
Scalable Generalized Linear Bandits: Online Computation and Hashing

Kwang-Sung Jun

UW-Madison

kjun@discovery.wisc.edu

Aniruddha Bhargava

UW-Madison

aniruddha@wisc.edu

Robert Nowak

UW-Madison

rdnowak@wisc.edu

Rebecca Willett

UW-Madison

willett@discovery.wisc.edu

Abstract

Generalized Linear Bandits (GLBs), a natural extension of the stochastic linear bandits, has been popular and successful in recent years. However, existing GLBs scale poorly with the number of rounds and the number of arms, limiting their utility in practice. This paper proposes new, scalable solutions to the GLB problem in two respects. First, unlike existing GLBs, whose per-time-step space and time complexity grow at least linearly with time t , we propose a new algorithm that performs online computations to enjoy a constant space and time complexity. At its heart is a novel Generalized Linear extension of the Online-to-confidence-set Conversion (GLOC method) that takes *any* online learning algorithm and turns it into a GLB algorithm. As a special case, we apply GLOC to the online Newton step algorithm, which results in a low-regret GLB algorithm with much lower time and memory complexity than prior work. Second, for the case where the number N of arms is very large, we propose new algorithms in which each next arm is selected via an inner product search. Such methods can be implemented via hashing algorithms (i.e., “hash-amenable”) and result in a time complexity sublinear in N . While a Thompson sampling extension of GLOC is hash-amenable, its regret bound for d -dimensional arm sets scales with $d^{3/2}$, whereas GLOC’s regret bound scales with d . Towards closing this gap, we propose a new hash-amenable algorithm whose regret bound scales with $d^{5/4}$. Finally, we propose a fast approximate hash-key computation (inner product) with a better accuracy than the state-of-the-art, which can be of independent interest. We conclude the paper with preliminary experimental results confirming the merits of our methods.

1 Introduction

This paper considers the problem of making generalized linear bandits (GLBs) scalable. In the stochastic GLB problem, a learner makes successive decisions to maximize her cumulative rewards. Specifically, at time t the learner observes a set of arms $\mathcal{X}_t \subseteq \mathbb{R}^d$. The learner then chooses an arm $\mathbf{x}_t \in \mathcal{X}_t$ and receives a stochastic reward y_t that is a noisy function of \mathbf{x}_t : $y_t = \mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*) + \eta_t$, where $\boldsymbol{\theta}^* \in \mathbb{R}^d$ is unknown, $\mu: \mathbb{R} \rightarrow \mathbb{R}$ is a known nonlinear mapping, and $\eta_t \in \mathbb{R}$ is some zero-mean noise. This reward structure encompasses generalized linear models [29]; e.g., Bernoulli, Poisson, etc.

The key aspect of the bandit problem is that the learner does not know how much reward she would have received, had she chosen another arm. The estimation on $\boldsymbol{\theta}^*$ is thus biased by the history of the selected arms, and one needs to mix in exploratory arm selections to avoid ruling out the optimal arm. This is well-known as the exploration-exploitation dilemma. The performance of a learner is evaluated by its *regret* that measures how much cumulative reward she would have gained additionally if she had known the true $\boldsymbol{\theta}^*$. We provide backgrounds and formal definitions in Section 2.

A linear case of the problem above ($\mu(z) = z$) is called the (stochastic) linear bandit problem. Since the first formulation of the linear bandits [7], there has been a flurry of studies on the problem [11,

34, 1, 9, 5]. In an effort to generalize the restrictive linear rewards, Filippi et al. [15] propose the GLB problem and provide a low-regret algorithm, whose Thompson sampling version appears later in Abeille & Lazaric [3]. Li et al. [27] evaluates GLBs via extensive experiments where GLBs exhibit lower regrets than linear bandits for 0/1 rewards. Li et al. [28] achieves a smaller regret bound when the arm set \mathcal{X}_t is finite, though with an impractical algorithm.

However, we claim that all existing GLB algorithms [15, 28] suffer from two scalability issues that limit their practical use: (i) under a large time horizon and (ii) under a large number N of arms.

First, existing GLBs require storing all the arms and rewards appeared so far, $\{(\mathbf{x}_s, y_s)\}_{s=1}^t$, so the space complexity grows linearly with t . Furthermore, they have to solve a batch optimization problem for the maximum likelihood estimation (MLE) at each time step t whose per-time-step time complexity grows at least linearly with t . While Zhang et al. [41] provide a solution whose space and time complexity do not grow over time, they consider a specific 0/1 reward with the logistic link function, and a generic solution for GLBs is not provided.

Second, existing GLBs have linear time complexities in N . This is impractical when N is very large, which is not uncommon in applications of GLBs such as online advertisements, recommendation systems, and interactive retrieval of images or documents [26, 27, 40, 21, 25] where arms are items in a very large database. Furthermore, the interactive nature of these systems requires prompt responses as users do not want to wait. This implies that the typical linear time in N is not tenable. Towards a *sublinear* time in N , locality sensitive hashings [18] or its extensions [35, 36, 30] are good candidates as they have been successful in fast similarity search and other machine learning problems like active learning [22], where the search time scales with N^ρ for some $\rho < 1$ (ρ is usually optimized and often ranges from 0.4 to 0.8 depending on the target search accuracy). Leveraging hashing in GLBs, however, relies critically on the objective function used for arm selections. The function must take a form that is readily optimized using *existing* hashing algorithms.¹ For example, algorithms whose objective function (a function of each arm $\mathbf{x} \in \mathcal{X}_t$) can be written as a distance or inner product between \mathbf{x} and a query \mathbf{q} are hash-amenable as there *exist* hashing methods for such functions.

To be scalable to a large time horizon, we propose a new algorithmic framework called Generalized Linear Online-to-confidence-set Conversion (GLOC) that takes in an online learning (OL) algorithm with a low ‘OL’ regret bound and turns it into a GLB algorithm with a low ‘GLB’ regret bound. The key tool is a novel generalization of the online-to-confidence-set conversion technique used in [2] (also similar to [14, 10, 16, 41]). This allows us to construct a confidence set for θ^* , which is then used to choose an arm \mathbf{x}_t according to the well-known optimism in the face of uncertainty principle. By relying on an online learner, GLOC inherently performs online computations and is thus free from the scalability issues in large time steps. While any online learner equipped with a low OL regret bound can be used, we choose the online Newton step (ONS) algorithm and prove a tight OL regret bound, which results in a practical GLB algorithm with almost the same regret bound as existing inefficient GLB algorithms. We present our proposed algorithms and their regret bounds in Section 3.

For large number N of arms, our proposed algorithm GLOC is not hash-amenable, to our knowledge, due to its nonlinear criterion for arm selection. As the first attempt, we derive a Thompson sampling [5, 3] extension of GLOC (GLOC-TS), which is hash-amenable due to its linear criterion. However, its regret bound scales with $d^{3/2}$ for d -dimensional arm sets, which is far from d of GLOC. Towards closing this gap, we propose a new algorithm Quadratic GLOC (QGLOC) with a regret bound that scales with $d^{5/4}$. We summarize the comparison of our proposed GLB algorithms in Table 1. In Section 4, we present GLOC-TS, QGLOC, and their regret bound.

Algorithm	Regret	Hash-amenable
GLOC	$\tilde{O}(d\sqrt{T})$	✗
GLOC-TS	$\tilde{O}(d^{3/2}\sqrt{T})$	✓
QGLOC	$\tilde{O}(d^{5/4}\sqrt{T})$	✓

Table 1: Comparison of GLBs algorithms for d -dimensional arm sets T is the time horizon. QGLOC achieves the smallest regret among hash-amenable algorithms.

Note that, while hashing achieves a time complexity sublinear in N , there is a nontrivial overhead of computing the projections to determine the hash keys. As an extra contribution, we reduce this overhead by proposing a new sampling-based approximate inner product method. Our proposed sampling method has smaller variance than the state-of-the-art sampling method proposed by [22, 24] when the vectors are normally distributed, which fits our setting where projection vectors are indeed normally distributed. Moreover, our method results in thinner tails in the distribution of estimation

¹ Without this designation, no *currently known* bandit algorithm achieves a sublinear time complexity in N .

error than the existing method, which implies a better concentration. We elaborate more on reducing the computational complexity of QOFUL in Section 5.

2 Preliminaries

We review relevant backgrounds here. \mathcal{A} refers to a GLB algorithm, and \mathcal{B} refers to an online learning algorithm. Let $\mathcal{B}_d(S)$ be the d -dimensional Euclidean ball of radius S , which overloads the notation \mathcal{B} . Let $\mathbf{A}_{\cdot i}$ be the i -th column vector of a matrix \mathbf{A} . Define $\|\mathbf{x}\|_{\mathbf{A}} := \sqrt{\mathbf{x}^\top \mathbf{A} \mathbf{x}}$ and $\text{vec}(\mathbf{A}) := [\mathbf{A}_{\cdot 1}; \mathbf{A}_{\cdot 2}; \dots; \mathbf{A}_{\cdot d}] \in \mathbb{R}^{d^2}$. Given a function $f : \mathbb{R} \rightarrow \mathbb{R}$, we denote by f' and f'' its first and second derivative, respectively. We define $[N] := \{1, 2, \dots, N\}$.

Generalized Linear Model (GLM) Consider modeling the reward y as one-dimensional exponential family such as Bernoulli or Poisson. When the feature vector \mathbf{x} is believed to correlate with y , one popular modeling assumption is the generalized linear model (GLM) that turns the *natural parameter* of an exponential family model into $\mathbf{x}^\top \boldsymbol{\theta}^*$ where $\boldsymbol{\theta}^*$ is a parameter [29]:

$$\mathbb{P}(y | z = \mathbf{x}^\top \boldsymbol{\theta}^*) = \exp \left(\frac{yz - m(z)}{g(\tau)} + h(y, \tau) \right), \quad (1)$$

where $\tau \in \mathbb{R}^+$ is a known scale parameter and m , g , and h are normalizers. It is known that $m'(z) = \mathbb{E}[y | z] =: \mu(z)$ and $m''(z) = \text{Var}(y | z)$. We call $\mu(z)$ the *inverse link* function. Throughout, we assume that the exponential family being used in a GLM has a *minimal representation*, which ensures that $m(z)$ is strictly convex [38, Prop. 3.1]. Then, the negative log likelihood (NLL) $\ell(z, y) := -yz + m(z)$ of a GLM is strictly convex. We refer to such GLMs as the *canonical* GLM. In the case of Bernoulli rewards $y \in \{0, 1\}$, $m(z) = \log(1 + \exp(z))$, $\mu(z) = (1 + \exp(-z))^{-1}$, and the NLL can be written as the logistic loss: $\log(1 + \exp(-y'(\mathbf{x}_t^\top \boldsymbol{\theta}^*)))$, where $y' = 2y - 1$.

Generalized Linear Bandits (GLB) Recall that \mathbf{x}_t is the arm chosen at time t by an algorithm. We assume that the arm set \mathcal{X}_t can be of an infinite cardinality, although we focus on finite arm sets in hashing part of the paper (Section 4). One can write down the reward model (1) in a different form:

$$y_t = \mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*) + \eta_t, \quad (2)$$

where η_t is conditionally R -sub-Gaussian given \mathbf{x}_t and $\{(\mathbf{x}_s, \eta_s)\}_{s=1}^{t-1}$. For example, Bernoulli reward model has η_t as $1 - \mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*)$ w.p. $\mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*)$ and $-\mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*)$ otherwise. Assume that $\|\boldsymbol{\theta}^*\|_2 \leq S$, where S is known. One can show that the sub-Gaussian scale R is determined by μ : $R = \sup_{z \in (-S, S)} \sqrt{\mu'(z)} \leq \sqrt{L}$, where L is the Lipschitz constant of μ . Throughout, we assume that each arm has ℓ_2 -norm at most 1: $\|\mathbf{x}\|_2 \leq 1, \forall \mathbf{x} \in \mathcal{X}_t, \forall t$. Let $\mathbf{x}_{t,*} := \max_{\mathbf{x} \in \mathcal{X}_t} \mathbf{x}^\top \boldsymbol{\theta}^*$. The performance of a GLB algorithm \mathcal{A} is analyzed by the expected cumulative regret (or simply *regret*): $\text{Regret}_{\mathcal{A}}^T := \sum_{t=1}^T \mu(\mathbf{x}_{t,*}^\top \boldsymbol{\theta}^*) - \mu((\mathbf{x}_t^{\mathcal{A}})^\top \boldsymbol{\theta}^*)$, where $\mathbf{x}_t^{\mathcal{A}}$ makes the dependence on \mathcal{A} explicit.

We remark that our results in this paper hold true for a strictly larger family of distributions than the canonical GLM, which we call the *non-canonical* GLM and explain below. The condition is that the reward model follows (2) where the R is now independent from μ that satisfies the following:

Assumption 1. μ is L -Lipschitz on $[-S, S]$ and continuously differentiable on $(-S, S)$. Furthermore, $\inf_{z \in (-S, S)} \mu'(z) = \kappa$ for some finite $\kappa > 0$ (thus μ is strictly increasing).

Define $\mu'(z)$ at $\pm S$ as their limits. Under Assumption 1, m is defined to be an integral of μ . Then, one can show that m is κ -strongly convex on $\mathcal{B}_1(S)$. An example of the non-canonical GLM is the probit model for 0/1 reward where μ is the Gaussian CDF, which is popular and competitive to the Bernoulli GLM as evaluated by Li et al. [27]. Note that canonical GLMs satisfy Assumption 1.

3 Generalized Linear Bandits with Online Computation

We describe and analyze a new GLB algorithm called Generalized Linear Online-to-confidence-set Conversion (GLOC) that performs online computations, unlike existing GLB algorithms.

GLOC employs the optimism in the face of uncertainty principle, which dates back to [7]. That is, we maintain a confidence set C_t (defined below) that traps the true parameter $\boldsymbol{\theta}^*$ with high probability (w.h.p.) and choose the arm with the largest feasible reward given C_{t-1} as a constraint:

$$(\mathbf{x}_t, \tilde{\boldsymbol{\theta}}_t) := \arg \max_{\mathbf{x} \in \mathcal{X}_t, \boldsymbol{\theta} \in C_{t-1}} \langle \mathbf{x}, \boldsymbol{\theta} \rangle \quad (3)$$

The main difference between GLOC and existing GLBs is in the computation of the C_t 's. Prior methods involve "batch" computations that involve all past observations, and so scale poorly with

t . In contrast, GLOC takes in an *online* learner \mathcal{B} , and uses \mathcal{B} as a co-routine instead of relying on a batch procedure to construct a confidence set. Specifically, at each time t GLOC feeds the loss function $\ell_t(\theta) := \ell(\mathbf{x}_t^\top \theta, y_t)$ into the learner \mathcal{B} which then outputs its parameter prediction θ_t . Let $\mathbf{X}_t \in \mathbb{R}^{t \times d}$ be the design matrix consisting of $\mathbf{x}_1, \dots, \mathbf{x}_t$. Define $\bar{\mathbf{V}}_t := \lambda \mathbf{I} + \mathbf{X}_t^\top \mathbf{X}_t$, where λ is the ridge parameter. Let $z_t := \mathbf{x}_t^\top \theta_t$ and $\mathbf{z}_t := [z_1; \dots; z_t]$. Let $\hat{\theta}_t := \bar{\mathbf{V}}_t^{-1} \mathbf{X}_t^\top \mathbf{z}_t$ be the ridge regression estimator taking \mathbf{z}_t as responses. Theorem 1 below is the key result for constructing our confidence set C_t , which is a function of the parameter predictions $\{\theta_s\}_{s=1}^t$ and the online (OL) regret bound B_t of the learner \mathcal{B} . All the proofs are in the supplementary material (SM).

Theorem 1. (*Generalized Linear Online-to-Confidence-Set Conversion*) Suppose we feed loss functions $\{\ell_s(\theta)\}_{s=1}^t$ into online learner \mathcal{B} . Let θ_s be the parameter predicted at time step s by \mathcal{B} . Assume that \mathcal{B} has an OL regret bound $B_t: \forall \theta \in \mathcal{B}_d(S), \forall t \geq 1$,

$$\sum_{s=1}^t \ell_s(\theta_s) - \ell_s(\theta) \leq B_t. \quad (4)$$

Let $\alpha(B_t) := 1 + \frac{4}{\kappa} B_t + \frac{8R^2}{\kappa^2} \log(\frac{2}{\delta} \sqrt{1 + \frac{2}{\kappa} B_t + \frac{4R^4}{\kappa^4 \delta^2}})$. Then, with probability (w.p.) at least $1 - \delta$,

$$\forall t \geq 1, \|\theta^* - \hat{\theta}_t\|_{\bar{\mathbf{V}}_t}^2 \leq \alpha(B_t) + \lambda S^2 - \left(\|\mathbf{z}_t\|_2^2 - \hat{\theta}_t^\top \mathbf{X}_t^\top \mathbf{z}_t \right) =: \beta_t. \quad (5)$$

Note that the center of the ellipsoid is the ridge regression estimator on the predicted natural parameters $z_s = \mathbf{x}_s^\top \theta_s$ rather than the rewards. Theorem 1 motivates the following confidence set:

$$C_t := \{\theta \in \mathbb{R}^d : \|\theta - \hat{\theta}_t\|_{\bar{\mathbf{V}}_t}^2 \leq \beta_t\} \quad (6)$$

which traps θ^* for all $t \geq 1$, w.p. at least $1 - \delta$. See Algorithm 1 for pseudocode. One way to solve the optimization problem (3) is to define the function $\theta(\mathbf{x}) := \max_{\theta \in C_{t-1}} \mathbf{x}^\top \theta$, and then use the Lagrangian method to write:

$$\mathbf{x}_t^{\text{GLOC}} := \arg \max_{\mathbf{x} \in \mathcal{X}_t} \mathbf{x}^\top \hat{\theta}_{t-1} + \sqrt{\beta_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}. \quad (7)$$

We prove the regret bound of GLOC in the following theorem.

Theorem 2. Let $\{\bar{\beta}_t\}$ be a nondecreasing sequence such that $\bar{\beta}_t \geq \beta_t$. Then, w.p. at least $1 - \delta$,

$$\text{Regret}_T^{\text{GLOC}} = O\left(L \sqrt{\bar{\beta}_T d T \log T}\right)$$

Although any low-regret online learner can be combined with GLOC, one would like to ensure that $\bar{\beta}_T$ is $O(\text{polylog}(T))$ in which case the total regret can be bounded by $\tilde{O}(\sqrt{T})$. This means that we must use online learners whose OL regret grows logarithmically in T such as [20, 31]. In this work, we consider the online Newton step (ONS) algorithm [20].

Online Newton Step (ONS) for Generalized Linear Models Note that ONS requires the loss functions to be α -exp-concave. One can show that $\ell_t(\theta)$ is α -exp-concave [20, Sec. 2.2]. Then, GLOC can use ONS and its OL regret bound to solve the GLB problem. However, motivated by the fact that the OL regret bound B_t appears in the radius $\sqrt{\beta_t}$ of the confidence set while a tighter confidence set tends to reduce the bandit regret in practice, we derive a tight data-dependent OL regret bound tailored to GLMs.

We present our version of ONS for GLMs (ONS-GLM) in Algorithm 2. $\ell'(z, y)$ is the first derivative w.r.t. z and the parameter ϵ is for inverting matrices conveniently (usually $\epsilon = 1$ or 0.1). The only difference from the original ONS [20] is that we rely on the strong convexity of $m(z)$ instead of the α -exp-concavity of the loss thanks to the GLM structure.² Theorem 3 states that we achieve the desired polylogarithmic regret in T .

Algorithm 1 GLOC

- 1: **Input:** $R > 0, \delta \in (0, 1), S > 0, \lambda > 0, \kappa > 0$, an online learner \mathcal{B} with known regret bounds $\{B_t\}_{t \geq 1}$.
 - 2: Set $\bar{\mathbf{V}}_0 = \lambda \mathbf{I}$.
 - 3: **for** $t = 1, 2, \dots$ **do**
 - 4: Compute \mathbf{x}_t by solving (3).
 - 5: Pull \mathbf{x}_t and then observe y_t .
 - 6: Receive θ_t from \mathcal{B} .
 - 7: Feed into \mathcal{B} the loss $\ell_t(\theta) = \ell(\mathbf{x}_t^\top \theta, y_t)$.
 - 8: Update $\bar{\mathbf{V}}_t = \bar{\mathbf{V}}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$ and $z_t = \mathbf{x}_t^\top \theta_t$.
 - 9: Compute $\hat{\theta}_t = \bar{\mathbf{V}}_t^{-1} \mathbf{X}_t^\top \mathbf{z}_t$ and β_t as in (5).
 - 10: Define C_t as in (6).
 - 11: **end for**
-

Algorithm 2 ONS-GLM

- 1: **Input:** $\kappa > 0, \epsilon > 0, S > 0$.
 - 2: $\mathbf{A}_0 = \epsilon \mathbf{I}$.
 - 3: Set $\theta_1 \in \mathcal{B}_d(S)$ arbitrarily.
 - 4: **for** $t = 1, 2, 3, \dots$ **do**
 - 5: Output θ_t .
 - 6: Observe \mathbf{x}_t and y_t .
 - 7: Incur loss $\ell(\mathbf{x}_t^\top \theta_t, y_t)$.
 - 8: $\mathbf{A}_t = \mathbf{A}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$
 - 9: $\theta'_{t+1} = \theta_t - \frac{\ell'(\mathbf{x}_t^\top \theta_t, y_t)}{\kappa} \mathbf{A}_t^{-1} \mathbf{x}_t$
 - 10: $\theta_{t+1} = \arg \min_{\theta \in \mathcal{B}_d(S)} \|\theta - \theta'_{t+1}\|_{\mathbf{A}_t}^2$
 - 11: **end for**
-

² A similar change to ONS has been applied in [16, 41].

Theorem 3. Define $g_s := \ell'(\mathbf{x}_s^\top \boldsymbol{\theta}_s, y_s)$. The regret of ONS-GLM satisfies, for any $\epsilon > 0$ and $t \geq 1$,

$$\sum_{s=1}^t \ell_s(\boldsymbol{\theta}_s) - \ell_s(\boldsymbol{\theta}^*) \leq \frac{1}{2\kappa} \sum_{s=1}^t g_s^2 \|\mathbf{x}_s\|_{\mathbf{A}_s^{-1}}^2 + 2\kappa S^2 \epsilon =: B_t^{\text{ONS}},$$

where $B_t^{\text{ONS}} = O(\frac{L^2 + R^2 \log(t)}{\kappa} d \log t)$, $\forall t \geq 1$ w.h.p. If $\max_{s \geq 1} |\eta_s|$ is bounded by \bar{R} w.p. 1, $B_t^{\text{ONS}} = O(\frac{L^2 + \bar{R}^2}{\kappa} d \log t)$.

We emphasize that the OL regret bound is data-dependent. A confidence set constructed by combining Theorem 1 and Theorem 3 directly implies the following regret bound of GLOC with ONS-GLM.

Corollary 1. Define β_t^{ONS} by replacing B_t with B_t^{ONS} in (5). With probability at least $1 - 2\delta$,

$$\forall t \geq 1, \boldsymbol{\theta}^* \in C_t^{\text{ONS}} := \left\{ \boldsymbol{\theta} \in \mathbb{R}^d : \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|_{\mathbf{V}_t}^2 \leq \beta_t^{\text{ONS}} \right\}. \quad (8)$$

Corollary 2. Run GLOC with C_t^{ONS} . Then, w.p. at least $1 - 2\delta$, $\forall T \geq 1$, $\text{Regret}_T^{\text{GLOC}} = \hat{O}\left(\frac{L(L+R)}{\kappa} d \sqrt{T} \log^{3/2}(T)\right)$ where \hat{O} ignores $\log \log(t)$. If $|\eta_t|$ is bounded by \bar{R} , $\text{Regret}_T^{\text{GLOC}} = \hat{O}\left(\frac{L(L+\bar{R})}{\kappa} d \sqrt{T} \log(T)\right)$.

We make regret bound comparisons ignoring $\log \log T$ factors. For generic arm sets, our dependence on d is optimal for linear rewards [34]. For the Bernoulli GLM, our regret has the same order as Zhang et al. [41]. One can show that the regret of Filippi et al. [15] has the same order as ours if we use their assumption that the reward y_t is bounded by R_{\max} . For unbounded noise, Li et al. [28] have regret $O((LR/\kappa) d \sqrt{T} \log T)$, which is $\sqrt{\log T}$ factor smaller than ours and has LR in place of $L(L+R)$. While $L(L+R)$ could be an artifact of our analysis, the gap is not too large for canonical GLMs. Let L be the smallest Lipschitz constant of μ . Then, $R = \sqrt{L}$. If $L \leq 1$, R satisfies $R > L$, and so $L(L+R) = O(LR)$. If $L > 1$, then $L(L+R) = O(L^2)$, which is larger than $LR = O(L^{3/2})$. For the Gaussian GLM with known variance σ^2 , $L = R = 1$.³ For finite arm sets, SupCB-GLM of Li et al. [28] achieves regret of $\tilde{O}(\sqrt{dT \log N})$ that has a better scaling with d but is not a practical algorithm as it wastes a large number of arm pulls. Finally, we remark that none of the existing GLB algorithms are scalable to large T . Zhang et al. [41] is scalable to large T , but is restricted to the Bernoulli GLM; e.g., theirs does not allow the probit model (non-canonical GLM) that is popular and shown to be competitive to the Bernoulli GLM [27].

Discussion The trick of obtaining a confidence set from an online learner appeared first in [13, 14] for the linear model, and then was used in [10, 16, 41]. GLOC is slightly different from these studies and rather close to Abbasi-Yadkori et al. [2] in that the confidence set is a function of a known regret bound. This generality frees us from re-deriving a confidence set for every online learner. Our result is essentially a nontrivial extension of Abbasi-Yadkori et al. [2] to GLMs.

One might have notice that C_t does not use $\boldsymbol{\theta}_{t+1}$ that is available before pulling \mathbf{x}_{t+1} and has the most up-to-date information. This is inherent to GLOC as it relies on the OL regret bound directly. One can modify the proof of ONS-GLM to have a tighter confidence set C_t that uses $\boldsymbol{\theta}_{t+1}$ as we show in SM Section E. However, this is now specific to ONS-GLM, which loses generality.

4 Hash-Amenable Generalized Linear Bandits

We now turn to a setting where the arm set is finite but very large. For example, imagine an interactive retrieval scenario [33, 25, 6] where a user is shown K images (e.g., shoes) at a time and provides relevance feedback (e.g., yes/no or 5-star rating) on each image, which is repeated until the user is satisfied. In this paper, we focus on showing one image (i.e., arm) at a time.⁴ Most existing algorithms require maximizing an objective function (e.g., (7)), the complexity of which scales linearly with the number N of arms. This can easily become prohibitive for large numbers of images. Furthermore, the system has to perform real-time computations to promptly choose which image to show the user in the next round. Thus, it is critical for a practical system to have a time complexity sublinear in N .

One naive approach is to select a subset of arms ahead of time, such as volumetric spanners [19]. However, this is specialized for an efficient exploration only and can rule out a large number of good arms. Another option is to use hashing methods. Locality-sensitive hashing and Maximum

³ The reason why R is not σ here is that the sufficient statistic of the GLM is y/σ , which is equivalent to dealing with the normalized reward. Then, σ appears as a factor in the regret bound.

⁴ One image at a time is a simplification of the practical setting. One can extend it to showing multiple images at a time, which is a special case of the combinatorial bandits of Qin et al. [32].

Inner Product Search (MIPS) are effective and well-understood tools but can only be used when the objective function is a distance or an inner product computation; (7) cannot be written in this form. In this section, we consider alternatives to GLOC which are compatible with hashing.

Thompson Sampling We present a Thompson sampling (TS) version of GLOC called GLOC-TS that chooses an arm $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}_t} \mathbf{x}^\top \dot{\boldsymbol{\theta}}_t$ where $\dot{\boldsymbol{\theta}}_t \sim \mathcal{N}(\hat{\boldsymbol{\theta}}_{t-1}, \beta_{t-1}^{\text{ONS}} \bar{\mathbf{V}}_{t-1}^{-1})$. TS is known to perform well in practice [8] and can solve the polytope arm set case in polynomial time⁵ whereas algorithms that solve an objective function like (3) (e.g., [1]) cannot since they have to solve an NP-hard problem [5]. We present the regret bound of GLOC-TS below. Due to space constraints, we present the pseudocode and the full version of the result in SM.

Theorem 4. (Informal) *If we run GLOC-TS with $\dot{\boldsymbol{\theta}}_t \sim \mathcal{N}(\hat{\boldsymbol{\theta}}_{t-1}, \beta_{t-1}^{\text{ONS}} \bar{\mathbf{V}}_{t-1}^{-1})$, $\text{Regret}_T^{\text{GLOC-TS}} = \hat{O}\left(\frac{L(L+R)}{\kappa} d^{3/2} \sqrt{T} \log^{3/2}(T)\right)$ w.h.p. If η_t is bounded by \bar{R} , then $\hat{O}\left(\frac{L(L+\bar{R})}{\kappa} d^{3/2} \sqrt{T} \log(T)\right)$.*

Notice that the regret now scales with $d^{3/2}$ as expected from the analysis of linear TS [4], which is higher than scaling with d of GLOC. This is concerning in the interactive retrieval or product recommendation scenario since the relevance of the shown items is harmed, which makes us wonder if one can improve the regret without losing the hash-amenability.

Quadratic GLOC We now propose a new hash-amenable algorithm called Quadratic GLOC (QGLOC). Recall that GLOC chooses the arm \mathbf{x}^{GLOC} by (7). Define $r = \min_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ and

$$\bar{m}_{t-1} := \min_{\mathbf{x}: \|\mathbf{x}\|_2 \in [r, 1]} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}, \quad (9)$$

which is r times the square root of the smallest eigenvalue of $\bar{\mathbf{V}}_{t-1}^{-1}$. It is easy to see that $\bar{m}_{t-1} \leq \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}$ for all $\mathbf{x} \in \mathcal{X}$ and that $\bar{m}_{t-1} \geq r/\sqrt{t+\lambda}$ using the definition of $\bar{\mathbf{V}}_{t-1}$. There is an alternative way to define \bar{m}_{t-1} without relying on r , which we present in SM.

Let $c_0 > 0$ be the exploration-exploitation tradeoff parameter (elaborated upon later). At time t , QGLOC chooses the arm

$$\mathbf{x}_t^{\text{QGLOC}} := \arg \max_{\mathbf{x} \in \mathcal{X}_t} \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 = \arg \max_{\mathbf{x} \in \mathcal{X}_t} \langle \mathbf{q}_t, \phi(\mathbf{x}) \rangle, \quad (10)$$

where $\mathbf{q}_t = [\hat{\boldsymbol{\theta}}_{t-1}; \text{vec}(\frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \bar{\mathbf{V}}_{t-1}^{-1})] \in \mathbb{R}^{d+d^2}$ and $\phi(\mathbf{x}) := [\mathbf{x}; \text{vec}(\mathbf{x}\mathbf{x}^\top)]$. The key property of QGLOC is that the objective function is now quadratic in \mathbf{x} , thus the name *Quadratic* GLOC, and can be written as an inner product. Thus, QGLOC is hash-amenable. We present the regret bound of QGLOC (10) in Theorem 5. The key step of the proof is that the QGLOC objective function (10) plus $c_0 \beta^{3/4} \bar{m}_{t-1}$ is a tight upper bound of the GLOC objective function (7).

Theorem 5. *Run QGLOC with C_t^{ONS} . Then, w.p. at least $1 - 2\delta$, $\text{Regret}_T^{\text{QGLOC}} = O\left(\left(\frac{1}{c_0} \left(\frac{L+R}{\kappa}\right)^{1/2} + c_0 \left(\frac{L+R}{\kappa}\right)^{3/2}\right) L d^{5/4} \sqrt{T} \log^2(T)\right)$. By setting $c_0 = \left(\frac{L+R}{\kappa}\right)^{-1/2}$, the regret bound is $O\left(\frac{L(L+R)}{\kappa} d^{5/4} \sqrt{T} \log^2(T)\right)$.*

Note that one can have a better dependence on $\log T$ when η_t is bounded (available in the proof). The regret bound of QGLOC is a $d^{1/4}$ factor improvement over that of GLOC-TS; see Table 1. Furthermore, in (10) c_0 is a free parameter that adjusts the balance between the exploitation (the first term) and exploration (the second term). Interestingly, the regret guarantee *does not break down* when adjusting c_0 in Theorem 5. Such a characteristic is not found in existing algorithms but is attractive to practitioners, which we elaborate in SM.

Maximum Inner Product Search (MIPS) Hashing While MIPS hashing algorithms such as [35, 36, 30] can solve (10) in time sublinear in N , these necessarily introduce an approximation error. Ideally, one would like the following guarantee on the error with probability at least $1 - \delta_H$:

Definition 1. *Let $\mathcal{X} \subseteq \mathbb{R}^{d'}$ satisfy $|\mathcal{X}| < \infty$. A data point $\tilde{\mathbf{x}} \in \mathcal{X}$ is called c_H -MIPS w.r.t. a given query \mathbf{q} if it satisfies $\langle \mathbf{q}, \tilde{\mathbf{x}} \rangle \geq c_H \cdot \max_{\mathbf{x} \in \mathcal{X}} \langle \mathbf{q}, \mathbf{x} \rangle$ for some $c_H < 1$. An algorithm is called c_H -MIPS if, given a query $\mathbf{q} \in \mathbb{R}^{d'}$, it retrieves $\mathbf{x} \in \mathcal{X}$ that is c_H -MIPS w.r.t. \mathbf{q} .*

Unfortunately, existing MIPS algorithms do not directly offer such a guarantee, and one must build a series of hashing schemes with varying hashing parameters like Har-Peled et al. [18]. Under the fixed budget setting T , we elaborate our construction that is simpler than [18] in SM.

⁵ConfidenceBall₁ algorithm of Dani et al. [11] can solve the problem in polynomial time as well.

Time and Space Complexity Our construction involves saving Gaussian projection vectors that are used for determining hash keys and saving the buckets containing pointers to the actual arm vectors. The time complexity for retrieving a c_H -MIPS solution involves determining hash keys and evaluating inner products with the arms in the retrieved buckets. Let $\rho^* < 1$ be an optimized value for the hashing (see [35] for detail). The time complexity for d' -dimensional vectors is $O\left(\log\left(\frac{\log(dT)}{\log(c_H^{-1})}\right) N^{\rho^*} \log(N)d'\right)$, and the space complexity (except the original data) is $O\left(\frac{\log(dT)}{\log(c_H^{-1})} N^{\rho^*} (N + \log(N)d')\right)$. While the time and space complexity grows with the time horizon T , the dependence is mild; $\log \log(T)$ and $\log(T)$, respectively. QGLOC uses $d' = d + d^2$,⁶ and GLOC-TS uses $d' = d'$. While both achieve a time complexity sublinear in N , the time complexity of GLOC-TS scales with d that is better than scaling with d^2 of QGLOC. However, GLOC-TS has a $d^{1/4}$ -factor worse regret bound than QGLOC.

Discussion While it is reasonable to incur small errors in solving the arm selection criteria like (10) and sacrifice some regret in practice, the regret bounds of QGLOC and GLOC-TS do not hold anymore. Though not the focus of our paper, we prove a regret bound under the presence of the hashing error in the fixed budget setting for QGLOC; see SM. Although the result therein has an inefficient space complexity that is linear in T , it provides the first low regret bound with time sublinear in N , to our knowledge.

5 Approximate Inner Product Computations with L1 Sampling

While hashing allows a time complexity sublinear in N , it performs an additional computation for determining the hash keys. Consider a hashing with U tables and length- k hash keys. Given a query \mathbf{q} and projection vectors $\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(Uk)}$, the hashing computes $\mathbf{q}^\top \mathbf{a}^{(i)}$, $\forall i \in [Uk]$ to determine the hash key of \mathbf{q} . To reduce such an overhead, approximate inner product methods like [22, 24] are attractive since hash keys are determined by discretizing the inner products; small inner product errors often do not alter the hash keys.

In this section, we propose an improved approximate inner product method called *L1 sampling* which we claim is more accurate than the sampling proposed by Jain et al. [22], which we call *L2 sampling*. Consider an inner product $\mathbf{q}^\top \mathbf{a}$. The main idea is to construct an unbiased estimate of $\mathbf{q}^\top \mathbf{a}$. That is, let $\mathbf{p} \in \mathbb{R}^d$ be a probability vector. Let

$$i_k \stackrel{\text{i.i.d.}}{\sim} \text{Multinomial}(\mathbf{p}) \quad \text{and} \quad G_k := q_{i_k} a_{i_k} / p_{i_k}, \quad k \in [m]. \quad (11)$$

It is easy to see that $\mathbb{E}G_k = \mathbf{q}^\top \mathbf{a}$. By taking $\frac{1}{m} \sum_{k=1}^m G_k$ as an estimate of $\mathbf{q}^\top \mathbf{a}$, the time complexity is now $O(mUk)$ rather than $O(d'Uk)$. The key is to choose the right \mathbf{p} . L2 sampling uses $\mathbf{p}^{(L2)} := [q_i^2 / \|\mathbf{q}\|_2^2]_i$. Departing from L2, we propose $\mathbf{p}^{(L1)}$ that we call L1 sampling and define as follows:

$$\mathbf{p}^{(L1)} := [|q_1|; \dots; |q_d|] / \|\mathbf{q}\|_1. \quad (12)$$

We compare L1 with L2 in two different point of view. Due to space constraints, we summarize the key ideas and defer the details to SM.

The first is on their concentration of measure. Lemma 1 below shows an error bound of L1 whose failure probability decays exponentially in m . This is in contrast to decaying polynomially of L2 [22], which is inferior.⁷

Lemma 1. Define G_k as in (11) with $\mathbf{p} = \mathbf{p}^{(L1)}$. Then, given a target error $\epsilon > 0$,

$$\mathbb{P}\left(\left|\frac{1}{m} \sum_{k=1}^m G_k - \mathbf{q}^\top \mathbf{a}\right| \geq \epsilon\right) \leq 2 \exp\left(-\frac{m\epsilon^2}{2\|\mathbf{q}\|_1^2 \|\mathbf{a}\|_{\max}^2}\right) \quad (13)$$

To illustrate such a difference, we fix \mathbf{q} and \mathbf{a} in 1000 dimension and apply L2 and L1 sampling 20K times each with $m = 5$ where we scale down the L2 distribution so its variance matches that of L1.

⁶ Note that this does not mean we need to store $\text{vec}(\mathbf{x}\mathbf{x}^\top)$ since an inner product with it is structured.

⁷ In fact, one can show a bound for L2 that fails with exponentially-decaying probability. However, the bound introduces a constant that can be arbitrarily large, which makes the tails thick. We provide details on this in SM.




Figure 2: Cumulative regrets with confidence intervals under the (a) logit and (b) probit model. (c) Cumulative regrets with confidence intervals of hash-amenable algorithms.

Figure 1(a) shows that L2 has thicker tails than L1. Note this is not a pathological case but a typical case for Gaussian \mathbf{q} and \mathbf{a} . This confirms our claim that L1 is safer than L2.

Another point of comparison is the variance of L2 and L1. We show that the variance of L1 may or may not be larger than L2 in SM; there is no absolute winner. However, if \mathbf{q} and \mathbf{a} follow a Gaussian distribution, then L1 induces smaller variances than L2 for large enough d ; see Lemma 9 in SM. Figure 1(b) confirms such a result. The actual gap between the variance of L2 and L1 is also nontrivial under the Gaussian assumption. For instance, with $d = 200$, the average variance of G_k induced by L2 is 0.99 whereas that induced by L1 is 0.63 on average. Although a stochastic assumption on the vectors being inner-producted is often unrealistic, in our work we deal with projection vectors \mathbf{a} that are truly normally distributed.

6 Experiments

We now show our experiment results comparing GLB algorithms and hash-amenable algorithms.

GLB Algorithms We compare GLOC with two different algorithms: UCB-GLM [28] and Online Learning for Logit Model (OL2M) [41].⁸ For each trial, we draw $\theta^* \in \mathbb{R}^d$ and N arms (\mathcal{X}) uniformly at random from the unit sphere. We set $d = 10$ and $\mathcal{X}_t = \mathcal{X}, \forall t \geq 1$. Note it is a common practice to scale the confidence set radius for bandits [8, 27]. Following Zhang et al. [41], for OL2M we set the squared radius $\gamma_t = c \log(\det(\mathbf{Z}_t)/\det(\mathbf{Z}_1))$, where c is a tuning parameter. For UCB-GLM, we set the radius as $\alpha = \sqrt{cd \log t}$. For GLOC, we replace β_t^{ONS} with $c \sum_{s=1}^t g_s^2 \|\mathbf{x}_s\|_{\mathbf{A}_s^{-1}}^2$. While parameter tuning in practice is nontrivial, for the sake of comparison we tune $c \in \{10^1, 10^{0.5}, \dots, 10^{-3}\}$ and report the best one. We perform 40 trials up to time $T = 3000$ for each method and compute confidence bounds on the regret.

We consider two GLM rewards: (i) the logit model (the Bernoulli GLM) and (ii) the probit model (non-canonical GLM) for 0/1 rewards that sets μ as the probit function. Since OL2M is for the logit model only, we expect to see the consequences of model mismatch in the probit setting. For GLOC and UCB-GLM, we specify the correct reward model. We plot the cumulative regret under the logit model in Figure 2(a). All three methods perform similarly, and we do not find any statistically significant difference based on paired t test. The result for the probit model in Figure 2(b) shows that OL2M indeed has higher regret than both GLOC and UCB-GLM due to the model mismatch in the probit setting. Specifically, we verify that at $t = 3000$ the difference between the regret of UCB-GLM and OL2M is statistically significant. Furthermore, OL2M exhibits a significantly higher variance in the regret, which is unattractive in practice. This shows the importance of being generalizable to *any* GLM reward. Note we observe a big increase in running time for UCB-GLM compared to OL2M and GLOC.

Hash-Amenable GLBs To compare hash-amenable GLBs, we use the logit model as above but now with $N=100,000$ and $T=5000$. We run QGLOC, QGLOC with hashing (QGLOC-Hash), GLOC-TS, and GLOC-TS with hashing (GLOC-TS-Hash), where we use the hashing to compute the objective function (e.g., (10)) on just 1% of the data points and save a significant amount of computation. Details on our hashing implementation is found in SM. Figure 2(c) summarizes the result. We observe that QGLOC-Hash and GLOC-TS-Hash increase regret from QGLOC and GLOC-TS, respectively, but only moderately, which shows the efficacy of hashing.

7 Future Work

In this paper, we have proposed scalable algorithms for the GLB problem: (i) for large time horizon T and (ii) for large number N of arms. There exists a number of interesting future work. First,

⁸We have chosen UCB-GLM over GLM-UCB of Filippi et al. [15] as UCB-GLM has a lower regret bound.

we would like to extend the GLM rewards to the single index models [23] so one does not need to know the function μ ahead of time under mild assumptions. Second, closing the regret bound gap between QGLOC and GLOC without loosing hash-amenability would be interesting: i.e., develop a hash-amenable GLB algorithm with $O(d\sqrt{T})$ regret. In this direction, a first attempt could be to design a hashing scheme that can directly solve (7) approximately.

Acknowledgments This work was partially supported by the NSF grant IIS-1447449 and the MURI grant 2015-05174-04. The authors thank Yasin Abbasi-Yadkori and Anshumali Shrivastava for providing constructive feedback and Xin Hunt for her contribution at the initial stage.

References

- [1] Abbasi-Yadkori, Yasin, Pal, David, and Szepesvari, Csaba. Improved Algorithms for Linear Stochastic Bandits. *Advances in Neural Information Processing Systems (NIPS)*, pp. 1–19, 2011.
- [2] Abbasi-Yadkori, Yasin, Pal, David, and Szepesvari, Csaba. Online-to-Confidence-Set Conversions and Application to Sparse Stochastic Bandits. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- [3] Abeille, Marc and Lazaric, Alessandro. Linear Thompson Sampling Revisited. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pp. 176–184, 2017.
- [4] Agrawal, Shipra and Goyal, Navin. Thompson Sampling for Contextual Bandits with Linear Payoffs. *CoRR*, abs/1209.3, 2012.
- [5] Agrawal, Shipra and Goyal, Navin. Thompson Sampling for Contextual Bandits with Linear Payoffs. In *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 127–135, 2013.
- [6] Ahukorala, Kumaripaba, Medlar, Alan, Ilves, Kalle, and Glowacka, Dorota. Balancing Exploration and Exploitation: Empirical Parameterization of Exploratory Search Systems. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, pp. 1703–1706, 2015.
- [7] Auer, Peter and Long, M. Using Confidence Bounds for Exploitation-Exploration Trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- [8] Chapelle, Olivier and Li, Lihong. An Empirical Evaluation of Thompson Sampling. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2249–2257, 2011.
- [9] Chu, Wei, Li, Lihong, Reyzin, Lev, and Schapire, Robert E. Contextual Bandits with Linear Payoff Functions. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 15, pp. 208–214, 2011.
- [10] Crammer, Koby and Gentile, Claudio. Multiclass Classification with Bandit Feedback Using Adaptive Regularization. *Mach. Learn.*, 90(3):347–383, 2013.
- [11] Dani, Varsha, Hayes, Thomas P, and Kakade, Sham M. Stochastic Linear Optimization under Bandit Feedback. In *Proceedings of the Conference on Learning Theory (COLT)*, pp. 355–366, 2008.
- [12] Datar, Mayur, Immorlica, Nicole, Indyk, Piotr, and Mirrokni, Vahab S. Locality-sensitive Hashing Scheme Based on P-stable Distributions. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, pp. 253–262, 2004.
- [13] Dekel, Ofer, Gentile, Claudio, and Sridharan, Karthik. Robust selective sampling from single and multiple teachers. In *Proceedings of the Conference on Learning Theory (COLT)*, 2010.
- [14] Dekel, Ofer, Gentile, Claudio, and Sridharan, Karthik. Selective sampling and active learning from single and multiple teachers. *Journal of Machine Learning Research*, 13:2655–2697, 2012.

- [15] Filippi, Sarah, Cappe, Olivier, Garivier, Aurélien, and Szepesvári, Csaba. Parametric Bandits: The Generalized Linear Case. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 586–594, 2010.
- [16] Gentile, Claudio and Orabona, Francesco. On Multilabel Classification and Ranking with Bandit Feedback. *Journal of Machine Learning Research*, 15:2451–2487, 2014.
- [17] Guo, Ruiqi, Kumar, Sanjiv, Choromanski, Krzysztof, and Simcha, David. Quantization based Fast Inner Product Search. *Journal of Machine Learning Research*, 41:482–490, 2016.
- [18] Har-Peled, Sariel, Indyk, Piotr, and Motwani, Rajeev. Approximate nearest neighbor: towards removing the curse of dimensionality. *Theory of Computing*, 8:321–350, 2012.
- [19] Hazan, Elad and Karnin, Zohar. Volumetric Spanners: An Efficient Exploration Basis for Learning. *Journal of Machine Learning Research*, 17(119):1–34, 2016.
- [20] Hazan, Elad, Agarwal, Amit, and Kale, Satyen. Logarithmic Regret Algorithms for Online Convex Optimization. *Mach. Learn.*, 69(2-3):169–192, 2007.
- [21] Hofmann, Katja, Whiteson, Shimon, and de Rijke, Maarten. Contextual Bandits for Information Retrieval. In *NIPS Workshop on Bayesian Optimization, Experimental Design and Bandits: Theory and Applications*, 2011.
- [22] Jain, Prateek, Vijayanarasimhan, Sudheendra, and Grauman, Kristen. Hashing Hyperplane Queries to Near Points with Applications to Large-Scale Active Learning. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 928–936, 2010.
- [23] Kalai, Adam Tauman and Sastry, Ravi. The Isotron Algorithm: High-Dimensional Isotonic Regression. In *Proceedings of the Conference on Learning Theory (COLT)*, 2009.
- [24] Kannan, Ravindran, Vempala, Santosh, and Others. Spectral algorithms. *Foundations and Trends in Theoretical Computer Science*, 4(3–4):157–288, 2009.
- [25] Konyushkova, Ksenia and Glowacka, Dorota. Content-based image retrieval with hierarchical Gaussian Process bandits with self-organizing maps. In *21st European Symposium on Artificial Neural Networks*, 2013.
- [26] Li, Lihong, Chu, Wei, Langford, John, and Schapire, Robert E. A Contextual-Bandit Approach to Personalized News Article Recommendation. *Proceedings of the International Conference on World Wide Web (WWW)*, pp. 661–670, 2010.
- [27] Li, Lihong, Chu, Wei, Langford, John, Moon, Taesup, and Wang, Xuanhui. An Unbiased Offline Evaluation of Contextual Bandit Algorithms with Generalized Linear Models. In *Proceedings of the Workshop on On-line Trading of Exploration and Exploitation 2*, volume 26, pp. 19–36, 2012.
- [28] Li, Lihong, Lu, Yu, and Zhou, Dengyong. Provable Optimal Algorithms for Generalized Linear Contextual Bandits. *CoRR*, abs/1703.0, 2017.
- [29] McCullagh, P and Nelder, J A. *Generalized Linear Models*. London, 1989.
- [30] Neyshabur, Behnam and Srebro, Nathan. On Symmetric and Asymmetric LSHs for Inner Product Search. *Proceedings of the International Conference on Machine Learning (ICML)*, 37: 1926–1934, 2015.
- [31] Orabona, Francesco, Cesa-Bianchi, Nicolo, and Gentile, Claudio. Beyond Logarithmic Bounds in Online Learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 22, pp. 823–831, 2012.
- [32] Qin, Lijing, Chen, Shouyuan, and Zhu, Xiaoyan. Contextual Combinatorial Bandit and its Application on Diversified Online Recommendation. In *SDM*, pp. 461–469, 2014.
- [33] Rui, Yong, Huang, T S, Ortega, M, and Mehrotra, S. Relevance feedback: a power tool for interactive content-based image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 8(5):644–655, 1998.

- [34] Rusmevichientong, Paat and Tsitsiklis, John N. Linearly Parameterized Bandits. *Math. Oper. Res.*, 35(2):395–411, 2010.
- [35] Shrivastava, Anshumali and Li, Ping. Asymmetric LSH (ALSH) for Sublinear Time Maximum Inner Product Search (MIPS). *Advances in Neural Information Processing Systems 27*, pp. 2321–2329, 2014.
- [36] Shrivastava, Anshumali and Li, Ping. Improved Asymmetric Locality Sensitive Hashing (ALSH) for Maximum Inner Product Search (MIPS). In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 812–821, 2015.
- [37] Slaney, Malcolm, Lifshits, Yury, and He, Junfeng. Optimal parameters for locality-sensitive hashing. *Proceedings of the IEEE*, 100(9):2604–2623, 2012.
- [38] Wainwright, Martin J and Jordan, Michael I. Graphical Models, Exponential Families, and Variational Inference. *Found. Trends Mach. Learn.*, 1(1-2):1–305, 2008.
- [39] Wang, Jingdong, Shen, Heng Tao, Song, Jingkuan, and Ji, Jianqiu. Hashing for Similarity Search: A Survey. *CoRR*, abs/1408.2, 2014.
- [40] Yue, Yisong, Hong, Sue Ann Sa, and Guestrin, Carlos. Hierarchical exploration for accelerating contextual bandits. *Proceedings of the International Conference on Machine Learning (ICML)*, pp. 1895–1902, 2012.
- [41] Zhang, Lijun, Yang, Tianbao, Jin, Rong, Xiao, Yichi, and Zhou, Zhi-hua. Online Stochastic Linear Optimization under One-bit Feedback. In *Proceedings of the International Conference on Machine Learning (ICML)*, volume 48, pp. 392–401, 2016.

Supplementary Material (Appendix)

A Proof of Theorem 1

We first describe the sketch of the proof. We perform a generalized version of the online-to-confidence-set conversion of Abbasi-Yadkori et al. [2, Theorem 1] that was strictly for linear rewards. Unlike that work, which deals with the squared loss, we now work with the negative log likelihood loss. The key is to use the strong-convexity of the loss function to turn the OL regret bound (4) into a quadratic equation in θ^* . Then, it remains to bound random quantities with a Martingale concentration bound.

Proof. Let $\ell'(z, y)$ and $\ell''(z, y)$ be the first and second derivative of the loss ℓ w.r.t. z . We lower bound the LHS of (4) using Taylor's theorem with some ξ_s between $\mathbf{x}_s^\top \theta^*$ and $\mathbf{x}_s^\top \theta_s$:

$$\begin{aligned} B_t &\geq \sum_{s=1}^t \ell_s(\theta_s) - \ell_s(\theta^*) = \sum_{s=1}^t \ell'(\mathbf{x}_s^\top \theta^*, y_s) \mathbf{x}_s^\top (\theta_s - \theta^*) + \frac{\ell''(\xi_s, y_s)}{2} (\mathbf{x}_s^\top (\theta_s - \theta^*))^2 \\ &\stackrel{(\text{Assumption 1})}{\geq} \sum_{s=1}^t \ell'(\mathbf{x}_s^\top \theta^*, y_s) \mathbf{x}_s^\top (\theta_s - \theta^*) + \frac{\kappa}{2} (\mathbf{x}_s^\top (\theta_s - \theta^*))^2. \end{aligned}$$

Since $\ell'(\mathbf{x}_s^\top \theta^*, y_s) = -y_s + \mu(\mathbf{x}_s^\top \theta^*) = -\eta_s$,

$$\sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2 \leq \frac{2}{\kappa} B_t + \frac{2}{\kappa} \sum_{s=1}^t \eta_s (\mathbf{x}_s^\top (\theta_s - \theta^*)).$$

Note that the second term in the RHS involves η_t that is unknown and random, which we bound using Abbasi-Yadkori et al. [2, Corollary 8]. That is, w.p. at least $1 - \delta$, for all $t \geq 1$,

$$\sum_{s=1}^t \eta_s (\mathbf{x}_s^\top (\theta_s - \theta^*)) \leq R \sqrt{\left(2 + 2 \sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2\right) \cdot \log \left(\frac{1}{\delta} \sqrt{1 + \sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2}\right)}.$$

Then,

$$\sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2 \leq \frac{2}{\kappa} B_t + \frac{2R}{\kappa} \sqrt{\left(2 + 2 \sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2\right) \cdot \log \left(\frac{1}{\delta} \sqrt{1 + \sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2}\right)}.$$

Define $q := \sqrt{1 + \sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2}$. Then, the inequality above can be written as $q^2 \leq 1 + \frac{2}{\kappa} B_t + \frac{2\sqrt{2}R}{\kappa} q \sqrt{\log(q/\delta)}$. The following Lemma is useful. See Section A.1 for a proof.

Lemma 2. Let $\delta \in (0, 1)$, $a \geq 0$, $f \geq 0$, $q \geq 1$. Then,

$$q^2 \leq a + f q \sqrt{\log \left(\frac{q}{\delta}\right)} \implies q^2 \leq 2a + f^2 \log \left(\frac{\sqrt{4a + f^4/(4\delta^2)}}{\delta}\right)$$

Applying Lemma 2 with $a := 1 + \frac{2}{\kappa} B_t$ and $f := 2\sqrt{2}R/\kappa$, we have

$$\sum_{s=1}^t (\mathbf{x}_s^\top (\theta_s - \theta^*))^2 \leq 1 + \frac{4}{\kappa} B_t + \frac{8R^2}{\kappa^2} \log \left(\frac{1}{\delta} \sqrt{4 + \frac{8}{\kappa} B_t + \frac{64R^4}{\kappa^4 \cdot 4\delta^2}}\right) = \beta'_t$$

Then, one can rewrite the above as

$$\|\mathbf{z}_t - \mathbf{X}_t \theta^*\|_2^2 \leq \beta'_t.$$

Let $\lambda > 0$. We add $\lambda \|\theta^*\|_2^2$ to the both sides.

$$\lambda \|\theta^*\|_2^2 + \|\mathbf{z}_t - \mathbf{X}_t \theta^*\|_2^2 \leq \lambda \|\theta^*\|_2^2 + \beta'_t \leq \lambda S^2 + \beta'_t.$$

Hereafter, we omit t from \mathbf{X}_t and \mathbf{z}_t for brevity. Since the LHS is quadratic in θ^* , we can rewrite it as an ellipsoid centered at $\hat{\theta}_t := \arg \min_{\theta} \lambda \|\theta\|_2^2 + \|\mathbf{z} - \mathbf{X}\theta\|_2^2 = \bar{\mathbf{V}}_t^{-1} \mathbf{X}^\top \mathbf{z}$ as follows:

$$\begin{aligned} \|\theta^* - \hat{\theta}_t\|_{\bar{\mathbf{V}}_t}^2 + \underbrace{\lambda \|\hat{\theta}_t\|_2^2 + \|\mathbf{z} - \mathbf{X}\hat{\theta}_t\|_2^2}_{= \|\mathbf{z}\|_2^2 - \hat{\theta}_t^\top \mathbf{X}^\top \mathbf{z}} &\leq \lambda S^2 + \beta'_t, \end{aligned}$$

which concludes the proof. \square

A.1 Proof of Lemma 2

Proof. Let $c := f\sqrt{\log(q/\delta)}$. Then, $q^2 \leq a + cq \implies q^2 - cq - a \leq 0$. Solving it for q , we get $q \leq \frac{c+\sqrt{c^2+4a}}{2}$. Then, using $(u+v)^2 \leq 2(u^2+v^2)$,

$$q^2 \leq \left(\frac{c + \sqrt{c^2 + 4a}}{2} \right)^2 \leq \frac{2(c^2 + c^2 + 4a)}{4} = c^2 + 2a \quad (14)$$

$$\iff q^2 \leq 2a + f^2 \log(q/\delta)$$

One might suspect that c has q in it, which might cause a problem. To be assured, one can prove the contrapositive: $q > \frac{c+\sqrt{c^2+4a}}{2} \implies q^2 - cq - a > 0$. To see this, $q^2 - cq - a = (q - \frac{c+\sqrt{c^2+4a}}{2})(q - \frac{c-\sqrt{c^2+4a}}{2})$ and since $q > \frac{c+\sqrt{c^2+4a}}{2}$ it suffices to show that $q - \frac{c-\sqrt{c^2+4a}}{2} > 0$. Then, since $q - \frac{c-\sqrt{c^2+4a}}{2} \geq q - \frac{c+\sqrt{c^2+4a}}{2} > 0$.

Using $\log u \leq \frac{1}{2}u$,

$$q^2 \leq 2a + f^2 \log(q/\delta) \leq 2a + \frac{f^2}{2\delta} q \iff q^2 - \frac{f^2}{2\delta} q - 2a \leq 0.$$

Solving the quadratic inequality for q , we have $q \leq \frac{f^2/(2\delta) + \sqrt{(f^4/(4\delta^2)) + 8a}}{2}$. This implies that

$$q^2 \leq \frac{2(f^4/(4\delta^2) + f^4/(4\delta^2) + 8a)}{4} = \frac{f^4}{4\delta^2} + 4a.$$

Now, applying this inequality on q in the RHS of (14),

$$q^2 \leq 2a + f^2 \log \left(\frac{\sqrt{4a + f^4/(4\delta^2)}}{\delta} \right)$$

□

B Proof of Theorem 2

Proof. Our proof closely follow a standard technique (cf. Abbasi-Yadkori et al. [1]). Define $\mathbf{x}_{t,*} = \arg \max_{\mathbf{x} \in \mathcal{X}_t} \langle \mathbf{x}, \boldsymbol{\theta}^* \rangle$. Let $r_t := \mu(\mathbf{x}_{t,*}^\top \boldsymbol{\theta}^*) - \mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*)$ be the instantaneous regret. Using $\mu(\mathbf{x}_{t,*}^\top \boldsymbol{\theta}^*) - \mu(\mathbf{x}_t^\top \boldsymbol{\theta}^*) \leq L(\mathbf{x}_{t,*}^\top \boldsymbol{\theta}^* - \mathbf{x}_t^\top \boldsymbol{\theta}^*)$,

$$\begin{aligned} \frac{r_t}{L} &\leq \mathbf{x}_{t,*}^\top \boldsymbol{\theta}^* - \mathbf{x}_t^\top \boldsymbol{\theta}^* \\ &\leq \mathbf{x}_t^\top \tilde{\boldsymbol{\theta}}_t - \mathbf{x}_t^\top \boldsymbol{\theta}^* \\ &= \mathbf{x}_t^\top (\tilde{\boldsymbol{\theta}}_t - \hat{\boldsymbol{\theta}}_{t-1}) + \mathbf{x}_t^\top (\hat{\boldsymbol{\theta}}_{t-1} - \boldsymbol{\theta}^*) \\ &\leq \|\mathbf{x}_t\|_{\mathbf{V}_{t-1}^{-1}} \|\tilde{\boldsymbol{\theta}}_t - \hat{\boldsymbol{\theta}}_{t-1}\|_{\mathbf{V}_{t-1}} + \|\mathbf{x}_t\|_{\mathbf{V}_{t-1}^{-1}} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_{t-1}\|_{\mathbf{V}_{t-1}} \\ &\leq 2\sqrt{\beta_t} \|\mathbf{x}_t\|_{\mathbf{V}_{t-1}^{-1}}. \end{aligned}$$

Note that

$$\sum_{t=1}^T \log \left(1 + \|\mathbf{x}_t\|_{\mathbf{V}_{t-1}^{-1}}^2 \right) = \log \left(\frac{\det(\mathbf{V}_T)}{\det(\lambda \mathbf{I})} \right) \leq d \log(1 + T/(d\lambda)), \quad (15)$$

which is due to Abbasi-Yadkori et al. [1, Lemma 11].

The following lemmas become useful.

Lemma 3. For any $q, x \geq 0$,

$$\min\{q, x\} \leq \max\{2, q\} \log(1 + x)$$

Proof. It is not hard to see that

$$x \in [0, a] \implies x \leq \frac{a}{\log(1+a)} \log(1+x) \quad (16)$$

We consider the following two cases.

Case 1. $q \leq 2$

If $x \leq 2$, by (16), $\min\{2, x\} = x \leq \frac{2}{\log(3)} \log(1+x) \leq 2 \log(1+x)$. If $x > 2$, $\min\{2, x\} = 2 \leq 2 \log(1+2) \leq 2 \log(1+x)$. Thus, for any x , $\min\{q, x\} \leq \min\{2, x\} \leq 2 \log(1+x)$.

Case 2. $q > 2$

If $x \leq q$, by (16), $\min\{q, x\} = x \leq \frac{q}{\log(1+q)} \log(1+x) < q \log(1+x)$. If $x > q$, $\min\{q, x\} = q \leq q \log(1+2) \leq q \log(1+x)$.

Combining both cases to complete the proof. \square

Lemma 4. If A is a value independent of t ,

$$\sum_{t=1}^T \min\{A, \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}^2\} \leq \max\{2, A\} d \log(1 + T/(d\lambda)) .$$

Proof. Combine Lemma 3 and (15). \square

Since r_t cannot be bigger than $2LS$,

$$\begin{aligned} \sum_{t=1}^T r_t &\leq \sum_{t=1}^T \min\{2LS, 2L\sqrt{\beta_t}\|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}\} \\ &\leq 2L\sqrt{\beta_T} \sum_{t=1}^T \min\{S/\sqrt{\beta_T}, \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}\} \\ &\stackrel{\text{(C.-S.)}}{\leq} 2L\sqrt{\beta_T} \sqrt{T \sum_{t=1}^T \min\left\{\frac{S^2}{\beta_T}, \|\mathbf{x}_t\|_{\mathbf{V}_t^{-1}}^2\right\}} \\ &\stackrel{\text{(Lem. 4)}}{\leq} 2L\sqrt{\beta_T} \sqrt{T \max\{2, S^2/\beta_T\} d \log(1 + T/(d\lambda))} \\ &= O\left(L\sqrt{\beta_T}\sqrt{T} \cdot \sqrt{d \log T}\right) \end{aligned}$$

where C.-S. stands for the Cauchy-Schwartz inequality. \square

C Proof of Theorem 3

Proof. We closely follow the proof of Hazan et al. [20]. Since $\ell(z, y)$ is κ -strongly convex w.r.t. $z \in \mathcal{B}_1(S)$,

$$\ell(\mathbf{x}_s^\top \boldsymbol{\theta}_s, y_s) - \ell(\mathbf{x}_s^\top \boldsymbol{\theta}^*, y_s) \leq \ell'(\mathbf{x}_s^\top \boldsymbol{\theta}_s, y_s) \cdot \mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*) - \frac{\kappa}{2} (\mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*))^2 . \quad (17)$$

Define $g_s := \ell'(\mathbf{x}_s^\top \boldsymbol{\theta}_s, y_s)$. Note that by the update rule of Algorithm 2,

$$\begin{aligned} \boldsymbol{\theta}'_{s+1} - \boldsymbol{\theta}^* &= \boldsymbol{\theta}_s - \boldsymbol{\theta}^* - \frac{g_s}{\kappa} \mathbf{A}_s^{-1} \mathbf{x}_s \\ \implies \|\boldsymbol{\theta}'_{s+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 &= \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 - \frac{2g_s}{\kappa} \mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*) + \frac{g_s^2}{\kappa^2} \|\mathbf{x}_s\|_{\mathbf{A}_s^{-1}}^2 \end{aligned} \quad (18)$$

By the property of the generalized projection (see Hazan et al. [20, Lemma 8])

$$\|\boldsymbol{\theta}'_{s+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 \geq \|\boldsymbol{\theta}_{s+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 .$$

Now, together with (18),

$$\begin{aligned} \|\boldsymbol{\theta}_{s+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 &\leq \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 - \frac{2g_s}{\kappa} \mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*) + \frac{g_s^2}{\kappa^2} \|\mathbf{x}_s\|_{\mathbf{A}_s^{-1}}^2 \\ \implies \sum_{s=1}^t g_s \mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*) &\leq \sum_{s=1}^t \frac{g_s^2}{2\kappa} \|\mathbf{x}_s\|_{\mathbf{A}_s^{-1}}^2 + \underbrace{\frac{\kappa}{2} \sum_{s=1}^t \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 - \|\boldsymbol{\theta}_{s+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2}_{=: D_1} . \end{aligned}$$

Note

$$\begin{aligned}
D_1 &= \|\theta_1 - \theta^*\|_{\mathbf{A}_1}^2 + \left(\sum_{s=2}^t \|\theta_s - \theta^*\|_{\mathbf{A}_s}^2 - \sum_{s=2}^t \|\theta_s - \theta^*\|_{\mathbf{A}_{s-1}}^2 \right) - \|\theta_{t+1} - \theta^*\|_{\mathbf{A}_t}^2 \\
&\leq \|\theta_1 - \theta^*\|_{\mathbf{A}_1}^2 + \left(\sum_{s=2}^t \|\theta_s - \theta^*\|_{\mathbf{A}_s}^2 - \sum_{s=2}^t \|\theta_s - \theta^*\|_{\mathbf{A}_{s-1}}^2 \right) \\
&\stackrel{(a)}{=} \|\theta_1 - \theta^*\|_{\mathbf{A}_1}^2 + \left(-\|\theta_1 - \theta^*\|_{\mathbf{x}_1 \mathbf{x}_1^\top}^2 + \sum_{s=1}^t \|\theta_s - \theta^*\|_{\mathbf{x}_s \mathbf{x}_s^\top}^2 \right) \\
&= \|\theta_1 - \theta^*\|_{\epsilon \mathbf{I}}^2 + \sum_{s=1}^t \|\theta_s - \theta^*\|_{\mathbf{x}_s \mathbf{x}_s^\top}^2 \leq 4\epsilon S^2 + \sum_{s=1}^t \|\theta_s - \theta^*\|_{\mathbf{x}_s \mathbf{x}_s^\top}^2
\end{aligned}$$

where (a) is due to $\mathbf{A}_s - \mathbf{A}_{s-1} = \mathbf{x}_s \mathbf{x}_s^\top$. Therefore,

$$\sum_{s=1}^t g_s \mathbf{x}_s^\top (\theta_s - \theta^*) \leq \sum_{s=1}^t \frac{g_s^2}{2\kappa} \|\mathbf{x}_s\|_{\mathbf{A}_{s-1}}^2 + 2\epsilon \kappa S^2 + \frac{\kappa}{2} \sum_{s=1}^t \|\theta_s - \theta^*\|_{\mathbf{x}_s \mathbf{x}_s^\top}^2.$$

Move the rightmost sum in the RHS to the LHS to see that the LHS now coincide with the RHS of (17). This leads to

$$\sum_{s=1}^t \ell(\mathbf{x}_s^\top \theta_s, y_s) - \ell(\mathbf{x}_s^\top \theta^*, y_s) \leq \frac{1}{2\kappa} \sum_{s=1}^t g_s^2 \|\mathbf{x}_s\|_{\mathbf{A}_{s-1}}^2 + 2\epsilon \kappa S^2 = B_t^{\text{ONS}}.$$

This yields the statement of the theorem.

For characterizing the order of B_t^{ONS} , notice that g_s is a random variable:

$$\begin{aligned}
g_s^2 &= (-y_s + \mu(\mathbf{x}_s^\top \theta_s))^2 = (-\mu(\mathbf{x}_s^\top \theta^*) - \eta_s + \mu(\mathbf{x}_s^\top \theta_s))^2 \\
&\leq 2(\mu(\mathbf{x}_s^\top \theta_s) - \mu(\mathbf{x}_s^\top \theta^*))^2 + 2\eta_s^2 \\
&\leq 2(L \cdot \mathbf{x}_s^\top (\theta_s - \theta^*))^2 + 2\eta_s^2 \\
&\leq 2L^2 \cdot 4S^2 + 2\eta_s^2
\end{aligned}$$

Let $\delta < 1$ be the target failure rate. By the sub-Gaussianity of η_s ,

$$\mathbb{P}(\forall s \geq 1, |\eta_s|^2 \geq 2R^2 \log(4s^2/\delta)) \leq \sum_{s \geq 1} \mathbb{P}(|\eta_s|^2 \geq 2R^2 \log(4s^2/\delta)) \leq \sum_{s \geq 1} \delta/(2s^2) \leq \delta.$$

Thus, w.p. at least $1 - \delta$, $\max_{s \leq t} g_s^2 \leq 8L^2 S^2 + 4R^2 \log(4t^2/\delta) = O(L^2 + R^2 \log(t/\delta))$. Furthermore, $\sum_{s=1}^t \|\mathbf{x}_s\|_{\mathbf{A}_{s-1}}^2 \leq d \log(1 + (t/\epsilon))$ by Hazan et al. [20, Lemma 11]. Thus, w.p. at least $1 - \delta$, $\forall t \geq 1, B_t^{\text{ONS}} = O\left(\frac{L^2 + R^2 \log(t/\delta)}{\kappa} d \log t\right)$.

For the case where $|\eta_s|$ is bounded by \bar{R} w.p. 1 (e.g., $\bar{R} = \frac{1}{2}$ for Bernoulli), $\max_{s \leq t} g_s^2 \leq 8L^2 S^2 + 2\bar{R}^2$, which leads to $B_t^{\text{ONS}} = O\left(\frac{L^2 + \bar{R}^2}{\kappa} d \log t\right)$. \square

D Proof of Corollaries 1 and 2

The proof of Corollary 1 a trivial consequence of combining Theorem 1 and Theorem 3.

Corollary 2 is simply a combination of Theorem 2 and Corollary 1. Note that $\bar{\beta}_t^{\text{ONS}} = \alpha(B_t^{\text{ONS}}) + \lambda S^2 \geq \beta_t^{\text{ONS}}$ by noticing that $\|\mathbf{z}_t\|_2^2 - \hat{\theta}_t^\top \mathbf{X}_t^\top \mathbf{z}_t$ is nonnegative (from the proof of Theorem 2). This concludes the proof.

E A Tighter Confidence Set

While the confidence set constructed by Theorem 2 is generic and allows us to rely on any online learner with a known regret bound, one can find a tighter confidence set by analyzing the online learner directly. We show one instance of such for ONS. A distinctive characteristic of our new confidence set, denoted by $C_t^{\text{ONS}+}$, is that it now depends on y_t (note that C_t^{ONS} depends on y_1, \dots, y_{t-1} only).

We deviate from the proof of Theorem 3. Recall that

$$D_1 = \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}^*\|_{\mathbf{A}_1}^2 + \left(\sum_{s=2}^t \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{A}_s}^2 - \sum_{s=2}^t \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{A}_{s-1}}^2 \right) - \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2$$

We previously dropped the term $\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2$. We now keep it, which leads to:

$$D_1 \leq 4\epsilon S^2 + \sum_{s=1}^t \|\boldsymbol{\theta}_s - \boldsymbol{\theta}^*\|_{\mathbf{x}_s \mathbf{x}_s^\top}^2 - \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2$$

Following the same argument,

$$\begin{aligned} & \left(\sum_{s=1}^t \ell_s(\mathbf{x}_s^\top \boldsymbol{\theta}_t) - \ell_s(\mathbf{x}_s^\top \boldsymbol{\theta}^*) \right) + \frac{\kappa}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2 \leq \sum_{s=1}^t \frac{g_s^2}{2\kappa} \|\mathbf{x}_s\|_{\mathbf{A}_{s-1}}^2 + 2\epsilon\kappa S^2 = B_t^{\text{ONS}} \\ & \implies \left(\sum_{s=1}^t \ell_s(\mathbf{x}_s^\top \boldsymbol{\theta}_t) - \ell_s(\mathbf{x}_s^\top \boldsymbol{\theta}^*) \right) \leq B_t^{\text{ONS}} - \frac{\kappa}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2, \end{aligned}$$

Combining the above with the proof of Theorem 1,

$$\begin{aligned} & \sum_{s=1}^t (\mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*))^2 \\ & \leq 1 + \frac{4}{\kappa} (B_t^{\text{ONS}} - \frac{\kappa}{2} \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2) + \frac{8R^2}{\kappa^2} \log \left(\frac{2}{\delta} \sqrt{1 + \frac{2}{\kappa} B_t^{\text{ONS}} + \frac{4R^4}{\kappa^4 \delta^2}} \right) \\ & \iff \sum_{s=1}^t (\mathbf{x}_s^\top (\boldsymbol{\theta}_s - \boldsymbol{\theta}^*))^2 + 2\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\mathbf{A}_t}^2 \leq 1 + \frac{4}{\kappa} B_t^{\text{ONS}} + \frac{8R^2}{\kappa^2} \log \left(\frac{2}{\delta} \sqrt{1 + \frac{2}{\kappa} B_t^{\text{ONS}} + \frac{4R^4}{\kappa^4 \delta^2}} \right) \end{aligned}$$

Define $\mathbf{z}_t = [\mathbf{x}_s^\top \boldsymbol{\theta}_s]_{s \in [t]}$, $\mathbf{z}_t^* = [\mathbf{x}_s^\top \boldsymbol{\theta}^*]_{s \in [t]}$ and $\mathbf{z}'_s = [\mathbf{x}_s^\top \boldsymbol{\theta}_{t+1}]_{s \in [t]}$. Then, the LHS above is

$$\begin{aligned} & \|\mathbf{z}_t - \mathbf{z}_t^*\|_2^2 + 2\|\mathbf{z}'_t - \mathbf{z}_t^*\|_2^2 + 2\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\epsilon \mathbf{I}}^2 \\ & = 3 \left\| \frac{\mathbf{z}_t + 2\mathbf{z}'_t}{3} - \mathbf{z}_t^* \right\|_2^2 - \frac{1}{3} \|\mathbf{z}_t + 2\mathbf{z}'_t\|_2^2 + \|\mathbf{z}_t\|_2^2 + 2\|\mathbf{z}'_t\|_2^2 + 2\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\epsilon \mathbf{I}}^2 \\ & = 3 \left\| \frac{\mathbf{z}_t + 2\mathbf{z}'_t}{3} - \mathbf{z}_t^* \right\|_2^2 + \frac{2}{3} \|\mathbf{z}_t - \mathbf{z}'_t\|_2^2 + 2\|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{\epsilon \mathbf{I}}^2. \end{aligned}$$

Let $\bar{\mathbf{z}}_s = (\mathbf{z}_s + 2\mathbf{z}'_s)/3$. Then,

$$\begin{aligned} \|\bar{\mathbf{z}}_t - \mathbf{z}_t^*\|_2^2 + \|\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*\|_{(2\epsilon/3)\mathbf{I}}^2 & \leq -\frac{2}{9} \|\mathbf{z}_t - \mathbf{z}'_t\|_2^2 + \frac{1}{3} + \frac{4}{3\kappa} B_t^{\text{ONS}} \\ & \quad + \frac{8R^2}{3\kappa^2} \log \left(\frac{2}{\delta} \sqrt{1 + \frac{2}{\kappa} B_t^{\text{ONS}} + \frac{4R^4}{\kappa^4 \delta^2}} \right) \end{aligned}$$

We lower bound the LHS with $\|\bar{\mathbf{z}}_t - \mathbf{z}_t^*\|_2^2 + \frac{2\epsilon}{3} \|\boldsymbol{\theta}^*\|_2^2 - \frac{4\epsilon}{3} S^2 =: D_2$. Define $\bar{\mathbf{W}}_t := \mathbf{X}_t^\top \mathbf{X}_t + (2\epsilon/3)\mathbf{I}$ and $\hat{\boldsymbol{\theta}}_t^+ = \bar{\mathbf{W}}_t^{-1} \mathbf{X}_t^\top \bar{\mathbf{z}}_t$. Then,

$$D_2 = \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t^+\|_{\bar{\mathbf{W}}_t}^2 + \underbrace{\|\bar{\mathbf{z}}_t - \mathbf{X}_t \hat{\boldsymbol{\theta}}_t^+\|_2^2 + \frac{2\epsilon}{3} \|\hat{\boldsymbol{\theta}}_t^+\|_2^2 - \frac{4\epsilon}{3} S^2}_{= \|\bar{\mathbf{z}}_t\|_2^2 - \bar{\mathbf{z}}_t^\top \mathbf{X}_t \hat{\boldsymbol{\theta}}_t^+}.$$

Thus,

$$\begin{aligned} \|\boldsymbol{\theta}^* - \hat{\boldsymbol{\theta}}_t^+\|_{\bar{\mathbf{W}}_t}^2 & \leq -\|\bar{\mathbf{z}}_t\|_2^2 + \bar{\mathbf{z}}_t^\top \mathbf{X}_t \hat{\boldsymbol{\theta}}_t^+ + \frac{4\epsilon}{3} S^2 - \frac{2}{9} \|\mathbf{z}_t - \mathbf{z}'_t\|_2^2 + \frac{1}{3} + \frac{4}{3\kappa} B_t^{\text{ONS}} \\ & \quad + \frac{8R^2}{3\kappa^2} \log \left(\frac{2}{\delta} \sqrt{1 + \frac{2}{\kappa} B_t^{\text{ONS}} + \frac{4R^4}{\kappa^4 \delta^2}} \right) \\ & =: \beta_t^{\text{ONS}+}. \end{aligned}$$

This leads to the following confidence set:

$$C_t^{\text{ONS}+} = \{\boldsymbol{\theta} \in \mathbb{R}^d : \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t^+\|_{\bar{\mathbf{W}}_t}^2 \leq \beta_t^{\text{ONS}+}\}.$$

F Details on GLOC-TS

For simplicity, we present the algorithm and the analysis when GLOC-TS is combined with the confidence set C_t^{ONS} . To present the algorithm, we use the definition β_t^{ONS} from Corollary 1. We use the notation $\beta_t^{\text{ONS}}(\delta)$ to show the dependence on δ explicitly. We present the algorithm of GLOC-TS in Algorithm 3.

Algorithm 3 GLOC-TS (GLOC - Thompson Sampling)

- 1: **Input:** time horizon T , $\delta \in (0, 1)$, $\lambda > 0$, $S > 0$, $\kappa > 0$.
 - 2: Let $\delta' = \delta/(8T)$.
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: Sample $\xi_t \sim \mathcal{N}(0, \mathbf{I})$.
 - 5: Compute parameter $\dot{\theta}_t = \hat{\theta}_{t-1} + \sqrt{\beta_{t-1}^{\text{ONS}}(\delta')} \cdot \bar{\mathbf{V}}_{t-1}^{-1/2} \xi_t$
 - 6: Solve $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \mathbf{x}^\top \dot{\theta}_t$.
 - 7: Pull \mathbf{x}_t and then observe y_t .
 - 8: Feed into \mathcal{B} the loss function $\ell_t(\theta) = \ell(\mathbf{x}_t^\top \theta, y_t)$.
 - 9: Update $\hat{\theta}_t$ and $\bar{\mathbf{V}}_t$.
 - 10: **end for**
-

Define $\gamma_t(\delta) = \beta_t(\delta)2d \log(2d/\delta)$ and $p = \frac{1}{4\sqrt{e\pi}}$. We present the full statement of Theorem 4 as follows.

Theorem 4 Let $\delta' = \delta/(8T)$. The cumulative regret of GLOC-TS over T steps is bounded as, w.p. at least $1 - \delta$,

$$\begin{aligned} \text{Regret}_T &\leq L \left(\sqrt{\beta_T(\delta')} + \sqrt{\gamma_T(\delta')(1 + 2/p)} \right) \sqrt{2Td \log(1 + T/\lambda)} + \frac{2L}{p} \sqrt{\gamma_T(\delta')} \sqrt{\frac{8T}{\lambda} \log(4/\delta)} \\ &= O \left(\frac{L(L + R)}{\kappa} d^{3/2} \sqrt{\log(d)T} \log^{3/2} T \right) \end{aligned}$$

Proof. Note that the proof is a matter of realizing that the proof of Abeille & Lazaric [3, Lemma 4] relies on a given confidence set C_t in a form of (6). To be specific, notice that the notations k_μ , c_μ , and $\beta_t(\delta')$ used in Abeille & Lazaric [3] is equivalent to L , κ , and $\kappa\sqrt{\beta_t(\delta')}$ in this paper, respectively. Furthermore, our result stated in (8) can replace the result of Abeille & Lazaric [3, Prop. 11]. Finally, one can verify that the proof of Abeille & Lazaric [3, Lemma 4] leads to the proof of the theorem above. \square

Although the theorem above is stated under the fixed-budget setting, one can easily change the failure rate δ' to $O(t^2/\delta)$ and enjoy an anytime regret bound.

G QGLOC without r

We present an alternative definition of m_t that does not rely on r :

$$\bar{m}'_{t-1} = \|\mathbf{x}^{\text{Greedy}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}, \text{ where } \mathbf{x}^{\text{Greedy}} := \arg \max_{\mathbf{x} \in \mathcal{X}_{t-1}} \langle \hat{\theta}_{t-1}, \mathbf{x} \rangle$$

We claim that $\bar{m}'_{t-1} \leq \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}$. To see this, suppose not: $\bar{m}'_{t-1} > \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}$. Then, $\mathbf{x}^{\text{Greedy}} \neq \mathbf{x}_t^{\text{QGLOC}}$. Furthermore, $\mathbf{x}^{\text{Greedy}}$ must be the maximizer of the QGLOC objective function while $\mathbf{x}_t^{\text{QGLOC}} \neq \mathbf{x}^{\text{Greedy}}$, which is a contradiction. The claim above allows all the proofs of our theorems on QGLOC to go through with \bar{m}'_{t-1} in place of \bar{m}_{t-1} .

However, computing \bar{m}'_{t-1} requires another hashing for finding the greedy arm defined above. Although this introduces a factor of 2 in the time complexity, it is quite cumbersome in practice, which is why we stick to m_{t-1} defined in (9) in the main text.

H Proof of Theorem 5

Let us first present the background. Denote by x_t^{GLOC} the solution of the optimization problem at line 3 of Algorithm 1. To find x_t^{GLOC} , one can fix \mathbf{x} and find the maximizer $\tilde{\boldsymbol{\theta}}(\mathbf{x}) := \max_{\boldsymbol{\theta} \in C_{t-1}} \mathbf{x}^\top \boldsymbol{\theta}$ in a closed form using the Lagrangian method:

$$\tilde{\boldsymbol{\theta}}(\mathbf{x}) = \hat{\boldsymbol{\theta}}_{t-1} + \sqrt{\beta_{t-1}} \cdot \frac{\bar{\mathbf{V}}_{t-1}^{-1} \mathbf{x}}{\|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}}, \quad (19)$$

which is how we obtained (7).

The following lemma shows that the objective function (10) plus $c_0 \beta^{3/4} \bar{m}_{t-1}$ is an upper bound of the GLOC's objective function (7). Note that this holds for any sequence $\{\beta_t\}$.

Lemma 5.

$$\begin{aligned} & \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle + \sqrt{\beta_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}} \\ & \leq \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \cdot \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1}. \end{aligned}$$

Furthermore,

$$\begin{aligned} & \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\text{GLOC}} \rangle + \sqrt{\beta_{t-1}} \|\mathbf{x}_t^{\text{GLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}} \\ & \leq \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\text{QGLOC}} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1}. \end{aligned} \quad (20)$$

Proof. Recall the GLOC optimization problem defined at line 3 of Algorithm 1. For a fixed \mathbf{x} , we need to solve

$$\begin{aligned} \tilde{\boldsymbol{\theta}}(\mathbf{x}) &= \arg \min_{\boldsymbol{\theta}} -\langle \boldsymbol{\theta}, \mathbf{x} \rangle \\ \text{s.t.} \quad & \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{t-1}\|_{\bar{\mathbf{V}}_{t-1}}^2 - \beta_{t-1} \leq 0 \end{aligned}$$

The Lagrangian is $\mathcal{L}(\boldsymbol{\theta}, \tau) := -\langle \boldsymbol{\theta}, \mathbf{x} \rangle + \tau(\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_{t-1}\|_{\bar{\mathbf{V}}_{t-1}}^2 - \beta_{t-1})$, and thus we need to solve

$$\max_{\tau \geq 0} \min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}, \tau).$$

According to the first-order optimality condition,

$$\begin{aligned} -\frac{\partial \mathcal{L}(\boldsymbol{\theta}, \tau)}{\partial \boldsymbol{\theta}} &= \mathbf{x} - \tau(2\bar{\mathbf{V}}_{t-1}\boldsymbol{\theta} - 2\bar{\mathbf{V}}_{t-1}\hat{\boldsymbol{\theta}}_{t-1}) = 0 \\ \boldsymbol{\theta} &= \hat{\boldsymbol{\theta}}_{t-1} + (2\tau)^{-1}\bar{\mathbf{V}}_{t-1}^{-1}\mathbf{x}, \end{aligned} \quad (21)$$

which results in $(\min_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}, \tau)) = -\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle - (4\tau)^{-1} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 - \tau\beta_{t-1}$. It remains to solve

$$\max_{\tau \geq 0} -\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle - (4\tau)^{-1} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 - \tau\beta_{t-1},$$

whose first-order optimality says that the solution is $\tau_* := (2\sqrt{\beta_{t-1}})^{-1} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}$. Plugging $\tau \leftarrow \tau_*$ in (21) leads to the solution $\tilde{\boldsymbol{\theta}}(\mathbf{x})$ defined in (19). Note that any choice of $\tau \geq 0$ leads to a lower bound on the Lagrangian $\mathcal{L}(\tilde{\boldsymbol{\theta}}(\mathbf{x}), \tau)$. Define

$$\tilde{\tau} := c_0 \beta_{t-1}^{-1/4} \bar{m}_{t-1}.$$

Then,

$$\begin{aligned} \mathcal{L}(\tilde{\boldsymbol{\theta}}(\mathbf{x}), \tilde{\tau}) &= -\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle - \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 - c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1} \\ &\leq \mathcal{L}(\tilde{\boldsymbol{\theta}}(\mathbf{x}), \tau_*) = -\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle - \sqrt{\beta_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}. \end{aligned}$$

This concludes the first part of the lemma. The second part of the lemma trivially follows from the first part by the definition (10). \square

Proof. Assume $\theta^* \in C_t^{\text{ONS}}$ for all $t \geq 1$, which happens w.p. at least $1 - \delta$. Suppose we pull $\mathbf{x}_t^{\text{QGLOC}}$ at every iteration. While we use the confidence set C_t^{ONS} , we omit the superscript ONS from β_t^{ONS} for brevity. Recall that $\bar{\beta}_t$ is an upper bound on β_t that is nondecreasing in t . We bound the instantaneous regret r_t as follows:

$$\begin{aligned}
\frac{r_t}{L} &= \frac{1}{L} \left(\mu(\langle \theta^*, \mathbf{x}_{t,*} \rangle) - \mu(\langle \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle) \right) \\
&\leq (\langle \theta^*, \mathbf{x}_{t,*} \rangle - \langle \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle) \\
&\leq \langle \tilde{\theta}(\mathbf{x}_t^{\text{GLOC}}, \mathbf{x}_t^{\text{GLOC}}) - \langle \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle \\
&= \langle \hat{\theta}_{t-1}, \mathbf{x}_t^{\text{GLOC}} \rangle + \sqrt{\bar{\beta}_{t-1}} \|\mathbf{x}_t^{\text{GLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}} - \langle \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle \\
&\stackrel{(\text{Lem. 5})}{\leq} \langle \hat{\theta}_{t-1}, \mathbf{x}_t^{\text{QGLOC}} \rangle + \frac{\bar{\beta}_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \bar{\beta}_{t-1}^{3/4} \bar{m}_{t-1} - \langle \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle \\
&= \langle \hat{\theta}_{t-1} - \theta^*, \mathbf{x}_t^{\text{QGLOC}} \rangle + \frac{\bar{\beta}_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \bar{\beta}_{t-1}^{3/4} \bar{m}_{t-1} \\
&\leq \underbrace{\|\hat{\theta}_{t-1} - \theta^*\|_{\bar{\mathbf{V}}_{t-1}^{-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}}_{=: A_1(t)} + \underbrace{\frac{\bar{\beta}_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2}_{=: A_2(t)} + \underbrace{c_0 \bar{\beta}_{t-1}^{3/4} \bar{m}_{t-1}}_{=: A_3(t)}.
\end{aligned}$$

Note that $\langle \theta^*, \mathbf{x} \rangle \in [-S, S]$ implies that $r_t \leq 2LS$. Then,

$$\begin{aligned}
r_t &\leq \min\{2LS, L \cdot A_1(t) + L \cdot A_2(t) + L \cdot A_3(t)\} \\
&\leq L \min\{2S, A_1(t)\} + \min\{2S, A_2(t)\} + \min\{2S, A_3(t)\}.
\end{aligned}$$

where the last inequality can be shown by a case-by-case analysis on each min operator.

Now, we consider computing $\sum_{t=1}^T r_t$. Using the same argument as the proof of Theorem 2 and Corollary 2,

$$L \sum_{t=1}^T \min\{2S, A_1(t)\} = O\left(\frac{L(L+R)}{\kappa} d\sqrt{T} \log^{3/2}(T)\right).$$

Then,

$$\begin{aligned}
&L \sum_{t=1}^T \min\left\{2S, \frac{\bar{\beta}_{T-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2\right\} \\
&\leq L \sum_{t=1}^T \min\left\{2S, \frac{\bar{\beta}_{T-1}^{1/4}}{4c_0 r} \sqrt{T + \lambda} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2\right\} \\
&\leq \frac{L \bar{\beta}_{T-1}^{1/4}}{4c_0 r} \sqrt{T + \lambda} \sum_{t=1}^T \min\left\{\frac{8c_0 r S}{\bar{\beta}_{T-1}^{1/4} \sqrt{T + \lambda}}, \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2\right\} \\
&\stackrel{(\text{Lem. 3})}{\leq} \frac{L \bar{\beta}_{T-1}^{1/4}}{4c_0 r} \sqrt{T + \lambda} \max\left\{2, \frac{8c_0 r S}{\bar{\beta}_{T-1}^{1/4} \sqrt{T + \lambda}}\right\} \sum_{t=1}^T \log(1 + \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2) \\
&\stackrel{(15)}{=} L \frac{1}{c_0} \cdot O\left(\left(\frac{L^2 + R^2}{\kappa^2} d \log^2 T\right)^{1/4}\right) \cdot O(\sqrt{T}) \cdot O(d \log T) \\
&= O\left(\frac{1}{c_0} L \left(\frac{L+R}{\kappa}\right)^{1/2} d^{5/4} \sqrt{T} \log^{3/2} T\right)
\end{aligned}$$

and

$$L \sum_{t=1}^T \min\{2S, c_0 \bar{\beta}_{t-1}^{3/4} \bar{m}_{t-1}\}$$

$$\begin{aligned}
&\stackrel{(9)}{\leq} L \sum_{t=1}^T \min\{2S, c_0 \bar{\beta}_{T-1}^{3/4} \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}\} \\
&\leq L c_0 \bar{\beta}_{T-1}^{3/4} \sum_{t=1}^T \min\left\{\frac{2S}{c_0 \bar{\beta}_{T-1}^{3/4}}, \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}\right\} \\
&\stackrel{(a)}{\leq} L c_0 \bar{\beta}_{T-1}^{3/4} \sqrt{T \sum_{t=1}^T \min\left\{\left(\frac{2S}{c_0 \bar{\beta}_{T-1}^{3/4}}\right)^2, \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2\right\}} \\
&\stackrel{(\text{Lem. 3})}{\leq} L c_0 \bar{\beta}_{T-1}^{3/4} \sqrt{T \max\left\{2, \left(\frac{2S}{c_0 \bar{\beta}_{T-1}^{3/4}}\right)^2\right\} \sum_{t=1}^T \log(1 + \|\mathbf{x}_t^{\text{QGLOC}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2)} \\
&\stackrel{(15)}{=} c_0 L \cdot O\left(\left(\frac{L^2 + R^2}{\kappa^2} d \log^2 T\right)^{3/4}\right) \cdot O(\sqrt{T d \log T}) \\
&= O\left(c_0 L \left(\frac{L+R}{\kappa}\right)^{3/2} d^{5/4} \sqrt{T} \log^2 T\right),
\end{aligned}$$

where (a) is due to the Cauchy-Schwartz inequality. Therefore, the regret is $O\left(L\left(\frac{1}{c_0}\left(\frac{L+R}{\kappa}\right)^{1/2} + c_0\left(\frac{L+R}{\kappa}\right)^{3/2}\right)d^{5/4}\sqrt{T}\log^2(T)\right)$. One can see that setting $c_0 = c'_0\left(\frac{L+R}{\kappa}\right)^{-1/2}$ leads to the stated regret bound. Note that one can improve $\log^2(T)$ in the regret bound to $\log^{7/4}(T)$ by making c_0 scale with $\log^{-1/4}(t)$, which is left as an exercise.

When the noise $|\eta_t|$ is bounded, we have a tighter β_t and thus we can replace $\log^2(T)$ in the regret bound to $\log^{5/4}(T)$. \square

I On c_0 of QGLOC

Observe that in (10) c_0 is a free parameter that adjusts the balance between the exploitation (the first term) and exploration (the second term). This is an interesting characteristic that is not available in existing algorithms but is attractive to practitioners. Specifically, in practice existing bandit algorithms like OFUL [1], LTS [5], and others [15, 41] usually perform exploration more than necessary, so one often enforces more exploitation by multiplying a small constant less than 1 to $\sqrt{\beta_t}$; e.g., see [40, 8]. Applying such a trick is theoretically not justified and foregoes the regret guarantee for existing algorithms, so a practitioner must take a leap of faith. In contrast, adjusting c_0 of QGLOC is exactly the common heuristic but now does not break the regret guarantee, which can assure practitioners.

J Details on Hashing

We first briefly introduce hashing methods for fast similarity search. Here, the similarity measure is often Euclidean distance [12] or inner product [35]. For a comprehensive review, we refer to Wang et al. [39]. The hashing methods build a hash table that consists of buckets where each bucket is identified by a unique hash key that is a sequence of k integers. At the hashing construction time, for each data point \mathbf{x} in the database we compute its hash key $h(\mathbf{x})$ using a function h (details shown below). We then organize the hash table so that each bucket contains pointers to the actual data points with the same hash key. The hash functions are decided at the construction time. Typically, for d' -dimensional hashing one draws k independent normally-distributed d' -dimensional vectors, which we call *projection vectors*. The hash function $h(\mathbf{x})$ outputs a discretized version of the inner product between these k vectors and the data point \mathbf{x} . When processing a query \mathbf{q} , we compute the hash key $h(\mathbf{q})$, retrieve the corresponding bucket, compute the similarities between the query and the data points therein, and pick the most similar one. It is important that one uses the same projection vectors for constructing the table and determining the hash key of the query. This means one needs to store the projection vectors. Finally, one typically constructs U independent hash tables to reduce the chance of missing very similar points (i.e., to increase the recall).

We now turn to operating QGLOC with hashing. Note that one would like to use an accurate hashing scheme since how accurately one solves (10) can impact the regret. However, a more accurate hashing have a higher space and time complexity. We characterize such a tradeoff in this section.

Recall that we aim to guarantee the accuracy of a MIPS hashing by the c_H -MIPS-ness. As we have mentioned in the main text, however, existing MIPS algorithms do not directly offer a c_H -MIPS guarantee. Instead of the c_H -MIPS guarantee, the standard MIPS algorithms provide a less convenient guarantee for an input parameter c and M as follows:

Definition 2. Let $\mathcal{X} \subseteq \mathbb{R}^{d'}$ such that $|\mathcal{X}| < \infty$. A data point $\tilde{\mathbf{x}} \in \mathcal{X}$ is called (c, M) -MIPS w.r.t. a given query \mathbf{q} if it satisfies $\langle \mathbf{q}, \tilde{\mathbf{x}} \rangle \geq cM$. An algorithm is called (c, M) -MIPS if, given a query $\mathbf{q} \in \mathbb{R}^{d'}$, it retrieves $\mathbf{x} \in \mathcal{X}$ that is (c, M) -MIPS w.r.t. \mathbf{q} whenever there exists $\mathbf{x}' \in \mathcal{X}$ such that $\langle \mathbf{q}, \mathbf{x}' \rangle \geq M$.

Note that when there is no such \mathbf{x} that $\langle \mathbf{q}, \mathbf{x} \rangle \geq M$, retrieving any arbitrary vector in \mathcal{X} is qualified as being (c, M) -MIPS. This also means that, by its contrapositive, if the hashing returns a vector that is not (c, M) -MIPS w.r.t. \mathbf{q} , then there is no \mathbf{x} such that $\langle \mathbf{q}, \mathbf{x} \rangle \geq M$ with high probability.

We emphasize that, to enjoy the (c, M) -MIPS-ness, a hashing must be built based on c and M . If we know $\max_{\mathbf{x} \in \mathcal{X}_t} \langle \mathbf{q}_t, \mathbf{x} \rangle$ ahead of time, then we can just set $M = \max_{\mathbf{x} \in \mathcal{X}_t} \langle \mathbf{q}_t, \mathbf{x} \rangle$ and build a hashing that is (c, M) -MIPS for some c , which gives a c -MIPS guarantee for the query \mathbf{q}_t . However, one does not have such information, and each query \mathbf{q}_t has its own “useful” value M .

To overcome such a difficulty, it seems natural to construct a c -MIPS hashing using multiple (c, M) -MIPS hashings with various M values covering a wide range. Indeed, such a construction is described in [18] for locality-sensitive hashing based on the Euclidean distance, which is complicated by multiple subroutines. We present here a streamlined construction of a c_H -MIPS hashing thanks to the existence of a high-probability upper and lower bound on the maximum of the QGLOC objective (10) as we show below. Note that one can derive a similar result for GLOC-TS.

Hereafter, we omit ONS from β_t^{ONS} . Note that there exists a simple upper bound $\bar{\beta}_t$ on β_t (see the proof of Corollary 2). Define $E_1(\delta)$ to be the event that θ^* belongs to C_t for all $t \geq 1$.

Lemma 6. Assume $E_1(\delta)$ and $\max_{\mathbf{x} \in \mathcal{X}_t} \langle \theta^*, \mathbf{x} \rangle \geq 1/2$. Suppose the target time horizon T is given. Then,

$$\begin{aligned} M_{\min} &:= 1/2 \leq \max_{t \in [T]} \max_{\mathbf{x} \in \mathcal{X}_t} \langle \hat{\theta}_{t-1}, \mathbf{x} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 \\ &\leq \sqrt{d}S + \bar{\beta}_{T-1}^{1/4} \cdot \frac{\sqrt{T+\lambda}}{4c_0 r \lambda} =: M_{\max}. \end{aligned}$$

Before presenting the proof, note that Lemma 6 assumes that $\max_{\mathbf{x} \in \mathcal{X}_t} \langle \theta^*, \mathbf{x} \rangle \geq 1/2$. In practice, this is not a restrictive assumption since it means that there exists at least one arm for which the reward is decently large in expectation. In interactive retrieval systems with binary user feedback, for example, it is reasonable to assume that there exists at least one item to which the user is likely to give a positive reward since otherwise any algorithm would work almost equally badly. One can change $1/2$ to any reasonable number v for which the reward $\mu(v)$ is considered high.

Proof. To show the lowerbound, recall that the objective function (10) is derived as an upperbound of the original GLOC objective function (7). Since the original GLOC objective function has θ^* as a feasible point by $E_1(\delta)$, $\max_{\mathbf{x} \in \mathcal{X}_t} \langle \theta^*, \mathbf{x} \rangle$ becomes a trivial lowerbound of the maximum of the GLOC objective and also of the maximum of QGLOC objective (10). This proves the lowerbound of the Lemma.

For the remaining part of the proof, we use notation $\mathbf{X}, \boldsymbol{\eta}, \mathbf{z}$ and $\bar{\mathbf{V}}$ in place of $\mathbf{X}_t, \boldsymbol{\eta}_t, \mathbf{z}_t$, and $\bar{\mathbf{V}}_t$ respectively (recall that \mathbf{z} is defined in Section 3).

It is easy to see that the eigenvalues of $\mathbf{X}\bar{\mathbf{V}}^{-2}\mathbf{X}^\top$ are all less than 1 (use the SVD of \mathbf{X}). Furthermore, $z_s = \mathbf{x}_s^\top \boldsymbol{\theta}_s \in [-S, S]$, and so $\|\mathbf{z}\|_2 \leq \sqrt{d}S$. Using these facts,

$$\begin{aligned} \|\hat{\boldsymbol{\theta}}_{t-1}\|_2 &= \|\bar{\mathbf{V}}^{-1}\mathbf{X}^\top\mathbf{z}\|_2 \\ &= \sqrt{(\mathbf{z})^\top \mathbf{X}\bar{\mathbf{V}}^{-2}\mathbf{X}^\top\mathbf{z}} \\ &= \|(\mathbf{X}\bar{\mathbf{V}}^{-2}\mathbf{X}^\top)^{1/2}\mathbf{z}\|_2 \\ &\leq \|\mathbf{z}\|_2 \leq \sqrt{d}S. \end{aligned}$$

Finally,

$$\begin{aligned} \max_{t \in [T]} \max_{\mathbf{x} \in \mathcal{X}_t} \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0\bar{m}_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 &\leq \max_{t \in [T]} \max_{\mathbf{x} \in \mathcal{X}_t} \|\hat{\boldsymbol{\theta}}_{t-1}\|_2 + \frac{\beta_{t-1}^{1/4}}{4c_0\bar{m}_{t-1}} \|\mathbf{x}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 \\ &\leq \max_{t \in [T]} \sqrt{d}S + \beta_{t-1}^{1/4} \cdot \frac{\sqrt{t+\lambda}}{4c_0r} \cdot \frac{1}{\lambda} \\ &\leq \sqrt{d}S + \beta_{T-1}^{1/4} \cdot \frac{\sqrt{T+\lambda}}{4c_0r} \cdot \frac{1}{\lambda}. \end{aligned}$$

□

Given a target approximation level $c_H < 1$, we construct a c_H -MIPS as follows. Define

$$J := \left\lceil \log_{1/\sqrt{c_H}} \frac{M_{\max}}{M_{\min}} \right\rceil = O(\log(dT)/\log(c_H^{-1})). \quad (22)$$

We build a series of J MIPS hashing schemes that are

$$(c_H^{1/2}, c_H^{j/2} M_{\max})\text{-MIPS for } j \in [J]. \quad (23)$$

We say that the MIPS hashing *succeeds (fails)* for a query \mathbf{q} if the retrieved vector is (not) (c, M) -MIPS w.r.t. \mathbf{q} . Theorem 6 shows that one can perform a binary search to find a vector $\mathbf{x} \in \mathcal{X}$ that is c_H -MIPS.

Theorem 6. *Upon given a query \mathbf{q} , perform a binary search over the J MIPS hashings (23) to find the smallest $j^* \in [J]$ for which the retrieved vector $\mathbf{x}^{(j^*)}$ from the j^* -th hashing succeeds. Then, $\mathbf{x}^{(j^*)}$ is c_H -MIPS w.r.t. \mathbf{q} with high probability.*

Proof. Assume the event that a retrieved vector $\mathbf{x}^{(j)}$ from j -th MIPS satisfies $(c_H^{1/2}, c_H^{j/2} M_{\max})$ -MIPS for $j \in [J]$, which happens with high probability.

Define the maximum $M^* := \max_{\mathbf{x} \in \mathcal{X}_t} \langle \mathbf{q}, \mathbf{x} \rangle$. The result of the binary search is that, for query \mathbf{q} , j^* -th MIPS succeeds but $(j^* - 1)$ -th MIPS fails. By the definition of $(c_H^{1/2}, c_H^{(j^*-1)/2} M_{\max})$ -MIPS, the fact that $(j^* - 1)$ -th hashing fails implies $M^* < c_H^{(j^*-1)/2} M_{\max}$. Then,

$$\langle \mathbf{q}, \mathbf{x}^{(j^*)} \rangle \geq c_H^{1/2} \cdot c_H^{j^*/2} M_{\max} = c_H \cdot c_H^{(j^*-1)/2} M_{\max} > c_H M^*.$$

□

Among various (c, M) -MIPS algorithms [35, 36, 30, 17], we adopt Shrivastava *et al.* [35]. Shrivastava *et al.* propose a reduction of MIPS to locality-sensitive hashing and present a result that their algorithm is (c, M) -MIPS with $O(N^{\rho^*} \log N)$ inner product computations and space $O(N^{1+\rho^*})$ for an optimized value ρ^* that is guaranteed to be less than 1; see [35, Theorem 5] for detail.

Let d' be the dimensionality of the projection vectors, where $d' = d^2 + d$ for QGLOC and $d' = d$ for GLOC-TS. One can recover the order of the space and time complexity stated in the main text by $O(\log(J)N^{\rho^*} \log(N)d')$ and $O(JN^{\rho^*}(N + \log(N)d'))$, respectively.

K The Regret Bound of QGLOC under Hashing Approximation Error

For most recommendation or interactive retrieval applications, it is reasonable to assume that the total number of steps T is bounded by a known constant (“fixed budget” in bandit terminology) since users do not interact with the system for too long. We present the regret of QGLOC combined with

MIPS hashing in Theorem 7. The theorem states that in the fixed budget setting T , we can set the target approximation level c_H as a function of T and enjoy the same order of regret as exactly solving the maximization problem (10). The proof is presented at the end of this section.

Theorem 7. *Let $T \geq 2$ and $c_H = \left(1 + \frac{\log(T)}{\sqrt{T}}\right)^{-1}$. Suppose we run QGLOC for T iterations where we invoke a c_H -MIPS hashing algorithm \mathcal{H} to approximately find the solution of (10) at each time $t \in [T]$. Denote this algorithm by $\text{QGLOC}(\mathcal{H})$. Assume that, w.p. at least $1 - \delta$, the hashing \mathcal{H} successfully retrieves a c_H -MIPS solution for every T queries made by $\text{QGLOC}(\mathcal{H})$. Then, w.p. at least $1 - 3\delta$, $\text{Regret}_T^{\text{QGLOC}(\mathcal{H})} = \hat{O}(\kappa^{-1}L(L+R)d^{5/4}\sqrt{T}\log^{7/4}(T))$.*

Note that we have made an assumption that, w.h.p., the hashing retrieves a c_H -MIPS solution for every T queries made by the algorithm. One naive way to construct such a hashing is to build T independent hashing schemes as follows, which is the first low-regret GLB algorithm with time sublinear in N , to our knowledge.

Corollary 3. *Let $T \geq 2$ and $c_H = \left(1 + \frac{\log(T)}{\sqrt{T}}\right)^{-1}$. Suppose we build T independent hashings where each hashing is c_H -MIPS w.p. at least $1 - \delta/T$. Suppose we run QGLOC for T iterations where at time t we use t -th MIPS hashing to solve (10). Then, w.p. at least $1 - 3\delta$, the regret bound is $\hat{O}(\kappa^{-1}L(L+R)d^{5/4}\sqrt{T}\log^{7/4}(T))$.*

Proof. The probability that at least one of T hashing schemes fails to output a c_H -MIPS solution is at most δ . Combining this with Theorem 7 completes the proof. \square

However, it is easy to see that its space complexity is $\Omega(T)$, which is not space-efficient. Of course, a better way would be to construct one hashing scheme that is, w.h.p., c_H -MIPS for any sequence of T queries. However, this is nontrivial for the following reason. In bandit algorithms, the second query depends on the result of the first query. Here, the result of the query \mathbf{q}_1 is obtained based on the hash keys computed using the *projection vectors*. Now, the second query \mathbf{q}_2 is based on the result of querying \mathbf{q}_1 , which means that \mathbf{q}_2 now correlates with the projection vectors. This breaks the *independence* of the query with the projection vectors, thus breaking the hashing guarantee. One way to get around the issue is the union bound. This is possible when there exists a finite set of possible queries, which is indeed how [18] manage to show a guarantee on such ‘adaptive’ queries. In our case, unfortunately, there are infinitely many possible queries, and thus the union bound does not apply easily. Resolving the issue above is of theoretical interest and left as future work.

Meanwhile, it is hardly the case that the correlation between \mathbf{q}_2 and the projection vectors has an malignant effect in practice. Indeed, many existing studies using hashing such as Jain et al. [22] ignore the correlation issue and do not provide any guarantee on the adaptive queries.

K.1 Proof of Theorem 7

Proof. Assume that $\boldsymbol{\theta}^*$ belongs to C_t^{ONS} , which happens w.h.p. In this proof, we drop QGLOC from $\mathbf{x}_t^{\text{QGLOC}}$ and use \mathbf{x}_t . Denote by $\mathbf{x}_t^{\mathcal{H}}$ the solution returned by the MIPS algorithm. We omit ONS from β_t^{ONS} to avoid clutter. Then, being c_H -MIPS guarantees that, $\forall t \in [T]$,

$$\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\mathcal{H}} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0\bar{m}_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 \geq c_H \left(\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0\bar{m}_{t-1}} \|\mathbf{x}_t\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 \right). \quad (24)$$

To avoid clutter, we use \mathbf{X} , \mathbf{y} , and $\boldsymbol{\eta}$ in place of \mathbf{X}_{t-1} , \mathbf{y}_{t-1} , and $\boldsymbol{\eta}_{t-1}$, respectively, when it is clear from the context. Note that, using the techniques in the proof of Theorem 2,

$$\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\mathcal{H}} \rangle = \langle \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle + \langle \hat{\boldsymbol{\theta}}_{t-1} - \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle \leq S + \sqrt{\beta_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}. \quad (25)$$

The instantaneous regret at time t divided by L is

$$\begin{aligned} \frac{r_t}{L} &\leq \langle \boldsymbol{\theta}^*, \mathbf{x}_{t,*} \rangle - \langle \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle \\ &\leq \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0\bar{m}_{t-1}} \|\mathbf{x}_t\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0\beta_{t-1}^{3/4}\bar{m}_{t-1} - \langle \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle \end{aligned}$$

$$\begin{aligned}
&\stackrel{(24)}{\leq} \frac{1}{c_H} \left(\langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\mathcal{H}} \rangle + \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 \right) + c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1} - \langle \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle \\
&= \left(\frac{1}{c_H} - 1 \right) \langle \hat{\boldsymbol{\theta}}_{t-1}, \mathbf{x}_t^{\mathcal{H}} \rangle + \langle \hat{\boldsymbol{\theta}}_{t-1} - \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle + \frac{1}{c_H} \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1} \\
&\stackrel{(25)}{\leq} \underbrace{\frac{\log T}{\sqrt{T}} S}_{A_1(t)} + \underbrace{\frac{\log T}{\sqrt{T}} \sqrt{\beta_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}}_{A_2(t)} + \underbrace{\langle \hat{\boldsymbol{\theta}}_{t-1} - \boldsymbol{\theta}^*, \mathbf{x}_t^{\mathcal{H}} \rangle + \frac{1}{c_H} \frac{\beta_{t-1}^{1/4}}{4c_0 \bar{m}_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}^2 + c_0 \beta_{t-1}^{3/4} \bar{m}_{t-1}}_{A_3(t)}.
\end{aligned}$$

Note that $r_t \leq 2LS$. Since $1/c_H \leq 2$, it is not hard to see that $\sum_{t=1}^T \min\{2LS, L \cdot A_3(t)\}$ is

$$O(\kappa^{-1} L(L+R) d^{5/4} \sqrt{T} \log^{7/4}(T))$$

using the same technique as the proof of Theorem 5. It remains to bound $\sum_{t=1}^T \min\{2LS, L \cdot A_1(t)\}$ and $\sum_{t=1}^T \min\{2S, L \cdot A_2(t)\}$:

$$\begin{aligned}
L \sum_{t=1}^T \min\{2S, A_1(t)\} &\leq L \sum_{t=1}^T \min\left\{2S, \frac{\log T}{\sqrt{T}} 2S\right\} \\
&\stackrel{(\log T \leq \sqrt{T})}{\leq} L \sum_{t=1}^T \frac{\log T}{\sqrt{T}} 2S \\
&= O(L \sqrt{T} \log T) \\
L \sum_{t=1}^T \min\{2S, A_2(t)\} &= L \sum_{t=1}^T \min\left\{2S, \frac{\log T}{\sqrt{T}} \sqrt{\beta_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}\right\} \\
&\leq L \sum_{t=1}^T \min\left\{2S, \sqrt{\beta_{t-1}} \|\mathbf{x}_t^{\mathcal{H}}\|_{\bar{\mathbf{V}}_{t-1}^{-1}}\right\}
\end{aligned}$$

Using the same argument as the proof of Theorem 2 and Corollary 2,

$$\sum_{t=1}^T \min\{2S, L \cdot A_2(t)\} = O\left(\frac{L(L+R)}{\kappa} d \sqrt{T} \log^{3/2}(T)\right).$$

Altogether, we notice that $\sum_{t=1}^T \min\{2S, L \cdot A_3(t)\}$ dominates the other terms. Notice that we have spent δ probability to control the event $\boldsymbol{\theta}^* \in C_t^{\text{ONS}}$, another δ for controlling the deviation of g_t which appears in the regret bound through β_t as shown in the proof of Theorem 2, and another δ for ensuring the hashing guarantee. This sums to 3δ and concludes the proof. \square

L Proof of Lemma 1

Proof. We first recall that for L1 sampling:

$$p_i = \frac{|q_i|}{\|\mathbf{q}\|_1}$$

Note that:

$$\begin{aligned}
|G_k| &= \left| \frac{q_{i_k} a_{i_k}}{p_{i_k}} \right| \\
&= \frac{|q_{i_k} a_{i_k}|}{|q_{i_k}| / \|\mathbf{q}\|_1} \\
&= \|\mathbf{q}\|_1 |a_{i_k}| \\
&\leq \|\mathbf{q}\|_1 \|\mathbf{a}\|_{\max} =: M
\end{aligned}$$

This means that G_k is a bounded random variable. Therefore, we can use Hoeffding's inequality to get a high-probability bound:

$$X_i = G_i - \mathbf{q}^\top \mathbf{a}$$

$$\begin{aligned}\mathbb{P}\left(\left|\frac{1}{m}\sum_{i=1}^m X_i\right| \geq \epsilon\right) &\leq 2\exp\left(-\frac{m\epsilon^2}{2M^2}\right) \\ &\leq 2\exp\left(-\frac{m\epsilon^2}{2\|\mathbf{q}\|_1^2 \|\mathbf{a}\|_{\max}^2}\right)\end{aligned}$$

□

M Comparison between L1 and L2 sampling

We first show the known high probability error bound of L2, which has a polynomially decaying probability of failure.

Lemma 7. [22, Lemma 3.4] Define G_k as in (11) with $\mathbf{p} = \mathbf{p}^{(L2)}$. Then, given a target error $\epsilon > 0$,

$$\mathbb{P}\left(\left|\frac{1}{m}\sum_{k=1}^m G_k - \mathbf{q}^\top \mathbf{a}\right| \geq \epsilon\right) \leq \frac{\|\mathbf{q}\|_2^2 \|\mathbf{a}\|_2^2}{m\epsilon^2}. \quad (26)$$

Proof. This proof is a streamlined version of the proof of Lemma 3.4 from [22]. Note that

$$\text{Var}G_k \leq \|\mathbf{q}\|^2 \|\mathbf{a}\|^2.$$

Define

$$X_m := \sum_{i=1}^m (G_i - \mathbf{q}^\top \mathbf{a})$$

Now, since $\mathbf{q}^\top \mathbf{a}$ is deterministic,

$$\text{Var}X_m = \sum_{i=1}^m \text{Var}G_i \leq m \|\mathbf{q}\|^2 \|\mathbf{a}\|^2$$

and,

$$\mathbb{E}X_m = 0$$

Let us apply Chebyshev's inequality to X_m :

$$\begin{aligned}\mathbb{P}(|X_m - \mathbb{E}X_m| \geq \alpha) &\leq \frac{\text{Var}X_m}{\alpha^2} \\ \Rightarrow \mathbb{P}(|X_m| \geq m\epsilon) &\leq \frac{m \|\mathbf{q}\|^2 \|\mathbf{a}\|^2}{m^2 \epsilon^2} \\ \Rightarrow \mathbb{P}\left(\left|\frac{1}{m}\sum_{i=1}^m G_i - \mathbf{q}^\top \mathbf{a}\right| \geq \epsilon\right) &\leq \frac{\|\mathbf{q}\|^2 \|\mathbf{a}\|^2}{m\epsilon^2} \\ \Rightarrow \mathbb{P}\left(\left|\frac{1}{m}\sum_{i=1}^m G_i - \mathbf{q}^\top \mathbf{a}\right| \geq \epsilon' \|\mathbf{q}\| \|\mathbf{a}\|\right) &\leq \frac{\|\mathbf{q}\|^2 \|\mathbf{a}\|^2}{m\epsilon'^2 \|\mathbf{q}\|^2 \|\mathbf{a}\|^2}, \text{ where } \epsilon = \epsilon' \|\mathbf{q}\| \|\mathbf{a}\| \\ \Rightarrow \mathbb{P}(|\tilde{\mathbf{q}}^\top \mathbf{a} - \mathbf{q}^\top \mathbf{a}| \geq \epsilon' \|\mathbf{q}\| \|\mathbf{a}\|) &\leq \frac{1}{m\epsilon'^2} \\ \Rightarrow \mathbb{P}(|\tilde{\mathbf{q}}^\top \mathbf{a} - \mathbf{q}^\top \mathbf{a}| \geq \epsilon' \|\mathbf{q}\| \|\mathbf{a}\|) &\leq \frac{1}{c}, \text{ where } m = c/\epsilon'^2 \\ \Rightarrow \mathbb{P}(|\tilde{\mathbf{q}}^\top \mathbf{a} - \mathbf{q}^\top \mathbf{a}| \geq \epsilon' \|\mathbf{q}\| \|\mathbf{a}\|) &\geq 1 - \frac{1}{c}\end{aligned}$$

□

To compare the concentration of measure of L1 and L2, we look at so-called ‘‘sample complexity’’ of L1 and L2. Let δ' be the target failure probability (set the RHS of (13) to δ'). Then, the sample complexity of L2 and L1 is $\frac{\|\mathbf{q}\|_2^2 \|\mathbf{a}\|_2^2}{\delta' \epsilon^2}$ and $\log(2/\delta') \frac{2\|\mathbf{q}\|_1^2 \|\mathbf{a}\|_{\max}^2}{\epsilon^2}$, respectively. Note L2 has a smaller scaling with \mathbf{q} since $\|\mathbf{q}\|_2 \leq \|\mathbf{q}\|_1$, but L1 has a better scaling with \mathbf{a} since $\|\mathbf{a}\|_{\max} \leq \|\mathbf{a}\|_2$. More importantly, the concentration bound of L2 decays polynomially with m whereas that of L1 decays exponentially. However, it is unclear whether or not the polynomial tail of L2 is just an artifact of analysis. In fact, we find that L2 has an exponential tail bound, but its scaling with \mathbf{q} can be very bad.

While we state the result below in Lemma 8, the key here is that the magnitude of G_k induced by L2 is at most $\|\mathbf{q}\|_2^2 \max_i |a_i|/|q_i|$ while that induced by L1 is at most $\|\mathbf{q}\|_1 \|\mathbf{a}\|_{\max}$. As q_i goes to zero, the support of the L2-based G_k can be arbitrarily large. Unless one knows that q_i is sufficiently bounded away from 0 (which is false in QGLOC), L2-based G_k raises a concern of having a “thick” tail.⁹

Lemma 8. Use $\mathbf{p} = \mathbf{p}^{(L2)}$. Then,

$$\mathbb{P} \left(\left| \frac{1}{m} \sum_{i=1}^m X_i \right| \geq \epsilon \right) \leq 2 \exp \left(- \frac{m\epsilon^2}{2 \|\mathbf{q}\|_2^4 \max_i |a_i/q_i|^2} \right)$$

Proof. From the results in Lemma 3.4 of [22] (restated as Lemma 7 in our paper), it seems that L2 sampling has polynomial tails, which is considered as heavy-tailed. Here we show that the tail of L2 sampling does not decay polynomially but exponentially by using Hoeffding’s inequality instead of Chebychev’s inequality. However, the scale that controls the tail thickness is quite bad, and the tail can be arbitrarily thick regardless of the norm or variance of \mathbf{q} .

We first recall that for L2 sampling:

$$p_i = \frac{q_i^2}{\|\mathbf{q}\|_2^2}$$

Let us first show an upper bound for $|G_k|$:

$$\begin{aligned} |G_k| &= \left| \frac{q_{i_k} a_{i_k}}{p_{i_k}} \right| \\ &= \frac{|q_{i_k} a_{i_k}|}{q_{i_k}^2 / \|\mathbf{q}\|_2^2} \\ &= \|\mathbf{q}\|_2^2 \left| \frac{a_{i_k}}{q_{i_k}} \right| \\ &\leq \|\mathbf{q}\|_2^2 \max_i \left| \frac{a_i}{q_i} \right| =: M' \end{aligned}$$

Recall that $\mathbb{E}G_k = 0$ and from above, $|G_k| \leq M'$. We can now apply Hoeffding’s inequality:

$$\begin{aligned} X_i &= G_i - \mathbf{q}^\top \mathbf{a} \\ \mathbb{P} \left(\left| \frac{1}{m} \sum_{i=1}^m X_i \right| \geq \epsilon \right) &\leq 2 \exp \left(- \frac{m\epsilon^2}{2M'^2} \right) \\ &\leq 2 \exp \left(- \frac{m\epsilon^2}{2 \|\mathbf{q}\|_2^4 \max_i \left| \frac{a_i}{q_i} \right|^2} \right) \end{aligned}$$

This is an exponential tail bound compared to the polynomial tail bound in Lemma 3.4 in [22]. We do note that the term $\max_i \left| \frac{a_i}{q_i} \right|$ could be very bad if $\min_i q_i$ is small and the corresponding a_i is non-zero. \square

The second comparison is on the variance of L1 and L2. The variance of G_k based on $\mathbf{p}^{(L2)}$ can be shown to be $\|\mathbf{q}\|_2^2 \|\mathbf{a}\|_2^2 - (\mathbf{q}^\top \mathbf{a})^2$. With L1, the variance of G_k is now $\|\mathbf{q}\|_1 (\sum_{i=1}^d |q_i| a_i^2) - (\mathbf{q}^\top \mathbf{a})^2$ whose first term can be upper-bounded by $\|\mathbf{q}\|_1 \|\mathbf{q}\|_2 \|\mathbf{a}\|_4^2$ using Cauchy-Schwartz. This means that the variance induced by L1 scales larger with \mathbf{q} than by L2 since $\|\mathbf{q}\|_2 \leq \|\mathbf{q}\|_1$ and scales smaller with \mathbf{a} since $\|\mathbf{a}\|_4^2 \leq \|\mathbf{a}\|_2^2$. Thus, neither is an absolute winner. However, if the vectors being inner-producted are normally distributed, then L1 has a smaller variance than L2 in most cases, for large enough d as we show in the lemma below. As mentioned in the main text, our projection vectors are truly normally distributed.

⁹ Still, the distribution does not belong to so-called “heavy-tailed” distributions.

Lemma 9. Suppose that $\mathbf{q} \sim \mathcal{N}(0, \mathbf{I}_d)$ and $\mathbf{a} \sim \mathcal{N}(0, \mathbf{I}_d)$ and that \mathbf{q} and \mathbf{a} are independent of each other. Then $\mathbb{E} \left(\|\mathbf{q}\|_1 \sum_{i=1}^d |q_i| a_i^2 \right) \leq \mathbb{E} \left(\|\mathbf{q}\|_2^2 \|\mathbf{a}\|_2^2 \right)$ for large enough d .

Proof. First note that if $\mathbf{x} \sim \mathcal{N}(0, I_d)$, then:

$$\mathbb{E} \|\mathbf{x}\|_2^2 = d.$$

Then, the RHS is

$$\begin{aligned} \mathbb{E} \|\mathbf{q}\|_2^2 \|\mathbf{a}\|_2^2 &= \mathbb{E} \|\mathbf{q}\|_2^2 \mathbb{E} \|\mathbf{a}\|_2^2 \\ &= d^2. \end{aligned}$$

For the LHS:

$$\begin{aligned} \mathbb{E} \left(\|\mathbf{q}\|_1 \sum_{i=1}^d |q_i| a_i^2 \right) &= \mathbb{E}_{\mathbf{q}} \left(\|\mathbf{q}\|_1 \sum_{i=1}^d |q_i| \mathbb{E}_{\mathbf{a}} a_i^2 \right) \\ &= \mathbb{E}_{\mathbf{q}} \left(\|\mathbf{q}\|_1 \sum_{i=1}^d |q_i| \right) \\ &= \mathbb{E}_{\mathbf{q}} \|\mathbf{q}\|_1^2 \\ &= \mathbb{E}_{\mathbf{q}} \left(\sum_{i=1}^d |q_i|^2 + 2 \sum_{i,j=1, i \neq j}^d |q_i| |q_j| \right) \\ &= \mathbb{E}_{\mathbf{q}} \sum_{i=1}^d |q_i|^2 + 2 \sum_{i,j=1, i \neq j}^d \mathbb{E}_{\mathbf{q}} |q_i| |q_j| \\ &= d + 2 \frac{2}{\pi} \frac{d(d-1)}{2} \\ &= d + \frac{2}{\pi} d(d-1) \\ &= \frac{2}{\pi} d^2 + \frac{\pi-2}{\pi} d. \end{aligned}$$

where we used the fact that $\sum_{i=1}^d |q_i|^2$ follows a Chi-squared distribution with d degrees of freedom and that $|q_i|$ follows a half-normal distribution, whose mean is $\sqrt{2/\pi}$.

Since $2/\pi < 1$, for large d , $\frac{2}{\pi} d^2 + \frac{\pi-2}{\pi} d < d^2$. □

N Hashing Implementation

We implement hashing for QGLOC and GLOC-TS based on python package OptimalLSH¹⁰. In practice, building a series of hashings with varying parameter shown in Section J can be burdensome. Departing from the theory, we use so-called multi-probe technique [37] that achieves a similar effect by probing nearby buckets. That is, upon determining a hash key and retrieving its corresponding bucket, one can also look at other hash keys that are different on only one component. We use $k = 12$ keys and $U = 24$ tables.

¹⁰<https://github.com/yahoo/Optimal-LSH>