

# A Multi-Agent Conversational Bandit Approach to Online Evaluation and Selection of User-Aligned LLM Responses

Xiangxiang Dai<sup>1</sup>, Yuejin Xie<sup>2</sup>, Maoli Liu<sup>1</sup>, Xuchuang Wang<sup>3</sup>,  
Zhuohua Li<sup>4, 1\*</sup>, Huanyu Wang<sup>5</sup>, John C.S. Lui<sup>1</sup>

<sup>1</sup>The Chinese University of Hong Kong

<sup>2</sup>Huazhong University of Science and Technology

<sup>3</sup>University of Massachusetts Amherst

<sup>4</sup>Guangzhou Institute of Technology, Xidian University

<sup>5</sup>Huawei Technologies Co., Ltd.

{xxdai23, mliu, csui}@cse.cuhk.edu.hk, yuejinxie@hust.edu.cn,  
xuchuangwang@umass.edu, lizhuohua@xidian.edu.cn, huanyuhello@zju.edu.cn

## Abstract

Prompt-based offline methods are commonly used to optimize large language model (LLM) responses, but evaluating these responses is computationally intensive and often fails to accommodate diverse response styles. This study introduces a novel online evaluation framework that employs a multi-agent conversational bandit model to select optimal responses while aligning with user preferences dynamically. To tackle challenges such as high-dimensional features, large response sets, adaptive conversational needs, and multi-device access, we propose MACO, Multi-Agent Conversational Online Learning, which comprises two key components: (1) MACO-A: Executed by local agents, it employs an online elimination mechanism to filter out low-quality responses. (2) MACO-S: Executed by the cloud server, it adaptively adjusts selection strategies based on aggregated preference data. An adaptive preference mechanism triggers asynchronous conversations to enhance alignment efficiency. Theoretical analysis demonstrates that MACO achieves near-optimal regret bounds, matching state-of-the-art performance in various degenerate cases. Extensive experiments utilizing Google and OpenAI text embedding models on the real-world datasets with different response styles, combined with Llama and GPT-4o, show that MACO consistently outperforms baseline methods by at least 8.29% across varying response set sizes and numbers of agents.

## 1 Introduction

Large Language Models (LLMs) have profoundly transformed the technological landscape of society (Zhu et al. 2025; Han et al. 2025; Zhang et al. 2024). A critical research direction is how to optimize responses from LLMs by exploring prompts (Liu et al. 2023). Currently, the research community is actively refining evaluation methodologies, such as prompt engineering (Guo et al. 2023; Par et al. 2023; Pryzant et al. 2023), to automatically enhance the accuracy of LLM responses and their alignment with user preferences. This process is collectively termed “**LLM response identification.**” In the offline phase, prompt-driven

approaches generate diverse response styles and formats. However, the complexity and diversity of response styles (Zhang et al. 2025a; Li et al. 2024) pose significant challenges. Currently, LLM response evaluation relies on offline pointwise scoring, assigning independent scores to each response (Liu et al. 2025b). This evaluation approach often requires generating numerous candidate responses and scoring them individually, which is time-consuming and computationally expensive due to the resource-intensive nature of LLM inference and evaluation metrics (Hu, Leung, and Farnia 2025; Zhang et al. 2025b). For instance, evaluating 205 zero-shot prompts on 784 GSM8K questions using Mistral-7B requires 78 Nvidia A6000 GPU hours (Zhou et al. 2025). In contrast, online evaluation restricts the candidate response set, enabling efficient identification of optimal responses.

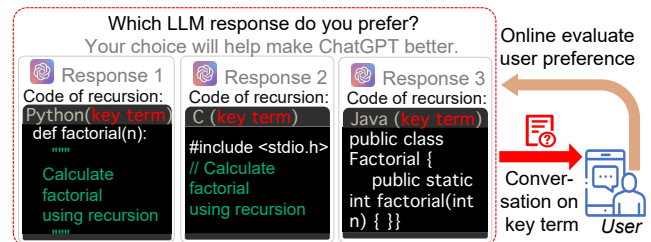


Figure 1: Evaluation aligned with online user feedback.

This study investigates the “*online*” evaluation of a curated set of candidate responses to select the optimal one. An attractive feature of online evaluation is its capability to adapt to personalized user preferences, such as humorous or formal tones, while ensuring high response quality. Although prior work (e.g., Rafailov et al. 2024) has explored preference optimization, it is often limited to binary “like/dislike” judgments. As preference demands grow increasingly diverse, LLM response identification must align with human preferences. To achieve alignment with human preferences, LLM response selection must prioritize both quality and user preference alignment. Traditional approaches rely on offline human quality control to build re-

\*Zhuohua Li is the corresponding author.

sponse pools with diverse styles, grouping responses by key terms (e.g., “C/C++” or “humorous tone”) (Zhang et al. 2020; Wu et al. 2021; Zhang et al. 2014), where these key terms can generalize to related responses for enabling preference inference. Our online frameworks take this further by allowing agents to directly query users about preferences on key terms, thereby advancing and improving the preference alignment process. For example, as shown in Fig. 1, when ChatGPT presents different code styles of recursion, user online feedback on programming style can optimize subsequent responses, enhancing preference alignment accuracy. However, multi-device access through distributed agents (e.g., queries on the Poe AI platform (Poe 2024) across smartphones, tablets, and desktops) generates fragmented and heterogeneous preference data, complicating the alignment process.

To adaptively identify suitable responses from an initial offline response set and align with human preferences, we adopt an “*online contextual conversational bandit framework*”, a dynamic and interactive response selection mechanism. Specifically, the server treats each potential LLM response as an “arm” and employs multi-armed bandit (MAB) algorithms to actively evaluate and select the next “arm” for response identification. Using this contextual conversational bandit framework, the server can also query users about preferences for key terms (e.g., humorous or serious tone) (Lei et al. 2020; Xie et al. 2021a), accelerating preference learning and alignment. However, existing conversational bandit methods fall short in addressing the unique challenges of LLM response identification, which include:

① Existing preference-aware bandit models, primarily used in recommendation systems (Wu et al. 2021; Liu et al. 2022), typically employ singular value decomposition (SVD) to extract low-dimensional feature vectors. However, the complex semantic information in LLM responses results in high-dimensional feature spaces (Cheng et al. 2025), significantly increasing computational complexity.

② Most conversational bandits assume an infinite arm framework (Liu et al. 2025a; Lei et al. 2020). In contrast, prompt engineering typically generates a large but finite set of responses as candidates. Current elimination-based contextual bandit algorithms can handle finite arm sets but rely on computationally intensive G-optimal designs (Huang et al. 2021; Lattimore and Szepesvári 2020; Li, Liu, and Lui 2024) to determine response selection distributions.

③ Existing conversational bandit studies (Wang et al. 2023; Zhang et al. 2020; Dai et al. 2024b) use predefined functions to control conversational frequency, which typically follow a fixed sequence of engagements to initiate a specific number of conversations. But this fixed interaction model struggles to adapt to the dynamic needs of LLM response identification, failing to adaptively adjust the direction and frequency of key term conversations.

④ As mentioned, the multi-device access inherent to LLMs significantly complicates conversational design. Existing conversational bandit methods are limited to single-agent scenarios and fail to address multi-device data access and processing, hindering efficient preference alignment. Conversely, existing multi-agent bandit studies multi-

agent bandit studies (Huang et al. 2021; Lin and Moothedath 2023; Wang et al. 2020) require local agents to upload user feedback or share identical arm sets, potentially leading to reduced flexibility and increased communication costs.

This paper makes the following contributions:

- **Model Formulation:** We design a multi-agent framework for online evaluation and selection of LLM responses, enhancing offline evaluation methods by incorporating user preferences for alignment, and implementing a centralized mechanism to achieve decision consensus among agents.

- **Algorithm Design:** We propose MACO, multi-agent conversational online learning, comprising MACO-A (executed by local agents) and MACO-S (executed by cloud server). MACO-A employs an online elimination mechanism to quickly filter out low-quality responses, while MACO-S dynamically adjusts selection strategies based on aggregated preference data. Our adaptive preference mechanism triggers conversations only when necessary, thus aligning with human preferences more efficiently.

- **Theoretical Analysis:** We rigorously prove regret upper and lower bounds for our proposed MACO, demonstrating near-optimal performance. The innovative conversational mechanism maintains performance while controlling communication costs at  $\mathcal{O}(d^2 M \log T)$ , eliminating the need for computationally intensive G-optimal designs used in traditional elimination-based bandits.

- **Empirical Validation:** Extensive experiments are conducted using two text embedding models from Google and OpenAI, on open-source LLM Llama and real-world datasets. Results demonstrate that the MACO significantly outperforms baselines across varying response set sizes and agent counts. Moreover, MACO reduces time overhead substantially, while preserving performance.

## 2 System Model

This section formulates the multi-agent conversational bandit for online LLM response evaluation and selection. For any vector  $\mathbf{x}$  and a positive semi-definite matrix  $\mathbf{M}$ , let  $\|\mathbf{x}\|_{\mathbf{M}} := \sqrt{\mathbf{x}^{\top} \mathbf{M} \mathbf{x}}$ . Denote the cardinality of a set  $\mathcal{A}$  as  $|\mathcal{A}|$  and  $[z] := \{1, \dots, z\}$  for  $\forall z \in \mathbb{N}^+$ .

### 2.1 Online LLM Response Identification

We define the set of local agents as  $\mathcal{M}$  with  $|\mathcal{M}| = M$ , which represent devices such as smartphones, laptops, and tablets. For any local agent  $m \in \mathcal{M}$ , the finite arm set of LLM responses is denoted as  $\mathcal{A}_m$ , which represents possible responses generated from various prompts. Given the heterogeneity of agents, different local agents may have different arm sets. As mentioned in Section 1, traditional offline techniques (e.g., prompt engineering) can help to construct a set of initial responses, but due to the diversity of LLM outputs and user preferences, it is essential to online evaluate LLM responses and select the optimal one, despite having an offline initiatory set of LLM responses. Our model adopts a time-slotted approach, denoted by discrete-time rounds  $\mathcal{T} = \{1, 2, 3, \dots, T\}$ , where each local agent selects one arm, i.e., LLM response, at each round  $t \in \mathcal{T}$ .

## 2.2 Multi-Agent User-Personalized Bandits

We consider a multi-agent conversational bandit setting involving  $M$  agents and a cloud server. At each round  $t \in \mathcal{T}$ , a local agent  $m \in \mathcal{M}$  selects an arm  $a_{m,t} \in \mathcal{A}_m$ , which denotes one possible LLM response, and receives reward feedback  $r_{m,t}$  reflects the user satisfaction. The user’s preference for LLM responses is represented by an “unknown” preference feature vector  $\theta^* \in \mathbb{R}^d$ , which all local agents aim to learn. For a local agent  $m \in \mathcal{M}$ , considering both the impact of the LLM response and the unknown user preference, the reward can be expressed as a linear combination with a noise term  $\eta_{m,t}$ :  $r_{m,t} = \langle \mathbf{x}_{a_{m,t}}, \theta^* \rangle + \eta_{m,t}$ , where  $\mathbf{x}_{a_{m,t}} \in \mathbb{R}^d$  is the embedding feature vector of the response, to capture the textual information (Bubeck et al. 2023; Liu et al. 2023). We will demonstrate the generalization of our model using two different open embedding approaches in Section 5. Our objective is to design a policy that evaluates and selects arms (i.e., LLM responses) each round to minimize cumulative regret, defined as the difference between the cumulative rewards of our policy and the best *unknown* policy across all local agents, tailored to personalized user preferences, which is defined as (Abbasi-Yadkori, Pál, and Szepesvári 2011; Dai et al. 2024c; Wu et al. 2021):

$$R_M(T) = \sum_{m=1}^M \sum_{t=1}^T \left( \mathbf{x}_{a_{m,t}^*}^\top \theta^* - \mathbf{x}_{a_{m,t}}^\top \theta^* \right), \quad (1)$$

where  $a_m^* \in \arg \max_{a \in \mathcal{A}_m} \mathbf{x}_a^\top \theta^*$  denotes the optimal arm with the highest expected reward at local agent  $m \in \mathcal{M}$ .

## 2.3 Conversational Contextual Mechanism

To adapt to and align with personalized user preferences, the cloud server also queries users’ local agents for feedback, beyond selecting arms. However, directly using all responses risks *information dispersion*, where ambiguous or overlapping groupings hinder clarity. To counter this, as shown in Fig. 2, we adopt the classical *key terms* concept (Zhang et al. 2020; Wu et al. 2021; Zhang et al. 2014; Li et al. 2024) to group related arms under core concepts, e.g., “C/C++” or “humorous tone”. Feedback on a key term applies only to its associated arms, enabling efficient preference inference with minimal interaction. In our multi-agent framework, local agent  $m$  queries a user with key term  $k_m \in \mathcal{K}_m \subseteq \mathcal{K}$  based on prompt  $y$ . Feedback, given user preference and noise  $\tilde{\eta}_{m,t}$ , is modeled as:  $\tilde{r}_{k_m,t} = \langle \tilde{\mathbf{x}}_{k_m,t}, \theta^* \rangle + \tilde{\eta}_{m,t}$ , where  $\tilde{\mathbf{x}}_k \in \mathbb{R}^d$  is the embedding vector of key term  $k$ . Unlike prior conversational bandits (Zhang et al. 2020; Wu et al. 2021; Zhao et al. 2022), which use fixed conversation schedules (e.g., linear or logarithmic in round  $t$ ) that may query unnecessarily, our algorithm (detailed in Section 3) queries *adaptively*, engaging users only when needed to refine preference estimates, improving efficiency and user experience (see Section 4 for detailed comparisons).

## 2.4 Distributed Communication Model

We consider a distributed model with  $M$  local agents and a cloud server, adopting a synchronous communication paradigm. In this setup, as shown in Fig. 2, each local

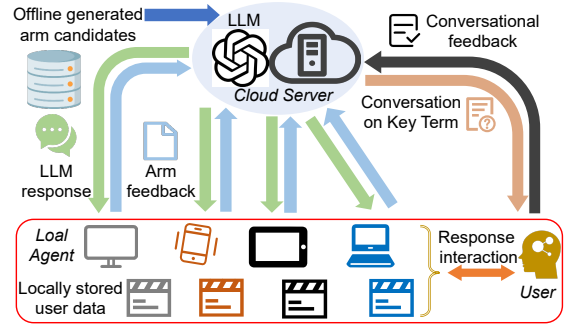


Figure 2: Multi-agent conversational bandit framework for online selecting LLM responses: Local agents handle response selection (arms), while a central server manages conversation flow through key term selection. Server aggregates interaction data across multiple agents to accelerate user preference learning.

agent communicates with the cloud server by uploading and downloading data with negligible latency. Moreover, the local agents do not directly communicate with each other. For simplicity, we focus on discrete-slot rounds solely for recording the selected arm. Querying key terms is interspersed with identifying LLM responses, allowing a key term to be queried and an arm to be pulled simultaneously. This aligns with the practical operations of conversational LLM systems. Consistent with (Huang et al. 2021), we define the standard communication cost as the cumulative count of scalar units transmitted between the cloud server and local agents, including both integers and real numbers.

## 3 Algorithm Design

We present the design of multi-agent conversational online learning (MACO) algorithms. Then, we compare our design to the traditional phase elimination-based online learning algorithms. At a high level, multi-agent conversational bandits face the issue of agent heterogeneity: naively aggregating data across agents does not necessarily lead to better user preference estimation. To mitigate this, the central algorithmic challenge lies in how to *guide exploration toward underrepresented directions in the feature space, ensuring accurate preference estimation across all relevant dimensions, with rigorous theoretical interpretability*.

### 3.1 MACO Algorithm on Local Agent

As outlined in Algorithm 1, which is executed by the local agents and referred to as MACO Agent (MACO-A), the online process of handling and updating information for LLM response evaluation within the multi-agent system operates as follows. Define  $\mathcal{T}_{m,a}^p$  as the set of rounds where local agent  $m$  selects arm  $a$  in phase  $p$ ,  $\tilde{\mathcal{T}}_{m,k}^p$  as the set of rounds when agent  $m$  conducts interaction on key term  $k$  in the same phase (like Fig. 1), and  $A \leq |\mathcal{A}|$  as the size of actually pulled arms from the LLM response set at each round. During each phase  $p$ , the local agent  $m \in \mathcal{M}$  computes the *information matrix*  $M_m^p$  from its *active arm set*  $\mathcal{A}_m^p$  (later updated in Line 13). Specifically,  $M_m^p :=$

---

**Algorithm 1: MACO on Local Agent (MACO-A)**

---

**Input:** Local agent count  $M$ , input dimension  $d$ , arm pool size  $A$ , confidence value  $\delta \in (0, 1]$ , phase  $p = 1$

```
1 while  $T$  has not been reached do
2   Calculate  $\mathbf{M}_m^p = \sum_{a \in \mathcal{A}_m^p} \frac{1}{|\mathcal{A}_m^p|} \mathbf{x}_a \mathbf{x}_a^\top$ ; Diagonalize
    $\mathbf{M}_m^p = \sum_{j=1}^d \lambda_{v_j} \mathbf{v}_j \mathbf{v}_j^\top$ 
3   Upload eigenvector  $\mathbf{v}_j$ , if its corresponding eigenvalue
   satisfies  $\lambda_{v_j} < h_p := \frac{3}{4(1-2^{-2p})d}$ 
4   Download  $\mathcal{K}_m^p$  and  $\{n_{m,k}^p\}_{k \in \mathcal{K}_m^p}$  from the cloud
   server
5   foreach  $k \in \mathcal{K}_m^p$  do ▷ Conduct
   conversations
6     Querying key term  $k$  for  $n_{m,k}^p$  times; Receive
     rewards  $\{\tilde{r}_{k,t}\}_{t \in \tilde{\mathcal{T}}_{m,k}^p}$  from conversational
     feedback
7   end
8   foreach  $a \in \mathcal{A}_m^p$  do ▷ Pull arms
9     Set  $n_{m,a}^p = \left\lceil \frac{d}{2(-2^{-2p}-1)|\mathcal{A}_m^p|} \log \frac{2AM \log T}{\delta} \right\rceil$ ; Pull  $a$ 
     for  $n_{m,a}^p$  times on the targeted LLM
10    Receive rewards  $\{r_{a,t}\}_{t \in \mathcal{T}_{m,a}^p}$  on the LLM
     response
11   end
12   Upload  $\mathbf{G}_m^p = \sum_{k \in \mathcal{K}_m^p} n_{m,k}^p \tilde{\mathbf{x}}_k \tilde{\mathbf{x}}_k^\top + \sum_{a \in \mathcal{A}_m^p} n_{m,a}^p \mathbf{x}_a \mathbf{x}_a^\top$ , and
      $\mathbf{W}_m^p = \sum_{t \in \bigcup_{k \in \mathcal{K}_m^p} \tilde{\mathcal{T}}_{m,k}^p} \tilde{r}_{k,t} \tilde{\mathbf{x}}_{k,t} + \sum_{t \in \bigcup_{a \in \mathcal{A}_m^p} \mathcal{T}_{m,a}^p} r_{a,t} \mathbf{x}_{a,t}$ 
13   Download  $\hat{\boldsymbol{\theta}}_p$  from the cloud server; Update the active
     LLM response set  $\mathcal{A}_m^{p+1}$  by eliminating sub-optimal
     LLM responses:  $\mathcal{A}_m^{p+1} =$ 
      $\left\{ a \in \mathcal{A}_m^p : \max_{a' \in \mathcal{A}_m^p} \langle \hat{\boldsymbol{\theta}}_p, \mathbf{x}_{a'} - \mathbf{x}_a \rangle \leq \frac{2^{-p+1}}{\sqrt{M}} \right\}$ ;
      $p = p + 1$ 
14 end
```

---

$\sum_{a \in \mathcal{A}_m^p} \frac{1}{|\mathcal{A}_m^p|} \mathbf{x}_a \mathbf{x}_a^\top$ , which captures principal directions in the feature space (Line 2). The eigenvalue  $\lambda_v$  of its eigenvector  $v$  reflects variance along that direction, with larger values aiding precise estimation of  $\boldsymbol{\theta}^*$ . Local agent  $m$  then diagonalizes its information matrix  $\mathbf{M}_m^p = \sum_{j=1}^d \lambda_{v_j} \mathbf{v}_j \mathbf{v}_j^\top$  to analyze all feature space directions (Line 2). If an eigenvalue  $\lambda_{v_j}$  falls below the threshold  $h_p := \frac{3}{4(1-2^{-2p})d}$  (determined by theoretical analysis), agent  $m$  uploads the corresponding eigenvector to the cloud server (Line 3), addressing under-explored feature space regions to improve LLM response selection accuracy.

The cloud server processes uploaded data and sends local agent  $m$  a set of key terms  $\mathcal{K}_m^p$  with required repetition times  $\{n_{m,k}^p\}_{k \in \mathcal{K}_m^p}$  (Line 4). Agent  $m$  then queries these key terms and pulls arms as specified, ensuring thorough exploration of LLM responses. During this process, arm pulls and key term queries can be interleaved flexibly (Lines 5-10), though shown sequentially for clarity. Agent  $m$  uploads

data on pulled arms, key terms, and rewards, stored in matrices  $\mathbf{G}_m^p$  and  $\mathbf{W}_m^p$  (Line 12). Next, it downloads the updated preference estimate  $\hat{\boldsymbol{\theta}}_p$  from the cloud server and updates its active arm set by eliminating suboptimal arms (Line 13). This process enables agent  $m$  to adaptively and accurately evaluate and select LLM responses tailored to user preferences while minimizing external data sharing by only uploading aggregated data ( $\mathbf{G}_m^p$ ,  $\mathbf{W}_m^p$ ) to the cloud server.

### 3.2 MACO Algorithm on Cloud Server

---

**Algorithm 2: MACO on Cloud Server (MACO-S)**

---

**Input:** Key term set  $\mathcal{K}$ , coverage parameter  $\beta$  in Condition 1.

**Initialization:** Let  $p = 1$ ,  $\mathbf{G} = \mathbf{0}$ ,  $\mathbf{W} = \mathbf{0}$

```
1 while  $T$  has not been reached do
2   foreach  $m \in \mathcal{M}$  do
3     Receive all eigenvectors uploaded by local agent
      $m$ , and denote this set as  $\mathcal{S}_m$ 
4     Initialize the set of key terms at phase  $p$  as
      $\mathcal{K}_m^p = \emptyset$ 
5     foreach  $v_j \in \mathcal{S}_m$  do
6        $k = \arg \max_{i \in \mathcal{K}} \tilde{\mathbf{x}}_i^\top \mathbf{v}_j$ ,  $\mathcal{K}_m^p = \mathcal{K}_m^p \cup \{k\}$ ;
        $n_{m,k}^p = \left\lceil \frac{3}{2(1-2^{-2p})\beta^2 2^{-2p}} \log \frac{2AM \log T}{\delta} \right\rceil$ 
7     end
8     Send  $\mathcal{K}_m^p$  and  $\{n_{m,k}^p\}_{k \in \mathcal{K}_m^p}$  to local agent  $m$ ;
     Receive  $\mathbf{G}_m^p$  and  $\mathbf{W}_m^p$  from local agent  $m$ 
9   end
10   $\mathbf{G} = \sum_{p \in [p]} \sum_{m \in \mathcal{M}} \mathbf{G}_m^p$ ,  $\mathbf{W} =$ 
      $\sum_{p \in [p]} \sum_{m \in \mathcal{M}} \mathbf{W}_m^p$ 
11  Broadcast  $\hat{\boldsymbol{\theta}}_p = \mathbf{G}^{-1} \mathbf{W}$  to all local agents;  $p = p + 1$ 
12 end
```

---

The MACO algorithm's cloud server component, MACO Server (MACO-S), tackles the challenge of local agent heterogeneity in the multi-agent conversational bandits model, as noted in Section 1. This diversity can hinder effective data aggregation, risking suboptimal estimation of the user preference  $\boldsymbol{\theta}^*$ . To counter this, the cloud server strategically uses key terms to probe and enhance information in underrepresented feature space directions, thus improving the accuracy of the estimation process.

In Algorithm 2, the cloud server receives eigenvectors from local agents indicating under-explored directions in the LLM response space (Line 6). It selects key terms  $k \in \mathcal{K}$  by maximizing the inner product with these directions and assigns repetition times  $n_{m,k}^p$ , sending them to local agents (Line 8). This enables targeted exploration of LLM responses related to these key terms. The server then aggregates data from all agents and estimates the preference parameter  $\boldsymbol{\theta}^*$  using linear regression, reducing uncertainty and improving tailored response predictions (Lines 10-11). To ensure invertibility, especially for large  $d$ ,  $\mathbf{G}$  can be initialized as an identity matrix.

### 3.3 Advantages over Phase Elimination Bandit

As noted in Section 1, LLM responses requiring online evaluation are finite, making phase elimination-based linear bandit algorithms (PE-Lin) preferable over classical conversational bandit frameworks (Zhang et al. 2020) due to better performance guarantees for finite arm sets. We enhance PE-Lin (Lattimore and Szepesvári 2020), where a learning agent estimates the preference vector  $\theta^*$  using least squares with *G-optimal design* to minimize prediction variance. This design optimizes a probability distribution  $\pi : \mathcal{X} \rightarrow [0, 1]$  over arm feature vectors  $\mathcal{X} \subset \mathbb{R}^d$ , satisfying:

$$\sum_{\mathbf{x} \in \mathcal{X}} \pi(\mathbf{x}) = 1, \quad M_m^p(\pi) = \sum_{\mathbf{x} \in \mathcal{X}} \pi(\mathbf{x}) \mathbf{x} \mathbf{x}^\top, \quad (2)$$

$$g(\pi) = \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_{M(\pi)^{-1}}^2 = d.$$

The agent plays arms per  $\pi$ , estimates  $\theta^*$ , and eliminates inferior arms. Computing *G-optimal design* in multi-agent settings is inefficient (Huang et al. 2021). Our algorithm, MACO, avoids this by leveraging multi-agent heterogeneity and adaptive conversational mechanisms, reducing computation costs. Running PE-Lin independently on each agent with server aggregation yields a regret bound of  $\tilde{O}(M\sqrt{dT})$ , equivalent to no communication. Our approach improves this to  $\tilde{O}(\sqrt{dMT})$  by enabling information sharing through lightweight conversations, as we will analyze in Section 4.

## 4 Performance Analysis

This section presents the theoretical results of MACO: regret, communication costs, and conversation frequency.

Following common practices in (Xie et al. 2021a; Dai et al. 2024b), we assume for any arm  $a$  and key term  $k$ ,  $\|\mathbf{x}_a\| = \|\tilde{\mathbf{x}}_k\| = 1$ . The length of preference vector  $\theta^*$  is bounded by 1, and the noise terms  $\eta_{m,t}$  and  $\tilde{\eta}_{m,t}$  are 1-subgaussian. We first present a “technical condition” that addresses general issues related to feature space coverage.

**Condition 1** (Feature Space Coverage). We say that a key term set  $\mathcal{K}$  as *sufficiently rich* for covering the feature space if, for any unit vector  $\mathbf{v} \in \mathbb{R}^d$ , there exists a key term  $k \in \mathcal{K}$  such that its feature vector  $\tilde{\mathbf{x}}_k$  satisfies  $\tilde{\mathbf{x}}_k^\top \mathbf{v} \geq \beta$ , where  $\beta \in (0, 1]$  is a coverage parameter close to 1.

**Remark 1.** Condition 1 ensures key terms in  $\mathcal{K}$  are well-distributed across the feature space, enabling effective uncertainty reduction for each agent. This is satisfied if  $\mathcal{K}$  contains an orthonormal basis of  $\mathbb{R}^d$ . It allows for bypassing the *G-optimal design* used in traditional elimination-based algorithms to minimize prediction variance.

For sufficiently rich key term sets, based on Condition 1, we provide the following theorems.

**Theorem 1** (Regret Bounds). *We have the following upper and lower regret bounds:*

1. **Upper Bound:** *With probability at least  $1 - \delta$ , the regret is bounded by  $\mathcal{O}(\sqrt{dMT} \log \frac{AM \log T}{\delta})$ .*
2. **Lower Bound:** *For any policy that selects at most one key term per round, there exists an instance where the policy incurs an expected regret of at least  $\Omega(\sqrt{dMT})$ .*

**Remark 2.** Theorem 1 highlights key performance insights for our approach: 1) For  $M = 1$ , the problem reduces to single-agent conversational bandits, with regret  $\tilde{O}(\sqrt{dT})$ , outperforming prior bounds of  $\tilde{O}(d\sqrt{T})$  (Wang et al. 2023; Zhang et al. 2020) by using phase elimination on finite arm sets, especially for high-dimensional LLM feature vectors. 2) In multi-agent settings, our upper bound matches the results of (Huang et al. 2021; Li, Liu, and Lui 2024) while avoiding computationally expensive *G-optimal design*, accelerating the online process. Moreover, unlike (Wang et al. 2020) which assumes all local agents share the same arm set, agents here have heterogeneous arm sets. 3) Collectively, the regret upper and lower bounds indicate MACO is minimax optimal up to a logarithmic factor, showing its tightness.

**Theorem 2** (Communication Cost). *The communication cost scales in  $\mathcal{O}(\log T)$ , independent of arm pool size  $A$ .*

**Remark 3.** The arm pool size  $A$  can reach thousands due to the diversity of LLM-generated responses. In contrast to (Huang et al. 2021), where the communication cost scales as  $\mathcal{O}(A \log T)$ , our approach significantly reduces overhead. Specifically, local agents process data independently and transmit only aggregated results to the cloud server, avoiding the need to upload full arm sets of size  $\mathcal{O}(A)$ .

**Theorem 3** (Bound on Conversation Frequency). *For any local agent  $m \in \mathcal{M}$  during phase  $p$ , let  $\gamma = \lambda_{\min}(M_m^p)$ , where  $\lambda_{\min}$  denotes the smallest eigenvalue, we have: 1) If  $\gamma \geq h_p$ , no conversations will be initiated. 2) If  $\gamma < h_p$ , the fraction of conversations relative to the total phase length is capped at  $\beta^{-2}(\frac{3}{4(1-2^{-2p})} - d\gamma)$ .*

**Remark 4.** MACO employs an adaptive conversation method, unlike deterministic schedules  $b(t)$  (e.g., linear or logarithmic) used in prior conversational bandit studies (Zhang et al. 2020; Wang et al. 2023; Xie et al. 2021b). Fixed-interval conversations can be inefficient, especially when user preferences are already well understood. In contrast, our model dynamically adjusts based on gaps in user preferences, enabling a more realistic interaction paradigm.

## 5 Performance Evaluation

### 5.1 Experimental Settings

**Embedding Models.** We demonstrate our framework’s generalization capabilities using different open embedding models from Google and OpenAI, which generate the embedding feature vector  $\mathbf{x}_a \in \mathbb{R}^d$  for the corresponding arm  $a$  (i.e., response) to capture text information:<sup>1</sup>

① **Text-embedding-preview-0409:** Google’s advanced embedding model, which streamlines synthetic training data creation by generating queries and task descriptions (Lee et al. 2024). ② **Text-embedding-3-large:** OpenAI’s generation embedding model, which surpasses its predecessor, though its technical details remain undisclosed (Muenighoff et al. 2023). ③ **Text-embedding-multilingual-002:** To evaluate the algorithm’s generalization across different languages, we also employ a multilingual embedding model of Google (Lee et al. 2024).

<sup>1</sup>Implementation Code: <https://github.com/TarferSoul/MACO>




Figure 3: Regret on embedding models from Google and OpenAI across different arm pool sizes  $A$ .

**Dataset Preprocessing.** We implement two response settings using the embedding models ① and ②, leveraging real-world datasets and the open-source Llama-3-8B-Instruct LLM (Ollama 2024; Sahoo et al. 2024). The first configuration generates 510 unique response arms from combinations of 13 key terms (e.g., “humorous,” “helpful”) (Köpf et al. 2023), while the second yields 455 arms by pairing five questions with two keyword styles. Both settings optimize response selection through cosine similarity between user preference and response feature vectors (Reimers and Gurevych 2019). Details are provided in the full version of our paper (Dai et al. 2025) for reference. In addition to Llama-generated responses, we evaluate MACO on the StyleEval (Li et al. 2024) and Gretel Multilingual (Van Segbroeck et al. 2024) datasets, which encompass diverse response styles (e.g., “polite,” “romantic”) and multilingual key-term variations for conversational LLM interfaces. The original multilingual dataset, limited to six languages, was expanded to 20 languages using GPT-4o to support larger-scale experiments (OpenAI 2024). As mentioned, the embedding model ③ is employed for multilingual support.

**Comparison Algorithms.** Results are averaged over five trials on a Linux Ubuntu machine (kernel 6.5.0) with a 5.40 GHz 13th Gen Intel Core i7-13700KF CPU and 32GB RAM. We set coverage parameter  $\beta = 1$  and confidence parameter  $\delta = 0.1$ , with an ablation study in the full version of our paper (Dai et al. 2025). Baselines, run independently on local agents, include: 1) TRIPLE-SH (Shi et al. 2024): Select optimal prompts for LLMs by adaptively eliminating arms with poor performance, where we directly set each arm as the corresponding LLM response. 2) LinUCB (Abbasi-Yadkori, Pál, and Szepesvári 2011): Online select arms and estimate user preference for *infinite* arm sets, excluding the conversational setting. 3) Arm-Con (Christakopoulou, Radlinski, and Hofmann 2016): Initiate conversations on user preference about arms, and use LinUCB for arm selection. 4) ConUCB (Zhang et al. 2020): Query key terms if conversations are allowed and utilize conversational feedback to accelerate learning. 5) ConLinUCB (Wang et al. 2023): Includes ConLinUCB-BS (uses barycentric spanner for conversations), ConLinUCB-MCR (selects terms with largest confidence radius), and ConLinUCB-UCB (uses LinUCB for term selection).

## 5.2 Evaluation Results

**Regret Across Different Arm Pool Sizes.** We compare MACO’s cumulative regret against seven baseline algorithms

in the first setting with  $M = 4$  local agents, using two embedding models. We vary arm pool sizes ( $A = 40, 50$ ), randomly selecting  $A$  arms from  $\mathcal{A}$  per local agent. Fig. 3 shows that non-conversational algorithms (LinUCB, Arm-Con) perform worst, while MACO outperforms all baselines, improving at least 8.29% over the best baseline, ConLinUCB-MCR. This stems from MACO’s multi-agent framework, where the cloud server aggregates local agent data to better estimate user preferences. Increasing  $A$  has minimal impact on MACO’s regret, supporting Theorem 1, which indicates regret grows at a square-root logarithmic rate with  $A$ .

**Regret Across Different Numbers of Local Agents.** We evaluate regret in the second setting with arm pool size  $A = 40$ , using the same embedding models, and varying the number of local agents ( $M = 8, 12$ ). Larger  $M$  reflects practical platforms grouping similar users to share learning, so we test MACO’s performance with increased  $M$ . Fig. 4 shows that without a multi-agent framework, baseline algorithms’ regrets grow linearly with  $M$ , following  $\tilde{O}(dM\sqrt{T})$ . In contrast, MACO leverages aggregated data from all agents, scaling regret as  $\tilde{O}(\sqrt{dMT})$ , significantly reducing regret growth. This highlights MACO’s effective multi-agent approach for online LLM response identification.

**Comparison of Overhead on Execution Time.** We compare the execution time (s) of our algorithm, MACO w/o G for emphasis, against ConLinUCB-BS (noted as the fastest in (Wang et al. 2023)) and MACO w/G (using traditional *G-optimal design*) under  $T = 5000$ , 6 phases ( $A = 40, M = 4$ ). The results of the two embedding models under the two response settings are labeled Settings (a), (b), (c), and (d). Section 5.2 (full version in (Dai et al. 2025)) shows MACO w/o G significantly reduces execution time by avoiding *G-optimal design* and leveraging multi-agent data aggregation to speed up learning. Additionally, MACO w/o G maintains the same average reward as MACO w/G, confirming that our conversational approach preserves performance while replacing *G-optimal design* with a more efficient, conversation-based design.

**Performance on Real-world Datasets.** We evaluate MACO’s cumulative regret against seven baseline algorithms on the StyleEval (Li et al. 2024) and Gretel Multilingual (Van Segbroeck et al. 2024) datasets with  $T = 1000$ , using Google and OpenAI text embedding models across varying numbers of local agents ( $M = 4, 8, 12, 16$ ). Table 1 presents results on the StyleEval dataset, where MACO consistently achieves the lowest regret, outperforming the best baseline




Figure 4: Regret on embedding models from Google and OpenAI across different agent counts  $M$ .

Algorithm	Text Embedding Model of Google				Text Embedding Model of OpenAI			
	$M = 4$	$M = 8$	$M = 12$	$M = 16$	$M = 4$	$M = 8$	$M = 12$	$M = 16$
TRIPLE-SH	5847.31±19.80	11478.73±23.16	17151.16±36.35	22673.76±21.64	7736.87±36.32	15302.63±62.52	22777.67±62.34	30138.45±25.60
LinUCB	495.67±10.87	1007.75±20.83	1510.77±31.77	2025.16±23.69	401.16±14.43	808.91±30.78	1214.77±30.88	1625.90±22.20
Arm-Con	425.89±12.13	871.32±15.63	1295.84±32.97	1742.96±53.81	319.82±14.35	645.34±30.84	953.95±38.38	1297.47±42.48
ConUCB	237.62±5.81	479.36±9.91	721.05±18.88	960.33±14.34	190.36±7.34	382.64±15.52	583.57±17.97	779.50±10.44
ConLinUCB-BS	991.73±27.14	2000.84±45.84	3005.28±67.18	4011.74±16.24	781.52±27.47	1583.90±58.27	2385.23±56.67	3177.74±49.60
ConLinUCB-MCR	439.70±13.92	895.52±15.70	1329.91±31.22	1793.32±16.30	382.10±15.50	769.74±29.30	1146.81±35.96	1536.44±34.82
ConLinUCB-UCB	1145.26±21.32	2313.41±33.67	3463.31±52.00	4634.67±43.57	937.34±31.69	1889.40±50.69	2836.83±46.72	3805.12±39.17
MACO (Ours)	<b>39.04±1.58</b>	<b>77.11±2.01</b>	<b>115.48±2.78</b>	<b>153.12±1.55</b>	<b>55.03±1.98</b>	<b>108.50±2.98</b>	<b>162.42±3.26</b>	<b>215.71±3.37</b>

Table 1: Cumulative regret for algorithms under two embeddings on the StyleEval dataset.

Setting	MACO (w/o G)		MACO (w/G)		ConLinUCB-BS	
	Time (s)	Reward	Time (s)	Reward	Time (s)	Reward
Setting (a)	2.576	61.849	9.766	61.847	18.124	59.811
Setting (b)	2.546	61.605	14.272	61.591	18.056	59.663
Setting (c)	2.576	47.405	6.369	47.381	17.926	46.104
Setting (d)	2.661	41.770	6.270	41.858	17.919	40.720

Table 2: Execution time and reward on different settings.

(ConUCB) by at least 79.56% for  $M = 4$  (Google embeddings) and 71.08% (OpenAI embeddings). Similarly, the results on the Gretel Multilingual dataset (with details in (Dai et al. 2025)), with MACO reducing regret by at least 80.07% (Google embeddings) and 69.23% (OpenAI embeddings) compared to ConUCB for  $M = 4$ . Across both datasets, MACO’s regret scales sublinearly with  $M$ , validating its robust multi-agent framework for online LLM response evaluation and selection even in diverse and multilingual settings.

## 6 Related Work

Research on prompt learning for automatically generating suitable LLM responses has made significant progress (Guo et al. 2023; Zhang et al. 2023; Poon et al. 2025). For example, Li et al. (2024) designs StyleChat for stylized response generation. However, offline generating methods face challenges like “data drift,” emphasizing the need for online approaches to optimize LLM responses (Bhardwaj et al. 2022; Chen, Zaharia, and Zou 2023). Xia et al. (2024) introduces an online non-stationary bandit method across different LLMs. Shi et al. (2024) proposes an online budget-limited LLM response optimization using various prompts. And Dai et al. (2024a) focuses on response identification over multiple LLM coordination. Nevertheless, these studies ignore the impact of user preferences and the natural multi-agent setting in LLM response identification.

Bandits tackle the exploitation-exploration tradeoff of online decision-making problems (Abbasi-Yadkori, Pál, and

Szepesvári 2011). Based on this, conversational contextual linear bandits, introduced by (Zhang et al. 2020), allow the cloud server to obtain user feedback on key terms to elicit preferences, in addition to arm selection. Later studies introduce clustering to avoid labeling efforts (Wu et al. 2021), integrate knowledge graphs for term selection (Zhao et al. 2022), and compute the barycentric spanner as an efficient exploration basis (Wang et al. 2023). Regarding the multi-agent bandit setting under finite arm sets, Wang et al. (2020) assumes homogeneous arm sets, Huang et al. (2021) requires the local agents to upload arm sets, increasing costs and privacy concerns, and (Li, Liu, and Lui 2024) utilizes the computationally intensive *G-optimal design*. Unlike existing works, we are the first to extend conversational bandits to multi-agent settings for online LLM response adaptation with reduced computation resources.

## 7 Conclusion and Future Work

This paper presents MACO, a novel multi-agent conversational bandit framework designed to online evaluate and select optimal responses from LLMs while minimizing cumulative regret and aligning with user preferences. Our MACO framework consists of local agents (MACO-A) that adaptively manage conversations and response selection, and a cloud server (MACO-S) that aggregates data to learn user preferences efficiently. We have proved that MACO achieves optimal regret bounds, reduces conversations, and enhances computational efficiency. Our extensive evaluations, using open LLMs such as Llama and GPT-4o, confirm that our approach significantly improves performance over traditional methods on real-world datasets. Future work could explore clustering similar user preferences, integrating privacy-preserving techniques, and extending the framework beyond linear reward models.

## Ethics Statement

This research raises no ethical concerns, with no involvement of human subjects or potential for harm.

## Acknowledgements

The work of Xiangxiang Dai was supported by the National Natural Science Foundation of China (625B2163). The work of John C.S. Lui was supported in part by the RGC GRF-14215722.

## References

- Abbasi-Yadkori, Y.; Pál, D.; and Szepesvári, C. 2011. Improved Algorithms for Linear Stochastic Bandits. In *Proceedings of the NeurIPS*.
- Bhardwaj, R.; Xia, Z.; Ananthanarayanan, G.; Jiang, J.; Shu, Y.; Karianakis, N.; Hsieh, K.; Bahl, P.; and Stoica, I. 2022. Ekya: Continuous learning of video analytics models on edge compute servers. In *Proceedings of the NSDI*, 119–135.
- Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y. T.; Li, Y.; Lundberg, S.; et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Chen, L.; Zaharia, M.; and Zou, J. 2023. FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance. *arXiv preprint arXiv:2305.05176*.
- Cheng, J.; Dai, X.; Ding, N.; Lui, J. C. S.; and Huang, J. 2025. Trading Vector Data in Vector Databases. *arXiv:2511.07139*.
- Christakopoulou, K.; Radlinski, F.; and Hofmann, K. 2016. Towards Conversational Recommender Systems. In *Proceedings of ACM SIGKDD International Conference*, 815–824.
- Dai, X.; Li, J.; Liu, X.; Yu, A.; and Lui, J. 2024a. Cost-Effective Online Multi-LLM Selection with Versatile Reward Models. *arXiv preprint arXiv:2405.16587*.
- Dai, X.; Wang, Z.; Xie, J.; Liu, X.; and Lui, J. C. 2024b. Conversational Recommendation with Online Learning and Clustering on Misspecified Users. *IEEE Transactions on Knowledge and Data Engineering*, 36(12): 7825–7838.
- Dai, X.; Wang, Z.; Xie, J.; Yu, T.; and Lui, J. C. 2024c. Online Learning and Detecting Corrupted Users for Conversational Recommendation Systems. *IEEE Transactions on Knowledge and Data Engineering*, 36(12): 8939–8953.
- Dai, X.; Xie, Y.; Liu, M.; Wang, X.; Li, Z.; Wang, H.; and Lui, J. 2025. Multi-Agent Conversational Online Learning for Adaptive LLM Response Identification. *arXiv preprint arXiv:2501.01849*.
- Guo, Q.; Wang, R.; Guo, J.; Li, B.; Song, K.; Tan, X.; Liu, G.; Bian, J.; and Yang, Y. 2023. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. *arXiv preprint arXiv:2309.08532*.
- Han, Z.; Wang, H.; Zhang, Z.; Dai, X.; Liu, X.; and Lui, J. 2025. HiLoRA: Adaptive Hierarchical LoRA Routing for Training-Free Domain Generalization. *arXiv e-prints*, arXiv–2510.
- Hu, X.; Leung, H.-f.; and Farnia, F. 2025. An Online Learning Approach to Prompt-based Selection of Generative Models and LLMs. In *Forty-second International Conference on Machine Learning*.
- Huang, R.; Wu, W.; Yang, J.; and Shen, C. 2021. Federated linear contextual bandits. *Advances in neural information processing systems*, 34: 27057–27068.
- Köpf, A.; Kilcher, Y.; von Rütte, D.; Anagnostidis, S.; Tam, Z.-R.; Stevens, K.; Barhoum, A.; Duc, N. M.; Stanley, O.; Nagyfi, R.; ES, S.; Suri, S.; Glushkov, D.; Dantluri, A.; Maguire, A.; Schuhmann, C.; Nguyen, H.; and Mattick, A. 2023. OpenAssistant Conversations – Democratizing Large Language Model Alignment. *arXiv preprint arXiv:2304.07327*.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit Algorithms*. Cambridge University Press.
- Lee, J.; Dai, Z.; Ren, X.; Chen, B.; Cer, D.; Cole, J. R.; Hui, K.; Boratko, M.; Kapadia, R.; Ding, W.; Luan, Y.; Duddu, S. M. K.; Abrego, G. H.; Shi, W.; Gupta, N.; Kusupati, A.; Jain, P.; Jonnalagadda, S. R.; Chang, M.-W.; and Naim, I. 2024. Gecko: Versatile Text Embeddings Distilled from Large Language Models. *arXiv preprint arXiv:2403.20327*.
- Lei, W.; He, X.; de Rijke, M.; and Chua, T.-S. 2020. Conversational recommendation: Formulation, methods, and evaluation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2425–2428.
- Li, J.; Zhang, Z.; Tu, Q.; Cheng, X.; Zhao, D.; and Yan, R. 2024. Stylechat: Learning recitation-augmented memory in llms for stylized dialogue generation. *arXiv preprint arXiv:2403.11439*.
- Li, Z.; Liu, M.; and Lui, J. C. S. 2024. FedConPE: Efficient Federated Conversational Bandits with Heterogeneous Clients. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 4533–4541. International Joint Conferences on Artificial Intelligence Organization.
- Lin, J.; and Moothedath, S. 2023. Federated Stochastic Bandit Learning with Unobserved Context. *arXiv preprint arXiv:2303.17043*.
- Liu, M.; Li, Z.; Dai, X.; and Lui, J. C. 2025a. Leveraging the Power of Conversations: Optimal Key Term Selection in Conversational Contextual Bandits. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 1833–1844.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35.
- Liu, X.; Zhao, H.; Yu, T.; Li, S.; and Lui, J. C. 2022. Federated online clustering of bandits. In *Proceedings of the UAI*, 1221–1231.
- Liu, Z.; Wang, P.; Xu, R.; Ma, S.; Ruan, C.; Li, P.; and Wu, Y. 2025b. Inference-Time Scaling for Generalist Reward Modeling. *arXiv preprint arXiv:2504.02495*.

- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2023. MTEB: Massive Text Embedding Benchmark. *arXiv preprint arXiv:2210.07316*.
- Ollama. 2024. <https://github.com/jmorganca/ollama/>.
- OpenAI. 2024. GPT. <https://platform.openai.com/docs/models/gpt-3-5>.
- Pan, R.; Xing, S.; Diao, S.; Liu, X.; Shum, K.; Zhang, J.; and Zhang, T. 2023. Plum: Prompt Learning using Metaheuristic. *arXiv preprint arXiv:2311.08364*.
- Poe. 2024. <https://poe.com/ChatGPT>.
- Poon, M.; Dai, X.; Liu, X.; Kong, F.; Lui, J.; and Zuo, J. 2025. Online Multi-LLM Selection via Contextual Bandits under Unstructured Context Evolution. *arXiv preprint arXiv:2506.17670*.
- Pryzant, R.; Iter, D.; Li, J.; Lee, Y. T.; Zhu, C.; and Zeng, M. 2023. Automatic prompt optimization with “gradient descent” and beam search. *arXiv preprint arXiv:2305.03495*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the EMNLP-IJCNLP*, 3982–3992.
- Sahoo, P.; Singh, A. K.; Saha, S.; Jain, V.; Mondal, S.; and Chadha, A. 2024. A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications. *arXiv preprint arXiv:2402.07927*.
- Shi, C.; Yang, K.; Yang, J.; and Shen, C. 2024. Best Arm Identification for Prompt Learning under a Limited Budget. *arXiv preprint arXiv:2402.09723*.
- Van Segbroeck, M.; Emadi, M.; Nathawani, D.; Ramaswamy, L.; Greco, J.; Boyd, K.; Grossman, M.; and Meyer, Y. 2024. Synthetic Multilingual LLM Prompts: A synthetic multilingual prompt dataset for prompting LLMs.
- Wang, Y.; Hu, J.; Chen, X.; and Wang, L. 2020. Distributed Bandit Learning: Near-Optimal Regret with Efficient Communication. In *Proceedings of the ICLR*.
- Wang, Z.; Liu, X.; Li, S.; and Lui, J. C. S. 2023. Efficient Explorative Key-Term Selection Strategies for Conversational Contextual Bandits. *Proceedings of the AAI*, 10288–10295.
- Wu, J.; Zhao, C.; Yu, T.; Li, J.; and Li, S. 2021. Clustering of Conversational Bandits for User Preference Learning and Elicitation. In *Proceedings of the ACM CIKM*, 2129–2139.
- Xia, Y.; Kong, F.; Yu, T.; Guo, L.; Rossi, R. A.; Kim, S.; and Li, S. 2024. Which LLM to Play? Convergence-Aware Online Model Selection with Time-Increasing Bandits. In *Proceedings of the ACM on Web Conference 2024*, 4059–4070.
- Xie, Z.; Yu, T.; Zhao, C.; and Li, S. 2021a. Comparison-Based Conversational Recommender System with Relative Bandit Feedback. In *Proceedings of the International ACM SIGIR Conference*, 1400–1409.
- Xie, Z.; Yu, T.; Zhao, C.; and Li, S. 2021b. Comparison-based Conversational Recommender System with Relative Bandit Feedback. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1400–1409.
- Zhang, C.; Zhong, H.; Zhang, K.; Chai, C.; Wang, R.; Zhuang, X.; Bai, T.; Jiantao, Q.; Cao, L.; Fan, J.; et al. 2025a. Harnessing Diversity for Important Data Selection in Pre-training Large Language Models. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, M.; Tang, J.; Zhang, X.; and Xue, X. 2014. Addressing cold start in recommender systems: A semi-supervised co-training algorithm. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, 73–82.
- Zhang, R.; Cheng, A.; Luo, Y.; Dai, G.; Yang, H.; Liu, J.; Xu, R.; Du, L.; Du, Y.; Jiang, Y.; et al. 2024. Decomposing the neurons: Activation sparsity via mixture of experts for continual test time adaptation. *arXiv preprint arXiv:2405.16486*.
- Zhang, R.; Dong, M.; Zhang, Y.; Heng, L.; Chi, X.; Dai, G.; Du, L.; Du, Y.; and Zhang, S. 2025b. Mole-vla: Dynamic layer-skipping vision language action model via mixture-of-layers for efficient robot manipulation. *arXiv preprint arXiv:2503.20384*.
- Zhang, X.; Xie, H.; Li, H.; and C.S. Lui, J. 2020. Conversational Contextual Bandit: Algorithm and Application. In *Proceedings of The Web Conference*, 662–672.
- Zhang, Z.; Wang, S.; Yu, W.; Xu, Y.; Iter, D.; Zeng, Q.; Liu, Y.; Zhu, C.; and Jiang, M. 2023. Auto-Instruct: Automatic Instruction Generation and Ranking for Black-Box Language Models. *arXiv preprint arXiv:2310.13127*.
- Zhao, C.; Yu, T.; Xie, Z.; and Li, S. 2022. Knowledge-Aware Conversational Preference Elicitation with Bandit Feedback. In *Proceedings of the ACM Web Conference 2022*, 483–492.
- Zhou, J. P.; Belardi, C. K.; Wu, R.; Zhang, T.; Gomes, C. P.; Sun, W.; and Weinberger, K. Q. 2025. On Speeding Up Language Model Evaluation. In *The Thirteenth International Conference on Learning Representations*.
- Zhu, Q.; Zhao, R.; Yan, H.; He, Y.; Chen, Y.; and Gui, L. 2025. Soft Reasoning: Navigating Solution Spaces in Large Language Models through Controlled Embedding Exploration. In *Forty-second International Conference on Machine Learning*.