

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CONTEXTUAL BANDITS WITH LLM-DERIVED PRIORS AND ADAPTIVE CALIBRATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) offer rich prior knowledge that can accelerate online decision-making, yet their use in contextual bandits lacks principled mechanisms for guiding exploration. We address this gap by proposing a lightweight framework that integrates LLM-derived priors with adaptive calibration in a multi-armed bandit setting. Our method first extracts prompt-based rewards from the LLM to provide task-specific supervision. We then construct an embedding-based estimator that quantifies uncertainty from the LLM’s representations, yielding a calibrated exploration signal. To remain robust under distribution shifts, we introduce an online contextual adapter that dynamically updates these uncertainty estimates during interaction. Experiments on LastFM and MovieLens-1M show that our method consistently outperforms both classical bandits and pure LLM-based agents, achieving higher cumulative rewards with significantly fewer LLM queries. Furthermore, we provide theoretical regret guarantees that establish improved sample efficiency compared to standard contextual bandits.

1 INTRODUCTION

In online decision-making tasks such as adaptive tutoring, information routing, and interactive recommendation, systems must continually balance exploration and exploitation to maximize long-term outcomes (Huang et al., 2024a; Chen et al., 2024b; Borchers & Shou, 2025). Traditional approaches to these problems are often framed as contextual bandits or reinforcement learning, where effective exploration strategies are crucial for data-efficient learning (Chen et al., 2024b). Large language models (LLMs) introduce new opportunities in this setting: beyond serving as inference engines, they may accelerate learning by reasoning about context, generalizing from limited feedback, and adapting decisions in real time.

However, for these applications to be truly reliable, LLMs need robust mechanisms to estimate uncertainty (Abbasi Yadkori et al., 2024; Ma et al., 2025). This is essential for balancing exploration and exploitation, ensuring the model’s robustness under distribution shifts, and supporting data-efficient learning (Zhang et al., 2023; Deng & Raffel, 2023). Despite their remarkable performance, current approaches to uncertainty estimation in LLMs often rely on heuristic methods such as temperature scaling, response entropy, or sampling-based diversity measures (Huang et al., 2024b). While these techniques can be useful in specific scenarios, they lack calibration, fail to generalize across different tasks, and offer limited insights into the model’s confidence levels. Other more sophisticated methods involve fine-tuning or ensemble approaches, which are computationally intensive and not suitable for general-purpose, pre-trained LLMs (Krishnan et al., 2024).

To address these limitations, we propose a lightweight and flexible framework for uncertainty-aware decision-making with LLMs. This framework includes two key components: (1) a prompt reward, which extracts supervision from the model’s outputs by leveraging internal knowledge, and (2) a last-layer embedding-based uncertainty estimator that captures semantic confidence. Additionally, our approach adapts to distribution shifts using an online contextual uncertainty adapter, allowing for dynamic calibration in changing environments. We demonstrate this framework within a multi-armed bandit (MAB) setting, integrating it with exploration strategies through calibrated Upper Confidence Bound/Lower Confidence Bound estimates.

The LLM and bandit integration is non-trivial due to the inherent challenges of mapping high-dimensional, task-agnostic LLM embeddings into actionable uncertainty estimates. These challenges

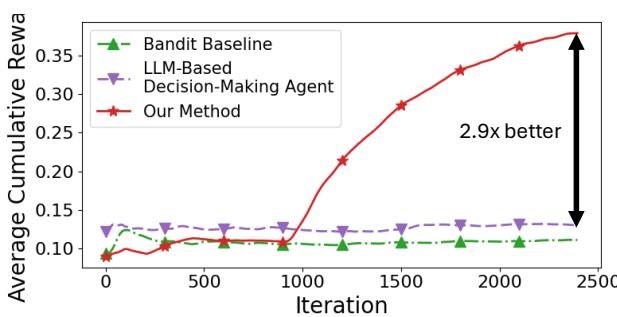


Figure 1: Performance comparison on the LastFM recommendation benchmark, comparing (i) a conventional bandit baseline, (ii) an LLM-based decision agent, and (iii) our calibrated LLM-MAB method. Further implementation details are provided in Section 5.1.

include resolving semantic misalignment, addressing representation drift, and ensuring consistent scale. Moreover, integrating these uncertainty estimates into structured decision-making frameworks, such as MABs, requires careful calibration to maintain theoretical guarantees while preserving the efficiency of pre-trained LLMs. Our method successfully addresses these complexities while remaining lightweight and compatible with zero-shot applications. To summarize, our key contributions are:

- We propose a lightweight LLM-calibrated agent that combines prompt-derived supervision with an embedding-based uncertainty estimator to guide multi-arm bandit exploration.
- We demonstrate an LLM-enabled enhancement to exploration strategies: dynamically adjusting conservative UCB estimates during early learning phases.
- We validate our framework on multiple benchmarks, achieving state-of-the-art cumulative reward while significantly reducing expensive LLM queries compared to pure LLM and classical bandit baselines (See Figure 1).
- We provide theoretical guarantees for the proposed method that quantifies the regret bound on online learning, which demonstrates improvement over pure LLM-based prediction or classic algorithms such as LinUCB.

Paper Outline: The remainder of this paper is organized as follows. Section 2 reviews related work on contextual bandits and bandits in LLMs. Section 3 introduces the preliminaries and formalizes our problem setup. Section 4 presents our methodology, including prompt reward extraction, embedding-based uncertainty quantification, and integration with multi-armed bandits, along with theoretical analysis. Section 5 describes the experimental setup and results, covering datasets, baselines, implementation details, ablation studies, and computational efficiency. Section 6 concludes with a summary of contributions, while the appendix provides additional case studies and theoretical proofs.

2 RELATED WORK

2.1 CONTEXTUAL BANDITS

Contextual bandits (CB) formalize sequential decision-making with side information under bandit feedback, balancing exploration and exploitation to maximize cumulative reward. Foundational reductions and algorithms established the modern landscape. [Agarwal et al. \(2014b\)](#) reduced CB to cost-sensitive classification with provable guarantees and practical policies; [Langford & Zhang \(2008\)](#) analyzed Epoch-Greedy as a simple exploration strategy. Linear models are central: [Li et al. \(2010a\)](#) introduced LinUCB with large-scale deployment on Yahoo news recommendation, while [Abbasi-Yadkori et al. \(2011a\)](#) provided tighter confidence-based analysis OFUL for stochastic linear rewards. Beyond linearity, [Filippi et al. \(2010\)](#) studied generalized linear bandits, and kernel methods such as GP- and RKHS-based UCB extended CBs to rich function classes ([Valko et al., 2013](#)). Bayesian and Thompson sampling perspectives offer complementary solutions. [Agrawal & Goyal \(2013\)](#) established regret guarantees for Thompson Sampling in linear bandits, and subsequent work refined practical implementations for large-scale problems ([Riquelme et al., 2018](#)). As representation

learning became critical, neural variants emerged: NeuralUCB and NeuralTS control optimism or posterior uncertainty over deep features to achieve sublinear regret in non-linear regimes while remaining computationally feasible (Zhou et al., 2020). Practical CB systems must also handle constraints and dynamics. Agarwal et al. (2014a) introduced Contextual Bandits with Knapsacks to manage resource and fairness constraints, while non-stationary environments have been studied via variation-budget and change-point models to maintain low regret under drift (Besbes et al., 2014).

2.2 LLM-BASED BANDITS

The integration of large language models (LLMs) with online learning and bandit algorithms has recently attracted increasing attention. Early studies recognized the synergy between bandits and LLMs: bandit methods provide exploration-exploitation guarantees, while LLMs offer contextual reasoning and prior knowledge. The survey of Bouneffouf & Feraud (2025) systematically outlines this emerging connection, in terms of both how bandit algorithms can enhance the efficiency and adaptability of LLMs, and how LLMs can contribute to the contextual and adaptive decision-making capabilities of bandit algorithms. Monea et al. (2024) found that LLMs exhibit in-context reinforcement-learning behavior, providing a mechanism by which few-shot prompts can induce sequential improvement without gradient updates. Almandari et al. (2024) showed empirically that LLM-generated prior knowledge can “jump-start” exploration by supplying informative priors, thereby accelerating learning in early rounds. Xia et al. (2024) studied the problem of in-context dueling bandits with LLM agents, demonstrating that preference-based feedback can be leveraged beyond numeric rewards to facilitate the combination. Beyond these methodological research, the combination of LLM with bandits also witnesses application-driven success, such as LLM-tailored health messaging (Song et al., 2025), online marketing (Ye et al., 2025), among others.

3 PRELIMINARIES

Notations. Let $n \in \mathbb{N}_+$ be a positive integer. $[n]$ denotes the set $\{1, \dots, n\}$. For any set \mathcal{S} , $|\mathcal{S}|$ denotes the number of elements in \mathcal{S} . For vector norms, $\|\mathbf{x}\|_p$ denotes the ℓ_p norm of vector \mathbf{x} .

Problem Formulation. We formulate the problem as a contextual decision-making task, where the expected reward of an action depends on both user and item contexts. The goal of the learning agent is to maximize the cumulative expected reward over time. Let $\mathcal{I} := \{1, \dots, I\}$ denote a finite set of items (e.g., movies, articles) to be recommended, referred to as arms, where $I \in \mathbb{N}_+$. Each item $i \in \mathcal{I}$ is associated with a context vector $\mathbf{c}_i \in \mathcal{C}$, where the context space \mathcal{C} represents a subspace of the language space (e.g., text descriptions or prompts). The learning agent interacts with users over $T \in \mathbb{N}_+$ rounds and has access to a large language model, denoted as LLM . In each round $t = 1, 2, \dots, T$, the agent observes the context $\mathbf{x}_t \in \mathcal{C}$ of the incoming user and a candidate subset of arms $\mathcal{I}_t \subseteq \mathcal{I}$. Based on the user context \mathbf{x}_t , item contexts $\{\mathbf{c}_i\}_{i \in \mathcal{I}_t}$, and past interactions, the agent constructs a prompt p_t to query the LLM and receives a textual response $\mathbf{o}_t \sim \text{LLM}(p_t)$. Using \mathbf{o}_t and potentially other relevant information (e.g., offline user-item profiling), the agent selects an item $i_t \in \mathcal{I}_t$ to recommend to the user.

The agent then receives a reward $r_t \in [0, 1]$, with expected value $\mathbb{E}[r_t] = f(\mathbf{x}_t, \mathbf{c}_{i_t})$. The reward function f is unknown and will be further specified below. Let $i_t^* = \arg \max_{i \in \mathcal{I}_t} f(\mathbf{x}_t, \mathbf{c}_i)$ denote the optimal arm at round t that yields the highest expected reward. The agent’s performance is measured by the cumulative regret:

$$R(T) = \sum_{t=1}^T f(\mathbf{x}_t, \mathbf{c}_{i_t^*}) - \sum_{t=1}^T f(\mathbf{x}_t, \mathbf{c}_{i_t}). \quad (1)$$

Reward Model. In this work, we assume a linear contextual reward model, where there exists an embedding function $e : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}^d$ and an unknown user preference vector $\boldsymbol{\theta}^* \in [0, 1]$, such that:

$$r_t = e(\mathbf{x}_t, \mathbf{c}_{i_t})^\top \boldsymbol{\theta}^* + \varepsilon_t, \quad (2)$$

where ε_t is a zero-mean 1-sub-Gaussian noise term. The function f is thus defined as $f(\mathbf{x}_t, \mathbf{c}_{i_t}) = e(\mathbf{x}_t, \mathbf{c}_{i_t})^\top \boldsymbol{\theta}^*$, and we assume $\|\boldsymbol{\theta}^*\|_2 \leq 1$. While our setting focuses on the linear model for clarity, our algorithm can be naturally extended to other settings such as logistic models (e.g.,

162 $f = \sigma(e(x_t, c_{i_t})^\top \theta^*)$), kernelized reward models, or general non-linear approximators (e.g., neural
 163 networks). With a little abuse of notation, we may use $x_{t,i} = [x_t, c_i]$ to denote the joint context of
 164 user t and item i .

4 METHODOLOGY

168 In this section, we describe our two-stage framework, comprising (i) an offline expert finetuning
 169 phase (Algorithm 1), and (ii) an online LLM–MAB integration phase (Algorithm 2). We detail
 170 our lightweight and flexible framework for uncertainty-aware decision-making with LLMs. Our
 171 approach comprises two complementary components: a *prompt reward* module that leverages internal
 172 LLM knowledge for supervisory signals, and a last-layer embedding-based uncertainty estimator
 173 that quantifies semantic confidence. We further introduce an online contextual uncertainty adapter,
 174 enabling dynamic calibration under distributional shifts. Finally, we describe how these elements
 175 integrate within a multi-armed bandit (MAB) setting, yielding calibrated exploration strategies.

4.1 PROMPT REWARD EXTRACTION

177 The first component of our framework harnesses the generative capabilities of pretrained LLMs to
 178 produce reward estimates without additional fine-tuning. Given a context $x \in \mathbb{R}^d$, we construct a
 179 natural language prompt that queries the LLM for an evaluation of the expected reward. The resulting
 180 probability or score is then normalized to fall in $[0, 1]$, forming the *prompt-derived reward* r_{prompt} .
 181 This mechanism exploits the LLM’s internal representations and world knowledge to provide zero-
 182 shot supervision, circumventing the need for task-specific training while capturing rich semantic
 183 cues.

4.2 EMBEDDING-BASED UNCERTAINTY QUANTIFICATION

186 To accompany prompt rewards with calibrated uncertainty estimates, we introduce a lightweight
 187 estimator based on the penultimate-layer embeddings of the LLM. For each context x , we extract
 188 the last-layer feature vector $e(x) \in \mathbb{R}^{d_{\text{last}}}$ from the LLM. We then learn a linear projection (logistic
 189 regression) head $\phi : \mathbb{R}^{d_{\text{last}}} \rightarrow [0, 1]$ on offline dataset predicting the observed reward. Simultaneously,
 190 we accumulate the empirical covariance matrix:

$$193 \quad \mathbf{V} = \sum_{i=1}^n e(x_i) e(x_i)^\top + \lambda \mathbf{I}, \quad (3)$$

196 where λ is a regularization constant. At inference time, we compute a confidence interval around the
 197 point estimate via:

$$199 \quad r_{\text{est}}(x) = \phi(e(x)), \quad u(x) = \beta \sqrt{e(x)^\top \mathbf{V}^{-1} e(x)}, \quad (4)$$

200 generating upper and lower confidence bounds $r_{\text{est}}(x) \pm u(x)$. This design captures semantic
 201 misalignment and representation drift by directly operating on high-dimensional embeddings, while
 202 remaining computationally efficient.

203 **Online Contextual Uncertainty Adapter.** While the expert module is initially trained on an offline
 204 dataset, real-world deployments often encounter distributional shifts that deviate from the training
 205 distribution. Such shifts can degrade the calibration quality of the reward estimator and uncertainty
 206 quantifier. To address this, we introduce an online contextual uncertainty adapter. To be specific,
 207 we deploy an online adapter that periodically updates ϕ and \mathbf{V} using newly observed $(e(x_t), r_t)$
 208 pairs. This adapter leverages incremental updates to the covariance matrix and online least-squares
 209 adjustment of ϕ , ensuring that uncertainty estimates remain accurate under distributional shifts.
 210 By blending offline pretraining with light online adaptation, our framework balances stability and
 211 flexibility.

4.3 INTEGRATION WITH MULTI-ARMED BANDITS

213 To effectively leverage both the semantic priors from prompt-based LLM supervision and the
 214 calibrated uncertainty estimates from embedding-based prediction, we integrate these signals into

216 a contextual multi-armed bandit (MAB) framework. This integration enables uncertainty-aware
 217 exploration that is both data-efficient in early rounds and adaptively grounded in observed feedback
 218 as more data becomes available. In particular, the LLM’s prompt-derived reward offers informative,
 219 zero-shot guidance at the start of learning, while the embedding-based estimator provides increasingly
 220 accurate posterior estimates with calibrated uncertainty as the environment is explored. By fusing
 221 these two perspectives, we design a composite reward function that smoothly transitions from
 222 prior-driven to data-driven decision-making. We define the adjusted reward for each arm as:

$$\tilde{r}_i = \text{clip}(r_{\text{prompt}}(x_i); r_{\text{est}}(x_i) - \nu(x_i), r_{\text{est}}(x_i) + \nu(x_i)). \quad (5)$$

223 During an initial exploration phase, actions are selected purely by maximizing r_{prompt} . In the
 224 subsequent joint phase, the clipped reward \tilde{r}_i drives selection, optionally augmented with classic UCB
 225 terms from a linear bandit on raw features. This calibrated fusion preserves theoretical exploration
 226 guarantees while capitalizing on pretrained LLM knowledge and uncertainty estimates. Experimental
 227 results in Section 5 demonstrate that this integration yields substantial improvements in cumulative
 228 reward and sample efficiency.

Algorithm 1 Offline LLM Expert Finetuning

233 1: **Input:** Offline dataset $\mathcal{D} = \{(x_t, r_t)\}_{t=1}^n$, foundation model LLM , regularization factor λ .
 234 2: **for** $t = 1, \dots, n$ **do**
 235 3: Input x_t into LLM and extract the second last layer’s embedding $e_t(x_t, r_t) \in \mathbb{R}^{d_{\text{last}}}$.
 236 4: **end for**
 237 5: Freeze all but the projection head ϕ (the last linear layer) of LLM , retrain ϕ through dataset \mathcal{D} .
 238 6: **Return:** Projection head $\phi : \mathbb{R}^{d_{\text{last}}} \rightarrow [0, 1]$, covariance matrix $\mathbf{V} = \sum_{t=1}^n e_t(x_t, r_t) e_t^\top(x_t, r_t) +$
 239 $\lambda \mathbf{I} \in \mathbb{R}^{d_{\text{last}} \times d_{\text{last}}}$.

Algorithm 2 Online LLM and MAB integration (A1): calibrate LLM with MAB

243 1: **Input:** Finetuned LLM , projection layer ϕ , covariance matrix \mathbf{V} , online decision rounds T ,
 244 exploration phase σT , coefficient $\beta_{\text{LLM}}, \beta_{\text{MAB}}$.
 245 2: **Initialize:** MAB covariance matrix $\mathbf{G}_t = \lambda \mathbf{I} \in \mathbb{R}^{d \times d}$, and regressand $\mathbf{b}_t = 0 \in \mathbb{R}^d$.
 246 3: **for** $t = 1, \dots, T$ **do**
 247 4: User u_t comes to the system with K items with contexts $x_{t,1}, \dots, x_{t,K} \in \mathbb{R}^d$
 248 5: Extract the second last layer’s embedding $e_t(x_{t,i}) \in \mathbb{R}^{d_{\text{last}}}$ for each item $i \in [K]$.
 249 6: Compute LLM reward prediction (1) estimated reward $r_{\text{LLM}}(i) = \text{Prompt}$, (2) UCB-style
 250 reward: $r_{\text{LLM}}(i) = \phi^\top e_t(x_{t,i}) + \beta_{\text{LLM}} \sqrt{e_t^\top(x_{t,i}) \mathbf{V}^{-1} e_t(x_{t,i})}$, (3) LCB-style reward: $r_{\text{LLM}}(i) =$
 251 $\phi^\top e_t(x_{t,i}) - \beta_{\text{LLM}} \sqrt{e_t^\top(x_{t,i}) \mathbf{V}^{-1} e_t(x_{t,i})}$.
 252 7: Compute UCB reward prediction for MAB: $r_{\text{UCB}}(i) = \hat{\theta}_t^\top x_{t,i} + \beta_{\text{MAB}} \sqrt{x_{t,i}^\top \mathbf{G}_t^{-1} x_{t,i}}$
 253 8: Compute LCB reward prediction for MAB: $r_{\text{LCB}}(i) = \hat{\theta}_t^\top x_{t,i} - \beta_{\text{MAB}} \sqrt{x_{t,i}^\top \mathbf{G}_t^{-1} x_{t,i}}$
 254 9: **if** $t \leq \sigma T$ **then**
 255 10: Select $i_t = \arg \max_{i \in [K]} r_{\text{LLM}}(i)$.
 256 11: **else**
 257 12: Select $i_t = \arg \max_{i \in [K]} \tilde{r}(i)$, where $\tilde{r}(i) = \text{CLIP}_{[r_{\text{LCB}}(i), r_{\text{UCB}}(i)]} r_{\text{LLM}}(i)$.
 258 13: **end if**
 259 14: Receive reward r_t for item i_t .
 260 15: Update MAB’s statistics $\mathbf{G}_{t+1} = \mathbf{G}_t + x_{t,i_t} x_{t,i_t}^\top$, $\mathbf{b}_{t+1} = \mathbf{b}_t + x_{t,i_t} r_t$.
 261 16: **end for**

 265 4.4 THEORETICAL ANALYSIS

266 To formally characterize the performance of our LLM–MAB integration, we establish theoretical
 267 assumptions and supporting lemmas.

270 **Assumption 1** (Linear reward model). Suppose the feature vectors $e_t(a) \in \mathbb{R}^d$ satisfy $\|\phi_t(a)\|_2 \leq L$.
 271 The reward $r_t \in [0, 1]$ of the played arm has the following form
 272

$$273 \quad r_t(a) = e_t(a)^\top \theta^* + \varepsilon_t(a), \quad \|\theta^*\|_2 \leq S,$$

274 and the noise $\varepsilon_t(a)$ are independent mean-zero subgaussian variables with parameter R .
 275

276 Denote the conditional mean of the reward $r_t(a)$ as $\mu_t(a)$.
 277

278 **Assumption 2** (Random LLM scores: biased subgaussian). For each round t and arm $a \in \mathcal{A}_t$, the
 279 LLM score $s_t(a)$ is revealed before choosing a_t , and conditioned on the history \mathcal{F}_{t-1} and the current
 280 context/candidate set,

$$280 \quad s_t(a) = \mu_t(a) + b_t(a) + \xi_t(a),$$

281 where the bias satisfies $|b_t(a)| \leq b$ for some $b \geq 0$, and the zero-mean noise $\xi_t(a)$ is σ_s -subgaussian
 282 conditionally on \mathcal{F}_{t-1} (i.e., $\mathbb{E}[e^{\lambda \xi_t(a)} | \mathcal{F}_{t-1}] \leq \exp(\frac{\lambda^2 \sigma_s^2}{2})$ for all $\lambda \in \mathbb{R}$). The sizes $A_t := |\mathcal{A}_t|$ are
 283 finite; define $A_{\max} := \max_{1 \leq t \leq T} A_t$.
 284

285 **Theorem 1** (Regret bound for Algorithm 2). Fix $\alpha \in (0, 1)$. Define
 286

$$287 \quad \Delta_\alpha := b + \sigma_s \sqrt{2 \log\left(\frac{A_{\max} T}{\alpha}\right)}.$$

288 Under Assumptions 1 and 2, define
 289

$$290 \quad \Psi := d \log\left(1 + \frac{(T - T_0)L^2}{d \lambda_{\min}(V_{T_0})}\right), \quad B_T := R \sqrt{2(d \log(1 + \frac{(T - T_0)L^2}{d \lambda_{\min}(V_{T_0})}) + \log \frac{1}{\delta_0})} + \sqrt{\lambda} S.$$

293 For any $\delta_0 \in (0, 1)$, with probability at least $1 - \delta_0 - \alpha$,

$$294 \quad R_T \leq \underbrace{2\Delta_\alpha T_0}_{\text{exploration}} + \underbrace{4 \min\left\{\Delta_\alpha(T - T_0), B_T \sqrt{2(T - T_0)\Psi}\right\}}_{\text{post exploration}}. \quad (6)$$

298 The proofs of the theorem are provided in Appendix C. Theorem 1 implies the following messages.
 299 First, the exploration phase induces a regret of $2\Delta_\alpha T_0$, which depends on the accuracy of the LLM
 300 measured by Δ_α . It depends on two aspects: the inherent bias b and the randomness of the LLM
 301 prediction σ_s . σ_s can usually be reduced by using a lower temperature. b depends on many factors,
 302 such as prompt, model, etc. If the LLM is predictive of the rewards, then Δ_α is small, and the
 303 exploration phase is leading to smaller error. Second, the post-exploration phase induces a regret of
 304 the minimum of two terms. The first term is the potential benefits of using LLM-based predictions,
 305 similar to the exploration phase. The second part has the form of the classic linear bandit regret
 306 bound. Nevertheless, it is better than the classic linear bandit regret bound because the factor B_T
 307 and Ψ are both smaller than the those in the classic linear regret bound, because the denominator
 308 term $\lambda_{\min}(V_{T_0})$ is larger than the denominator term λ in the classic linear regret bound. Overall, if
 309 $\Delta_\alpha(T - T_0) \leq B_T \sqrt{2(T - T_0)\Psi}$, or, equivalently, $T - T_0 \leq B_T^2 / \Delta_\alpha^2 \Psi$, then the post-exploration
 310 phase is leading to smaller errors. This says that, while the long-term rate of the regret of Algorithm
 311 2 coincides with the classic linear bandit regret bound, the short-term rate is better than the classic
 312 linear bandit regret bound. Moreover, if the LLM is powerful in predicting the rewards with delicate
 313 fine-tuning and prompt engineering, we can benefit a lot from a small Δ_α . Third, the quality of
 314 the embedding-based prediction matters a lot, which is captured by the parameter R . When the
 315 embedding has high quality for encoding the semantic information of the task, then R is small, and
 316 the embedding-based prediction is more accurate with a smaller B_T . Otherwise, the embedding-based
 317 prediction is less accurate with a larger B_T .

5 EXPERIMENTS

320 In this section, we aim to answer these questions:
 321

- 322 RQ1: Does the integrated LLM–MAB agent consistently outperform both classical bandit
 323 methods (e.g., LinUCBs) and pure LLM-based strategies (e.g., CoRRAL) across diverse
 324 recommendation benchmarks?

324

- 325 RQ2: How do different settings of the embedding-based uncertainty weight β_{LLM} and LLM
- 326 sampling temperature affect the exploration–exploitation trade-off and final cumulative
- 327 reward?
- 328 RQ3: How many expensive LLM queries can be avoided by delegating decision rounds to a
- 329 lightweight MAB component, while still preserving or improving cumulative reward?

330

331 **5.1 EXPERIMENTAL SETUPS**

332 **5.1.1 DATASETS.**

334 The experiments are conducted using two publicly available datasets, LastFM (Cantador et al., 2011)
 335 and MovieLens-1M (Harper & Konstan, 2015), and to evaluate the proposed method in diverse settings.
 336 Following Chen et al. (2024c), we choose the LastFM and MovieLens-1M datasets. LastFM (Cantador
 337 et al., 2011): A dataset collected from a radio listening application, containing user play history
 338 and metadata of songs. This dataset is particularly useful for evaluating recommendation systems.
 339 MovieLens-1M (Harper & Konstan, 2015): A dataset derived from the MovieLens project, which
 340 contains 1 million movie ratings from 6,040 users on 3,952 movies. It serves as a benchmark for
 341 collaborative filtering and recommendation algorithms. We only use them for research purposes.
 342

343 **5.1.2 BASELINES**

344 To provide a comprehensive comparison, several state-of-the-art baselines are selected:

345

- 346 Single LinUCB (Li et al., 2010b): A linear upper confidence bound algorithm that provides
 347 a simple baseline for the bandit setting.
- 348 LinUCB (Li et al., 2010b): The classic linear UCB algorithm with improved performance
 349 compared to the single version.
- 350 Linear Bandit (AdamLinear) (Foster & Rakhlin, 2020): An adaptive gradient method applied
 351 in the context of linear bandits, known for its robustness and efficiency.
- 352 LLM baseline (LLama 3.2 3B (Grattafiori et al., 2024)): A generative pre-trained model
 353 used as a basic comparison point for natural language processing tasks integrated into
 354 decision-making agents.
- 355 Uncertainty-based LLM (Tanneru et al., 2024) (varying temperature hyperparameters): The
 356 same GPT baseline but with different temperature hyperparameters to explore the impact of
 357 uncertainty in decision-making.
- 358 CoRRAL (Chen et al., 2024a): A recent algorithm that leverages adaptive mechanisms for
 359 efficient exploration and exploitation, serving as a strong comparative baseline.

360

361 **6 IMPLEMENTATION DETAILS**

362 For conducting experiments, we use a single NVIDIA A100 GPU equipped with 80GB of GPU
 363 memory. For MovieLens dataset, we set the total rounds $T = 6,500$, exploration phase $\sigma T = 1,000$,
 364 regularization $\lambda = 1.0$, and LLM-uncertainty weight $\beta_{LLM} = 0.8$ (selected via validation). All
 365 methods are run for 5 seeds, and results are averaged. For LastFM dataset, we set the total rounds
 366 $T = 2,500$, exploration phase $\sigma T = 1,000$, regularization $\lambda = 1.0$, and LLM-uncertainty weight
 367 $\beta_{LLM} = 0.8$ (selected via validation). All methods are run for 10 seeds, and results are averaged. For
 368 prompt-based reward extraction, we employ the Llama 3.2-3B-Instruct model to compute
 369 a reward score for each sample, where each sample consists of a context-item pair. The details
 370 of prompt design refer to Section 6.1. For embedding-based uncertainty quantification, we use
 371 logistic regression configured with L2 penalty, regularization strength $\lambda = 1.0$, the saga solver, and
 372 a maximum of 100 iterations. All input embeddings are first standardized to zero mean and unit
 373 variance via StandardScaler. For fair comparison, in all our experiments, we select LLM backbones
 374 from the Llama 3.2-3B-Instruct model (Grattafiori et al., 2024).

375 **6.1 PROMPT DESIGN**

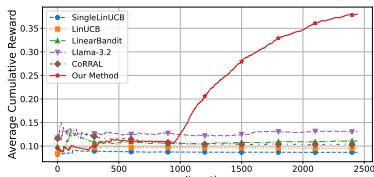
376 For the MovieLens dataset:

378 • **System message:** You are a movie recommender system specialized on the MovieLens
 379 dataset.
 380 • **User message:**
 381 User watching history: {query}
 382 Candidate movies: {candidates}
 383 Recommend exactly one movie by outputting only its name (no explanations).

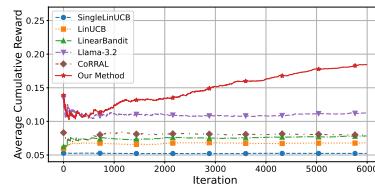
384 For the LastFM dataset:

385 • **System message:** You are a music recommender system specialized on the Last.FM dataset.
 386 • **User message:**
 387 User listening history: {query}
 388 Candidate tracks: {candidates}
 389 Recommend exactly one track by outputting only its name (no explanations).

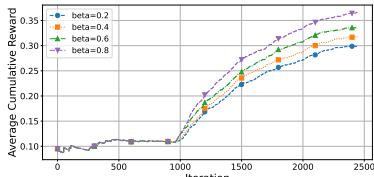
390 **Evaluation metrics.** To comprehensively evaluate the performance of our proposed model, we
 391 measure its effectiveness using the *Average Cumulative Rewards*: The mean total reward accumulated
 392 per episode (or round) across all runs, reflecting both the quality of individual decisions and the
 393 long-term performance stability.



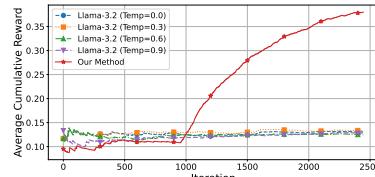
404 Figure 2: LastFM Cumulative Reward.



405 Figure 3: MovieLens Cumulative Reward.



406 Figure 4: LastFM β_{LLM} Ablation.



407 Figure 5: LastFM Uncertainty.

408 Table 1: Average Cumulative Reward on the ML-1M and LastFM benchmarks. The symbol \uparrow
 409 indicates that higher values are better.

Method	Average Cumulative Reward (\uparrow)	
	LastFM	ML-1M
SingleLinUCB	0.0913	0.0528
LinUCB	0.0993	0.0683
LinearBandit	0.1269	0.0783
Llama 3.2	0.1386	0.1350
CoRRAL	0.1511	0.0843
<i>Our Method</i>	0.3801 (+151.6%)	0.1846 (+36.7%)

410 As shown in Table 1, our method dramatically outperforms all baselines on the LastFM and MovieLens-1M benchmarks. Also, Figure 2 and Figure 3 illustrate that *our method* consistently achieves the highest average cumulative reward after about 1000 rounds on both LastFM and MovieLens-1M from the per-iteration cumulative reward curves, outperforming classical bandit baselines as well as LLM-based strategies. The first 1000 rounds can be viewed as a warm-up phase, during which the offline-trained model adapts to the online environment using observed feedback. Notably, our method not only accumulates rewards more rapidly but also maintains greater stability throughout the exploitation phase after this initial period. This is attributed to the early elimination of suboptimal actions and the effective integration of LLM-derived prompt rewards with embedding-based uncertainty estimates.

432 Together, these components enable informed exploration in the early stages and efficient exploitation
 433 thereafter, resulting in robust and consistent performance across both recommendation tasks.
 434

435 436 6.2 UNCERTAINTY COMPARISON (RQ2)

437 All experiments in this section follow the online integration procedure from Algorithm 2. We
 438 focus on two aspects of uncertainty modeling: the weight of the embedding-based uncertainty term
 439 β_{LLM} , and the effect of LLM sampling temperature when used standalone as a decision agent
 440 compared to our method. Figure 3 plots the average cumulative reward on LastFM for different
 441 values of β_{LLM} . When β_{LLM} is very small, the agent under-explores, relying almost entirely on
 442 the prompt-derived reward prior. As β_{LLM} increases, the learned uncertainty estimator kicks in,
 443 guiding more balanced exploration and yielding substantial gains. Beyond the optimal regime,
 444 further increasing β_{LLM} yields diminishing returns, as over-emphasis on uncertainty can lead to
 445 excessive exploration. This ablation confirms that a properly calibrated uncertainty weight is crucial
 446 for robust performance. Figure 4 compares *our method* against the LLaMA 3.2 3B baseline run
 447 with four different temperature settings ($\tau \in \{0, 0.3, 0.6, 0.9\}$). While temperature tuning slightly
 448 affects the baseline’s exploration behavior, all LLaMA variants plateau early and fail to match the
 449 upward trajectory of *our method*. In contrast, by explicitly combining prompt rewards with an
 450 embedding-based uncertainty signal, our approach continues to drive accurate exploration throughout
 451 the decision rounds, achieving increasingly higher cumulative rewards afterwards. This demonstrates
 452 that static temperature adjustments alone cannot substitute for a principled uncertainty estimator in
 453 bandit decision-making.

454 455 6.3 COMPUTATIONAL EFFICIENCY (RQ3)

456 457 458 Table 2: LLM and MAB call statistics on MovieLens-1M and LastFM. *CP* means call percentage.

459 Statistic	460 CoRRAL		461 <i>Our Method</i>	
	462 <i>Number</i>	463 <i>CP(%)</i>	464 <i>Number</i>	465 <i>CP(%)</i>
466 ML-1M LLM calls	467 3000	468 49.67	469 1000	470 16.56
471 LastFM LLM calls	472 1500	473 60.12	474 1000	475 40.08

476 Table 2 demonstrates that, unlike CoRRAL, which splits decision rounds roughly equally between
 477 LLM and MAB, *our method* confines LLM calls to an initial exploration phase and delegates the
 478 vast majority of subsequent decisions to the lightweight MAB component. By reducing LLM usage
 479 by over 20% compared to CoRRAL, we cut API latency and cost substantially, while still achieving
 480 superior cumulative reward.

481 482 7 CONCLUSION

483 In this paper, we introduce a lightweight framework that fuses prompt-derived priors from an LLM
 484 with an embedding-based uncertainty estimator in a contextual bandit. On LastFM and ML-1M, our
 485 method achieves state-of-the-art cumulative reward, cuts LLM queries by over 20%, and maintains a
 486 strong exploration-exploitation balance.

487 488 489 LIMITATIONS

490 Our approach assumes access to an open-source LLM that exposes internals such as embeddings. In
 491 cases where only closed-source or black-box APIs are available, obtaining these representations may
 492 not be straightforward. Adapting our method to work with limited API access, for example, using
 493 proxy embeddings or lightweight adapter modules, could introduce additional engineering overhead
 494 but is unlikely to affect the core algorithmic insights.

486
487

ETHICAL CONSIDERATIONS

488
489
490

By leveraging a large pretrained language model, our method inherits potential biases present in its training data, which could lead to unfair or harmful recommendations.

491
492

REFERENCES

493
494
495

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *NeurIPS*, 2011a.

496
497
498
499
500
501

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011b.

502

Yasin Abbasi-Yadkori, Ilja Kuzborskij, András György, and Csaba Szepesvári. To believe or not to believe your llm: Iterative prompting for estimating epistemic uncertainty. *Advances in Neural Information Processing Systems*, 37:58077–58117, 2024.

503
504

Alekh Agarwal, Ashwinkumar Badanidiyuru, John Langford, Aleksandrs Slivkins Chourasia, and Thorsten Joachims. Contextual bandits with knapsacks. In *COLT*, 2014a.

505
506

Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert E. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *ICML*, 2014b.

507
508
509

Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *ICML*, 2013.

510
511

Parand A Alamdari, Yanshuai Cao, and Kevin H Wilson. Jump starting bandits with llm-generated prior knowledge. *arXiv preprint arXiv:2406.19317*, 2024.

512
513
514

Omar Besbes, Yonatan Gur, and Assaf Zeevi. Stochastic multi-armed bandit problem with non-stationary rewards. In *NeurIPS*, 2014.

515
516
517

Conrad Borchers and Tianze Shou. Can large language models match tutoring system adaptivity? a benchmarking study. *arXiv preprint arXiv:2504.05570*, 2025.

518
519

Djallel Bouneffouf and Raphael Feraud. Multi-armed bandits meet large language models. *arXiv preprint arXiv:2505.13355*, 2025.

520
521
522
523

Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. Second workshop on information heterogeneity and fusion in recommender systems (hetrec2011). In *Proceedings of the fifth ACM conference on Recommender systems*, pp. 387–388, 2011.

524
525

Dingyang Chen, Qi Zhang, and Yinglun Zhu. Efficient sequential decision making with large language models. *arXiv preprint arXiv:2406.12125*, 2024a.

526
527
528
529

Yulin Chen, Ning Ding, Hai-Tao Zheng, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. Empowering private tutoring by chaining large language models. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 354–364, 2024b.

530
531
532

Yuxin Chen, Junfei Tan, An Zhang, Zhengyi Yang, Leheng Sheng, Enzhi Zhang, Xiang Wang, and Tat-Seng Chua. On softmax direct preference optimization for recommendation. *arXiv preprint arXiv:2406.09215*, 2024c.

533
534
535

Haikang Deng and Colin Raffel. Reward-augmented decoding: Efficient controlled text generation with a unidirectional reward model. *arXiv preprint arXiv:2310.09520*, 2023.

536
537
538

Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *NeurIPS*, 2010.

539

Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International conference on machine learning*, pp. 3199–3210. PMLR, 2020.

540 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 541 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
 542 models. *arXiv preprint arXiv:2407.21783*, 2024.

543 F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm
 544 transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.

545 Chengkai Huang, Tong Yu, Kaige Xie, Shuai Zhang, Lina Yao, and Julian McAuley. Foundation mod-
 546 els for recommender systems: A survey and new perspectives. *arXiv preprint arXiv:2402.11143*,
 547 2024a.

548 Hsiao-Yuan Huang, Yutong Yang, Zhaoxi Zhang, Sanwoo Lee, and Yunfang Wu. A survey of
 549 uncertainty estimation in llms: Theory meets practice. *arXiv preprint arXiv:2410.15326*, 2024b.

550 Ranganath Krishnan, Piyush Khanna, and Omesh Tickoo. Enhancing trust in large language models
 551 with uncertainty-aware fine-tuning. *arXiv preprint arXiv:2412.02904*, 2024.

552 John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side
 553 information. In *NeurIPS*, 2008.

554 Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to
 555 personalized news article recommendation. In *WWW*, 2010a.

556 Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to
 557 personalized news article recommendation. In *Proceedings of the 19th international conference on
 558 World wide web*, pp. 661–670, 2010b.

559 Huan Ma, Jingdong Chen, Guangyu Wang, and Changqing Zhang. Estimating llm uncertainty with
 560 logits. *arXiv preprint arXiv:2502.00290*, 2025.

561 Giovanni Monea, Antoine Bosselut, Kianté Brantley, and Yoav Artzi. Llms are in-context reinforce-
 562 ment learners. 2024.

563 Carlos Riquelme, George Tucker, and Jasper Snoek. Deep bayesian bandits showdown: An empirical
 564 comparison of thompson sampling and UCB with deep neural networks. In *ICLR*, 2018.

565 Haochen Song, Dominik Hofer, Rania Islambouli, Laura Hawkins, Ananya Bhattacharjee, Meredith
 566 Franklin, and Joseph Jay Williams. Investigating the relationship between physical activity and
 567 tailored behavior change messaging: Connecting contextual bandit with large language models.
 568 *arXiv preprint arXiv:2506.07275*, 2025.

569 Sree Harsha Tanneru, Chirag Agarwal, and Himabindu Lakkaraju. Quantifying uncertainty in
 570 natural language explanations of large language models. In *International Conference on Artificial
 571 Intelligence and Statistics*, pp. 1072–1080. PMLR, 2024.

572 Michal Valko, Rémi Munos, and Hilbert Kappen. Finite-time analysis of kernelised contextual
 573 bandits. In *UAI*, 2013.

574 Fanzeng Xia, Hao Liu, Yisong Yue, and Tongxin Li. Beyond numeric rewards: In-context dueling
 575 bandits with llm agents. In *Findings of ACL*, 2024.

576 Zikun Ye, Hema Yoganarasimhan, and Yufeng Zheng. Lola: Llm-assisted online learning algorithm
 577 for content experiments. *Marketing Science*, 2025.

578 Hongxin Zhang, Weihua Du, Jiaming Shan, Qinrong Zhou, Yilun Du, Joshua B Tenenbaum, Tianmin
 579 Shu, and Chuang Gan. Building cooperative embodied agents modularly with large language
 580 models. *arXiv preprint arXiv:2307.02485*, 2023.

581 Zhengyuan Zhou, Lihong Li, Jian Li, and Quanquan Wang. Neural contextual bandits with ucb-based
 582 exploration. In *ICML*, 2020.

583

584

585

586

587

588

589

590

591

592

593

594 **A CASE STUDY**
595

596 In this case study, we compare three recommendation strategies for a 25-year-old male user whose
 597 watch history includes over twenty films such as Son in Law, Beetlejuice, Aliens, Escape from L.A.,
 598 and Lord of the Flies. Presented with the same set of twenty candidate movies, the pure LLaMA 3.2
 599 model mistakenly selects Bloody Child, and the classic LinUCB algorithm chooses Thinner, whereas
 600 our LLM-calibrated LinUCB approach correctly ranks Star Wars: Episode VI – Return of the Jedi at
 601 the top. This choice is accurate for two complementary reasons. First, Star Wars: Episode VI shares
 602 core sci-fi and adventure elements with several of the user’s past favorites (Aliens and Escape from
 603 L.A.), so the LLM’s semantic scoring naturally favors it. Second, during the online exploration phase,
 604 our method’s UCB adjustment identifies that Star Wars consistently yields higher real-time reward
 605 feedback than other high-scoring candidates, confirming its true preference alignment. In contrast, a
 606 standalone LLM can only rely on static semantic similarity (and so can be misled by less relevant
 607 titles), and a standalone bandit method lacks that rich prior, leading to suboptimal initial choices. By
 608 combining LLM priors with principled exploration, our hybrid method both recognizes the user’s
 609 genre tastes and rigorously validates them through interaction.

600 **Comparison at LLama 3.2, LinUCB, and our method**601 **Query:**

602 This user’s age is 25, gender is Male; History:

603 Son in Law | Beetlejuice | Aliens | Escape from L.A.|... | Cutting Edge, The | Young Poisoner’s
 604 Handbook, The | Lord of the Flies

605 **Candidates:**

- 606 • Full Speed
- 607 • My Boyfriend’s Back
- 608 • Shooting Fish
- 609 • Mr. Jealousy
- 610 • Gold Diggers: The Secret of Bear Mountain
- 611 • Star Wars: Episode VI – Return of the Jedi
- 612 • Little Odessa
- 613 • Twin Peaks: Fire Walk with Me
- 614 • Second Best
- 615 • Wisdom
- 616 • Dazed and Confused
- 617 • American Werewolf in London, An
- 618 • Ulee’s Gold
- 619 • Children of a Lesser God
- 620 • D3: The Mighty Ducks
- 621 • The Siege
- 622 • Out of Sight
- 623 • Cutting Edge, The
- 624 • The Young Poisoner’s Handbook
- 625 • Lord of the Flies

626 **Ground Truth:**

627 Star Wars: Episode VI – Return of the Jedi

628 **Choices:**

- 629 1) Llama 3.2 output: Bloody Child
- 630 2) LinUCB output: Thinner
- 631 3) Our Method output: **Star Wars: Episode VI – Return of the Jedi**

644 **B THE USE OF LARGE LANGUAGE MODELS (LLMs)**
645

646 In preparing this manuscript, we leveraged large language models (LLMs) as writing assistants
 647 to improve readability and clarity. Specifically, LLMs were employed to polish sentences and

648 adjust phrasing to align with the conventions of academic writing. Importantly, the scientific ideas,
649 methodological details, and results presented in this paper are entirely original and were conceived,
650 designed, and validated by the authors. The use of LLMs was limited to language refinement rather
651 than content generation, ensuring that all technical contributions remain the authors' own.
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 **C THEORETICAL ANALYSIS**
 703

704 **Lemma 1** (Uniform score envelope). *Fix $\alpha \in (0, 1)$. Define*

706
$$\Delta_\alpha := b + \sigma_s \sqrt{2 \log\left(\frac{A_{\max} T}{\alpha}\right)}.$$

 707

708 *Then, with probability at least $1 - \alpha$,*

710
$$\max_{1 \leq t \leq T} \max_{a \in \mathcal{A}_t} |s_t(a) - \mu_t(a)| \leq \Delta_\alpha.$$

 711

712 *Proof of Lemma 1.* For any fixed (t, a) , conditional subgaussianity and Hoeffding-type tails give
 713 $\Pr(|\xi_t(a)| > \tau \mid \mathcal{F}_{t-1}) \leq 2e^{-\tau^2/(2\sigma_s^2)}$. Adding bias $b_t(a)$ with $|b_t(a)| \leq b$ yields
 714 $\Pr(|s_t(a) - \mu_t(a)| > b + \tau \mid \mathcal{F}_{t-1}) \leq 2e^{-\tau^2/(2\sigma_s^2)}$. Union bound over all (t, a) and set
 715 $\tau = \sigma_s \sqrt{2 \log(A_{\max} T / \alpha)}$ to obtain the claim. \blacksquare
 716

717 Lemma 1 quantifies the accuracy of the LLM prediction. It depends on two aspects: the inherent
 718 bias b and the randomness of the LLM prediction σ_s . σ_s can usually be reduced by using lower
 719 temperature. b depends on many factors, such as prompt, model, etc.

720 *Proof of Theorem 1. Exploration Phase.* Conditioned on the event \mathcal{E}_α (Lemma 1), at any $t \leq T_0$,
 721 since $a_t = \arg \max s_t(\cdot)$ and $|s_t - \mu_t| \leq \Delta_\alpha$,

722
$$\mu_t(a_t^*) - \mu_t(a_t) \leq s_t(a_t^*) + \Delta_\alpha - (s_t(a_t) - \Delta_\alpha) \leq 2\Delta_\alpha.$$

 723

724 Summing yields the term $2\Delta_\alpha T_0$.

725 **Post-exploration Phase.** We will use a similar argument as Abbasi-Yadkori et al. (2011b). Let

726
$$\beta_t = R \sqrt{2 \log \frac{\det(V_{t-1})^{1/2}}{\det(V_{T_0})^{1/2}} \frac{1}{\delta_0}} + \sqrt{\lambda} S,$$

 727

728 and

729
$$w_t(a) = \|e_t(a)\|_{V_{t-1}^{-1}}.$$

 730

731 Abbasi-Yadkori et al. (2011b) proved that

732
$$\beta_t \leq B_T = R \sqrt{2 \left(d \log\left(1 + \frac{(T-T_0)L^2}{d\lambda_{\min}(V_{T_0})}\right) + \log \frac{1}{\delta_0} \right)} + \sqrt{\lambda} S.$$

 733

734 By the self-normalized inequality with prior V_{T_0} (Abbasi-Yadkori et al., 2011b), with probability \geq
 735 $1 - \delta_0$, for all $t > T_0$ and all a , $|e_t(a)^\top (\hat{\theta}_{t-1} - \theta^*)| \leq \beta_t w_t(a)$, hence $\mu_t(a) \in [\text{LCB}_t(a), \text{UCB}_t(a)]$.
 736 By the selection rule, if a_t is selected, $r_t(a_t) \geq r_t(a_t^*)$.

737 On one hand, if $\Delta_\alpha \geq \beta_t w_t(a_t)$, we have

738
$$\mu_t(a_t^*) - \mu_t(a_t) \leq r_t(a_t^*) - r_t(a_t) + 4\beta_t w_t(a_t) \leq 4\beta_t w_t(a_t) \quad \text{for all } t > T_0. \quad (7)$$

 739

740 On the other hand, if the LLM prediction accuracy Δ_α is smaller than the bound $\beta_t w_t(a_t)$, we can
 741 prove that

742
$$|r_t(a_t) - \mu_t(a_t)| = |r_t(a_t) - s_t(a_t) + s_t(a_t) - \mu_t(a_t)| \quad (8)$$

 743

744 If $r_t(a_t) = s_t(a_t)$, then $|r_t(a_t) - \mu_t(a_t)| \leq \Delta_\alpha$ by Lemma 1. If $r_t(a_t) < s_t(a_t)$, then $0 < s_t(a_t) - r_t(a_t) \leq \mu_t(a_t) + \Delta_\alpha - r_t(a_t) \leq \Delta_\alpha$, hence $|r_t(a_t) - \mu_t(a_t)| \leq \Delta_\alpha$. Similarly, when $r_t(a_t) > s_t(a_t)$, we can prove the same bound. Therefore, in this case, we always have $|r_t(a_t) - \mu_t(a_t)| \leq \Delta_\alpha$.
 745 By the selection rule of Algorithm 2, we have

746
$$r_t(a_t) \geq \text{clip}(s_t(a_t^*), \text{LCB}_t(a_t^*), \text{UCB}_t(a_t^*)).$$

 747

- If $\text{clip}(s_t(a_t^*), \text{LCB}_t(a_t^*), \text{UCB}_t(a_t^*)) = s_t(a_t^*)$, then $r_t(a_t) \geq s_t(a_t^*)$.
- If $\text{clip}(s_t(a_t^*), \text{LCB}_t(a_t^*), \text{UCB}_t(a_t^*)) = \text{LCB}_t(a_t^*)$, then $r_t(a_t) \geq \text{LCB}_t(a_t^*) \geq s_t(a_t^*)$.
- If $\text{clip}(s_t(a_t^*), \text{LCB}_t(a_t^*), \text{UCB}_t(a_t^*)) = \text{UCB}_t(a_t^*)$, then $r_t(a_t) \geq \text{UCB}_t(a_t^*) \geq \mu_t(a_t^*) \geq s_t(a_t^*) - \Delta_\alpha$.

To conclude, if a_t is selected and Δ_α is smaller than the bound $\beta_t w_t(a_t)$, then we have

$$s_t(a_t^*) - r_t(a_t) \leq \Delta_\alpha, \quad \mu_t(a_t^*) - \mu_t(a_t) \leq 3\Delta_\alpha.$$

Therefore, the per-round regret is bounded by

$$\mu_t(a_t^*) - \mu_t(a_t) \leq 4 \min(\Delta_\alpha, \beta_t w_t(a_t)).$$

The above analysis gives

$$\begin{aligned} \sum_{t=T_0+1}^T (\mu_t(a_t^*) - \mu_t(a_t)) &\leq \sum_{t=T_0+1}^T 4 \min(\Delta_\alpha, \beta_t w_t(a_t)) \\ &\leq \min \left\{ \sum_{t=T_0+1}^T 4\Delta_\alpha, \sum_{t=T_0+1}^T 4\beta_t w_t(a_t) \right\}. \end{aligned}$$

The first summation in the minimum is simply $4\Delta_\alpha(T - T_0)$. The second summation in the minimum is bounded by $4B_T \sqrt{(T - T_0)\Psi}$, following the classic argument in Abbasi-Yadkori et al. (2011b). Therefore, the per-round regret is bounded by

$$\begin{aligned} \sum_{t=T_0+1}^T (\mu_t(a_t^*) - \mu_t(a_t)) &\leq \min \left\{ \sum_{t=T_0+1}^T 4\Delta_\alpha, \sum_{t=T_0+1}^T 4\beta_t w_t(a_t) \right\} \\ &\leq 4 \min \left\{ \Delta_\alpha(T - T_0), B_T \sqrt{2(T - T_0)\Psi} \right\}. \end{aligned}$$

Taking the minimum of the two post-exploration controls completes equation 6. ■