of Reward Function In this work, we propose an LLM-driven reward, where the policy is trained to maximize the log-likelihood that the LLM assigns to the reference response. Formally, given a user profile \mathcal{P} , a query x , and a reference personalized response y , we define the reward as:

$$\begin{aligned} R(\text{LLM}(\mathcal{P} \parallel x), y) &= \log p_\phi(y | \mathcal{P}, x) \\ &= \sum_{j=1}^{|y|} \log p_\phi(y_j | \mathcal{P}, x, y_{<j}), \end{aligned} \quad (4)$$

where ϕ are the parameters of the LLM, and $p_\phi(\cdot)$ denotes its next-token distribution. Unlike task-specific metrics such as accuracy or ROUGE-1, which only provide coarse-grained feedback on the final answer (Liu et al., 2025), the log-likelihood of the reference response serves as a semantically richer reward signal. The increased granularity allows the policy to distinguish between plausible and optimal profiles even when the final metrics might be identical. Empirical results in Section 5.2 demonstrate the advantages of this reward over metric-based formulations.

Moreover, we show in Appendix B that, with the log-likelihood reward, the objective defined in Equation 1 corresponds to maximizing the evidence lower bound (ELBO) of the marginalization-based RAG formulation (Lewis et al., 2020). Directly optimizing this marginal likelihood is intractable in our setting due to the combinatorial space of possible profiles, whereas our objective provides a tractable surrogate with rich learning signals. In the next section, we empirically demonstrate that this log-likelihood-based reward is robust across different types of downstream tasks.

4 Experimental Setup

4.1 Dataset and Evaluation

We evaluate the performance of PURPLE using Phi-4-Mini-Instruct (Microsoft et al., 2025) and Llama-3-8B-Instruct (Grattafiori et al., 2024) as the frozen LLM, and further scale to Llama-3-70B-Instruct. Our experiments span

Table 1: Results on the LaMP benchmark. The best and second-best results in each column are highlighted in **bold** and underlined, respectively. We report the mean and standard deviation over three runs with different random seeds for LLM generation. For GPT-5-nano, we report only a single run due to API cost constraints.

Task	Citation		Movie	Rating	News			Scholar	Tweet
Metric	Acc. / F1	Acc. / F1	MAE / RMSE	RG1 / RGL / MT	RG1 / RGL / MT	RG1 / RGL / MT	RG1 / RGL / MT	RG1 / RGL / MT	
<i>With Phi-4-Mini-Instruct (3.84B)</i>									
BM25	63.9 _{0.53} / 63.8 _{0.69}	34.5 _{0.65} / 29.6 _{0.61}	0.438 _{0.00} / 0.852 _{0.01}	14.4 _{0.16} / 12.8 _{0.16} / 11.9 _{0.12}	39.7 _{0.08} / 33.2 _{0.04} / 42.2 _{0.08}	38.3 _{0.04} / 33.6 _{0.08} / 35.2 _{0.04}			
Contriever	64.6 _{0.33} / 64.5 _{0.20}	36.0 _{0.37} / 31.1 _{0.24}	<u>0.424</u> _{0.01} / <u>0.830</u> _{0.03}	14.6 _{0.04} / 13.1 _{0.00} / 12.2 _{0.04}	39.7 _{0.00} / 33.3 _{0.08} / 42.0 _{0.08}	38.6 _{0.08} / 33.7 _{0.08} / 35.8 _{0.04}			
IC-RALM	62.0 _{0.20} / 61.7 _{0.37}	33.6 _{0.12} / 28.8 _{0.16}	0.471 _{0.01} / 0.857 _{0.02}	13.4 _{0.04} / 11.9 _{0.04} / 11.0 _{0.00}	37.6 _{0.12} / 31.0 _{0.12} / 41.0 _{0.33}	38.1 _{0.12} / 33.3 _{0.16} / 35.2 _{0.12}			
REPLUG	61.2 _{0.94} / 60.2 _{0.78}	<u>37.4</u> _{0.49} / <u>32.1</u> _{0.37}	0.486 _{0.01} / 0.891 _{0.02}	14.1 _{0.04} / 12.7 _{0.00} / 11.4 _{0.08}	38.1 _{0.63} / 32.8 _{0.90} / 40.4 _{1.51}	42.2 _{0.08} / 37.3 _{0.00} / 38.6 _{0.04}			
RankGPT (Llama-3-8B-Instruct)	63.9 _{0.82} / 63.7 _{0.65}	34.2 _{0.94} / 28.9 _{1.18}	0.446 _{0.00} / 0.863 _{0.01}	14.7 _{0.33} / 13.1 _{0.24} / 12.3 _{0.24}	39.7 _{0.00} / 33.3 _{0.04} / 42.0 _{0.04}	38.2 _{0.04} / 33.4 _{0.02} / 35.2 _{0.08}			
RankGPT (GPT-5-nano)	<u>65.9</u> / 65.6	35.5 / 31.4	0.444 / 0.865	14.6 / 13.0 / 12.1	<u>39.8</u> / <u>33.4</u> / <u>42.3</u>	38.5 / 33.7 / 35.5			
ICR (Llama-3-8B-Instruct)	65.2 _{0.53} / <u>65.0</u> _{0.45}	34.1 _{0.73} / 29.8 _{1.02}	<u>0.424</u> _{0.00} / <u>0.830</u> _{0.02}	<u>15.0</u> _{0.00} / <u>13.4</u> _{0.04} / <u>12.2</u> _{0.00}	39.5 _{0.08} / 33.1 _{0.08} / 42.0 _{0.04}	38.6 _{0.12} / 33.6 _{0.24} / 35.5 _{0.16}			
PURPLE (Ours)	66.0 _{0.20} / 65.6 _{0.16}	38.6 _{0.33} / 34.2 _{0.53}	0.419 _{0.01} / 0.808 _{0.02}	15.1 _{0.04} / 13.5 _{0.00} / 12.6 _{0.08}	40.0 _{0.04} / 33.5 _{0.04} / 42.4 _{0.00}	<u>39.0</u> _{0.04} / <u>34.0</u> _{0.04} / <u>35.9</u> _{0.00}			
<i>With Llama-3-8B-Instruct (8.03B)</i>									
BM25	56.9 _{0.61} / 56.5 _{0.53}	45.2 _{0.37} / 36.5 _{0.98}	0.328 _{0.01} / 0.664 _{0.02}	16.7 _{0.33} / 15.1 _{0.37} / 14.7 _{0.33}	41.0 _{0.02} / 35.2 _{0.04} / <u>40.6</u> _{0.02}	31.4 _{0.16} / 26.6 _{0.16} / 27.6 _{0.24}			
Contriever	58.5 _{0.20} / <u>58.1</u> _{0.41}	47.2 _{0.37} / 39.1 _{0.29}	0.314 _{0.00} / <u>0.631</u> _{0.01}	<u>17.2</u> _{0.04} / <u>15.6</u> _{0.04} / <u>15.1</u> _{0.04}	41.1 _{0.08} / 35.4 _{0.12} / 40.5 _{0.04}	<u>32.1</u> _{0.16} / <u>27.2</u> _{0.24} / <u>28.5</u> _{0.16}			
IC-RALM	<u>59.3</u> _{0.04} / 56.6 _{0.29}	38.0 _{0.82} / 29.9 _{0.41}	0.351 _{0.01} / 0.671 _{0.01}	14.1 _{0.24} / 12.5 _{0.24} / 12.1 _{0.08}	37.5 _{0.08} / 31.5 _{0.04} / 39.2 _{0.04}	29.2 _{0.69} / 24.6 _{0.57} / 25.6 _{0.49}			
REPLUG	55.5 _{1.06} / 47.5 _{2.00}	42.0 _{1.35} / 31.8 _{1.14}	0.314 _{0.00} / 0.633 _{0.00}	14.8 _{0.04} / 13.2 _{0.00} / 11.8 _{0.04}	42.5 _{0.08} / 37.1 _{0.06} / 40.8 _{0.12}	30.4 _{0.20} / 26.1 _{0.12} / 25.8 _{0.37}			
RankGPT (Llama-3-8B-Instruct)	57.0 _{0.20} / 56.3 _{0.12}	47.2 _{0.90} / 38.4 _{0.53}	0.318 _{0.01} / 0.637 _{0.01}	16.9 _{0.20} / 15.3 _{0.16} / 14.8 _{0.29}	41.0 _{0.04} / 35.4 _{0.08} / 40.5 _{0.16}	31.0 _{0.12} / 26.3 _{0.08} / 27.3 _{0.16}			
RankGPT (GPT-5-nano)	59.5 / 58.0	45.1 / 36.2	0.321 / 0.638	17.1 / 15.4 / 15.0	41.0 / 35.3 / 40.5	31.5 / 26.5 / 27.8			
ICR (Llama-3-8B-Instruct)	58.4 _{0.29} / 57.3 _{0.37}	<u>48.0</u> _{0.37} / <u>39.3</u> _{0.73}	<u>0.312</u> _{0.01} / <u>0.631</u> _{0.02}	17.1 _{0.04} / 15.4 _{0.00} / 14.9 _{0.00}	41.3 _{0.16} / 35.5 _{0.29} / 40.8 _{0.24}	31.3 _{0.33} / 26.8 _{0.24} / 28.1 _{0.29}			
PURPLE (Ours)	59.2 _{0.82} / 58.8 _{0.78}	49.6 _{0.65} / 41.6 _{0.53}	0.307 _{0.01} / 0.624 _{0.01}	17.6 _{0.04} / 15.9 _{0.00} / 15.3 _{0.08}	<u>41.4</u> _{0.49} / <u>35.8</u> _{0.73} / 40.3 _{0.37}	32.5 _{0.04} / 27.5 _{0.00} / 28.8 _{0.06}			
<i>With Llama-3-70B-Instruct (70.6B)</i>									
BM25	71.3 _{0.37} / 70.8 _{0.37}	54.5 _{0.37} / 47.2 _{0.41}	0.245 _{0.01} / 0.544 _{0.01}	18.1 _{0.33} / 16.5 _{0.29} / 14.9 _{0.37}	43.8 _{0.53} / 38.2 _{0.41} / 40.5 _{0.49}	36.3 _{0.16} / 30.9 _{0.20} / 33.0 _{0.12}			
Contriever	71.2 _{0.78} / 70.8 _{0.73}	56.6 _{0.16} / 49.5 _{0.29}	<u>0.238</u> _{0.00} / 0.528 _{0.00}	18.6 _{0.08} / <u>17.0</u> _{0.08} / <u>15.6</u> _{0.08}	44.1 _{0.08} / <u>38.6</u> _{0.12} / 41.1 _{0.00}	36.5 _{0.00} / <u>31.4</u> _{0.00} / <u>33.3</u> _{0.00}			
IC-RALM	66.5 _{0.04} / 66.4 _{0.00}	49.0 _{0.29} / 41.8 _{0.12}	0.261 _{0.00} / 0.555 _{0.00}	14.8 _{0.02} / 13.3 _{0.04} / 12.2 _{0.00}	39.7 _{0.00} / 34.1 _{0.24} / 38.6 _{0.16}	32.0 _{0.20} / 27.4 _{0.00} / 28.8 _{0.20}			
REPLUG	65.8 _{0.33} / 65.4 _{0.41}	51.7 _{0.04} / 44.0 _{0.08}	0.258 _{0.01} / 0.554 _{0.02}	15.2 _{0.04} / 13.9 _{0.04} / 12.1 _{0.04}	41.8 _{0.12} / 36.6 _{0.33} / 38.6 _{0.35}	32.1 _{0.08} / 27.6 _{0.12} / 27.7 _{0.08}			
RankGPT (Llama-3-8B-Instruct)	69.7 _{0.16} / 69.1 _{0.12}	<u>56.9</u> _{0.16} / 48.8 _{0.41}	0.247 _{0.00} / 0.546 _{0.01}	18.0 _{0.24} / 16.5 _{0.29} / 15.2 _{0.29}	44.0 _{0.04} / <u>38.6</u> _{0.08} / 41.5 _{0.12}	35.8 _{0.00} / 30.6 _{0.00} / 32.4 _{0.08}			
RankGPT (GPT-5-nano)	73.8 / 73.5	55.3 / 48.2	0.240 / <u>0.523</u>	<u>18.7</u> / <u>17.0</u> / 15.8	44.6 / 38.8 / 41.5	<u>36.6</u> / 31.3 / 33.2			
ICR (Llama-3-8B-Instruct)	71.6 _{0.16} / 71.0 _{0.16}	56.8 _{0.24} / <u>49.6</u> _{0.57}	<u>0.238</u> _{0.00} / 0.531 _{0.00}	18.3 _{0.00} / 16.8 _{0.04} / 15.3 _{0.16}	44.3 _{0.06} / 38.5 _{0.04} / <u>41.2</u> _{0.20}	36.3 _{0.16} / 31.1 _{0.24} / 33.2 _{0.20}			
PURPLE (Ours)	<u>72.2</u> _{0.49} / <u>71.8</u> _{0.57}	57.0 _{0.08} / 49.9 _{0.41}	0.236 _{0.02} / 0.520 _{0.01}	18.8 _{0.04} / 17.1 _{0.00} / <u>15.6</u> _{0.08}	<u>44.4</u> _{0.02} / 38.8 _{0.06} / 41.0 _{0.04}	37.3 _{0.01} / 32.1 _{0.04} / 34.0 _{0.06}			

a range of personalization settings, including personalized classification, regression, and both short- and long-text generation from the LaMP (Salemi et al., 2024) and LongLaMP (Kumar et al., 2024) benchmarks. We follow the prompt templates of Salemi et al. (2024) and Kumar et al. (2024) to incorporate user profiles into the original queries.

Specifically, we evaluate PURPLE on **nine personalization tasks**, including two classification tasks — *Personalized Citation Identification* (Citation) and *Personalized Movie Tagging* (Movie) — evaluated with Accuracy and F1; one regression task — *Personalized Product Rating* (Rating) — evaluated with MAE and RMSE; and six generation tasks, evaluated with ROUGE-1 (RG1), ROUGE-L (RGL) (Lin, 2004), and METEOR (MT) (Banerjee and Lavie, 2005). The generation tasks are further divided into short-text generation — *Personalized News Headline Generation* (News), *Personalized Scholarly Title Generation* (Scholar), and *Personalized Tweet Paraphrasing* (Tweet) — and long-text generation — *Personalized Abstract Generation* (Abstract), *Personalized Topic Generation* (Topic), and *Personalized Product Review Generation* (Review). In all experiments, we first use Contriever (Izacard et al., 2022) to retrieve 20 candidate records as the user history \mathcal{H} , and then select 5 of them to construct the user profile \mathcal{P} . Appendix D further elaborates implementation details.

4.2 Baseline Methods

We focus on the setting where the LLM is kept frozen and no ground-truth profile is available for training the reranker. Therefore, we compare against three categories of prior methods that likewise neither fine-tune the LLM nor rely on supervision from ground-truth profiles. Methods that operate under different assumptions or settings are not included in our main comparisons, as these differences prevent a fair and controlled evaluation. We discuss these methods in detail in Appendix C.

The baselines we compare with include (i) **Zero-Shot Rerankers**, represented by RankGPT (Sun et al., 2023) and ICR (Chen et al., 2025). For both methods, we adopt Llama-3-8B-Instruct as the reranker LLM. We also report the performance of RankGPT with GPT-5-nano to reflect an upper bound of the methods that distill from the ranking results of state-of-the-art proprietary LLMs (Pradeep et al., 2023a,b; Tamber et al., 2023; Gangi Reddy et al., 2024). (ii) **In-Context Retrieval-Augmented Language Models (RALMs)**, represented by IC-RALM (Ram et al., 2023) and REPLUG (Shi et al., 2024). Both methods consider each record in isolation when generating a response. They incorporate multiple records from the user profile either through context switching (IC-RALM) or marginalization (REPLUG). Additionally, we include (iii) **Efficient**

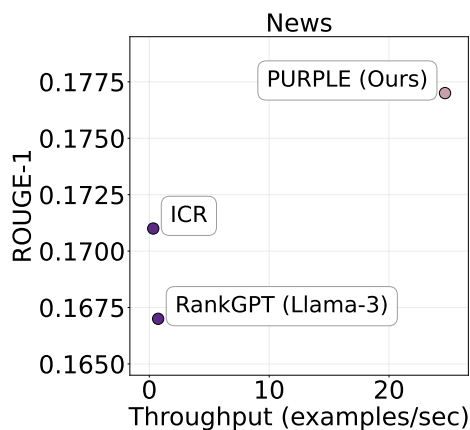


Figure 3: Performance–throughput tradeoff on News.

Sparse and Dense Retrievers, applied directly as rerankers. We include BM25 (Robertson and Zaragoza, 2009) for the sparse retriever and Contriever (Izacard et al., 2022) for the dense retriever. These methods represent the efficiency-oriented side of the efficiency–performance trade-off.

5 Results and Analysis

5.1 Overall Performance Comparison

Table 1 presents the results of PURPLE and the baselines on the LaMP benchmark. Table 3 in the Appendix presents the results on the LongLaMP benchmark. The main findings are as follows:

PURPLE consistently outperforms strong baselines across LLM scales. Across all tasks and LLMs of varying sizes, PURPLE achieves consistent improvements over existing methods. Compared with Contriever, which is of comparable model size, PURPLE’s propensity scores provide more effective ranking signals than raw relevance. Compared with zero-shot rerankers, which use much larger backbone LLMs and incur higher inference cost, PURPLE achieves stronger personalization with a much smaller model, since reinforcement learning training help better capture the utility of multiple records. Compared with in-context RALMs, which provide user records one at a time to the LLM and combine multiple records post hoc, our single-stage modeling more effectively captures personalized signals, highlighting the advantage of treating user profiles holistically.

PURPLE outperforms strong baselines with high computational throughput. Figure 3 shows that, on a representative LaMP dataset, PURPLE outperforms existing methods while being

more efficient. Specifically, PURPLE achieves higher performance than ICR and RankGPT (Llama-3) while maintaining high computational throughput. Additional results across the six LaMP tasks are provided in Figure 6 in the Appendix.

PURPLE is effective across task types, including regression. Since our reward is based on the log-likelihood that the LLM assigns to the reference response, it does not directly reflect numerical distances between regression targets. Nevertheless, PURPLE still achieves strong gains on the regression task (Rating). This demonstrates that log-likelihood provides a principled and broadly applicable reward signal across diverse task forms.

5.2 Ablation Studies on Model Architecture and Reward Modeling

Table 2 presents the ablation results of PURPLE using Phi-4-Mini-Instruct. Overall, we examine three key design choices. First, instead of performing token-level cross-attention, we test a simplified variant that pools the query into a single embedding vector, which is then prepended as an extra token to the Transformer encoder (w/o CA). This approach is less effective, indicating that fine-grained token-level interactions between the query and the user records are crucial for accurate personalization. Second, we remove the Transformer encoder entirely, resulting in a point-wise scoring model where each record is processed independently (w/o RDM). This variant shows the largest performance drop across tasks. Although the model can still leverage individually informative records, it fails to model dependencies such as redundancy and complementarity among records. Third, we examine an alternative reward design that uses task-specific evaluation metrics as the reward (Accuracy for classification, MAE for regression, and ROUGE-1 for generation). This metric-based reward is substantially more coarse-grained, providing weaker learning signals and making optimization more challenging. This variant consistently underperforms PURPLE across all tasks, thereby strongly motivating our log-likelihood-based reward.

These results highlight that token-level cross-attention, inter-record dependency modeling, and the log-likelihood-based reward are all indispensable, validating our design choices of treating user profiles as structured contexts rather than isolated records, and providing semantically richer reward signals via reference log-likelihoods.

Table 2: Ablation studies of PURPLE using Phi-4-Mini-Instruct. CA and RDM stand for cross-attention and inter-record dependency modeling, respectively. MR refers to using task-specific metrics as the reward.

Task	Citation	Movie	Rating	News	Scholar	Tweet
Metric	Acc. / F1	Acc. / F1	MAE / RMSE	RG1 / RGL / MT	RG1 / RGL / MT	RG1 / RGL / MT
PURPLE	66.2 / 65.8	38.2 / 33.6	0.405 / 0.788	15.2 / 13.5 / 12.5	40.0 / 33.5 / 42.4	39.1 / 34.0 / 35.9
w/o CA	64.8 / 64.5	35.1 / 29.7	0.440 / 0.816	14.8 / 13.2 / 12.4	40.0 / 33.5 / 42.2	39.1 / 34.1 / 36.0
w/o RDM	61.3 / 60.6	35.0 / 31.1	0.449 / 0.850	14.5 / 12.8 / 11.9	39.7 / 33.1 / 41.9	39.0 / 34.0 / 36.1
w/ MR	64.8 / 64.8	38.0 / 33.3	0.433 / 0.854	15.0 / 13.2 / 12.4	39.5 / 32.9 / 41.8	38.7 / 33.7 / 35.6

5.3 Analysis on TopK Selection at Inference

To further justify our inference strategy of selecting the top- K records based on the propensity scores, we compare the top- K selections of PURPLE, ICR, RankGPT, and Contriever on the six LaMP datasets. For each test set example, we consider the top-5 records proposed by each method and randomly sample five orderings from the $5! = 120$ permutations as controls. These six orderings are ranked based on downstream generation quality using Llama-3-8B-Instruct, as measured by task-specific evaluation metrics. As shown in Figure 4, orderings induced by PURPLE’s propensity scores are more frequently ranked as the best among the six candidates. This indicates that PURPLE’s scoring function better captures relative dependencies between records, highlighting that our method not only identifies useful records but also arranges them in an order that maximizes downstream personalization utility. We attribute this advantage of PURPLE to its holistic modeling of the user profile, and emphasize that our method adopts a profile-level optimization strategy rather than relying solely on local, pairwise relevance.

5.4 Human Evaluation and Case Study

To complement our automatic evaluation and verify the practical utility of the selected profiles, we conduct a rigorous side-by-side human evaluation. While automatic metrics measure lexical overlap, they often fail to capture the subtle stylistic nuances, such as voice, stance, and informal phrasing, that define true personalization. We focus our human evaluation on *Personalized Tweet Paraphrasing* (Tweet), which requires significant stylistic adaptation to user personas. In this study, we randomly sample 250 instances. For each instance, human evaluators are presented with the input query, a snippet of the user’s history, and two anonymized model outputs: one generated using the recent base-

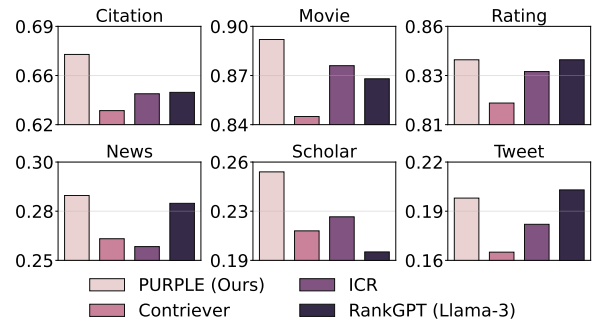


Figure 4: Ranking accuracy across the six LaMP tasks.

line, ICR, and the other using PURPLE. Evaluators perform a blind, forced-choice comparison to determine which output better reflects the user’s specific persona while maintaining semantic fidelity. The results demonstrate a strong evaluator preference for PURPLE, surpassing ICR by an absolute margin of 14.4% (57.2% vs. 42.8%). This confirms that our method effectively captures personal stylistic features, such as slang usage and informal tone, where heuristic relevance retrieval often fails. We provide detailed qualitative case studies analyzing these behaviors in Appendix E.

5.5 Sensitivity to Profile Size

While the main experiments employ a profile size of $K = 5$, we further investigate the impact of varying profile sizes $K \in \{5, 10, 15\}$ on *Personalized Tweet Paraphrasing* (Tweet) using Phi-4-Mini-Instruct. Throughout this analysis, the size of the candidate pool remains fixed at 20. As illustrated in Figure 5, PURPLE consistently outperforms ICR across all settings. However, we observe that increasing K beyond 5 generally leads to a slight performance degradation. This finding aligns with prior results in the RAG literature, which suggest that larger contexts may introduce redundancy or interference rather than additional utility (Yu et al., 2024; Salemi et al., 2024).




Figure 5: Effect of varying profile sizes on Tweet.

6 Conclusion

In this work, we study the problem of retrieval-augmented personalization for LLMs. Our key intuitions are twofold: (i) *record relevance does not reliably predict personalization utility*, and (ii) *utility is non-monotonic across records*, making greedy aggregation suboptimal. To address these limitations, we propose PURPLE, a contextual bandit framework that optimizes user profiles by leveraging downstream performance as feedback. PURPLE jointly models query–record interactions and inter-record dependencies, enabling adaptive selection of user profiles beyond static heuristics. Instead of task-specific metrics, PURPLE relies on the log-likelihood that the LLM assigns to the reference response as the reward, providing a semantically richer reward signal that facilitates more effective training. Extensive experiments on nine real-world personalization tasks spanning classification, regression, and generation show that PURPLE consistently outperforms heuristic retrievers, LLM-based rerankers, and in-context RALMs, while being more efficient. These results establish contextual bandit-based retrieval as a principled and scalable paradigm for personalizing LLMs.

Limitations

While PURPLE demonstrates strong effectiveness across personalization tasks, it has several limitations. First, the method relies on the log-likelihood of reference personalized responses under a frozen LLM as the training reward, which assumes access to high-quality personalized supervision. In practical deployment, such explicit supervision may be sparse or unavailable, as user feedback often takes implicit or noisy forms (e.g., engagement signals). Extending PURPLE to effectively leverage

these weaker signals remains an open challenge. Second, although using response log-likelihood as the reward provides a unified optimization signal, PURPLE is still trained separately for each task in our current study. We do not explicitly evaluate the extent to which a policy learned under this unified objective can generalize to unseen tasks or domains. Exploring more transferable training and deployment paradigms is thus left for future work.

References

- 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, Sigir '18*, pages 135–144, New York, NY, USA. Association for Computing Machinery.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-RAG: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Shijie Chen, Bernal Jimenez Gutierrez, and Yu Su. 2025. [Attention in large language models yields efficient zero-shot re-rankers](#). In *The Thirteenth International Conference on Learning Representations*.
- Revanth Gangi Reddy, JaeHyeok Doo, Yifei Xu, Md Arafat Sultan, Deevya Swain, Avirup Sil, and Heng Ji. 2024. [FIRST: Faster improved listwise reranking with single token decoding](#). In *Proceedings of the 2024 Conference on Empirical Methods*

- in *Natural Language Processing*, pages 8642–8652, Miami, Florida, USA. Association for Computational Linguistics.
- Michael Glass, Gaetano Rossiello, Md Faisal Mahub Chowdhury, Ankita Naik, Pengshan Cai, and Alfio Gliozzo. 2022. [Re2G: Retrieve, rerank, generate](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2701–2715, Seattle, United States. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. [Retrieval augmented language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3929–3938. PMLR.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Transactions on Machine Learning Research*.
- Bowen Jiang, Zhuoqun Hao, Young Min Cho, Bryan Li, Yuan Yuan, Sihao Chen, Lyle Ungar, Camillo Jose Taylor, and Dan Roth. 2025. [Know me, respond to me: Benchmarking LLMs for dynamic user profiling and personalized responses at scale](#). In *Second Conference on Language Modeling*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Omar Khattab and Matei Zaharia. 2020. [ColBERT: Efficient and effective passage search via contextualized late interaction over BERT](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, Sigir '20*, pages 39–48, New York, NY, USA. Association for Computing Machinery.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations*.
- Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A. Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka, Chien Van Nguyen, Thien Huu Nguyen, and Hamed Zamani. 2024. [LongLaMP: A benchmark for personalized long-form text generation](#). *Preprint*, arXiv:2407.11016.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with PagedAttention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, Sosp '23*, pages 611–626, New York, NY, USA. Association for Computing Machinery.
- John Langford and Tong Zhang. 2007. [The epoch-greedy algorithm for multi-armed bandits with side information](#). In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Xinyu Li, Ruiyang Zhou, Zachary C. Lipton, and Liu Leqi. 2024. [Personalized language modeling from personalized human feedback](#). *Preprint*, arXiv:2402.05133.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Wei Liu, Siya Qi, Xinyu Wang, Chen Qian, Yali Du, and Yulan He. 2025. [NOVER: Incentive training for language models via verifier-free reinforcement learning](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7439–7458, Suzhou, China. Association for Computational Linguistics.
- Microsoft, :, Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, Dong Chen, Dongdong Chen,

- Junkun Chen, Weizhu Chen, Yen-Chun Chen, Yiling Chen, Qi Dai, and 57 others. 2025. [Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-LoRAs](#). *Preprint*, arXiv:2503.01743.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. [Document ranking with a pre-trained sequence-to-sequence model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, Online. Association for Computational Linguistics.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. [Multi-stage document ranking with BERT](#). *Preprint*, arXiv:1910.14424.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Liang Pang, Jun Xu, Qingyao Ai, Yanyan Lan, Xueqi Cheng, and Jirong Wen. 2020. [SetRank: Learning a permutation-invariant ranking model for information retrieval](#). In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, Sigir '20*, pages 499–508, New York, NY, USA. Association for Computing Machinery.
- Emiliano Penalzoza, Olivier Gouvert, Haolun Wu, and Laurent Charlin. 2025. [TEARS: Text representations for scrutable recommendations](#). In *Proceedings of the ACM on Web Conference 2025, Www '25*, pages 4949–4968, New York, NY, USA. Association for Computing Machinery.
- Ronak Pradeep, Rodrigo Nogueira, and Jimmy Lin. 2021. [The expando-mono-duo design pattern for text ranking with pretrained sequence-to-sequence models](#). *Preprint*, arXiv:2101.05667.
- Ronak Pradeep, Sahel Sharifmoghammad, and Jimmy Lin. 2023a. [RankVicuna: Zero-shot listwise document reranking with open-source large language models](#). *Preprint*, arXiv:2309.15088.
- Ronak Pradeep, Sahel Sharifmoghammad, and Jimmy Lin. 2023b. [RankZephyr: Effective and robust zero-shot listwise reranking is a breeze!](#) *Preprint*, arXiv:2312.02724.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#). *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Stephen Robertson and Hugo Zaragoza. 2009. [The Probabilistic Relevance Framework: BM25 and Beyond](#). *Foundations and Trends in Information Retrieval*, 4(1-2):1–174.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. [LaMP: When large language models meet personalization](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392, Bangkok, Thailand. Association for Computational Linguistics.
- Xiaoteng Shen, Rui Zhang, Xiaoyan Zhao, Jieming Zhu, and Xi Xiao. 2024. [PMG: Personalized multimodal generation with large language models](#). In *Proceedings of the ACM Web Conference 2024, Www '24*, pages 3833–3843, New York, NY, USA. Association for Computing Machinery.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [REPLUG: Retrieval-augmented black-box language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8371–8384, Mexico City, Mexico. Association for Computational Linguistics.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? Investigating large language models as re-ranking agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937, Singapore. Association for Computational Linguistics.
- Yihang Sun, Tao Feng, Ge Liu, and Jiaxuan You. 2025. [PREMIUM: LLM personalization with individual-level preference feedback](#).
- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. [Policy gradient methods for reinforcement learning with function approximation](#). In *Advances in Neural Information Processing Systems*, volume 12. MIT Press.
- Manveer Singh Tamber, Ronak Pradeep, and Jimmy Lin. 2023. [Scaling down, LiTting up: Efficient zero-shot listwise reranking with seq2seq encoder-decoder models](#). *Preprint*, arXiv:2312.16098.
- Zhaoxuan Tan, Zheyuan Liu, and Meng Jiang. 2024. [Personalized pieces: Efficient personalized large language models through collaborative efforts](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6459–6475, Miami, Florida, USA. Association for Computational Linguistics.

Ronald J. Williams. 1992. [Simple statistical gradient-following algorithms for connectionist reinforcement learning](#). *Machine Learning*, 8(3):229–256.

Bin Wu, Zhengyan Shi, Hossein A. Rahmani, Varsha Ramineni, and Emine Yilmaz. 2024. [Understanding the role of user profile in the personalization of large language models](#). *Preprint*, arXiv:2406.17803.

Shangyu Wu, Ying Xiong, Yufei Cui, Haolun Wu, Can Chen, Ye Yuan, Lianming Huang, Xue Liu, Tei-Wei Kuo, Nan Guan, and Chun Jason Xue. 2025. [Retrieval-augmented generation for natural language processing: A survey](#). *Preprint*, arXiv:2407.13193.

Soyoung Yoon, Eunbi Choi, Jiyeon Kim, Hyeongu Yun, Yireun Kim, and Seung-won Hwang. 2024. [ListT5: Listwise reranking with fusion-in-decoder improves zero-shot retrieval](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2287–2308, Bangkok, Thailand. Association for Computational Linguistics.

Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2024. [Making retrieval-augmented language models robust to irrelevant context](#). In *The Twelfth International Conference on Learning Representations*.

Yue Yu, Wei Ping, Zihan Liu, Boxin Wang, Jiaxuan You, Chao Zhang, Mohammad Shoeybi, and Bryan Catanzaro. 2024. [RankRAG: Unifying context ranking with retrieval-augmented generation in llms](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 121156–121184. Curran Associates, Inc.

Ye Yuan, Mohammad Amin Shabani, and Siqi Liu. 2026. [Embedding-based context-aware reranker](#). In *The Fourteenth International Conference on Learning Representations*.

Lemei Zhang, Peng Liu, Marcus Tiedemann Oekland Henriksboe, Even W. Lauvraak, Jon Atle Gulla, and Heri Ramampiaro. 2024. [PersonalSum: A user-subjective guided personalized summarization dataset for large language models](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 99333–99346. Curran Associates, Inc.

A Details of Gradient Estimation

To estimate the gradient in Equation 2, we first draw a batch of examples $\{(\mathcal{H}_b, x_b, y_b)\}_{b=1}^B$. For each example, we sample M profiles $\mathcal{P}_b^1, \dots, \mathcal{P}_b^M$ from $\pi_\theta(\cdot | \mathcal{C}_b)$, and compute the empirical mean of the gradient estimates. This learning procedure corresponds to the REINFORCE algorithm (Sutton et al., 1999), with the Monte Carlo gradient

estimate:

$$\begin{aligned} \nabla_{\theta} \mathcal{J}(\theta) &\approx \\ &\frac{1}{B} \sum_{b=1}^B \frac{1}{M} \sum_{m=1}^M \nabla_{\theta} \log \pi_{\theta}(\mathcal{P}_b^m | \mathcal{C}_b) \tilde{r}_b^m. \end{aligned} \quad (5)$$

To reduce variance in gradient estimation, we apply reward normalization over the M profiles sampled for each example. Concretely, for each example with rewards $\mathbf{r}_b = [r_b^1, \dots, r_b^M]^{\top}$, where $r_b^m = R(\phi(\mathcal{P}_b^m, x_b), y_b)$, the normalized reward is computed as $\tilde{r}_b^m = \frac{r_b^m - \text{mean}(\mathbf{r}_b)}{\text{std}(\mathbf{r}_b)}$.

B Motivating our Reward

The choice of using response log-likelihood as the reward is grounded in the generative modeling formulation of RAG (Lewis et al., 2020), where the user profile is treated as a latent variable and the response likelihood is obtained by marginalizing over all possible profiles. Applying Jensen’s inequality to the maximum likelihood objective gives:

$$\begin{aligned} &\mathbb{E}_{(\mathcal{H}, x, y) \sim \mathcal{D}} \\ &\left[\log \sum_{\mathcal{P} \in \text{Perm}_K(\mathcal{H})} \pi_{\theta}(\mathcal{P} | \mathcal{C}) p_{\phi}(y | \mathcal{P}, x) \right] \quad (6) \\ &\geq \mathbb{E}_{(\mathcal{H}, x, y) \sim \mathcal{D}, \mathcal{P} \sim \pi_{\theta}(\cdot | \mathcal{C})} [\log p_{\phi}(y | \mathcal{P}, x)]. \end{aligned}$$

Therefore, maximizing the expected log-likelihood is equivalent to optimizing the evidence lower bound (ELBO), with p_{ϕ} modeled by a frozen LLM.

C Discussion of Alternative Personalization Methods

This section provides a brief discussion of personalization methods that fall outside our experimental setting. These approaches are not included in our empirical comparisons because they operate under different assumptions or training regimes, which prevents a fair and controlled evaluation against our method. We review them here to contextualize our design choices and to clarify the scope of the methods considered in this work.

In their benchmark paper, Salemi et al. (2024) propose two baseline LLM personalization strategies: In-Prompt Augmentation (IPA) and Fusion-in-Decoder (FiD). IPA represents the standard processing applied in our method, as well as in all compared baselines: it integrates all retrieved records into a single prompt and feeds them to a frozen

Table 3: Results on the LongLaMP benchmark. The best and second-best results in each column are highlighted in **bold** and underlined, respectively. We report the mean and standard deviation over three runs with different random seeds for LLM generation. For GPT-5-nano, we report only a single run due to API cost constraints.

Task	Abstract	Topic	Review
Metric	R1 / RL / MT	R1 / RL / MT	R1 / RL / MT
<i>With Phi-4-Mini-Instruct (3.84B)</i>			
BM25	38.6 _{0.12} / 22.0 _{0.16} / 26.2 _{0.08}	25.2 _{0.41} / 12.7 _{0.20} / <u>17.1</u> _{0.16}	25.8 _{1.39} / 13.2 _{0.53} / 16.5 _{0.16}
Contriever	38.5 _{0.04} / 21.9 _{0.12} / 26.1 _{0.08}	<u>25.7</u> _{1.80} / <u>13.0</u> _{0.73} / 16.6 _{0.20}	25.2 _{1.92} / 13.1 _{0.57} / 15.9 _{0.78}
IC-RALM	37.2 _{0.04} / 21.1 _{0.02} / 25.1 _{0.04}	24.8 _{1.47} / 12.4 _{0.78} / 15.9 _{0.20}	25.0 _{1.35} / 12.7 _{0.49} / 15.8 _{0.16}
REPLUG	36.0 _{0.24} / 21.2 _{0.20} / 23.8 _{0.20}	23.5 _{1.06} / 12.1 _{0.20} / 13.1 _{0.41}	20.8 _{2.94} / 11.2 _{1.14} / 12.5 _{1.63}
RankGPT (Llama-3-8B-Instruct)	38.6 _{0.12} / 22.0 _{0.08} / 26.2 _{0.08}	25.1 _{0.53} / 12.7 _{0.29} / 17.1 _{0.18}	25.1 _{1.63} / 12.9 _{0.53} / 15.8 _{0.45}
RankGPT (GPT-5-nano)	39.1 / 22.4 / 26.9	24.9 / 12.5 / 17.5	<u>27.1</u> / <u>13.7</u> / <u>16.6</u>
ICR (Llama-3-8B-Instruct)	38.5 _{0.24} / 22.0 _{0.16} / 26.2 _{0.16}	<u>25.7</u> _{1.71} / <u>13.0</u> _{0.73} / 16.5 _{0.24}	25.8 _{1.67} / 13.2 _{0.53} / 16.4 _{0.45}
PURPLE (Ours)	<u>38.8</u> _{0.08} / <u>22.1</u> _{0.16} / <u>26.4</u> _{0.12}	26.2 _{1.14} / 13.2 _{0.61} / 17.5 _{0.12}	27.9 _{0.04} / 14.4 _{0.08} / 17.1 _{0.04}
<i>With Llama-3-8B-Instruct (8.03B)</i>			
BM25	42.2 _{0.04} / <u>24.2</u> _{0.04} / 31.8 _{0.08}	29.8 _{0.73} / 14.6 _{0.24} / <u>20.2</u> _{0.12}	<u>30.9</u> _{2.00} / <u>15.3</u> _{0.82} / <u>20.8</u> _{0.41}
Contriever	<u>42.3</u> _{0.24} / <u>24.2</u> _{0.29} / 31.7 _{0.24}	<u>30.6</u> _{1.39} / 15.2 _{0.53} / 20.0 _{0.32}	30.1 _{2.49} / 14.9 _{1.02} / 19.7 _{0.90}
IC-RALM	38.3 _{0.90} / 20.4 _{0.78} / 29.1 _{0.37}	25.0 _{0.90} / 11.4 _{1.06} / 15.5 _{1.92}	29.8 _{1.27} / 13.6 _{0.98} / 18.3 _{0.98}
REPLUG	38.2 _{0.45} / 20.7 _{0.33} / 28.4 _{0.20}	20.2 _{1.22} / 10.8 _{0.57} / 11.9 _{1.27}	15.5 _{2.04} / 8.8 _{0.94} / 8.8 _{1.02}
RankGPT (Llama-3-8B-Instruct)	42.2 _{0.04} / <u>24.2</u> _{0.04} / 31.9 _{0.04}	30.8 _{1.39} / 15.2 _{0.57} / 20.4 _{0.36}	30.5 _{2.45} / 15.1 _{1.02} / 20.4 _{0.82}
RankGPT (GPT-5-nano)	42.5 / 24.5 / 32.1	28.7 / 14.2 / <u>20.2</u>	33.6 / 16.5 / 21.4
ICR (Llama-3-8B-Instruct)	41.9 _{0.24} / 23.9 _{0.20} / 31.5 _{0.16}	<u>30.6</u> _{1.35} / <u>15.1</u> _{0.57} / <u>20.2</u> _{0.20}	30.7 _{1.96} / 15.2 _{0.78} / 20.3 _{0.41}
PURPLE (Ours)	<u>42.3</u> _{0.08} / <u>24.2</u> _{0.12} / <u>32.0</u> _{0.24}	<u>30.4</u> _{1.63} / <u>15.0</u> _{0.73} / <u>19.8</u> _{0.24}	<u>30.2</u> _{1.61} / <u>14.9</u> _{0.83} / <u>20.0</u> _{0.94}

LLM. FiD, on the other hand, encodes multiple inputs separately within the encoder, producing representations that are attended to by the decoder; as a result, it relies on encoder-decoder architectures and requires fine-tuning of the LLM.

PREMIUM (Sun et al., 2025) utilizes a pre-defined tag library to construct candidate tag sets from user queries using a neural network. The network is trained by leveraging users’ preference feedback on LLM responses generated from different candidate tag sets. While a suitable tag set can be readily identified for certain personalization tasks, such as *Personalized Movie Tagging* (Movie), this is generally not the case for open-ended personalized generation. Moreover, the tag sets considered in their work correspond to static topical features (e.g., Investment, Bakery, Comedy), whereas personalization may involve more nuanced attributes such as writing style, tone, and language use. These attributes are difficult to capture with a fixed set of tags but are more naturally reflected in user history records. We consider a more general setting in which only a collection of history records (and no online user feedback) is available, and we aim to select a useful subset for retrieval-augmented personalized generation.

Personalized Pieces (Tan et al., 2024) trains LoRA adapters for representative users and com-

bins them for new users based on user profile similarity, while Li et al. (2024) incorporates user information by fine-tuning the LLM under a personalized training objective. As a result, these works fall outside the class of RAG-based methods considered in our comparisons.

D Implementation Details

We employ a frozen pre-trained Contriever to first encode both queries and user history records into token embeddings. The only trainable components are the remaining modules of the user records encoder. These include a cross-attention layer that integrates query information into record embeddings, a Transformer encoder that captures inter-record dependencies, and an MLP decoder that maps the updated record encodings into scalar propensity scores. We set the number of Transformer encoder layers to $l = 12$, resulting in a parameter size roughly twice that of Contriever, while still being substantially smaller than the zero-shot rerankers and in-context RALMs used as baselines. For gradient estimation, we use a batch size of $B = 16$ and sample $M = 32$ user profiles for each example. We train the model for 10 epochs using the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a learning rate of 1×10^{-4} . During training, we apply a gradient clipping norm of




Figure 6: Performance–throughput tradeoff across the six LaMP tasks.

1.0. The checkpoint achieving the best validation performance is selected for testing.

In all experiments, we use frozen LLMs both to generate personalized responses and to evaluate the log-likelihood of reference responses conditioned on the query and user profiles (i.e., our reward). For generation, we set the temperature to $T = 0.7$ and employ nucleus sampling (Holtzman et al., 2020) with $top_p = 0.8$. For Phi-4-Mini-Instruct and Llama-3-8B-Instruct, we deploy them on a single NVIDIA H100 GPU. For Llama-3-70B-Instruct, we deploy the model across four NVIDIA H100 GPUs using vLLM (Kwon et al., 2023). All LLMs are deployed in BF16 precision.

E Qualitative Case Studies

In this section, we present a detailed case study from our human evaluation where PURPLE and the baseline ICR diverge significantly, illustrating the qualitative superiority of our method. In the Tweet task, preserving the logic of the original sentence is as critical as matching the user’s informal voice. The example below demonstrates a failure in logical reasoning by ICR. The input tweet contains a comparative structure (“at least we attended... unlike someone”). The baseline fails to parse this dependency, resulting in a paraphrase that states the exact opposite of the truth. PURPLE successfully disentangles the complex sentence structure and preserves the correct meaning.

User Persona / History:

- @tommoger thanks tom, and glad it all went well
- @callamKAPOW I tried to change it but the pic was too big and I can’t find the old one
- Thinks I can pass English, I just wish I could be as confident about geography

Input Tweet:

Paraphrase the following tweet without any explanation before or after it: @picklesticks6 and I am in the same boat, but at least we attended the math lessons unlike someone we both know who is taking their exam today and feels sad about it. =(

Reference Response:

@picklesticks6 same, but at least we turned up to the maths lessons unlike someone we know whos taking their exam today =(

ICR (Baseline) Generation:

@picklesticks6 we both missed math classes, and it’s making our exam prep tougher, unlike someone who went and feels down about it today. =(

PURPLE (Ours) Generation:

@picklesticks6 feels the same about math exams as me, but at least we went to the lessons, unlike someone else who feels sad today about their exam. =(

Analysis:

This case highlights a critical semantic failure in the baseline. ICR generates “we both missed math classes,” which directly contradicts the input stating “at least we attended.” It effectively swaps the roles of the speaker and the third party. PURPLE demonstrates superior natural language understanding by correctly preserving the fact that the speaker *did* attend the lesson. Furthermore, PURPLE better matches the user’s persona (which discusses academic confidence in English/Geography) by maintaining the conversational flow without losing the factual premise.

F LLM Usage

During the preparation of this manuscript, we employed LLMs to assist with English writing refinement and style polishing. All technical content, including the design of PURPLE, theoretical formulations, experimental setup, and reported results, was conceived, implemented, and validated by the authors. The LLMs were used solely for linguistic improvement and did not contribute to the research methodology or experimental findings.