Communication Efficient Federated Learning for Generalized Linear Bandits

Anonymous Author(s)
Affiliation
Address
email

Abstract

Contextual bandit algorithms have been recently studied under the federated learning setting to satisfy the demand of keeping data decentralized and pushing the learning of bandit models to the client side. But limited by the required communication efficiency, existing solutions are restricted to linear models to exploit their closed-form solutions for parameter estimation. Such a restricted model choice greatly hampers these algorithms' practical utility. In this paper, we take the first step to addressing this challenge by studying generalized linear bandit models under the federated learning setting. We propose a communication-efficient solution framework that employs online regression for local update and offline regression for global update. We rigorously proved, though the setting is more general and challenging, our algorithm can attain sub-linear rate in both regret and communication cost, which is also validated by our extensive empirical evaluations.

1 Introduction

As a classic model for sequential decision making problems, contextual bandit has been widely used for a variety of real-world applications, including recommender systems [19], display advertisement [21] and clinical trials [7]. While most existing bandit solutions are designed under a centralized setting (i.e., data is readily available at a central server), in response to the increasing application scale and public concerns of privacy, there is increasing research effort on federated bandit learning lately [28, 6, 27, 11, 17], where N clients collaborate with limited communication bandwidth to minimize the overall cumulative regret incurred over a finite time horizon T, while keeping each client's raw data local. Compared with standard federated learning [23, 13] that works with fixed datasets, federated bandit learning is characterized by its online interactions with the environment, which continuously provides new data samples to the clients over time. This brings in new challenges in addressing the conflict between the need of timely data/model aggregation for regret minimization and the need of communication efficiency with decentralized data. A carefully designed model update method and communication strategy become vital to strike this balance.

Existing federated bandit learning solutions only partially addressed this challenge by considering simple bandit models, like context-free bandit [27] and contextual linear bandit [28, 6, 17], where closed-form solution for both local and global model update exists. Therefore, efficient communication for global bandit model update is realized by directly aggregating local sufficient statistics, such that the only concern left is how to control the communication frequency over time horizon T. However, such a solution framework does not apply to the more complicated bandit models that are often preferred in practice, such as generalized linear bandit (GLB) [8] or neural bandit [30], where only iterative solutions exist for parameter estimation (e.g., gradient-based optimization). To enable joint model estimation, now the learning system needs to solve distributed optimization for multiple times as new data is collected from the environment, and each requires iterative gradient/model aggregation

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

among clients. This is much more expensive compared with linear models, and it naturally leads to the question: whether a communication efficient solution to this challenging problem is still possible?

In this paper, we answer this question affirmatively by proposing the first provably communication efficient algorithm for federated GLB that only requires $\tilde{O}(\sqrt{T})$ communication cost, while still attaining the optimal order of regret. Our proposed algorithm employs a combination of online and offline regression, with online regression adjusting each client's model using its newly collected data, and offline (distributed) regression occasionally soliciting local gradients from all N clients for joint model estimation when sufficient amount of new data has been accumulated. In order to balance exploration and exploitation in arm selection, we propose a novel way to construct the confidence set based on the sequence of *offline-and-online* model updates that each client has received. The initialization of online regression with offline regression introduces dependencies that break the standard martingale argument, which requires proof techniques unique to this paper.

We also explored other non-trivial solution ideas to further justify our current design. Specifically, in practice, a common way to update the deployed model for applications with streaming data is to set a schedule and periodically re-train the model using iterative optimization methods. For comparison, we propose and rigorously analyze a federated GLB algorithm designed based on this idea, as well as a variant that further enables online updates on the clients. We also consider another solution idea motivated by distributed/batched online convex optimization, which is characterized by lazy online updates over batches of data. Moreover, extensive empirical evaluations on both synthetic and real-world datasets are performed to validate the effectiveness of our algorithm.

2 Related Work

GLB, as an important extension of linear bandit models, has demonstrated encouraging performance in modeling binary rewards (such as clicks) that are ubiquitous in real-world applications [18]. The study of GLB under a centralized setting dates back to Filippi et al. [8], who proposed a UCB-type algorithm that achieved $\tilde{O}(d\sqrt{T})$ regret. Li et al. [20] later proposed two improvements: a similar UCB-type algorithm that improves the result of [8] by a factor of $O(\log T)$, which has been popularly used in practice as it avoids the projection step needed in [8]; and another impractical algorithm that further improves the result by a factor of $O(\sqrt{d})$ assuming fixed number of arms. To improve the time and space complexity of the aforementioned GLB algorithms, followup works adopted online regression methods. In particular, motivated by the online-to-confidence-set conversion technique from [2], Jun et al. [12] proposed both UCB and Thompsan sampling algorithms with online Newton step, and Ding et al. [4] proposed a Thompson sampling algorithm with online gradient descent, which, however, requires an additional context regularity assumption to obtain a sub-linear regret.

GLB under federated/distributed setting still remains under-explored. The most related works are the federated/distributed linear bandits [16, 28, 6, 11, 17]. In these works, thanks to the existence of closed-form solution for linear models, the clients only communicate their local sufficient statistics for global model update. Korda et al. [16] considered a peer-to-peer (P2P) communication network and assumed the clients form clusters, i.e., each cluster is associated with a unique bandit problem. But as they only focused on reducing *per-round* communication, the communication cost is still linear over time. Huang et al. [11] considered a star-shaped communication network as in our paper, but their proposed phase-based elimination algorithm only works in fixed arm set setting. The closest works to ours are [28, 6, 17], which uses event-triggered communication protocols to obtain sub-linear communication cost over time for federated linear bandit with a time-varying arm set.

Another related line of research is the standard federated learning that considers offline supervised learning problems [13]. Since its debut in [23], FedAvg has become the most popularly used algorithm for offline federated learning. However, despite its popularity, several works [22, 14, 24] identified that FedAvg suffers from a *client-drift* problem when the clients' data are non-IID (which is an important signature of our case), i.e., local iterates in each client drift towards their local minimum. This leads to a sub-optimal convergence rate of FedAvg: for example, one has to suffer a sub-linear convergence rate for strongly convex and smooth losses, though a linear convergence rate is expected under a centralized setting. To alleviate this, Pathak and Wainwright [26] proposed an operator splitting procedure to guarantee linear convergence to a neighborhood of the global minimum. Later, Mitra et al. [24] introduced variance reduction techniques to guarantee exact linear convergence to the global minimum.

3 Preliminaries

In this section, we first introduce the general problem formulation of federated bandit learning, and discuss the existing solutions under the linear reward assumption. Then we formulate the federated GLB problem considered in this paper, followed by detailed discussions about the new challenges compared with its linear counterpart.

3.1 Federated Bandit Learning

Consider a learning system with 1) N clients responsible for taking actions and receiving corresponding reward feedback from the environment, e.g., each client being an edge device directly interacting with a user, and 2) a central server responsible for coordinating the communication between the clients for joint model estimation.

At each time step $t = 1, 2, ..., T$, all N clients interact with the environment in a round-robin manner, i.e., each client $i \in [N]$ chooses an arm $\mathbf{x}_{t,i}$ from its time-varying candidate set $\mathcal{A}_{t,i} = \{\mathbf{x}_{t,i}^{(1)}, \mathbf{x}_{t,i}^{(2)}, \dots, \mathbf{x}_{t,i}^{(K)}\}$, where $\mathbf{x}_{t,i}^{(a)} \in \mathbb{R}^d$ denotes the context vector associated with the a-th arm for client i at time t. Without loss of generality, we assume $||\mathbf{x}_{t,i}^{(a)}||_2 \leq 1, \forall i, a, t$. Then client i receives the corresponding reward $y_{t,i} \in \mathbb{R}$ from the environment, which is drawn from the reward distribution governed by an unknown parameter $\theta_\star \in \mathbb{R}^d$ (assume $\|\theta_\star\| \leq S$), i.e., $y_{t,i} \sim p_{\theta_\star}(y|\mathbf{x}_{t,i}^{(a)})$. The interaction between the learning system and the environment repeats itself, and the goal of the learning system is to minimize the cumulative (pseudo) regret over all N clients in the finite time horizon T, i.e., $R_T = \sum_{t=1}^{T} \sum_{i=1}^{N} r_{t,i}$, where $r_{t,i} = \max_{\mathbf{x} \in \mathcal{A}_{t,i}} \mathbf{E}[y|\mathbf{x}] - \mathbf{E}[y_{t,i}|\mathbf{x}_{t,i}]$.

In a federated learning setting, the clients cannot directly communicate with each other, but through the central server, i.e., a star-shaped communication network. Raw data collected by each client $i \in [N]$, i.e., $\{(\mathbf{x}_{s,i}, y_{s,i})\}_{s \in [T]}$, is stored locally and cannot be shared with anyone else. Instead, the clients can only communicate the parameters of the learning algorithm, e.g., models, gradients, or sufficient statistics; and the communication cost is measured by the total number of times data being transferred across the system up to time T, which is denoted as C_T.

3.2 Federated Linear Bandit

Prior works have studied communication-efficient federated linear bandit [28, 6], i.e., the reward function is a linear model $y_{t,i} = \mathbf{x}_{t,i}^\top \theta_\star + \eta_{t,i}$, where $\eta_{t,i}$ denotes zero-mean sub-Gaussian noise. Consider an imaginary centralized agent that has direct access to the data of all clients, so that it can compute the global sufficient statistics $A_t = \sum_{i \in [N]} \sum_{s \in [t]} \mathbf{x}_{s,i} \mathbf{x}_{s,i}^\top, b_t = \sum_{i \in [N]} \sum_{s \in [t]} \mathbf{x}_{s,i} y_{s,i}$. Then the cumulative regret incurred by this distributed learning system can match that under a centralized setting, if all N clients select arms based on the global sufficient statistics $\{A_t, b_t\}$. However, it requires $N^2 T$ communication cost for the immediate sharing of each client's update to the sufficient statistics with all other clients, which is expensive for most applications.

To ensure communication efficiency, prior works like DisLinUCB [28] let each client i maintain a local copy $\{A_{t-1,i}, b_{t-1,i}\}$ for arm selection, which receives immediate local update using each newly collected data sample, i.e., $A_{t,i} = A_{t-1,i} + \mathbf{x}_{t,i} \mathbf{x}_{t,i}^\top, b_{t,i} = b_{t-1,i} + \mathbf{x}_{t,i} y_{t,i}$. Then client i checks whether the event $(t - t_{\text{last}}) \log(\frac{\det A_{t,i}}{\det A_{t_{\text{last}}}}) > D$ is true, where t_{last} denotes the time step of last global update. If true, a new global update is triggered, such that the server will collect all clients' local update since t_{last}, aggregate them to compute $\{A_t, b_t\}$, and then synchronize the local sufficient statistics of all clients, i.e., set $\{A_{t,i}, b_{t,i}\} = \{A_t, b_t\}, \forall i \in [N]$.

3.3 Federated Generalized Linear Bandit

In this paper, we study federated bandit learning with generalized linear models, i.e., the conditional distribution of reward y given context vector $\mathbf{x}$ is drawn from the exponential family [8, 20]:

$$p_{\theta_\star}(y|\mathbf{x}) = \exp\left(\frac{y\mathbf{x}^\top \theta_\star - m(\mathbf{x}^\top \theta_\star)}{g(\tau)} + h(y, \tau)\right) \tag{1}$$

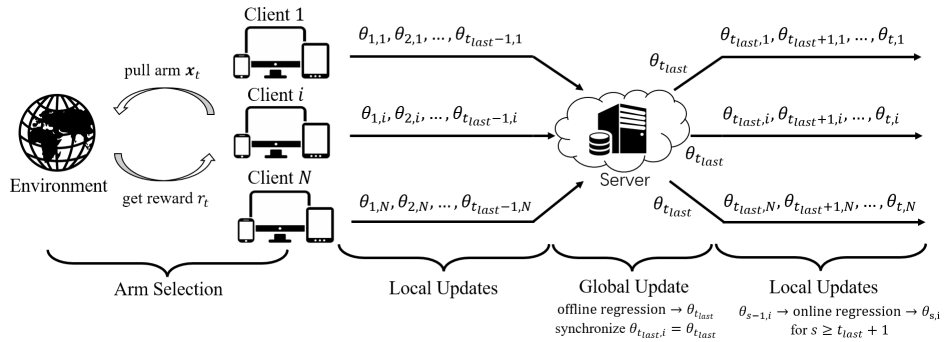

Figure 1: Illustration of FedGLB-UCB algorithm, which uses online regression for local update, i.e., immediately update each client's local model $\theta_{t,i}$ using its newly collected data sample, and uses offline regression for global update, i.e., synchronize all N clients to a globally updated model $\theta_{t_{\text{last}}}$ using all the data samples collected so far.

where $\tau \in \mathbb{R}^+$ is a known scale parameter. Given a function $f : \mathbb{R} \to \mathbb{R}$, we denote its first and second derivatives by $\dot{f}$ and $\ddot{f}$, respectively. It is known that $\dot{m}(\mathbf{x}^\top \theta_\star) = \mathbb{E}[y|\mathbf{x}] := \mu(\mathbf{x}^\top \theta_\star)$, which is called the inverse link function, and $\ddot{m}(\mathbf{x}^\top \theta_\star) = \mathbb{V}(y|\mathbf{x}^\top \theta_\star)$. Based on Eq.(1), the reward $y_{t,i}$ observed by client i at time t can be equivalently represented as $y_{t,i} = \mu(\mathbf{x}_{t,i}^\top \theta_\star) + \eta_{t,i}$, where $\eta_{t,i}$ denotes the sub-Gaussian noise. Then we denote the negative log-likelihood of $y_{i,t}$ given $\mathbf{x}_{i,t}$ as $l(\mathbf{x}_{t,i}^\top \theta_\star, y_{t,i}) = -\log p_{\theta_\star}(y_{t,i}|\mathbf{x}_{t,i}) = -y_{t,i}\mathbf{x}_{t,i}^\top \theta_\star + m(\mathbf{x}_{t,i}^\top \theta_\star)$. In addition, we adopt the following two assumptions about the reward, which are standard for GLB [8].

Assumption 1. *The link function μ is continuously differentiable on $(-S, S)$, k_μ-Lipschitz on $[-S, S]$, and $\inf_{z \in [-S,S]} \dot{\mu}(z) = c_\mu > 0$.*

Assumption 2. $\mathbb{E}[\eta_{t,i}|\mathcal{F}_{t,i}] = 0, \forall t, i$, *where $\mathcal{F}_{t,i} = \sigma\{\mathbf{x}_{t,i}, [\mathbf{x}_{s,j}, y_{s,j}]_{(s,j):s<t \cap j=i}\}$ denotes the σ-algebra generated by client i's previously pulled arms and observed rewards, and $\max_{t,i} |\eta_{t,i}| \leq R_{\max}$ for some constant $R_{\max} > 0$.*

New Challenges Compared with federated linear bandit discussed in Section 3.2, new challenges arise in designing a communication-efficient algorithm for federated GLB due to the absence of a closed form solution:

- *Iterative communication for global update:* compared with the global update for federated linear bandit that only requires one round of communication to share the sufficient statistics, now it takes multiple iterations of gradient aggregation to obtain converged global optimization. Moreover, as the clients collect more data samples over time during bandit learning, the required number of iterations for convergence also increases.
- *Drifting issue with local update:* during local model update, iterative optimization using only local gradient can push the updated model away from the global model, i.e., forget the knowledge gained during previous communications [15].

4 Methodology

In this section we propose the first algorithm for federated GLB that addresses the aforementioned challenges. We rigorously prove that it attains sub-linear rate in T for both regret and communication cost. In addition, we propose and analyze different variants of our algorithm to facilitate understanding of our algorithm design.

4.1 FedGLB-UCB Algorithm

To ensure communication-efficient model updates for federated GLB, we propose to use online regression for local update, i.e., update each client's local model only with its newly collected data samples, and use offline regression for global update, i.e., solicit all clients' local gradients for joint model estimation. Based on the resulting sequence of *offline-and-online* model updates, the confidence ellipsoid for $\theta_\star$ is constructed for each client to select arms using the OFUL principle.

We name this algorithm Federated Generalized Linear Bandit with Upper Confidence Bound, or FedGLB-UCB for short. We illustrate its key components in Figure 1 and describe its procedures in Algorithm 1. In the following, we discuss about each component of FedGLB-UCB in details.

Algorithm 1 FedGLB-UCB

1: **Input:** threshold D, regularization parameter $\lambda > 0$, $\delta \in (0, 1)$ and c_μ.
2: **Initialize** $\forall i \in [N]$: $A_{0,i} = \frac{\lambda}{c_\mu}\mathbf{I} \in \mathbb{R}^{d\times d}, b_{0,i} = \mathbf{0} \in \mathbb{R}^d, \theta_{0,i} = \mathbf{0} \in \mathbb{R}^d, \Delta A_{0,i} = \mathbf{0} \in \mathbb{R}^{d\times d}$; $A_0 = \frac{\lambda}{c_\mu}\mathbf{I} \in \mathbb{R}^{d\times d}, b_0 = \mathbf{0} \in \mathbb{R}^d, \theta_0 = \mathbf{0} \in \mathbb{R}^d, t_{\text{last}} = 0$
3: **for** $t = 1, 2, ..., T$ **do**
4: **for** client $i = 1, 2, ..., N$ **do**
5: Observe arm set $\mathcal{A}_{t,i}$ for client i
6: Select arm $\mathbf{x}_{t,i} \in \mathcal{A}_{t,i}$ by Eq.(5), and observe reward $y_{t,i}$
7: Update client i: $A_{t,i} = A_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top$, $\Delta A_{t,i} = \Delta A_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top$
8: **if** $(t - t_{\text{last}}) \log \frac{\det(A_{t,i})}{\det(A_{t,i}-\Delta A_{t,i})} < D$ **then**
9: **Client** i: perform local update $\theta_{t,i}$ = ONS-Update$(\theta_{t-1,i}, A_{t,i}, \nabla l(\mathbf{x}_{t,i}^\top\theta_{t-1,i}, y_{t,i}))$, $b_{t,i} = b_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top\theta_{t-1,i}$
10: **else**
11: **Clients** $\forall i \in [N]$: send $\Delta A_{t,i}$ to server, and reset $\Delta A_{t,i} = \mathbf{0}$
12: **Server**: compute $A_t = A_{t_{\text{last}}} + \sum_{i=1}^N \Delta A_{t,i}$
13: **Server**: perform global update θ_t = AGD-Update$(\theta_{t_{\text{last}}}, J_t)$ (see Eq.(3) for the choice of J_t), $b_t = b_{t_{\text{last}}} + \sum_{i=1}^N \Delta A_{t,i}\theta_t$, and set $t_{\text{last}} = t$
14: **Clients** $\forall i \in [N]$: set $\theta_{t,i} = \theta_t, A_{t,i} = A_t, b_{t,i} = b_t$

• **Local update.** As mentioned earlier, iterative optimization over local dataset $\{(\mathbf{x}_{s,i}, y_{s,i})\}_{s\in[t]}$ leads to the drifting issue that pushes the updated model to the local optimum. Due to the small size of this local dataset, the confidence ellipsoid centered at the converged model has increased width, which leads to increased regret in bandit learning. However, as we will prove in Section 4.2, completely disabling local update and restricting all clients to use the previous globally updated model for arm selection is also a bad choice, because the learning system will then need more frequent global updates to adapt to the growing dataset.

To enable local update while alleviating the drifting issue, we adopt online regression in each client, such that the local model estimation $\theta_{t,i}$ is only updated for one step using the sample $(\mathbf{x}_{t,i}, y_{t,i})$ collected at time t. Prior works [2, 12] showed that UCB-type algorithms with online regression can attain comparable cumulative regret to the standard UCB-type algorithms [1, 20], as long as the selected online regression method guarantees logarithmic online regret. As the negative log-likelihood loss defined in Section 3.3 is exp-concave and online Newton step (ONS) is known to attain logarithmic online regret in this case [10, 12], ONS is chosen for the local update of FedGLB-UCB and its description is given in Appendix A. At time step t, after client i pulls an arm $\mathbf{x}_{t,i} \in \mathcal{A}_{t,i}$ and observes the reward $y_{t,i}$, its model $\theta_{t-1,i}$ is immediately updated by the ONS update rule (line 9 in Algorithm 1), where $\nabla l(\mathbf{x}_{t,i}^\top\theta_{t-1,i}, y_{t,i})$ denotes the gradient w.r.t. $\theta_{t-1,i}$, and $A_{t,i}$ denotes the covariance matrix for client i at time t.

• **Global update** The global update of FedGLB-UCB requires communication among the N clients, which imposes communication cost in two aspects: 1) each global update for federated GLB requires multiple rounds of communication among N clients, i.e., iterative aggregation of local gradients; and 2) global update needs to be performed for multiple times over time horizon T, in order to adapt to the growing dataset collected by each client during bandit learning. Consider a particular time step $t \in [T]$ when global update happens, the distributed optimization objective is:

$$\min_{\theta\in\Theta} F_t(\theta) := \frac{1}{N}\sum_{i=1}^N F_{t,i}(\theta) \tag{2}$$

where $F_{t,i}(\theta) = \frac{1}{t}\sum_{s=1}^t l(\mathbf{x}_{s,i}^\top\theta, y_{s,i}) + \frac{\lambda}{2t}||\theta||_2^2$ denotes the *average* regularized negative log-likelihood loss for client $i \in [N]$, and $\lambda > 0$ denotes the regularization parameter. Based on Assumption 1, $\{F_{t,i}(\theta)\}_{i\in[N]}$ are $\frac{\lambda}{Nt}$-strongly-convex and $(k_\mu + \frac{\lambda}{Nt})$-smooth in θ (proof in Appendix

B), and we denote the unique minimizer of Eq.(2) as $\hat{\theta}_t^{\text{MLE}}$. In this case, it is known that the number of communication rounds J_t required to attain a specified sub-optimality ϵ_t, such that $F_t(\theta) - \min_{\theta\in\Theta} F_t(\theta) \leq \epsilon_t$, has a lower bound $J_t = \Omega\left(\sqrt{(k_\mu Nt)/\lambda + 1}\log\frac{1}{\epsilon_t}\right)$ [3], which means J_t increases at least at the rate of $\sqrt{Nt}$. This lower bound is matched by the distributed version of accelerated gradient descent (AGD) [25]:

$$J_t \leq 1 + \sqrt{(k_\mu Nt)/\lambda + 1}\log\frac{(k_\mu + \frac{2\lambda}{Nt})\|\theta_t^{(1)} - \hat{\theta}_t^{\text{MLE}}\|_2^2}{2\epsilon_t} \tag{3}$$

where the superscript (i) denotes the i-th iteration of AGD.

In order to minimize the number of communication rounds in one global update, AGD is chosen as the offline regression method for FedGLB-UCB, and its description is given in Appendix A. However, other federated/distributed optimization methods can be readily used in place of AGD, as our analysis only requires the convergence result of the adopted method. We should note that ϵ_t is essential to the regret-communication trade-off during the global update at time t: a larger ϵ_t leads to a wider confidence ellipsoid, which increases regret, while a smaller ϵ_t requires more communication rounds J_t, which increases communication cost. In Section 4.2, we will discuss about the proper choice of ϵ_t to attain desired trade-off between the two conflicting objectives.

To reduce the total number of global updates over time horizon T, we adopt the event-triggered communication from [28], such that global update is triggered if the following event is true for any client $i \in [N]$ (line 8):

$$(t - t_{\text{last}})\log\frac{\det(A_{t,i})}{\det(A_{t,i} - \Delta A_{t,i})} > D \tag{4}$$

where $\Delta A_{t,i}$ denotes client i's local update to its covariance matrix since last global update at t_{last}, and $D > 0$ is the chosen threshold for the event-trigger. During the global update, the model estimation $\theta_{t,i}$, covariance matrix $A_{t,i}$ and vector $b_{t,i}$ for all clients $i \in [N]$ will be updated (line 11-14). We should note that the LHS of Eq.(4) is essentially an upper bound of the cumulative regret that client i's locally updated model has incurred since t_{last}. Therefore, this event-trigger guarantees that a global update only happens when effective regret reduction is possible.

• **Arm selection** To balance exploration and exploitation during bandit learning, FedGLB-UCB uses the OFUL principle for arm selection [1], which requires the construction of a confidence ellipsoid for each client i. We propose a novel construction of the confidence ellipsoid based on the sequence of model updates that each client i has received up to time t: basically, there are 1) one global update at t_{last}, i.e., the joint offline regression across all clients' accumulated data till t_{last}: $\{(\mathbf{x}_{s,i}, y_{s,i})\}_{s\in[t_{\text{last}}],i\in[N]}$, which resets all clients' local models to $\theta_{t_{\text{last}}}$; and 2) multiple local updates from $t_{\text{last}} + 1$ to t, i.e., the online regression on client i's own data sequence $\{(\mathbf{x}_{s,i}, y_{s,i})\}_{s\in[t_{\text{last}}+1,t]}$ to get $\{\theta_{s,i}\}_{s\in[t_{\text{last}}+1,t]}$ step by step. This can be more easily understood by the illustration in Figure 1. The resulting confidence ellipsoid is centered at the ridge regression estimator $\hat{\theta}_{t,i} = A_{t,i}^{-1}b_{t,i}$ [2, 12], which is computed using the predicted rewards given by the past sequence of model updates $\{\theta_{t_{\text{last}}}\} \cup \{\theta_{s,i}\}_{s\in[t_{\text{last}}+1,t]}$ (see the update of $b_{t,i}$ in line 9 and 13 of Algorithm 1). Then at time step t, client i selects the arm that maximizes the UCB score:

$$\mathbf{x}_{t,i} = \arg\max_{\mathbf{x}\in\mathcal{A}_{t,i}} \mathbf{x}^\top\hat{\theta}_{t-1,i} + \alpha_{t-1,i}||\mathbf{x}||_{A_{t-1,i}^{-1}} \tag{5}$$

where $\alpha_{t-1,i}$ is the parameter of the confidence ellipsoid given in Lemma 4.2. Note that compared with standard federated/distributed learning where clients only need to communicate gradients for joint model estimation, in our problem, due to the time-varying arm set, it is also necessary to communicate the confidence ellipsoid among clients, i.e., $A_t \in \mathbb{R}^{d\times d}$ and $b_t \in \mathbb{R}^d$ (line 14 in Algorithm 1), as the clients need to be prepared for all possible arms $\mathbf{x} \in \mathbb{R}^d$ that may appear in future for the sake of regret minimization.

4.2 Theoretical Analysis

In this section, we construct the confidence ellipsoid based on the *offline-and-online* estimators described in Section 4.1. Then we analyze the cumulative regret and communication cost of FedGLB-UCB, followed by theoretical comparisons with its different variants.

• **Construction of confidence ellipsoid** Compared with prior works that convert a sequence of online regression estimators to confidence ellipsoid [2, 12], our confidence ellipsoid is built on the combination of an offline regression estimator $\theta_{t_{\text{last}}}$ for global update, and the subsequent online regression estimators $\{\theta_{s-1,i}\}_{s\in[t_{\text{last}}+1,t]}$ for local updates on each client i. This construction is new and requires proof techniques unique to our proposed solution. In the following, we highlight the key steps, and refer our readers to the appendix for details.

To simplify the use of notations, we assume without loss of generality that the global update at t_{last} is triggered by the N-th client, such that no more new data will be collected at t_{last}, i.e., the first data sample obtained after the global update has index $t_{\text{last}}+1$. We start our construction by considering the following loss difference introduced by the global and local model updates: $\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}},y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big] + \sum_{s=t_{\text{last}}+1}^{t}\big[l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big]$, where the first term is the loss difference between the globally updated model $\theta_{t_{\text{last}}}$ and $\theta_\star$, and the second term is between the sequence of locally updated models $\{\theta_{s-1,i}\}_{s\in[t_{\text{last}}+1,t]}$ and $\theta_\star$. This extends the definition of online regret used in the construction in [2, 12]; and due to the existence of offline regression, the obtained upper bounds in Lemma 4.1 are unique to our solution.

Lemma 4.1 (Upper Bound of Loss Difference)**.** *Denote the sub-optimality of the global model update procedure at time step t_{last} as $\epsilon_{t_{last}}$, such that $F_{t_{last}}(\theta) - \min_{\theta\in\mathcal{B}_d(S)} F_{t_{last}}(\theta) \le \epsilon_{t_{last}}$. Then under Assumption 1 and 2, we have*

$$\sum_{s=1}^{t_{last}}\sum_{i=1}^{N}\big[l(\mathbf{x}_{s,i}^\top\theta_{t_{last}},y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big] \le B_1 \tag{6}$$

where $B_1 = Nt_{last}\epsilon_{t_{last}} + \frac{\lambda}{2}S^2$, and with probability at least $1-\delta$,

$$\sum_{s=t_{last}+1}^{t}\big[l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big] \le B_2 \tag{7}$$

where $B_2 = \frac{1}{2c_\mu}\sum_{s=t_{last}+1}^{t}\|\nabla l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})\|^2_{A^{-1}_{s,i}} + \frac{c_\mu}{2}\Big[\frac{1}{c_\mu}R_{\max}\sqrt{d\log(1+Nt_{last}c_\mu/d\lambda)+2\log(1/\delta)} + Nt_{last}\sqrt{\frac{2k_\mu}{\lambda c_\mu}+\frac{2}{Nt_{last}c_\mu}}\sqrt{\epsilon_{t_{last}}} + \sqrt{\frac{\lambda}{c_\mu}}S\Big]^2$, respectively.

Specifically, B_1 corresponds to the convergence of the offline (distributed) optimization in previous global update; B_2 is essentially the online regret upper bound of ONS, with the major difference that it is initialized using the globally updated model $\theta_{t_{\text{last}}}$, instead of an arbitrary model as in standard ONS. Then due to the c_μ-strongly-convexity of $l(z,y)$ w.r.t. z, i.e., $l(\mathbf{x}_s^\top\theta,y_s) - l(\mathbf{x}_s^\top\theta_\star,y_s) \ge \big[\mu(\mathbf{x}_s^\top\theta_\star)-y_s\big]\mathbf{x}_s^\top(\theta-\theta_\star) + \frac{c_\mu}{2}\big[\mathbf{x}_s^\top(\theta-\theta_\star)\big]^2$, and by rearranging terms in Eq.(6) and Eq.(7), we have: $\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2 \le \frac{2}{c_\mu}B_1 + \frac{2}{c_\mu}\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)$, and $\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)\big]^2 \le \frac{2}{c_\mu}B_2 + \frac{2}{c_\mu}\sum_{s=t_{\text{last}}+1}^{t}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)$, whose LHS is quadratic in $\theta_\star$. To further upper bound the RHS, we should note that the term $\frac{2}{c_\mu}\sum_{s=t_{\text{last}}+1}^{t}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)$ is standard in [2, 12] as $\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)$ is $\mathcal{F}_{s,i}$-measurable for online estimator $\theta_{s-1,i}$. However, this is not true for the term $\frac{2}{c_\mu}\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)$ as the offline regression estimator $\theta_{t_{\text{last}}}$ depends on all data samples collected till t_{last}; and thus we have to develop a different approach to bound it. This leads to Lemma 4.2 below, which provides the confidence ellipsoid for $\theta_\star$.

Lemma 4.2 (Confidence Ellipsoid of FedGLB-UCB)**.** *With probability at least $1-2\delta$, for all $t\in[T], i\in[N]$,*

$$\|\hat{\theta}_{t,i}-\theta_\star\|^2_{A_{t,i}} \le \beta_{t,i} + \frac{\lambda}{c_\mu}S^2 - \|\boldsymbol{z}_{t,i}\|_2^2 + \hat{\theta}_{t,i}^\top b_{t,i} := \alpha_{t,i}^2$$

where $\boldsymbol{z}_{t,i}$ denotes the vector of predicted rewards $[\mathbf{x}_{1,1}^\top\theta_{t_{last}},\mathbf{x}_{1,2}^\top\theta_{t_{last}},\ldots,\mathbf{x}_{t_{last},N-1}^\top\theta_{t_{last}},\mathbf{x}_{t_{last},N}^\top\theta_{t_{last}},\mathbf{x}_{t_{last}+1,i}^\top\theta_{t_{last},i},\mathbf{x}_{t_{last}+2,i}^\top\theta_{t_{last}+1,i},\ldots,\mathbf{x}_{t,i}^\top\theta_{t-1,i}]^\top$, and $\beta_{t,i} = \frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{last}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big) + B_1 + \frac{4R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{last}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}\Big(\|\theta_{t_{last}}\|_2 + \|\theta_\star\|_2 + \sqrt{B_1}\Big) + \frac{4B_2}{c_\mu} + \frac{8R_{\max}^2}{c_\mu^2}\log\Big(\frac{N}{\delta}\sqrt{4+\frac{8}{c_\mu}B_2+\frac{64R_{\max}^4}{c_\mu^4\cdot 4\delta^2}}\Big) + 1.$

Table 1: Comparison between FedGLB-UCB and its variants with different design choices.

Global Upd.	Local Upd.	Setting	R_T	C_T
AGD	ONS	$D = \frac{T}{Nd\log(NT)}$	$\frac{k_\mu(k_\mu+R_{\max})}{c_\mu} d\sqrt{NT}\log(NT)$	$dN^2\sqrt{T}\log^2(NT)$
AGD	no update	$B = \sqrt{NT}$	$\frac{k_\mu R_{\max}}{c_\mu} d\sqrt{NT}\log(NT)$	$N^2 T\log(NT)$
AGD	ONS	$B = d^2 N\log(NT)$	$\frac{k_\mu(k_\mu+R_{\max})}{c_\mu} d\sqrt{NT}\log(NT)\log(T)$	$d^2N^{2.5}\sqrt{T}\log^2(NT)$
ONS	ONS	$B = \sqrt{NT}$	$\frac{k_\mu(k_\mu+R_{max})}{c_\mu} d(NT)^{3/4}\log(NT)$	$N^{1.5}\sqrt{T}$

• **Regret and communication cost** From Lemma 4.2, we can see that $\alpha_{t,i}$ grows at a rate of $Nt_{\text{last}}\sqrt{\epsilon_{t_{\text{last}}}}$ through its dependence on the B_2 term. To make sure the growth rate of $\alpha_{t,i}$ matches that in standard GLB algorithms [20, 12], we set $\epsilon_{t_{\text{last}}} = \frac{1}{N^2 t_{\text{last}}^2}$, which leads to the following corollary.

Corollary 4.2.1 (Order of $\beta_{t,i}$)**.** *With* $\epsilon_{t_{last}} = \frac{1}{N^2 t_{last}^2}$, $\beta_{t,i} = O\big(\frac{d\log NT}{c_\mu^2}[k_\mu^2 + R_{\max}^2]\big)$.

Then using a similar argument as the proof for Theorem 4 of [28], we obtain the following upper bounds on R_T and C_T for FedGLB-UCB (proof in Appendix E).

Theorem 4.3 (Regret and Communication Cost Upper Bound of FedGLB-UCB)**.** *Under Assumption 1, 2, and by setting* $\epsilon_t = \frac{1}{N^2t^2}, \forall t$ *and* $D = \frac{T}{Nd\log(NT)}$*, the cumulative regret* R_T *of FedGLB-UCB has an upper bound* $R_T = O\left(\frac{k_\mu(k_\mu+R_{\max})}{c_\mu} d\sqrt{NT}\log(NT/\delta)\right)$*, with probability at least* $1-2\delta$*. The corresponding cumulative communication cost has an upper bound* $C_T = O\Big(dN^2\sqrt{T}\log^2(NT)\Big)$.

Theorem 4.3 shows that FedGLB-UCB recovers the standard $O\big(d\sqrt{NT}\log(NT)\big)$ rate in regret as in the centralized setting, while only incurring a communication cost that is sub-linear in T. Note that, to obtain $O\big(d\sqrt{NT}\log(NT)\big)$ regret for federated linear bandit, the DisLinUCB algorithm incurs a communication cost of $O(dN^{1.5}\log(NT))$ [28], which is smaller than that of FedGLB-UCB by a factor of $\sqrt{NT}\log(NT)$. As the frequency of global updates is the same for both algorithms (due to their use of the same event-trigger), this additional communication cost is caused by the iterative optimization procedure for the global update, which is required for GLB model estimation. Moreover, as we mentioned in Section 4.1, there is not much room for improvement here as the use of AGD already matches the lower bound up to a logarithmic factor.

To facilitate the understanding of our algorithm design and investigate the impact of different components of FedGLB-UCB on its regret and communication efficiency trade-off, we propose and analyze three variants, which are also of independent interest, and report the results in Table 1. Detailed descriptions, as well as proof for these results can be found in Appendix F. Note that all three variants perform global update according to a fixed schedule $\mathcal{S} = \{t_1 := \lfloor\frac{T}{B}\rfloor, t_2 := 2\lfloor\frac{T}{B}\rfloor, \ldots, t_B := B\lfloor\frac{T}{B}\rfloor\}$, where B denotes the total number of global updates specified in advance to trade-off between R_T and C_T, and these variants differ in their global and local update strategies.

5 Experiments

We performed extensive empirical evaluations of FedGLB-UCB on both synthetic and real-world datasets, and the results (averaged over 10 runs) are reported in Figure 2 and Figure 3, respectively. We included the three variants of FedGLB-UCB (listed in Table 1), One-UCB-GLM, N-UCB-GLM [20] and N-ONS-GLM [12] as baselines, where One-UCB-GLM learns a shared bandit model across all clients, and N-UCB-GLM and N-ONS-GLM learn a separated bandit model for each client with no communication. Additional results and discussions about experiments can be found in Appendix G.

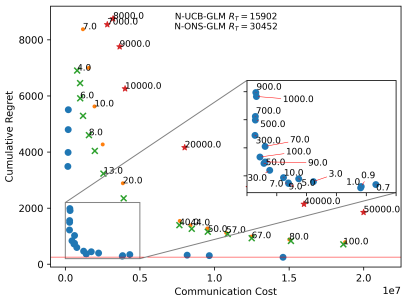

Figure 2: Synthetic dataset

• **Synthetic Dataset** We simulated the federated GLB setting defined in Section 3.3, with $T = 2000, N = 200, d = 10, S = 1$, $\mathcal{A}_t$ ($K = 25$) uniformly sampled from a ℓ_2 unit sphere, and reward

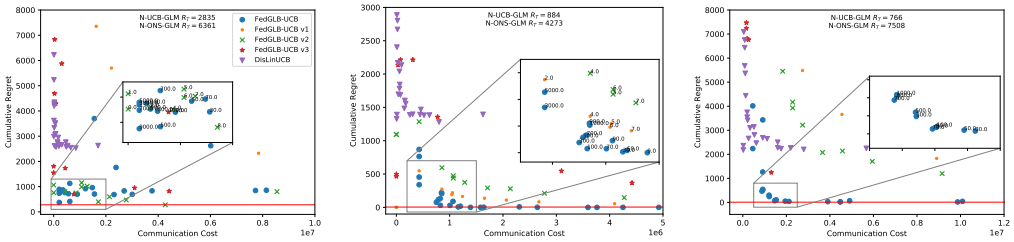

(a) CoverType (b) MagicTelescope (c) Mushroom

Figure 3: Experiment results on real world datasets.

$y_{t,,i} \sim \text{Bernoulli}(\mu(\mathbf{x}_{t,,i}^\top \theta_\star))$, with $\mu(z) = (1 + \exp(-z))^{-1}$. To compare the algorithms' R_T and C_T under different trade-off settings, we run FedGLB-UCB with different threshold value D (logarithmically spaced between 10^{-1} and 10^3) and its variants with different number of global updates B. Note that each dot in the result figure illustrates the C_T (x-axis) and R_T (y-axis) that a particular instance of FedGLB-UCB or its variants obtained by time T, and the corresponding value for D or B is labeled next to the dot. R_T of One-UCB-GLM is illustrated as the red horizontal line, and R_T of N-UCB-GLM and N-ONS-GLM are labeled on the top of the figure. We can observe that for FedGLB-UCB and its variants, R_T decreases as C_T increases, interpolating between the two extreme cases: independently learned bandit models by N-UCB-GLM, N-ONS-GLM; and the jointly learned bandit model by One-UCB-GLM. FedGLB-UCB significantly reduces C_T, while attaining low R_T, i.e., its regret is even comparable with One-UCB-GLM that requires at least $C_T = N^2T$ (8×10^7 in this simulation) for gradient aggregation at each time step.

• **Real-world Dataset** The results above demonstrate the effectiveness of FedGLB-UCB when data is generated by a well-specified generalized linear model. To evaluate its performance in a more challenging and practical scenario, we performed experiments using real-world datasets: CoverType, MagicTelescope and Mushroom from the UCI Machine Learning Repository [5]. To convert them to contextual bandit problems, we pre-processed these datasets following the steps in prior works [8], with $T = 2000$ and $N = 20$. Moreover, to demonstrate the advantage of GLB over linear model, we included DisLinUCB [28] as an additional baseline. Since the parameters being communicated in DisLinUCB and FedGLB-UCB are different, to ensure a fair comparison of C_T in this experiment, we measure communication cost (x-axis) by the number of integers or real numbers transferred across the learning system (instead of the frequency of communications). Note that DisLinUCB has no $C_T \geq 3 \times 10^6$in Figure 3 because its global update is already happening in every round and cannot be increased further. As mentioned earlier, due to the difference in messages being sent, the communication in DisLinUCB's *per global update* is much smaller than that in FedGLB-UCB. However, because linear models failed to capture the complicated reward mappings in these three datasets, we can see that DisLinUCB is clearly outperformed by FedGLB-UCB and its variants. This shows that, by offering a larger variety of modeling choices, e.g., linear, Poisson, logistic regression, etc., FedGLB-UCB has more potential in dealing with the complicated data in real-world applications.

Conclusion

In this paper, we take the first step to address the new challenges in communication efficient federated bandit learning beyond linear models, where closed-form solutions do not exist, and propose a solution framework for federated GLB that employs online regression for local update and offline regression for global update. For arm selection, we propose a novel confidence ellipsoid construction based on the sequence of *offline-and-online* model estimations. We rigorously prove that the proposed algorithm attains sub-linear rate for both regret and communication cost, and also analyze the impact of each component of our algorithm via theoretical comparison with different variants. In addition, extensive empirical evaluations are performed to validate the effectiveness of our algorithm.

An important further direction of this work is the lower bound analysis for the communication cost, analogous to the communication lower bound for standard distributed optimization by Arjevani and Shamir [3]. Moreover, in our algorithm, clients' locally updated models are not utilized for global model update, so that another interesting direction is to investigate whether using such knowledge, e.g., by model aggregation, can further improve communication efficiency.

References

[1] Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.

[2] Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2012.

[3] Yossi Arjevani and Ohad Shamir. Communication complexity of distributed convex learning and optimization. *arXiv preprint arXiv:1506.01900*, 2015.

[4] Qin Ding, Cho-Jui Hsieh, and James Sharpnack. An efficient algorithm for generalized linear bandit: Online stochastic gradient descent and thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 1585–1593. PMLR, 2021.

[5] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.

[6] Abhimanyu Dubey and AlexSandy' Pentland. Differentially-private federated linear bandits. *Advances in Neural Information Processing Systems*, 33, 2020.

[7] Audrey Durand, Charis Achilleos, Demetris Iacovides, Katerina Strati, Georgios D Mitsis, and Joelle Pineau. Contextual bandits for adapting treatment in a mouse model of de novo carcinogenesis. In *Machine learning for healthcare conference*, pages 67–82. PMLR, 2018.

[8] Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *NIPS*, volume 23, pages 586–594, 2010.

[9] Elad Hazan. Introduction to online convex optimization. *arXiv preprint arXiv:1909.05207*, 2019.

[10] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

[11] Ruiquan Huang, Weiqiang Wu, Jing Yang, and Cong Shen. Federated linear contextual bandits. *Advances in Neural Information Processing Systems*, 34, 2021.

[12] Kwang-Sung Jun, Aniruddha Bhargava, Robert Nowak, and Rebecca Willett. Scalable generalized linear bandits: Online computation and hashing. *arXiv preprint arXiv:1706.00136*, 2017.

[13] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.

[14] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.

[15] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[16] Nathan Korda, Balazs Szorenyi, and Shuai Li. Distributed clustering of linear bandits in peer to peer networks. In *International conference on machine learning*, pages 1301–1309. PMLR, 2016.

[17] Chuanhao Li and Hongning Wang. Asynchronous upper confidence bound algorithms for federated linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 6529–6553. PMLR, 2022.

[18] Lihong Li, Wei Chu, John Langford, Taesup Moon, and Xuanhui Wang. 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*, pages 19–36. JMLR Workshop and Conference Proceedings, 2012.

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

[20] Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *International Conference on Machine Learning*, pages 2071–2080. PMLR, 2017.

[21] Wei Li, Xuerui Wang, Ruofei Zhang, Ying Cui, Jianchang Mao, and Rong Jin. Exploitation and exploration in a performance based contextual advertising system. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 27–36, 2010.

[22] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019.

[23] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.

[24] Aritra Mitra, Rayana Jaafar, George J Pappas, and Hamed Hassani. Achieving linear convergence in federated learning under objective and systems heterogeneity. *arXiv preprint arXiv:2102.07053*, 2021.

[25] Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.

[26] Reese Pathak and Martin J Wainwright. Fedsplit: an algorithmic framework for fast federated optimization. *Advances in Neural Information Processing Systems*, 33:7057–7066, 2020.

[27] Chengshuai Shi, Cong Shen, and Jing Yang. Federated multi-armed bandits with personalization. In *International Conference on Artificial Intelligence and Statistics*, pages 2917–2925. PMLR, 2021.

[28] Yuanhao Wang, Jiachen Hu, Xiaoyu Chen, and Liwei Wang. Distributed bandit learning: Near-optimal regret with efficient communication. In *International Conference on Learning Representations*, 2019.

[29] Lijun Zhang, Tianbao Yang, Rong Jin, Yichi Xiao, and Zhi-Hua Zhou. Online stochastic linear optimization under one-bit feedback. In *International Conference on Machine Learning*, pages 392–401. PMLR, 2016.

[30] Dongruo Zhou, Lihong Li, and Quanquan Gu. Neural contextual bandits with ucb-based exploration. In *International Conference on Machine Learning*, pages 11492–11502. PMLR, 2020.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] , see Section 4.2
   (c) Did you discuss any potential negative societal impacts of your work? [No]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...

(a) Did you state the full set of assumptions of all theoretical results? [Yes] , see Section 3.3

(b) Did you include complete proofs of all theoretical results? [Yes] , the complete proof is given in appendix

3. If you ran experiments...

(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

(a) If your work uses existing assets, did you cite the creators? [Yes] , in Section 5

(b) Did you mention the license of the assets? [N/A]

(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Description of AGD and ONS

For the sake of completeness, we first provide the descriptions for AGD [25] and ONS [10, 12], which are used by FedGLB-UCB in Section 4.1 as sub-routines for global and local model updates.

Algorithm 2 ONS-Update

1: **Input:** $\theta_{t-1,i}, A_{t,i}, \nabla l(\mathbf{x}_{t,i}^\top \theta_{t-1,i}, y_{t,i})$
2: $\theta'_{t,i} = \theta_{t-1,i} - \frac{1}{c_\mu} A_{t,i}^{-1} \nabla l(\mathbf{x}_{t,i}^\top \theta_{t-1,i}, y_{t,i})$
3: $\theta_{t,i} = \arg\min_{\theta \in \mathcal{B}_d(S)} ||\theta - \theta'_{t,i}||^2_{A_{t,i}}$
4: **Output:** $\theta_{t,i}$

Algorithm 3 AGD-Update

1: **Input** : initial θ, number of inner iterations J
2: **Initialization**: set $\theta^{(1)} = \vartheta^{(1)} = \theta$, and define the sequences $\{\upsilon_j := \frac{1+\sqrt{1+4\upsilon_{j-1}^2}}{2}\}_{j\in[J]}$ (with $\upsilon_0 = 0$), and $\{\gamma_j = \frac{1-\upsilon_j}{\upsilon_{j+1}}\}_{j\in[J]}$
3: **for** $j = 1, 2, \ldots, J$ **do**
4: **Clients** compute and send local gradient $\{\nabla F_i(\theta^{(j)})\}_{i\in[N]}$ to the server
5: **Server** aggregates local gradients $\nabla F(\theta^{(j)}) = \frac{1}{N}\sum_{i=1}^N \nabla F_i(\theta^{(j)})$, and execute the following update rule to get $\theta^{(j+1)}$:
- $\vartheta^{(j+1)} = \theta^{(j)} - \frac{1}{k_\mu + \frac{\lambda}{Nt}} \nabla F(\theta^{(j)})$
- $\theta^{(j+1)} = (1-\gamma_j)\vartheta^{(j+1)} + \gamma_j \vartheta^{(j)}$

6: **Output:** $\theta^{(J+1)}$

B Technical Lemmas

When applying the following self-normalized bound in the analysis of federated bandit algorithm with event-trigger, a subtle difference from the analysis of standard bandit algorithm is that the sequence of data points used to update each client is controlled by the data-dependent event-trigger, e.g. Eq (4), which introduces dependencies on the future data, and thus breaks the standard argument. This problem also exists in prior works of distributed linear bandit, but was not addressed rigorously (see Lemma H.1. of [28]). Specifically, each client i observes the sequence of data points in a different order, i.e., it first observes each newly collected local data points from the environment, and then observes (in the form of their gradients) the batch of new data points that other clients have collected at the end of the epoch. Then, if we consider a data point that is contained in the batch of new data collected by other clients, the index of this data point (as observed by client i) has dependency all the way to the end of this batch, i.e., its index is only determined after some client triggers the global update.

Therefore, in order to avoid this dependency on future data points, when constructing the filtration, we should avoid including the σ-algebra that 'cuts a batch in half', but instead only include the σ-algebra generated by the sequence of data points up to the end of each batch, where we consider each locally observed data point as a batch as well. Denote the sequence of time indices corresponding to these data points as $\{t_p\}_{p\in[P]}$ for some $P > 0$. Then the constructed filtration $\{\mathcal{F}_{t_p}\}_{p\in[P]}$ is essentially a batched version of the standard $\{\mathcal{F}_t\}_{t=1}^\infty$. The self-normalized bound below still holds, i.e., by changing the stopping time construction from $\cup_{t\geq 1} B_t(\delta)$ to $\cup_{t\in\{t_p\}_{p\in[P]}} B_t(\delta)$ in the proof of Theorem 1 in [1], where $B_t(\delta)$ denotes the bad event that the bound does not hold. Therefore, instead of holding uniformly over all t, the self-normalized bound now only holds for all $t \in \{t_p\}_{p\in[P]}$, i.e., the sequence of time steps when client i gets updated, which is also what we need.

Lemma B.1 (Vector-valued self-normalized bound (Theorem 1 of [1])). *Let $\{\mathcal{F}_t\}_{t=1}^\infty$ be a filtration. Let $\{\eta_t\}_{t=1}^\infty$ be a real-valued stochastic process such that η_t is $\mathcal{F}_{t+1}$-measurable, and η_t is conditionally zero mean R-sub-Gaussian for some $R \geq 0$. Let $\{X_t\}_{t=1}^\infty$ be a $\mathbb{R}^d$-valued stochastic process such that X_t is $\mathcal{F}_t$-measurable. Assume that V is a $d \times d$ positive definite matrix. For any $t > 0$,*

define

$$V_t = V + \sum_{\tau=1}^{t} X_\tau X_\tau^\top \quad \mathcal{S}_t = \sum_{\tau=1}^{t} \eta_\tau X_\tau$$

Then for any $\delta > 0$, with probability at least $1-\delta$,

$$||\mathcal{S}_t||_{V_t^{-1}} \leq R\sqrt{2\log\frac{\det(V_t)^{1/2}}{\det(V)^{1/2}\delta}}, \quad \forall t \geq 0$$

Lemma B.2 (Corollary 8 of [2])**.** *Under the same assumptions as Lemma B.1, consider a sequence of real-valued variables $\{Z_t\}_{t=1}^{\infty}$ such that Z_t is $\mathcal{F}_t$-measurable. Then for any $\delta > 0$, with probability at least $1-\delta$,*

$$|\sum_{\tau=1}^{t} \eta_\tau Z_\tau| \leq R\sqrt{2(V + \sum_{\tau=1}^{t} Z_\tau^2)\log\left(\frac{\sqrt{V + \sum_{\tau=1}^{t} Z_\tau^2}}{\delta\sqrt{V}}\right)}, \forall t \geq 0$$

Lemma B.3. *Under Assumption 1, $F_{i,t}(\theta)$ for $i = 1, 2, \ldots, N$ is smooth with constant $k_\mu + \frac{\lambda}{Nt}$*

Proof. By Assumption 1, $\mu(\cdot)$ is Lipschitz continuous with constant k_μ, i.e., $|\mu(\mathbf{x}^\top\theta_1) - \mu(\mathbf{x}^\top\theta_2)| \leq k_\mu|\mathbf{x}^\top(\theta_1 - \theta_2)|$. Then we can show that

$$\begin{aligned}
&||\nabla F_{i,t}(\theta_1) - \nabla F_{i,t}(\theta_2)|| \\
&= ||\frac{1}{t}\sum_{s=1}^{t} \mathbf{x}_{s,i}[\mu(\mathbf{x}_{s,i}^\top\theta_1) - \mu(\mathbf{x}_{s,i}^\top\theta_2)] + \frac{\lambda}{Nt}(\theta_1 - \theta_2)|| \\
&\leq \frac{1}{t}\sum_{s=1}^{t} ||\mathbf{x}_{s,i}[\mu(\mathbf{x}_{s,i}^\top\theta_1) - \mu(\mathbf{x}_{s,i}^\top\theta_2)]|| + \frac{\lambda}{Nt}||\theta_1 - \theta_2|| \\
&\leq \frac{1}{t}\sum_{s=1}^{t} |\mu(\mathbf{x}_{s,i}^\top\theta_1) - \mu(\mathbf{x}_{s,i}^\top\theta_2)| + \frac{\lambda}{Nt}||\theta_1 - \theta_2|| \\
&\leq \frac{k_\mu}{t}\sum_{s=1}^{t} |\mathbf{x}_{s,i}^\top(\theta_1 - \theta_2)| + \frac{\lambda}{Nt}||\theta_1 - \theta_2|| \leq (k_\mu + \frac{\lambda}{Nt})||\theta_1 - \theta_2||
\end{aligned}$$

Therefore, $\nabla F_{i,t}(\theta)$ is Lipschitz continuous with constant $k_\mu + \frac{\lambda}{Nt}$, and $\nabla F_t(\theta) = \frac{1}{N}\sum_{i=1}^{N} \nabla F_{i,t}(\theta)$ is Lipschitz continuous with constant $k_\mu + \frac{\lambda}{Nt}$ as well. □

Lemma B.4 (Matrix Weighted Cauchy-Schwarz)**.** *If $A \in \mathbb{R}^{d\times d}$ is a PSD matrix, then $x^T Ay \leq \sqrt{x^T Ax \cdot y^T Ay}$ holds for any vectors $x, y \in \mathbb{R}^d$.*

Proof. Consider a quadratic function $(x+ty)^T A(x+ty) = x^T Ax + 2(x^T Ay)t + (y^T Ay)t^2$ for some variable $t \in \mathbb{R}$, where $x, y \in \mathbb{R}^d$ are arbitrary vectors. Since A is PSD, the value of this quadratic function $(x+ty)^T A(x+ty) = x^T Ax + 2(x^T Ay)t + (y^T Ay)t^2 \geq 0, \forall t$, which means there can be at most one root. This is equivalent to saying the discriminant of this quadratic function $4(x^T Ay)^2 - 4x^T Ax \cdot y^T Ay \leq 0$, which finishes the proof. □

Lemma B.5 (Confidence Ellipsoid Centered at Global Model)**.** *Consider time step $t \in [T]$ when a global update happens, such that the distributed optimization over N clients is performed to get the globally updated model θ_t. Denote the sub-optimality of the final iteration as ϵ_t, such that $F_t(\theta_t) - F_t(\hat{\theta}_t^{MLE}) \leq \epsilon_t$; then with probability at least $1-\delta$, for all $t \in [T]$,*

$$||\theta_t - \theta_\star||_{A_t} \leq \alpha_t$$

where $\alpha_t = Nt\sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{Ntc_\mu}}\sqrt{\epsilon_t} + \frac{R_{max}}{c_\mu}\sqrt{d\log(1 + Ntc_\mu/(d\lambda)) + 2\log(1/\delta)} + \sqrt{\frac{\lambda}{c_\mu}}S$, and $A_t = \frac{\lambda}{c_\mu}I + \sum_{i=1}^{N}\sum_{s=1}^{t} \mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top$.

Proof. Recall that the unique minimizer of Eq.(2) is denoted as $\hat{\theta}_t^{\text{MLE}}$, so by taking gradient w.r.t. θ we have, $g_t(\hat{\theta}_t^{\text{MLE}}) = \sum_{i=1}^N \sum_{s=1}^t \mathbf{x}_{s,i} y_{s,i}$, where we define $g_t(\theta) = \lambda\theta + \sum_{i=1}^N \sum_{s=1}^t \mu(\mathbf{x}_{s,i}^\top \theta)\mathbf{x}_{s,i}$. First, we start with standard arguments [8, 20] to show that $\|\theta_t - \theta_\star\|_{A_t} \leq \frac{1}{c_\mu}\|g_t(\theta_t) - g_t(\theta_\star)\|_{A_t^{-1}}$. Specifically, by Assumption 1 and the Fundamental Theorem of Calculus, we have

$$g_t(\theta_t) - g_t(\theta_\star) = G_t(\theta_t - \theta_\star)$$

where $G_t = \int_0^1 \nabla g_t(a\theta_t + (1-a)\theta_\star)da$. Again by Assumption 1, $\nabla g_t(\theta) = \lambda I + \sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top \dot{\mu}(\mathbf{x}_{s,i}^\top \theta)$ is continuous, and $\nabla g_t(\theta) \succcurlyeq \lambda I + c_\mu \sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top \succ 0$, so $G_t \succ 0$, i.e., G_t is invertible. Therefore, we have

$$\theta_t - \theta_\star = G_t^{-1}[g_t(\theta_t) - g_t(\theta_\star)]$$

Note that $G_t \succcurlyeq \lambda I + c_\mu \sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top = c_\mu A_t$, so $G_t^{-1} \preccurlyeq \frac{1}{c_\mu} A_t^{-1}$. Hence,

$$\begin{aligned}
\|\theta_t - \theta_\star\|_{A_t} &= \|G_t^{-1}[g_t(\theta_t) - g_t(\theta_\star)]\|_{A_t} \leq \|\frac{1}{c_\mu} A_t^{-1}[g_t(\theta_t) - g_t(\theta_\star)]\|_{A_t} = \frac{1}{c_\mu}\|g_t(\theta_t) - g_t(\theta_\star)\|_{A_t^{-1}} \\
&\leq \frac{1}{c_\mu}\|g_t(\theta_t) - g_t(\hat{\theta}_t^{\text{MLE}})\|_{A_t^{-1}} + \frac{1}{c_\mu}\|g_t(\hat{\theta}_t^{\text{MLE}}) - g_t(\theta_\star)\|_{A_t^{-1}} \qquad (8)
\end{aligned}$$

where the first term depends on the sub-optimality of the offline regression estimator θ_t to the unique minimizer $\hat{\theta}_t^{(\text{MLE})}$, and the second term is standard for GLB [20].

For the first term, based on the definition of $g_t(\cdot)$, we have

$$\begin{aligned}
&||g_t(\theta_t) - g_t(\hat{\theta}_t^{(\text{MLE})})||_{A_t^{-1}} \\
&\leq ||\lambda\theta_t + \sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}\mu(\mathbf{x}_{s,i}^\top \theta_t) - \sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i} y_{s,i}||_{A_t^{-1}} \\
&= ||\sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}[-y_{i,s} + \mu(\mathbf{x}_{s,i}^\top \theta_t)] + \lambda\theta_t||_{A_t^{-1}} \\
&= ||Nt\nabla F_t(\theta_t)||_{A_t^{-1}}
\end{aligned}$$

where the last equality is due to the definition of $F_t(\theta)$ in Eq.(2). We can further bound it using the property of Rayleigh quotient and the fact that $A_t \succcurlyeq \frac{\lambda}{c_\mu} I$, which gives us

$$||g_t(\theta_t) - g_t(\hat{\theta}_t^{(\text{MLE})})||_{A_t^{-1}} \leq \frac{Nt||\nabla F_t(\theta_t)||_2}{\sqrt{\lambda_{\min}(A_t)}} \leq \frac{Nt||\nabla F_t(\theta_t)||_2}{\sqrt{\lambda/c_\mu}}$$

Based on Lemma B.3, $F_t(\theta)$ is $(k_\mu + \frac{\lambda}{Nt})$-smooth, which means

$$\frac{1}{2k_\mu + 2\lambda/(Nt)}\|\nabla F_t(\theta_t)\|^2 \leq F_t(\theta_t) - F_t(\hat{\theta}_t^{\text{MLE}}) \leq \epsilon_t$$

where the second inequality is by definition of ϵ_t. Putting everything together, we have the following bound for the first term

$$\frac{1}{c_\mu}\|g_t(\theta_t) - g_t(\hat{\theta}_t^{\text{MLE}})\|_{A_t^{-1}} \leq Nt\sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{Ntc_\mu}}\sqrt{\epsilon_t}$$

For the second term, similarly, based on the definition of $g_t(\cdot)$, we have

$$\begin{aligned}
&\frac{1}{c_\mu}||g_t(\hat{\theta}_t^{\text{MLE}}) - g_t(\theta_\star)||_{A_t^{-1}} \\
&= \frac{1}{c_\mu}||\sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i} y_{s,i} - \sum_{s=1}^t \sum_{i=1}^N \mu(\mathbf{x}_{s,i}^\top \theta_\star)\mathbf{x}_{s,i} - \lambda\theta_\star||_{A_t^{-1}} \\
&\leq \frac{1}{c_\mu}||\sum_{s=1}^t \sum_{i=1}^N \mathbf{x}_{s,i}\eta_{s,i}||_{A_t^{-1}} + \sqrt{\frac{\lambda}{c_\mu}}S
\end{aligned}$$

Then based on the self-normalized bound in Lemma B.1 (Theorem 1 of [1]), we have $||\sum_{s=1}^{t}\sum_{i=1}^{N}\mathbf{x}_{s,i}\eta_{s,i}||_{A_t^{-1}} \leq R_{max}\sqrt{d\log(1+Ntc_\mu/d\lambda)+2\log(1/\delta)}, \forall t$, with probability at least $1-\delta$.

Substituting the upper bounds for these two terms back into Eq.(8), we have, with probability at least $1-\delta$,

$$\begin{aligned}\|\theta_t-\theta_\star\|_{A_t} &\leq \frac{1}{c_\mu}\|g_t(\theta_t)-g_t(\hat{\theta}_t^{(\text{MLE})})\|_{A_t^{-1}}+\frac{1}{c_\mu}\|g_t(\hat{\theta}_t^{(\text{MLE})})-g_t(\theta_\star)\|_{A_t^{-1}}\\ &\leq Nt\sqrt{\frac{2k_\mu}{\lambda c_\mu}+\frac{2}{Ntc_\mu}}\sqrt{\epsilon_t}+\frac{R_{max}}{c_\mu}\sqrt{d\log(1+Ntc_\mu/(d\lambda))+2\log(1/\delta)}+\sqrt{\frac{\lambda}{c_\mu}}S\end{aligned}$$

which finishes the proof. □

C Proof of Lemma 4.1

Proof. Denote the two terms for loss difference as $A_1 = \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}},y_{s,i})-l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big]$, and $A_2=\sum_{s=t_{\text{last}}+1}^{t}\big[l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})-l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big]$. We can upper bound the term A_1 by

$$\begin{aligned}A_1 &= \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}},y_{s,i})-l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\big]-\frac{\lambda}{2}\|\theta_\star\|_2^2+\frac{\lambda}{2}\|\theta_\star\|_2^2\\ &\leq \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}},y_{s,i})-\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}l(\mathbf{x}_{s,i}^\top\hat{\theta}_{t_{\text{last}}}^{\text{MLE}},y_{s,i})-\frac{\lambda}{2}\|\hat{\theta}_{t_{\text{last}}}^{\text{MLE}}\|_2^2+\frac{\lambda}{2}\|\theta_\star\|_2^2\\ &\leq \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}},y_{s,i})+\frac{\lambda}{2}\|\theta_{t_{\text{last}}}\|_2^2-\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}l(\mathbf{x}_{s,i}^\top\hat{\theta}_{t_{\text{last}}}^{\text{MLE}},y_{s,i})-\frac{\lambda}{2}\|\hat{\theta}_{t_{\text{last}}}^{\text{MLE}}\|_2^2+\frac{\lambda}{2}S^2\\ &\leq Nt_{\text{last}}\epsilon_{t_{\text{last}}}+\frac{\lambda}{2}S^2 := B_1\end{aligned}$$

where the first inequality is because $\hat{\theta}_{t_{\text{last}}}^{\text{MLE}}$ minimizes Eq.(2), such that $\sum_{s=1}^{t_{\text{last}}}l(\mathbf{x}_{s,i}^\top\hat{\theta}_{t_{\text{last}}}^{\text{MLE}},y_{s,i})+\frac{\lambda}{2}\|\hat{\theta}_{t_{\text{last}}}^{\text{MLE}}\|_2^2\leq\sum_{s=1}^{t_{\text{last}}}l(\mathbf{x}_{s,i}^\top\theta,y_{s,i})+\frac{\lambda}{2}\|\theta\|_2^2$ for any $\theta\in\mathcal{B}_d(S)$, and the last inequality is because $F_{t_{\text{last}}}(\theta_{t_{\text{last}}})-F_{t_{\text{last}}}(\hat{\theta}_{t_{\text{last}}}^{\text{MLE}})\leq\epsilon_{t_{\text{last}}}$ by definition.

Now we start with standard arguments [12, 29] in order to bound the term A_2, which is essentially the online regret of ONS, except that its initial model is the globally updated model $\theta_{t_{\text{last}}}$. First, since $l(z,y)$ is c_μ-strongly-convex w.r.t. z, we have

$$l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})-l(\mathbf{x}_{s,i}^\top\theta_\star,y_{s,i})\leq[\mu(\mathbf{x}_{s,i}^\top\theta_{s-1,i})-y_{s,i}]\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)-\frac{c_\mu}{2}||\theta_{s-1,i}-\theta_\star||^2_{\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top} \tag{9}$$

To further bound the RHS of Eq.(9), recall from the ONS local update rule in Algorithm 2 that, for each client $i\in[N]$ at the end of each time step $s\in[t_{\text{last}}+1,t]$,

$$\begin{aligned}\theta'_{s,i} &= \theta_{s-1,i}-\frac{1}{c_\mu}A_{s,i}^{-1}\nabla l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})\\ \theta_{s,i} &= \underset{\theta\in\mathcal{B}_d(S)}{\arg\min}\,||\theta'_{s,i}-\theta||^2_{A_{s,i}}\end{aligned}$$

Then due to the property of generalized projection (Lemma 8 of [10]), we have

$$\begin{aligned}&\|\theta_{s,i}-\theta_\star\|^2_{A_{s,i}}\\ &\leq\|\theta_{s-1,i}-\theta_\star-\frac{1}{c_\mu}A_{s,i}^{-1}\nabla l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})\|^2_{A_{s,i}}\\ &\leq\|\theta_{s-1,i}-\theta_\star\|^2_{A_{s,i}}-\frac{2}{c_\mu}\nabla l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})^\top(\theta_{s-1,i}-\theta_\star)+\frac{1}{c_\mu^2}\|\nabla l(\mathbf{x}_{s,i}^\top\theta_{s-1,i},y_{s,i})\|^2_{A_{s,i}^{-1}}\end{aligned}$$

By rearranging terms, we have

$$
\begin{aligned}
&\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})^\top (\theta_{s-1,i} - \theta_\star) \\
&\leq \frac{1}{2c_\mu} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,j}^{-1}}^2 + \frac{c_\mu}{2} \left(\|\theta_{s-1,i} - \theta_\star\|_{A_{s,i}}^2 - \|\theta_{s,i} - \theta_\star\|_{A_{s,i}}^2 \right) \\
&= \frac{1}{2c_\mu} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 + \frac{c_\mu}{2} \|\theta_{s-1,i} - \theta_\star\|_{A_{s-1,i}}^2 \\
&\quad + \frac{c_\mu}{2} \left(\|\theta_{s-1,i} - \theta_\star\|_{A_{s,i}}^2 - \|\theta_{s-1,i} - \theta_\star\|_{A_{s-1,i}}^2 \right) - \frac{c_\mu}{2} \|\theta_{s,i} - \theta_\star\|_{A_{s,i}}^2 \\
&= \frac{1}{2c_\mu} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 + \frac{c_\mu}{2} \|\theta_{s-1,i} - \theta_\star\|_{A_{s-1,i}}^2 + \frac{c_\mu}{2} \|\theta_{s-1,i} - \theta_\star\|_{\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top}^2 - \frac{c_\mu}{2} \|\theta_{s,i} - \theta_\star\|_{A_{s,i}}^2
\end{aligned}
$$

Note that $\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i}) = \mathbf{x}_{s,i}[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - y_{s,i}]$, so with the inequality above, we can further bound the RHS of Eq.(9):

$$
\begin{aligned}
&l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i}) - l(\mathbf{x}_{s,i}^\top \theta_\star, y_{s,i}) \leq [\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - y_{s,i}] \mathbf{x}_{s,i}^\top (\theta_{s-1,i} - \theta_\star) - \frac{c_\mu}{2} ||\theta_{s-1,i} - \theta_\star||_{\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top}^2 \\
&\leq \frac{1}{2c_\mu} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 + \frac{c_\mu}{2} \|\theta_{s-1,i} - \theta_\star\|_{A_{s-1,i}}^2 - \frac{c_\mu}{2} \|\theta_{s,i} - \theta_\star\|_{A_{s,i}}^2
\end{aligned}
$$

Then summing over $s \in [t_{\text{last}} + 1, t]$, we have

$$
A_2 \leq \frac{1}{2c_\mu} \sum_{s=t_{\text{last}}+1}^{t} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 + \frac{c_\mu}{2} \|\theta_{t_{\text{last}},i} - \theta_\star\|_{A_{t_{\text{last}},i}}^2 - \frac{c_\mu}{2} \|\theta_{t,i} - \theta_\star\|_{A_{t,i}}^2
$$

where $A_{t_{\text{last}},i} = A_{t_{\text{last}}}, \theta_{t_{\text{last}},i} = \theta_{t_{\text{last}}}, \forall i \in [N]$ due to the global update (line 15 in Algorithm 1).

We should note that the second term above itself essentially corresponds to a confidence ellipsoid centered at the globally updated model $\theta_{t_{\text{last}}}$, and its appearance in the upper bound for the loss difference (online regret) of local updates is because the local update is initialized by $\theta_{t_{\text{last}}}$. And based on Lemma B.5, with probability at least $1 - \delta$,

$$
\begin{aligned}
\|\theta_{t_{\text{last}},i} - \theta_\star\|_{A_{t_{\text{last}},i}} \leq\ & N t_{\text{last}} \sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{N t_{\text{last}} c_\mu}} \sqrt{\epsilon_{t_{\text{last}}}} \\
& + \frac{1}{c_\mu} R_{\max} \sqrt{d \log\left(1 + N t_{\text{last}} c_\mu / d\lambda\right) + 2 \log\left(1/\delta\right)} + \sqrt{\frac{\lambda}{c_\mu}} S
\end{aligned}
$$

Therefore, with probability at least $1 - \delta$,

$$
\begin{aligned}
A_2 \leq\ & \frac{1}{2c_\mu} \sum_{s=t_{\text{last}}+1}^{t} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 + \frac{c_\mu}{2} \Big[N t_{\text{last}} \sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{N t_{\text{last}} c_\mu}} \sqrt{\epsilon_{t_{\text{last}}}} \\
& + \frac{1}{c_\mu} R_{\max} \sqrt{d \log\left(1 + N t_{\text{last}} c_\mu / d\lambda\right) + 2 \log\left(1/\delta\right)} + \sqrt{\frac{\lambda}{c_\mu}} S \Big]^2 := B_2
\end{aligned}
$$

which finishes the proof for Lemma 4.1.

□

D Proof of Lemma 4.2 and Corollary 4.2.1

Proof of Lemma 4.2. Due to c_μ-strongly convexity of $l(z, y)$ w.r.t. z, we have $l(\mathbf{x}_{s,i}^\top \theta, y_{s,i}) - l(\mathbf{x}_{s,i}^\top \theta_\star, y_{s,i}) \geq \left[\mu(\mathbf{x}_{s,i}^\top \theta_\star) - y_{s,i}\right] \mathbf{x}_{s,i}^\top (\theta - \theta_\star) + \frac{c_\mu}{2} \left[\mathbf{x}_{s,i}^\top (\theta - \theta_\star)\right]^2$. Substituting this to the LHS

of Eq.(6) and Eq.(7), we have

$$
\begin{aligned}
B_1 &\geq \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[l(\mathbf{x}_{s,i}^\top\theta_{t_{\text{last}}}, y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star, y_{s,i})\big] \\
&\geq \sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mu(\mathbf{x}_{s,i}^\top\theta_\star) - y_s\big]\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star) + \frac{c_\mu}{2}\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star)\big]^2 \\
B_2 &\geq \sum_{s=t_{\text{last}}+1}^{t}\big[l(\mathbf{x}_{s,i}^\top\theta_{s-1,i}, y_{s,i}) - l(\mathbf{x}_{s,i}^\top\theta_\star, y_{s,i})\big] \\
&\geq \sum_{s=t_{\text{last}}+1}^{t}\big[\mu(\mathbf{x}_{s,i}^\top\theta_\star) - y_s\big]\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star) + \frac{c_\mu}{2}\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2
\end{aligned}
$$

By rearranging the terms, we have

$$
\begin{aligned}
\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star)\big]^2 &\leq \frac{2}{c_\mu}B_1 + \frac{2}{c_\mu}\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star) \\
\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2 &\leq \frac{2}{c_\mu}B_2 + \frac{2}{c_\mu}\sum_{s=t_{\text{last}}+1}^{t}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)
\end{aligned}
$$

where the LHS is quadratic in $\theta_\star$. For the RHS, we will further upper bound the second term as shown below.

• **Upper Bound for** $\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2$ Note that $\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)$ is $\mathcal{F}_{s,i}$-measurable, and $\eta_{s,i}$ is $\mathcal{F}_{s+1,i}$-measurable and conditionally R_{max}-sub-Gaussian. By applying Lemma B.2 (Corollary 8 of [2]) w.r.t. client i's filtration $\{\mathcal{F}_{s,i}\}_{s=t_{\text{last}}+1}^{\infty}$, where $\mathcal{F}_{s,i} = \sigma\big([\mathbf{x}_{k,j}, \eta_{k,j}]_{k,j:k\leq t_{\text{last}}\cap j\leq N}, [\mathbf{x}_{k,j}, \eta_{k,j}]_{k,j:t_{\text{last}}+1\leq k\leq s-1\cap j=i}, \mathbf{x}_{s,i}\big)$, and taking union bound over all $i \in [N]$, with probability at least $1-\delta$, for all $t \in [T], i \in [N]$,

$$
\sum_{s=t_{\text{last}}+1}^{t}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star) \leq
$$

$$
R_{\max}\sqrt{2\big(1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2\big)\cdot\log\big(\frac{N}{\delta}\sqrt{1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2}\big)}
$$

Therefore,

$$
\begin{aligned}
1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2 &\leq 1 + \frac{2}{c_\mu}B_2 \\
&+ \frac{2R_{\max}}{c_\mu}\sqrt{2\big(1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2\big)\cdot\log\big(\frac{N}{\delta}\sqrt{1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2}\big)}
\end{aligned} \tag{10}
$$

Then by applying Lemma 2 of [12], i.e., if $q^2 \leq a + fq\sqrt{\log(\frac{q}{\delta/N})}$ then $q^2 \leq 2a + f^2\log(\frac{\sqrt{4a+f^4/(4\delta^2)}}{\delta/N})$ (for $a, f \geq 0, q \geq 1$). And by setting $q = \sqrt{1 + \sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2}$, $a = 1 + \frac{2}{c_\mu}B_2$, $f = \frac{2\sqrt{2}R_{\max}}{c_\mu}$, we have

$$
\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i} - \theta_\star)\big]^2 \leq 1 + \frac{4B_2}{c_\mu} + \frac{8R_{\max}^2}{c_\mu^2}\log\left(\frac{N}{\delta}\sqrt{4 + \frac{8}{c_\mu}B_2 + \frac{64R_{\max}^4}{c_\mu^4\cdot 4\delta^2}}\right), \forall t, i \tag{11}
$$

with probability at least $1-\delta$.

• **Upper Bound for** $\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star)\big]^2$ Note that $\theta_{t_{\text{last}}}$ depends on all data samples in $\{(\mathbf{x}_{s,i}, y_{s,i})\}_{s\in[t_{\text{last}}]}$ as a result of the offline regression method, and therefore $\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}} - \theta_\star)$ is no

longer $\mathcal{F}_{s,i}$-measurable for $s \in [1, t_{\text{last}})$. Hence, we cannot use Lemma B.2 as before. Instead, we have

$$
\begin{aligned}
&\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star) = \big(\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}\big)^\top(\theta_{t_{\text{last}}}-\theta_\star)\\
&= \big(\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}\big)^\top(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)^{-1}(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)(\theta_{t_{\text{last}}}-\theta_\star)\\
&\le \sqrt{\big(\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}\big)^\top(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)^{-1}\big(\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}\big)\cdot(\theta_{t_{\text{last}}}-\theta_\star)^\top(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)(\theta_{t_{\text{last}}}-\theta_\star)}\\
&= \sqrt{\|\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\eta_{s,i}\mathbf{x}_{s,i}\|^2_{(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)^{-1}}\cdot\|\theta_{t_{\text{last}}}-\theta_\star\|^2_{(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}}\\
&\le R_{\max}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)\cdot\|\theta_{t_{\text{last}}}-\theta_\star\|^2_{(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}},
\end{aligned}
$$

with probability at least $1-\delta$, where the first inequality is due to the matrix-weighted Cauchy-Schwarz inequality in Lemma B.4, such that $x^\top A^{-1}Ay \le \sqrt{x^\top A^{-1}x\cdot y^\top A^\top A^{-1}Ay} = \sqrt{x^\top A^{-1}x\cdot y^\top Ay}$ for symmetric PD matrix A, and the second inequality is obtained by applying the self-normalized bound in Lemma B.1 w.r.t. the filtration $\{\mathcal{F}_s\}_{s\in\{t_p\}_{p=1}^B}$, where $\mathcal{F}_s = \sigma\big([\mathbf{x}_{k,j},\eta_{k,j}]_{k,j:k\le s-1\cap j\le N},[\mathbf{x}_{k,j},\eta_{k,j}]_{k,j:k=s\cap j\le N-1},\mathbf{x}_{s,N}\big)$ and $\{t_p\}_{p=1}^B$ denotes the sequence of time steps when global update happens, and B denotes the total number of global updates. By substituting it back, we have

$$
\begin{aligned}
&\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2\\
&\le \frac{2}{c_\mu}B_1+\frac{2R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)\cdot\|\theta_{t_{\text{last}}}-\theta_\star\|^2_{I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top}}\\
&\le \frac{2}{c_\mu}B_1+\frac{2R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)\cdot\big(\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2+\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2\big)}
\end{aligned}
\tag{12}
$$

Then by applying the Proposition 9 of [2], i.e. if $z^2 \le a+bz$ then $z \le b+\sqrt{a}$ (for $a,b\ge 0$), and setting $z=\sqrt{\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2+\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2}, a=\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2+\frac{2}{c_\mu}B_1, b=\frac{2R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}$,we have

$$
\begin{aligned}
&\sqrt{\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2+\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2}\\
\le&\frac{2R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}+\sqrt{\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2+B_1}
\end{aligned}
\tag{13}
$$

Taking square on both sides, and rearranging terms, we have

$$
\begin{aligned}
&\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2\\
\le&\frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)+B_1\\
&+\frac{4R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}\sqrt{\|\theta_{t_{\text{last}}}-\theta_\star\|_2^2+B_1}\\
\le&\frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)+B_1\\
&+\frac{4R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}(\|\theta_{t_{\text{last}}}-\theta_\star\|_2+\sqrt{B_1})\\
\le&\frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)+B_1\\
&+\frac{4R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}(\|\theta_{t_{\text{last}}}\|_2+\|\theta_\star\|_2+\sqrt{B_1})
\end{aligned}
\tag{14}
$$

Now putting everything together, we have the following confidence region for $\theta_\star$,

$$
P\big(\forall t,i,\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\big[\mathbf{x}_{s,i}^\top(\theta_{t_{\text{last}}}-\theta_\star)\big]^2+\sum_{s=t_{\text{last}}+1}^{t}\big[\mathbf{x}_{s,i}^\top(\theta_{s-1,i}-\theta_\star)\big]^2\le\beta_{t,i}\big)\ge 1-2\delta \tag{15}
$$

where $\beta_{t,i} = \frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big) + B_1 + \frac{4R_{\max}}{c_\mu}\sqrt{2\log\big(\frac{1}{\delta}\sqrt{\det(I+\sum_{s=1}^{t_{\text{last}}}\sum_{i=1}^{N}\mathbf{x}_{s,i}\mathbf{x}_{s,i}^\top)}\big)}(\|\theta_{t_{\text{last}}}\|_2 + \|\theta_\star\|_2 + \sqrt{B_1}) + 1 + \frac{4B_2}{c_\mu} + \frac{8R_{\max}^2}{c_\mu^2}\log\big(\frac{N}{\delta}\sqrt{4+\frac{8}{c_\mu}B_2+\frac{64R_{\max}^4}{c_\mu^4\cdot 4\delta^2}}\big)$.

Denote $\mathbf{X}_{t,i} = \begin{bmatrix}\mathbf{x}_{1,1}^\top\\ \cdots\\ \mathbf{x}_{t_{\text{last}},N}^\top\\ \mathbf{x}_{i,t_{\text{last}}+1}^\top\\ \cdots\\ \mathbf{x}_{i,t}^\top\end{bmatrix}\in\mathbb{R}^{(Nt_{\text{last}}+t-t_{\text{last}})\times d}$, and $\mathbf{z}_{t,i} = \begin{bmatrix}\mathbf{x}_{1,1}^\top\theta_{t_{\text{last}}}\\ \cdots\\ \mathbf{x}_{t_{\text{last}},N}^\top\theta_{t_{\text{last}}}\\ \mathbf{x}_{i,t_{\text{last}}+1}^\top\theta_{t_{\text{last}},i}\\ \cdots\\ \mathbf{x}_{i,t}^\top\theta_{t-1,i}\end{bmatrix}\in\mathbb{R}^{Nt_{\text{last}}+t-t_{\text{last}}}$. We can rewrite the inequality above as

$$
\begin{aligned}
&\|\mathbf{z}_{t,i}-\mathbf{X}_{t,i}\theta_\star\|_2^2+\frac{\lambda}{c_\mu}\|\theta_\star\|_2^2\le\beta_{t,i}+\frac{\lambda}{c_\mu}\|\theta_\star\|_2^2\le\beta_{t,i}+\frac{\lambda}{c_\mu}S^2\\
\Leftrightarrow&\|\mathbf{z}_{t,i}-\mathbf{X}_{t,i}\theta_\star\|_2^2+\frac{\lambda}{c_\mu}\|\theta_\star\|_2^2-\|\mathbf{z}_{t,i}-\mathbf{X}_{t,i}\hat{\theta}_{t,i}\|_2^2-\frac{\lambda}{c_\mu}\|\hat{\theta}_{t,i}\|_2^2+\|\mathbf{z}_{t,i}-\mathbf{X}_{t,i}\hat{\theta}_{t,i}\|_2^2+\frac{\lambda}{c_\mu}\|\hat{\theta}_{t,i}\|_2^2\\
&\le\beta_{t,i}+\frac{\lambda}{c_\mu}S^2
\end{aligned}
$$

where $\hat{\theta}_{t,i} = A_{t,i}^{-1}\mathbf{X}_{t,i}^\top\mathbf{z}_{t,i}$ denotes the Ridge regression estimator based on the predicted rewards given by the past sequence of model updates, and the regularization parameter is $\frac{\lambda}{c_\mu}$. Note that by expanding $\hat{\theta}_{t,i}$, we can show $\hat{\theta}_{t,i}^\top A_{i,t}\hat{\theta}_{t,i} = \mathbf{z}_{i,t}^\top\mathbf{X}_{i,t}\hat{\theta}_{t,i}$, and $\hat{\theta}_{t,i}^\top A_{i,t}\theta_\star = \mathbf{z}_{i,t}^\top\mathbf{X}_{i,t}\theta_\star$. Therefore, we have

$$
\|\hat{\theta}_{t,i}-\theta_\star\|_{A_{t,i}}^2\le\beta_{t,i}+\frac{\lambda}{c_\mu}S^2-(\|\mathbf{z}_{t,i}\|_2^2-\hat{\theta}_{t,i}^\top\mathbf{X}_{t,i}^\top\mathbf{z}_{t,i})
$$

which finishes the proof of Lemma 4.2. □

Proof of Corollary 4.2.1. Under the condition that $\epsilon_{t_{\text{last}}} \leq \frac{1}{N^2 t_{\text{last}}^2}$,

$$
\begin{aligned}
B_1 &\leq \frac{1}{N t_{\text{last}}} + \frac{\lambda}{2} S^2 = O(1) \\
B_2 &\leq \frac{1}{2c_\mu} \sum_{s=t_{\text{last}}+1}^{t} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 \\
&\quad + \frac{c_\mu}{2} \Big[\sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{N t_{\text{last}} c_\mu}} + \frac{1}{c_\mu} R_{\max} \sqrt{d \log(1 + N t_{\text{last}} c_\mu / d\lambda) + 2\log(1/\delta)} + \sqrt{\frac{\lambda}{c_\mu}} S \Big]^2
\end{aligned}
$$

Note that $\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i}) = \mathbf{x}_{s,i}[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - y_{s,i}]$. We can upper bound the squared prediction error by

$$
\begin{aligned}
&\big[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - y_{s,i}\big]^2 \\
&= \big[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - \mu(\mathbf{x}_{s,i}^\top \theta_\star) - \eta_{s,i}\big]^2 \\
&\leq 2\big[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - \mu(\mathbf{x}_{s,i}^\top \theta_\star)\big]^2 + 2\eta_{s,i}^2 \\
&\leq 2k_\mu^2 \big[\mathbf{x}_{s,i}^\top (\theta_{s-1,i} - \theta_\star)\big]^2 + 2\eta_{s,i}^2 \\
&\leq 8k_\mu^2 S^2 + 2\eta_{s,i}^2
\end{aligned}
$$

where the first inequality is due to AM-QM inequality, and the second inequality is due to the k_μ-Lipschitz continuity of $\mu(\cdot)$ according to Assumption 1. Since $|\eta_{s,i}| \leq R_{\max}$, $\big[\mu(\mathbf{x}_{s,i}^\top \theta_{s-1,i}) - y_{s,i}\big]^2 \leq k_\mu^2 S^2 + R_{\max}^2$. In addition, due to Lemma 11 of [1], i.e., $\sum_{s=t_{\text{last}}+1}^{t} \|\mathbf{x}_{s,i}\|_{A_{s,i}^{-1}}^2 \leq 2\log(\frac{\det(A_{t,i})}{\det(\lambda I)})$ Therefore,

$$
\frac{1}{2c_\mu} \sum_{s=t_{\text{last}}+1}^{t} \|\nabla l(\mathbf{x}_{s,i}^\top \theta_{s-1,i}, y_{s,i})\|_{A_{s,i}^{-1}}^2 = O\big(\frac{d \log NT}{c_\mu}[k_\mu^2 S^2 + R_{\max}^2]\big)
$$

so $B_2 = O\big(\frac{d \log NT}{c_\mu}[k_\mu^2 S^2 + R_{\max}^2]\big)$. Hence,

$$
\beta_{t,i} = O(d\frac{R_{\max}^2}{c_\mu^2} \log NT + d\frac{k_\mu^2}{c_\mu^2} \log NT + d\frac{R_{\max}^2}{c_\mu^2} \log NT) = O(\frac{d \log NT}{c_\mu^2}[k_\mu^2 + R_{\max}^2])
$$

which finishes the proof. □

E Proof of Theorem 4.3

Proof. Since $\mu(\cdot)$ is k_μ-Lipschitz continuous, we have $\mu(\mathbf{x}_{t,\star}^\top \theta_\star) - \mu(\mathbf{x}_{t,i}^\top \theta_\star) \leq k_\mu(\mathbf{x}_{t,\star}^\top \theta_\star - \mathbf{x}_{t,i}^\top \theta_\star)$. Then we have the following upper bound on the instantaneous regret,

$$
\begin{aligned}
\frac{r_{t,i}}{k_\mu} &\leq \mathbf{x}_{t,\star}^\top \theta_\star - \mathbf{x}_{t,i}^\top \theta_\star \leq \mathbf{x}_{t,i}^\top \tilde{\theta}_{t-1,i} - \mathbf{x}_{t,i}^\top \theta_\star \\
&= \mathbf{x}_{t,i}^\top (\tilde{\theta}_{t-1,i} - \hat{\theta}_{t-1,i}) + \mathbf{x}_{t,i}^\top (\hat{\theta}_{t-1,i} - \theta_\star) \\
&\leq \|\mathbf{x}_{t,i}\|_{A_{t-1,i}^{-1}} \|\tilde{\theta}_{t-1,i} - \hat{\theta}_{t-1,i}\|_{A_{t-1,i}} + \|\mathbf{x}_{t,i}\|_{A_{t-1,i}^{-1}} \|\hat{\theta}_{t-1,i} - \theta_\star\|_{A_{t-1,i}} \\
&\leq 2\alpha_{t-1,i} \cdot \|\mathbf{x}_{t,i}\|_{A_{t-1,i}^{-1}}
\end{aligned}
$$

which holds for all $i \in [N], t \in [T]$, with probability at least $1 - 2\delta$. And $\tilde{\theta}_{t-1.i}$ denotes the optimistic estimate in the confidence ellipsoid that maximizes the UCB score when client i selects arm at time step t.

Now consider an imaginary centralized agent that has direct access to all clients' data, and we denote its covariance matrix as $\tilde{A}_{t,i} = \frac{\lambda}{c_\mu} I + \sum_{s=1}^{t-1} \sum_{j=1}^{N} \mathbf{x}_{s,j} \mathbf{x}_{s,j} + \sum_{j=1}^{i} \mathbf{x}_{t,j} \mathbf{x}_{t,j}^\top$, i.e., $\tilde{A}_{t,i}$ is immediately

updated after any client obtains a new data sample from the environment. Then we can obtain the following upper bound for $r_{t,i}$, which is dependent on the determinant ratio between the covariance matrix of the imaginary centralized agent and that of client i, i.e., $\det(\tilde{A}_{t-1,i})/\det(A_{t-1,i})$.

$$r_{t,i} \leq 2k_\mu \alpha_{t-1,i}\sqrt{\mathbf{x}_{t,i}^\top A_{t-1,i}^{-1}\mathbf{x}_{t,i}} \leq 2k_\mu \alpha_{t-1,i}\sqrt{\mathbf{x}_{t,i}^\top \tilde{A}_{t-1,i}^{-1}\mathbf{x}_{t,i} \cdot \frac{\det(\tilde{A}_{t-1,i})}{\det(A_{t-1,i})}}$$

We refer to the time period in-between two consecutive global updates as an epoch, and denote the total number of epochs as $B \in \mathbb{R}$, i.e., the p-th epoch refers to the period from $t_{p-1}+1$ to t_p, for $p \in [B]$, where t_p denotes the time step when the p-th global update happens. Then the p-th epoch is called a 'good' epoch if the determinant ratio $\frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} \leq 2$, where A_{t_p} is the aggregated sufficient statistics computed at the p-th global update. Otherwise, it is called a 'bad' epoch. In the following, we bound the cumulative regret in 'good' and 'bad' epochs separately.

Suppose the p-th epoch is a good epoch, then for any client $i \in [N]$, and time step $t \in [t_{p-1}+1, t_p]$, we have $\frac{\det(\tilde{A}_{t-1,i})}{\det(A_{t-1,i})} \leq \frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} \leq 2$, because $A_{t-1,i} \succcurlyeq A_{t_{p-1}}$ and $\tilde{A}_{t-1,i} \preccurlyeq A_{t_p}$. Therefore, the instantaneous regret incurred by any client i at any time step t of a good epoch can be bounded by

$$r_{t,i} \leq 2\sqrt{2}k_\mu \alpha_{t-1,i}\sqrt{\mathbf{x}_{t,i}^\top \tilde{A}_{t-1,i}^{-1}\mathbf{x}_{t,i}}$$

with probability at least $1-2\delta$. Therefore, using standard arguments for UCB-type algorithms, e.g., Theorem 2 in [20], the cumulative regret for all the 'good epochs' is

$$\begin{aligned} REG_{good} &\leq 2\sqrt{2}k_\mu \alpha_{t-1,i} \sum_{t=1}^{T}\sum_{i=1}^{N} \|\mathbf{x}_{t,i}\|_{\tilde{A}_{t-1,i}^{-1}} \\ &= O\left(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu} d\sqrt{NT}\log NT\right) \end{aligned}$$

which matches the regret upper bound of GLOC [12].

Now suppose the p-th epoch is bad. Then the cumulative regret incurred by all N clients during this 'bad epoch' can be upper bounded by:

$$\begin{aligned} &\sum_{t=t_{p-1}+1}^{t_p}\sum_{i=1}^{N} r_{t,i} \\ &\leq O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}\sqrt{d\log(NT)}) \sum_{t=t_{p-1}+1}^{t_p}\sum_{i=1}^{N} \min(1, ||\mathbf{x}_{t,i}||_{A_{t-1,i}^{-1}}) \\ &\leq O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}\sqrt{d\log(NT)}) \sum_{i=1}^{N}\sqrt{(t_p - t_{p-1})\log\frac{\det(A_{t_p-1,i})}{\det(A_{t_p-1,i} - \Delta A_{t_p-1,i})}} \\ &\leq O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu} N\sqrt{d\log(NT)D}) \end{aligned}$$

where the last inequality is due to the event-trigger design in Algorithm 1. Following the same argument as [28], there can be at most $R = \lceil d\log(1+\frac{NTc_\mu}{\lambda d})\rceil = O\big(d\log(NT)\big)$ 'bad epochs', because $\det(A_{t_B}) \leq \det(\tilde{A}_{T,N}) \leq (\frac{\lambda}{c_\mu} + \frac{NT}{d})^d$. Therefore, the cumulative regret for all the 'bad epochs' is

$$REG_{bad} = O\left(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu} d^{1.5}\log^{1.5}(NT)ND^{0.5}\right)$$

Combining the regret upper bound for 'good' and 'bad' epochs, the cumulative regret

$$R_T = O\left(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}(d\sqrt{NT}\log(NT) + d^{1.5}\log^{1.5}(NT)ND^{0.5})\right).$$

To obtain upper bound for the communication cost C_T, we first upper bound the total number of epochs B. Denote the length of an epoch, i.e., the number of time steps between two consecutive global updates, as $\alpha > 0$, so that there can be at most $\lceil \frac{T}{\alpha} \rceil$ epochs with length longer than α. For a particular epoch p with less than α time steps, we have $t_p - t_{p-1} < \alpha$. Moreover, due to the event-trigger design in Algorithm 1, we have $(t_p - t_{p-1}) \log \frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} > D$, which means $\log \frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} > \frac{D}{\alpha}$. Since $\sum_{p=1}^{B} \log \frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} \leq R$, the number of epochs with less than α time steps is at most $\lceil \frac{R\alpha}{D} \rceil$. Therefore, the total number of epochs.

$$B \leq \lceil \frac{T}{\alpha} \rceil + \lceil \frac{R\alpha}{D} \rceil$$

which is minimized it by choosing $\alpha = \sqrt{\frac{DT}{R}}$, so $B \leq \sqrt{\frac{TR}{D}} = O(d^{0.5} \log^{0.5}(NT) T^{0.5} D^{-0.5})$.

At the end of each epoch, FedGLB-UCB has a global update step that executes AGD among all N clients. As mentioned in Section 4.1, the number of iterations required by AGD has upper bound

$$J_t \leq 1 + \sqrt{\frac{k_\mu}{\lambda} Nt + 1} \log \frac{(k_\mu + \frac{2\lambda}{Nt}) \|\theta_t^{(1)} - \hat{\theta}_t^{\text{MLE}}\|_2^2}{2\epsilon_t},$$

and under the condition that $\epsilon_t = \frac{1}{N^2 t^2}, \forall t \in [T]$, we have $J_t = O\big(\sqrt{NT} \log(NT)\big), \forall t \in [T]$. Moreover, each iteration of AGD involves communication with N clients, so the communication cost

$$C_T = O(d^{0.5} \log^{1.5}(NT) T N^{1.5} D^{-0.5})$$

In order to match the regret under centralized setting, we set the threshold $D = \frac{T}{Nd \log(NT)}$, which gives us $R_T = O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu} d\sqrt{NT} \log(NT))$, and $C_T = O(dN^2 \sqrt{T} \log^2(NT))$. □

F Theoretical Analysis for Variants of FedGLB-UCB

In this section, we describe and analyze the variants of FedGLB-UCB listed in Table 1. The first variant, FedGLB-UCB$_1$, completely disables local update, and we can see that it requires a linear communication cost in T to attain the $O(d\sqrt{NT} \log(NT))$ regret. As we mentioned in Section 4.1, this is because in the absence of local update, FedGLB-UCB$_1$ requires more frequent global updates, i.e., $\sqrt{NT}$ in total, to control the sub-optimality of the employed bandit model w.r.t the growing training set. The second variant, denoted as FedGLB-UCB$_2$, is exactly the same as FedGLB-UCB, except for its fixed communication schedule. This leads to additional $d\sqrt{N}$ global updates, as fixed update schedule cannot adapt to the actual quality of collected data. The third variant, denoted as FedGLB-UCB$_3$, uses ONS for both local and global update, such that only one round of gradient aggregation among N clients is performed for each global update, i.e., lazy ONS update over batched data. It incurs the least communication cost among all variants, but its regret grows at a rate of $(NT)^{3/4}$ due to the inferior quality of its lazy ONS update.

F.1 FedGLB-UCB$_1$: scheduled communication + no local update

Though many real-world applications are online problems in nature, i.e., the clients continuously collect new data samples from the users, standard federated/distributed learning methods do not provide a principled solution to adapt to the growing datasets. A common practice is to manually set a fixed global update schedule in advance, i.e., periodically update and deploy the model.

To demonstrate the advantage of FedGLB-UCB over this straightforward solution, we present and analyze the first variant FedGLB-UCB$_1$, which completely disables local update, and performs global update according to a fixed schedule $\mathcal{S} = \{t_1 := \lfloor \frac{T}{B} \rfloor, t_2 := 2\lfloor \frac{T}{B} \rfloor, \ldots, t_B := B\lfloor \frac{T}{B} \rfloor\}$, where B is the total number of global updates up to time step T. The description of FedGLB-UCB$_1$ is presented in Algorithm 4.

In FedGLB-UCB$_1$, each client stores a local model $\theta_{t-1,i}$, and the corresponding covariance matrix $A_{t-1,i}$. Note that $\{\theta_{t-1,i}, A_{t-1,i}\}_{i \in [N]}$ are only updated at time steps $t \in \mathcal{S}$, and remain unchanged

Algorithm 4 FedGLB-UCB$_1$

Input: communication schedule $\mathcal{S}$, regularization parameter $\lambda > 0$, $\delta \in (0,1)$ and c_μ.
Initialize $\forall i \in [N]$: $\theta_{0,i} = \mathbf{0} \in \mathbb{R}^d, A_{0,i} = \frac{\lambda}{c_\mu}\mathbf{I} \in \mathbb{R}^{d\times d}, \mathbf{X}_{0,i} = \mathbf{0} \in \mathbb{R}^{0\times d}, \mathbf{y}_{0,i} = \mathbf{0} \in \mathbb{R}^0$, $t_{\text{last}} = 0$
for $t = 1, 2, ..., T$ **do**
  for client $i = 1, 2, ..., N$ **do**
    Observe arm set $\mathcal{A}_{t,i}$ for client i
    Select arm $x_{t,i} \in \mathcal{A}_{t,i}$ according to Eq. (16) and observe reward $y_{t,i}$
    Update client i: $\mathbf{X}_{t,i} = \begin{bmatrix} \mathbf{X}_{t-1,i} \\ \mathbf{x}_{t,i}^\top \end{bmatrix}$,$\mathbf{y}_{t,i} = \begin{bmatrix} \mathbf{y}_{t-1,i} \\ y_{t,i} \end{bmatrix}$
  if $t \notin \mathcal{S}$ **then**
    Clients: set $\theta_{t,i} = \theta_{t-1,i}, A_{t,i} = A_{t-1,i}, \forall i \in [N]$
  else
    Clients: send $\{\mathbf{X}_{t,i}^\top \mathbf{X}_{t,i}\}_{i\in[N]}$ to server
    Server compute $A_t = \frac{\lambda}{c_\mu}\mathbf{I} + \sum_{i=1}^N \mathbf{X}_{t,i}^\top \mathbf{X}_{t,i}$ and send A_t to all clients.
    Clients: set $A_{t,i} = A_t$, for $i \in [N]$
    Server update global model $\theta_t = \text{AGD-Update}(\theta_{t_{\text{last}}}, J_t)$, and set $t_{\text{last}} = t$
    Clients set local models $\theta_{t,i} = \theta_t, \forall i \in [N]$

for $t \notin \mathcal{S}$. At time t, client i selects the arm that maximizes the following UCB score:

$$\mathbf{x}_{t,i} = \arg\max_{\mathbf{x}\in\mathcal{A}_{t,i}} \mathbf{x}^\top \theta_{t-1,i} + \alpha_{t-1,i} ||\mathbf{x}||_{A_{t-1,i}^{-1}} \tag{16}$$

where $\alpha_{t-1,i}$ is given in Lemma B.5. The regret and communication cost of FedGLB-UCB$_1$ is given in the following theorem.

Theorem F.1 (Regret and Communication Cost Upper Bound of FedGLB-UCB$_1$)**.** *Under the condition that $\epsilon_t = \frac{1}{N^2t^2}$, and the total number of global synchronizations $B = \sqrt{NT}$, the cumulative regret R_T has upper bound*

$$R_T = O\left(\frac{k_\mu R_{\max} d}{c_\mu}\sqrt{NT}\log(NT/\delta)\right)$$

with probability at least $1-\delta$. The cumulative communication cost has upper bound

$$C_T = O(N^2 T \log(NT))$$

Proof. First, based on Lemma B.5 and under the condition that $\epsilon_t = \frac{1}{N^2t^2}$, we have

$$\|\theta_t - \theta_\star\|_{A_t} \le \alpha_t$$

holds $\forall t$, where $\alpha_t = \sqrt{\frac{2k_\mu}{\lambda c_\mu} + \frac{2}{Ntc_\mu}} + \frac{R_{max}}{c_\mu}\sqrt{d\log(1 + Ntc_\mu/(d\lambda)) + 2\log(1/\delta)} + \sqrt{\frac{\lambda}{c_\mu}}S = O(\frac{R_{max}}{c_\mu}\sqrt{d\log(Nt)})$, which matches the order in [20].

Similar to the proof of Theorem 4.3, we decompose all B epochs into 'good' and 'bad' epochs according to the log-determinant ratio: the p-th epoch, for $p \in [B]$, is a 'good' epoch if the determinant ratio $\frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} \le 2$. Otherwise, it is a 'bad' epoch. In the following, we bound the cumulative regret in 'good' and 'bad' epochs separately.

Suppose epoch p is a good epoch, then for any client $i \in [N]$, and time step $t \in [t_{p-1}+1, t_p]$, we have $\frac{\det(\tilde{A}_{t-1,i})}{\det(A_{t-1,i})} \le \frac{\det(A_{t_p})}{\det(A_{t_{p-1}})} \le 2$, because $A_{t-1,i} = A_{t_{p-1}}$ and $\tilde{A}_{t-1,i} \preccurlyeq A_{t_p}$. Therefore, the instantaneous regret incurred by any client i at any time step t of a good epoch p can be bounded by

$$\begin{aligned} r_{t,i} &\le 2k_\mu \alpha_{t_{p-1}} \sqrt{\mathbf{x}_{t,i}^\top A_{t-1,i}\mathbf{x}_{t,i}} \le 2k_\mu \alpha_{t_{p-1}} \sqrt{\mathbf{x}_{t,i}^\top A_{t-1}^{-1}\mathbf{x}_{t,i} \cdot \frac{\det(\tilde{A}_{t-1,i})}{\det(A_{t-1,i})}} \\ &\le 2\sqrt{2}k_\mu \alpha_T \sqrt{\mathbf{x}_{t,i}^\top A_{t-1}^{-1}\mathbf{x}_{t,i}} \end{aligned}$$

By standard arguments [1, 20], the cumulative regret incurred in all good epochs can be upper bounded by $O(\frac{k_\mu R_{\max}}{c_\mu} d\sqrt{NT}\log(NT/\delta))$ with probability at least $1-\delta$.

By Assumption 1, $\mu(\cdot)$ is Lipschitz continuous with constant k_μ, i.e., $|\mu(\mathbf{x}^\top\theta_1) - \mu(\mathbf{x}^\top\theta_2)| \leq k_\mu|\mathbf{x}^\top(\theta_1-\theta_2)|$, so the instantaneous regret $r_{t,i}$ is uniformly bounded $\forall t \in [T], i \in [N]$ by $2k_\mu S$. Now suppose epoch p is bad, then we can upper bound the cumulative regret in this bad epoch by $2k_\mu S\frac{NT}{B}$, where $\frac{NT}{B}$ is the number of time steps in each epoch. Since there can be at most $O(d\log NT)$ bad epochs, the cumulative regret incurred in all bad epochs can be upper bounded by $O(\frac{NT}{B}k_\mu Sd\log(NT))$. Combining both parts together, the cumulative regret upper bound is

$$R_T = O\left(\frac{NT}{B}k_\mu Sd\log(NT) + \frac{k_\mu R_{\max} d}{c_\mu}\sqrt{NT}\log(NT)\right)$$

To recover the regret under centralized setting, we set $B=\sqrt{NT}$, so

$$R_T = O\left(\frac{k_\mu R_{max}}{c_\mu} d\sqrt{NT}\log(NT)\right)$$

Note that FedGLB-UCB$_1$ has $B=\sqrt{NT}$ global updates in total, and during each global update, there are J_t rounds of communications, for $t\in\mathcal{S}$. As mentioned earlier, for AGD to attain $\epsilon_t = \frac{1}{N^2t^2}$ sub-optimality, the required number of inner iterations

$$J_t \leq 1 + \sqrt{\frac{k_\mu + \frac{\lambda}{Nt}}{\frac{\lambda}{Nt}}}\log\frac{(k_\mu + \frac{\lambda}{Nt} + \frac{\lambda}{Nt})\|\theta_t^{(0)} - \hat{\theta}_t^{\text{MLE}}\|_2^2}{2\epsilon_t} = O\left(\sqrt{Nt}\log(Nt)\right)$$

Therefore, the communication cost over time horizon T is

$$\begin{aligned}
C_T &= N\cdot\sum_{t\in\mathcal{S}} J_t \\
&= N\cdot[\sqrt{\sqrt{NT}}\log(\sqrt{NT}) + \sqrt{2\sqrt{NT}}\log(2\sqrt{NT}) + \cdots + \sqrt{\sqrt{NT}\cdot\sqrt{NT}}\log(\sqrt{NT}\cdot\sqrt{NT})] \\
&\leq N^{5/4}T^{1/4}\log(NT)[\sqrt{1}+\sqrt{2}+\cdots+\sqrt{\sqrt{NT}}] \\
&\leq N^{5/4}T^{1/4}\log(NT)\cdot\frac{3}{2}(\sqrt{NT}+\frac{1}{2})^{3/2} \\
&= O(N^2T\log(NT))
\end{aligned}$$

which finishes the proof. □

F.2 FedGLB-UCB$_2$: scheduled communication

For the second variant FedGLB-UCB$_2$, we enabled local update on top of FedGLB-UCB$_1$. Therefore, compared with the original algorithm FedGLB-UCB, the only difference is that FedGLB-UCB$_2$ uses scheduled communication instead of event-triggered communication. Its description is given in Algorithm 5.

The regret and communication cost of FedGLB-UCB$_2$ is given in the following theorem.

Theorem F.2 (Regret and Communication Cost Upper Bound of FedGLB-UCB$_2$)**.** *Under the condition that $\epsilon_t = \frac{1}{N^2t^2}$, and the total number of global synchronizations $B = d^2N\log(NT)$, the cumulative regret R_T has upper bound*

$$R_T = O\left(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu} d\sqrt{NT}\log(NT/\delta)\sqrt{\log\frac{T}{d^2N\log NT}}\right)$$

with probability at least $1-\delta$. The cumulative communication cost has upper bound

$$C_T = O(d^2N^{2.5}\sqrt{T}\log^2(NT))$$

Proof. Compared with the analysis for FedGLB-UCB, the main difference in the analysis for FedGLB-UCB$_2$ is how we bound the regret incurred in 'bad epochs'. Using the same argument, the cumulative regret for the 'good epochs' is $REG_{good} = O(\frac{k_\mu(k_\mu+R_{\max})}{c_\mu} d\sqrt{NT}\log NT/\delta)$.

Algorithm 5 FedGLB-UCB$_2$

Input: communication schedule $\mathcal{S}$, regularization parameter $\lambda > 0$, $\delta \in (0,1)$ and c_μ.
Initialize $\forall i \in [N]$: $A_{0,i} = \frac{\lambda}{c_\mu}\mathbf{I} \in \mathbb{R}^{d\times d}, b_{0,i} = \mathbf{0} \in \mathbb{R}^d, \theta_{0,i} = \mathbf{0} \in \mathbb{R}^d, \Delta A_{0,i} = \mathbf{0} \in \mathbb{R}^{d\times d}$; $A_0 = \frac{\lambda}{c_\mu}\mathbf{I} \in \mathbb{R}^{d\times d}, b_0 = \mathbf{0} \in \mathbb{R}^d, \theta_0 = \mathbf{0} \in \mathbb{R}^d$, $t_{\text{last}} = 0$
for $t = 1, 2, ..., T$ **do**
  for client $i = 1, 2, ..., N$ **do**
    Observe arm set $\mathcal{A}_{t,i}$ for client i
    Select arm $\mathbf{x}_{t,i} \in \mathcal{A}_{t,i}$ by Eq.(5), and observe reward $y_{t,i}$
    Update client i: $A_{t,i} = A_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top$, $\Delta A_{t,i} = \Delta A_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top$
  if $t \notin \mathcal{S}$ **then**
    Clients $\forall i \in [N]$: $\theta_{t,i} = \text{ONS-Update}(\theta_{t-1,i}, A_{t,i}, \nabla l(\mathbf{x}_{t,i}^\top\theta_{t-1,i}, y_{t,i}))$, $b_{t,i} = b_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top\theta_{t-1,i}$
  else
    Clients $\forall i \in [N]$: send $\Delta A_{t,i}$ to server, and reset $\Delta A_{t,i} = \mathbf{0}$
    Server compute $A_t = A_{t_{\text{last}}} + \sum_{i=1}^N \Delta A_{t,i}$
    Server perform global model update $\theta_t = \text{AGD-Update}(\theta_{t_{\text{last}}}, J_t)$ (see Eq.(3) for choice of J_t), $b_t = b_{t_{\text{last}}} + \sum_{i=1}^N \Delta A_{t,i}\theta_t$, and set $t_{\text{last}} = t$
    Clients $\forall i \in [N]$: set $\theta_{t,i} = \theta_t, A_{t,i} = A_t, b_{t,i} = b_t$

Now consider a particular bad epoch $p \in [B]$. Then the cumulative regret incurred by all N clients during this 'bad epoch' can be upper bounded by:

$$
\begin{aligned}
&\sum_{t=t_{p-1}+1}^{t_p}\sum_{i=1}^N r_{t,i} \\
&\le O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}\sqrt{d\log(NT/\delta)})\sum_{t=t_{p-1}+1}^{t_p}\sum_{i=1}^N \min(1, ||\mathbf{x}_{t,i}||_{A_{t-1,i}^{-1}}) \\
&\le O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}\sqrt{d\log(NT/\delta)})\sum_{i=1}^N\sqrt{(t_p - t_{p-1})\log\frac{\det(A_{t_p-1,i})}{\det(A_{t_p-1,i} - \Delta A_{t_p-1,i})}} \\
&\le O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}dN\sqrt{\log(NT/\delta)}\sqrt{\frac{T}{B}\log(\frac{T}{B})})
\end{aligned}
$$

where the last inequality is because all epochs has length $\frac{T}{B}$ as defined by $\mathcal{S}$. Again, since there can be at most $O(d\log NT)$ 'bad epochs', the cumulative regret for the 'bad epochs' is upper bounded by

$$
REG_{bad} = O(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}d^2\log^{1.5}(NT/\delta)N\sqrt{\frac{T}{B}\log(\frac{T}{B})}).
$$

Combining the cumulative regret for both 'good' and 'bad' epochs, and setting $B = d^2N\log(NT)$, we have

$$
R_T = O\left(\frac{k_\mu(k_\mu + R_{\max})}{c_\mu}d\sqrt{NT}\log(NT/\delta)\sqrt{\log(\frac{T}{d^2N\log NT})}\right)
$$

Now that FedGLB-UCB$_2$ has $B = d^2N\log(NT)$ global updates in total, and during each global update, there are $J_t = O(\sqrt{NT}\log(NT))$ rounds of communications, for $t \in \mathcal{S}$. Therefore, the communication cost over time horizon T is

$$
\begin{aligned}
C_T = N\cdot\sum_{t\in\mathcal{S}} J_t &= O(N\cdot d^2N\log(NT)\cdot\sqrt{NT}\log(NT)) \\
&= O(d^2N^{2.5}\sqrt{T}\log^2(NT))
\end{aligned}
$$

which finishes the proof. □

F.3 FedGLB-UCB$_3$: scheduled communication + ONS for global update

The previous two variants both adopt iterative optimization method, i.e., AGD, for the global update, which introduces a $\sqrt{NT}\log(NT)$ factor in the communication cost. In this section, we try to avoid this by studying the third variant FedGLB-UCB$_3$ that adopts ONS for both local and global update, such that only one step of ONS is performed (based on all new data samples N clients collected in this epoch). It can be viewed as the ONS-GLM algorithm [12] with lazy batch update.

Algorithm 6 FedGLB-UCB$_3$

Input: communication schedule $\mathcal{S}$, regularization parameter $\lambda > 0$, $\delta \in (0,1)$ and c_μ
Initialize $\forall i \in [N]$: $\theta_{0,i} = \mathbf{0} \in \mathbb{R}^d, A_{0,i} = \lambda\mathbf{I} \in \mathbb{R}^{d\times d}, V_{0,i} = \lambda\mathbf{I} \in \mathbb{R}^{d\times d}, b_{0,i} = \mathbf{0} \in \mathbb{R}^d$; $\theta_0 = \mathbf{0} \in \mathbb{R}^d, A_0 = \lambda\mathbf{I} \in \mathbb{R}^{d\times d}, V_0 = \lambda\mathbf{I} \in \mathbb{R}^{d\times d}, b_0 = \mathbf{0} \in \mathbb{R}^d$, $t_{\text{last}} = 0$
for $t = 1, 2, ..., T$ **do**
  for client $i = 1, 2, ..., N$ **do**
    Observe arm set $\mathcal{A}_{t,i}$ for client $i \in [N]$
    Select arm $\mathbf{x}_{t,i} = \arg\max_{\mathbf{x}\in\mathcal{A}_{t,i}} \mathbf{x}^\top \hat{\theta}_{t-1,i} + \alpha_{t-1,i}\|\mathbf{x}\|_{V_{t-1,i}^{-1}}$, where $\hat{\theta}_{t-1,i} = V_{t-1,i}^{-1} b_{t-1,i}$ and $\alpha_{t-1,i}$ is given in Lemma F.4; and then observe reward $y_{t,i}$
    Compute loss $l(z_{t,i}, y_{t,i})$, where $z_{t,i} = \mathbf{x}_{t,i}^\top \theta_{t-1,i}$
    Update client i: $A_{t,i} = A_{t-1,i} + \nabla l(z_{t,i}, y_{t,i})\nabla l(z_{t,i}, y_{t,i})^\top$, $V_{t,i} = V_{t-1,i} + \mathbf{x}_{t,i}\mathbf{x}_{t,i}^\top$
  if $t \notin \mathcal{S}$ **then**
    Clients $\forall i \in [N]$: $\theta_{t,i} = \text{ONS-Update}(\theta_{t-1,i}, A_{t,i}, \nabla l(z_{t,i}, y_{t,i}))$, $b_{t,i} = b_{t-1,i} + \mathbf{x}_{t,i} z_{t,i}$
  else
    Clients $\forall i \in [N]$: send gradient $\nabla F_{t,i}(\theta_{t_{\text{last}}}) = \sum_{s=t_{\text{last}}+1}^{t} \nabla l(\mathbf{x}_{s,i}^\top \theta_{t_{\text{last}}}, y_{s,i})$ and $\Delta V_{t,i} = V_{t,i} - V_{t_{\text{last}},i}$ to server
    Server $A_t = A_{t_{\text{last}}} + (\sum_{i=1}^N \nabla F_{t,i}(\theta_{t_{\text{last}}}))(\sum_{i=1}^N \nabla F_{t,i}(\theta_{t_{\text{last}}}))^\top$, $V_t = V_{t_{\text{last}}} + \sum_{i=1}^N \Delta V_{t,i}$, $b_t = b_{t_{\text{last}}} + \sum_{i=1}^N \Delta V_{t,i}\theta_{t_{\text{last}}}$, $\theta_t = \text{ONS-Update}(\theta_{t_{\text{last}}}, A_t, \sum_{i=1}^N \nabla F_{t,i}(\theta_{t_{\text{last}}}))$
    Clients $\forall i \in [N]$: $\theta_{t,i} = \theta_t, A_{t,i} = A_t, V_t = V_t, b_{t,i} = b_t$
    Set $t_{\text{last}} = t$

Recall that the update schedule is denoted as $\mathcal{S} = \{t_1 := \lfloor \frac{T}{B} \rfloor, t_2 := 2\lfloor \frac{T}{B} \rfloor, \ldots, t_q := q\lfloor \frac{T}{B} \rfloor, \ldots, t_B := B\lfloor \frac{T}{B} \rfloor\}$, where B denotes the total number of global updates up to T. Compared with [12], the main difference in our construction is that the loss function in the online regression problem may contain multiple data samples, i.e., for global update, or one single data sample, i.e., for local update. Then for a client $i \in [N]$ at time step t (suppose t is in the $(q+1)$-th epoch, so $t \in [t_q + 1, t_{q+1}]$), the sequence of loss functions observed by the online regression estimator till time t is:

$$\underbrace{\sum_{s=1}^{t_1}\sum_{i=1}^{N} l(\mathbf{x}_{s,i}^\top\theta_0, y_{s,i}), \sum_{s=t_1+1}^{t_2}\sum_{i=1}^{N} l(\mathbf{x}_{s,i}^\top\theta_{t_1}, y_{s,i}), \ldots, \sum_{s=t_{q-1}+1}^{t_q}\sum_{i=1}^{N} l(\mathbf{x}_{s,i}^\top\theta_{t_{q-1}}, y_{s,i})}_{\text{global updates at } t_1,\, t_2,\, \ldots,\, t_q}, \underbrace{l(\mathbf{x}_{t_q+1,i}^\top\theta_{t_q}, y_{t_q+1,i}), \ldots, l(\mathbf{x}_{t,i}^\top\theta_{t-1,i}, y_{t,i})}_{\text{local updates at } t_q+1,\, \ldots,\, t}$$

We can see that the first q terms correspond to the global ONS updates that are computed using the whole batch of data collected by N clients in each epoch, and the remaining $t - t_q$ terms are local ONS updates that are computed using each new data sample collected by client i in the $(q+1)$-th epoch.

To facilitate further analysis, we introduce a new set of indices for the data samples, so that we can unify the notations for the loss functions above. Imagine all the arm pulls are performed by an imaginary centralized agent, such that, in each time step $t \in [T]$, it pulls an arm for clients $1, 2, \ldots, N$ one by one. Therefore, the sequence of data sample obtained by this imaginary agent can be denoted as $(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_s, y_s), \ldots, (\mathbf{x}_{NT}, y_{NT})$. Moreover, we denote n_p as the total number of data samples collected by all N clients till the p-th ONS update (including both global and local ONS update), and denote the updated model as θ_p, for $p \in [P]$. Note that P denotes the total number of updates up to time t (total number of terms in the sequence above), such that $P = q + t - t_q$. Then this sequence of loss functions can be rewritten as:

$$\underbrace{F_1(\theta_0), F_2(\theta_1), \ldots, F_q(\theta_{q-1})}_{\text{global updates}}, \underbrace{F_{q+1}(\theta_q), \ldots, F_P(\theta_{P-1})}_{\text{local updates}}$$

where $F_p(\theta_{p-1}) = \sum_{s=n_{p-1}+1}^{n_p} l(\mathbf{x}_s^\top \theta_{p-1}, y_s)$, for $p \in [P]$.

• **Online regret upper bound for lazily-updated ONS** To construct the confidence ellipsoid based on this sequence of global and local ONS updates, we first need to upper bound the online regret that ONS incurs on this sequence of loss functions, which is given in Lemma F.3.

Lemma F.3 (Online regret upper bound)**.** *Under the condition that the learning rate of ONS is set to* $\gamma = \frac{1}{2}\min(\frac{1}{4S\sqrt{k_\mu^2 S^2 + R_{\max}^2}}, \frac{c_\mu}{(k_\mu^2 S^2 + R_{\max}^2)\max_{p\in[P]}(n_p - n_{p-1})})$*, then the cumulative online regret over* P *steps*

$$\sum_{p=1}^{P} F_p(\theta_{p-1}) - F_p(\theta_\star) \leq B_P$$

where $B_P = \frac{1}{2\gamma}\sum_{p=1}^{P} ||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + 2\gamma\lambda S^2$.

Proof of Lemma F.3. Recall from the proof of Corollary 4.2.1 that $|\mu(\mathbf{x}_s^\top \theta_{p-1}) - y_s| \leq \sqrt{k_\mu^2 S^2 + R_{\max}^2} := G, \forall s$. First, we need to show that $F_p(\theta_{p-1}) = \sum_{s=n_{p-1}+1}^{n_p} l(\mathbf{x}_s^\top \theta_{p-1}, y_s)$ is $\frac{c_\mu}{(n_p - n_{p-1})G^2}$-exp-concave, or equivalently, $\nabla^2 F_p(\theta_{p-1}) \succcurlyeq \frac{c_\mu}{(n_p - n_{p-1})G^2}\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top$ (Lemma 4.2 in [9]). Taking first and second order derivative of $F_p(\theta_{p-1})$ w.r.t. θ_{p-1}, we have

$$\nabla F_p(\theta_{p-1}) = \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s[-y_s + \mu(\mathbf{x}_s^\top \theta_{p-1})] = \mathbf{X}_p^\top[\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p],$$

$$\nabla^2 F_p(\theta_{p-1}) = \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s\mathbf{x}_s^\top \dot{\mu}(\mathbf{x}_s^\top \theta_{p-1})$$

where $\mathbf{X}_p = [\mathbf{x}_{n_{p-1}+1}, \mathbf{x}_{n_{p-1}+2}, \ldots, \mathbf{x}_{n_p}]^\top \in \mathbb{R}^{(n_p - n_{p-1})\times d}$, and $\mathbf{y}_p = [y_{n_{p-1}+1}, y_{n_{p-1}+2}, \ldots, y_{n_p}]^\top \in \mathbb{R}^{n_p - n_{p-1}}$. Then due to Assumption 1, we have $\nabla^2 F_p(\theta_{p-1}) \succcurlyeq c_\mu \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s\mathbf{x}_s^\top = c_\mu \mathbf{X}_p^\top \mathbf{X}_p$. For any vector $u \in \mathbb{R}^d$, we can show that,

$$\begin{aligned}
&u^\top \nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top u \\
&= u^\top \mathbf{X}_p^\top[\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p][\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p]^\top \mathbf{X}_p u \\
&= \left[(\mathbf{X}_p u)^\top[\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p]\right]^2 \\
&\leq \|\mathbf{X}_p u\|_2^2 \cdot \|\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p\|_2^2 \\
&\leq u^\top \mathbf{X}_p^\top \mathbf{X}_p u \cdot (n_p - n_{p-1})G^2
\end{aligned}$$

where the first inequality is due to Cauchy-Schwarz inequality, and the second inequality is because $\|\mu(\mathbf{X}_p\theta_{p-1}) - \mathbf{y}_p\|_2^2 = \sum_{s=n_{p-1}+1}^{n_p}[-y_s + \mu(\mathbf{x}_s^\top \theta_{p-1})]^2 \leq (n_p - n_{p-1})G^2$. Therefore, $\mathbf{X}_p^\top \mathbf{X}_p \succcurlyeq \frac{1}{(n_p - n_{p-1})G^2}\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top$, which gives us

$$\nabla^2 F_p(\theta_{p-1}) \succcurlyeq \frac{c_\mu}{(n_p - n_{p-1})G^2}\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top$$

Then due to Lemma 4.3 of [9], under the condition that $\gamma_p \leq \frac{1}{2}\min(\frac{1}{4GS}, \frac{c_\mu}{(n_p - n_{p-1})G^2})$, we have

$$\begin{aligned}
&F_p(\theta_{p-1}) - F_p(\theta_\star) \\
&\quad \leq \nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) - \frac{\gamma_p}{2}(\theta_{p-1} - \theta_\star)^\top \nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) \qquad (17)
\end{aligned}$$

Then we start to upper bound the RHS of the inequality above. Recall that the ONS update rule is:

$$\theta_p' = \theta_{p-1} - \frac{1}{\gamma}A_p^{-1}\nabla F_p(\theta_{p-1})$$

$$\theta_p = \arg\min_{\theta\in\Theta} ||\theta_p' - \theta||_{A_p}^2$$

where $A_p = \sum_{\rho=1}^{p} \nabla F_\rho(\theta_{\rho-1}) \nabla F_\rho(\theta_{\rho-1})^\top$, and γ is set to $\min_{p\in[P]} \gamma_p = \frac{1}{2}\min(\frac{1}{4GS}, \frac{c_\mu}{G^2 \max_{p\in[P]}(n_p - n_{p-1})})$. So we have

$$\theta_p' - \theta_\star = \theta_{p-1} - \theta_\star - \frac{1}{\gamma} A_p^{-1} \nabla F_p(\theta_{p-1})$$

Then due to the property of the generalized projection, and by substituting into the update rule, we have

$$||\theta_p - \theta_\star||_{A_p}^2 \leq ||\theta_p' - \theta_\star||_{A_p}^2 \leq ||\theta_{p-1} - \theta_\star||_{A_p}^2 - \frac{2}{\gamma}(\theta_{p-1} - \theta_\star)^\top \nabla F_p(\theta_{p-1}) + \frac{1}{\gamma^2}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2$$

By rearranging terms,

$$\nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) \leq \frac{1}{2\gamma}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + \frac{\gamma}{2}\big(||\theta_{p-1} - \theta_\star||_{A_p}^2 - ||\theta_p - \theta_\star||_{A_p}^2\big)$$

After summing over P steps, we have

$$\sum_{p=1}^{P} \nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) \leq \frac{1}{2\gamma}\sum_{p=1}^{P}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + \frac{\gamma}{2}\sum_{p=1}^{P}\big(||\theta_{p-1} - \theta_\star||_{A_p}^2 - ||\theta_p - \theta_\star||_{A_p}^2\big)$$

The second term can be simplified,

$$\begin{aligned}
&\sum_{p=1}^{P}\big(||\theta_{p-1} - \theta_\star||_{A_p}^2 - ||\theta_p - \theta_\star||_{A_p}^2\big) \\
&= ||\theta_0 - \theta_\star||_{A_1}^2 + \sum_{p=2}^{P}\big(||\theta_{p-1} - \theta_\star||_{A_p}^2 - ||\theta_{p-1} - \theta_\star||_{A_{p-1}}^2\big) - ||\theta_P - \theta_\star||_{A_P}^2 \\
&\leq ||\theta_0 - \theta_\star||_{A_1}^2 + \sum_{p=2}^{P}\big(||\theta_{p-1} - \theta_\star||_{A_p}^2 - ||\theta_{p-1} - \theta_\star||_{A_{p-1}}^2\big) \\
&= ||\theta_0 - \theta_\star||_{A_1}^2 + \sum_{p=2}^{P}||\theta_{p-1} - \theta_\star||_{\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top}^2 \\
&= ||\theta_0 - \theta_\star||_{A_1}^2 + \sum_{p=1}^{P}||\theta_{p-1} - \theta_\star||_{\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top}^2 - ||\theta_0 - \theta_\star||_{\nabla F_1(\theta_0)\nabla F_1(\theta_0)^\top}^2 \\
&= 4\lambda S^2 + \sum_{p=1}^{P}||\theta_{p-1} - \theta_\star||_{\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top}^2
\end{aligned}$$

which leads to

$$\begin{aligned}
\sum_{p=1}^{P} \nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) \leq \frac{1}{2\gamma}\sum_{p=1}^{P}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + 2\gamma\lambda S^2 \\
+ \frac{\gamma}{2}\sum_{p=1}^{P}||\theta_{p-1} - \theta_\star||_{\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top}^2
\end{aligned}$$

By rearranging terms, we have

$$\begin{aligned}
&\sum_{p=1}^{P}\big[\nabla F_p(\theta_{p-1})^\top(\theta_{p-1} - \theta_\star) - \frac{\gamma}{2}||\theta_{p-1} - \theta_\star||_{\nabla F_p(\theta_{p-1})\nabla F_p(\theta_{p-1})^\top}^2\big] \\
&\leq \frac{1}{2\gamma}\sum_{p=1}^{P}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + 2\gamma\lambda S^2
\end{aligned}$$

Combining with Eq.(17), we obtain the following upper bound for the P-step online regret

$$\sum_{p=1}^{P} F_p(\theta_{p-1}) - F_p(\theta_\star) \leq \frac{1}{2\gamma}\sum_{p=1}^{P}||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 + 2\gamma\lambda S^2$$

where $A_p = \sum_{\rho=1}^{p} \nabla F_\rho(\theta_{\rho-1})\nabla F_\rho(\theta_{\rho-1})^\top$. □

Corollary F.3.1 (Order of B_P)**.** *Under the condition that* $\gamma = \frac{1}{2}\min(\frac{1}{4S\sqrt{k_\mu^2 S^2 + R_{\max}^2}}, \frac{c_\mu}{(k_\mu^2 S^2 + R_{\max}^2)\max_{p\in[P]}(n_p - n_{p-1})})$*, the online regret upper bound* $B_P = O(\frac{k_\mu^2 + R_{max}^2}{c_\mu} d \log(n_P) \max_{p\in[P]}(n_p - n_{p-1}))$.

Proof of Corollary F.3.1. Recall that $A_p = \sum_{\rho=1}^{p} \nabla F_\rho(\theta_{\rho-1}) \nabla F_\rho(\theta_{\rho-1})^\top$. Therefore, we have

$$\sum_{p=1}^{P} ||\nabla F_p(\theta_{p-1})||_{A_p^{-1}}^2 \leq \log \frac{\det(A_P)}{\det(\lambda I)} = \log \frac{\det(\lambda I + \sum_{p=1}^{P} \nabla F_p(\theta_{p-1}) \nabla F_p(\theta_{p-1})^\top)}{\det(\lambda I)}$$
$$\leq d \log\big(1 + \frac{1}{d\lambda} \sum_{p=1}^{P} \|\nabla F_p(\theta_{p-1})\|_2^2\big)$$

where the first inequality is due to Lemma 11 of [1], and the second due to the determinant-trace inequality (Lemma 10 of [1]), i.e., $\det(\lambda I + \sum_{p=1}^{P} \nabla F_p(\theta_{p-1}) \nabla F_p(\theta_{p-1})^\top) \leq \big(\frac{tr(\lambda I + \sum_{p=1}^{P} \nabla F_p(\theta_{p-1}) \nabla F_p(\theta_{p-1})^\top)}{d}\big)^d = \big(\frac{d\lambda + \sum_{p=1}^{P} \|\nabla F_p(\theta_{p-1})\|_2^2}{d}\big)^d$. Also note that $\nabla F_p(\theta_{p-1}) = \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s \big[\mu(\mathbf{x}_s^\top \theta_{p-1}) - y_s\big]$, so we have

$$\sum_{p=1}^{P} ||\nabla F_p(\theta_{p-1})||_2^2 = \sum_{p=1}^{P} || \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s \big[\mu(\mathbf{x}_s^\top \theta_{p-1}) - y_s\big] ||_2^2$$
$$\leq G^2 \sum_{p=1}^{P} || \sum_{s=n_{p-1}+1}^{n_p} \mathbf{x}_s ||_2^2 \leq G^2 \sum_{p=1}^{P} (n_p - n_{p-1})^2 \leq G^2 n_P^2$$

where the second inequality is due to Jensen's inequality and the assumption that $\|\mathbf{x}_s\| \leq 1, \forall s$. Substituting this back gives us

$$\sum_{p=1}^{P} F_p(\theta_{p-1}) - F_p(\theta_\star) \leq \frac{1}{2\gamma} d \log\big(1 + \frac{1}{d\lambda} G^2 n_P^2\big) + 2\gamma\lambda S^2$$
$$= \frac{(k_\mu^2 S^2 + R_{\max}^2)\max_{p\in[P]}(n_p - n_{p-1})}{c_\mu} d \log\big(1 + \frac{1}{d\lambda}(k_\mu^2 S^2 + R_{\max}^2) n_P^2\big)$$
$$+ \frac{c_\mu}{(k_\mu^2 S^2 + R_{\max}^2)\max_{p\in[P]}(n_p - n_{p-1})} \lambda S^2$$

where the equality is because $\max_{p\in[P]}(n_p - n_{p-1})$ dominates $\gamma = \frac{1}{2}\min(\frac{1}{4GS}, \frac{c_\mu}{G^2 \max_{p\in[P]}(n_p - n_{p-1})})$. □

• **Construct Confidence Ellipsoid for FedGLB-UCB$_3$** With the online regret bound B_P in Lemma F.3, the steps to construct the confidence ellipsoid largely follows that of Theorem 1 in [12], with the main difference in our batch update. We include the full proof here for the sake of completeness.

Lemma F.4 (Confidence Ellipsoid for FedGLB-UCB$_3$)**.** *Under the condition that the learning rate of ONS* $\gamma = \frac{1}{2}\min(\frac{1}{4S\sqrt{k_\mu^2 S^2 + R_{\max}^2}}, \frac{c_\mu}{(k_\mu^2 S^2 + R_{\max}^2)\max_{p\in[P]}(n_p - n_{p-1})})$*, we have* $\forall t \in [T], i \in [N]$

$$\|\theta_\star - \hat{\theta}_{t,i}\|_{V_{t,i}}^2 \leq \lambda S^2 + 1 + \frac{4}{c_\mu} B_P + \frac{8R_{max}^2}{c_\mu^2} \log\big(\frac{N}{\delta}\sqrt{4 + \frac{8}{c_\mu} B_P + \frac{64 R_{max}^2}{c_\mu^4 \cdot 4\delta^2}}\big)\}$$
$$- \hat{\theta}_{t,i}^\top b_{t,i} - \sum_{s=1}^{n_P} z_s^2 := \alpha_{t,i}^2$$

with probability at least $1 - \delta$.

Proof of Lemma F.4. Due to c_μ-strongly convexity of $l(z,y)$ w.r.t. z, we have $l(\mathbf{x}_s^\top\theta_{p-1}, y_s) - l(\mathbf{x}_s^\top\theta_\star, y_s) \geq \big[\mu(\mathbf{x}_s^\top\theta_\star) - y_s\big]\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) + \frac{c_\mu}{2}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2$. Therefore,

$$
\begin{aligned}
F_p(\theta_{p-1}) - F_p(\theta_\star) &= \sum_{s=n_{p-1}+1}^{n_p} l(\mathbf{x}_s^\top\theta_{p-1}, y_s) - l(\mathbf{x}_s^\top\theta_\star, y_s) \\
&\geq \sum_{s=n_{p-1}+1}^{n_p} \big[\mu(\mathbf{x}_s^\top\theta_\star) - y_s\big]\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) + \frac{c_\mu}{2}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2 \\
&= -\sum_{s=n_{p-1}+1}^{n_p} \eta_s\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) + \frac{c_\mu}{2}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2
\end{aligned}
$$

where η_s is the R-sub-Gaussian noise in the reward y_s. Summing over P steps we have

$$
B_P \geq \sum_{p=1}^{P} F_p(\theta_{p-1}) - F_p(\theta_\star) \geq \sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p} \eta_s\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) + \frac{c_\mu}{2}\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2
$$

By rearranging terms, we have

$$
\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2 \leq \frac{2}{c_\mu}\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\eta_s\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) + \frac{2}{c_\mu}B_P
$$

Then as $\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)$ for $s \in [n_{p-1}+1, n_p]$ is $\mathcal{F}_s$-measurable for lazily updated online estimator θ_{p-1}, we can use Corollary 8 from [2], which leads to

$$
\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\eta_s\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star) \leq
$$

$$
R_{max}\sqrt{\big(2+2\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}(\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star))^2\big)\cdot\log\big(\frac{1}{\delta}\sqrt{1+\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}(\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star))^2}\big)}
$$

Then we have

$$
\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2 \leq \frac{2}{c_\mu}B_P
$$

$$
+\frac{2R_{max}}{c_\mu}\sqrt{\big(2+2\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}(\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star))^2\big)\cdot\log\big(\frac{1}{\delta}\sqrt{1+\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}(\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star))^2}\big)}
$$

Then by applying Lemma 2 from [12], we have

$$
\sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2 \leq 1 + \frac{4}{c_\mu}B_P + \frac{8R_{max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{4+\frac{8}{c_\mu}B_P+\frac{64R_{max}^2}{c_\mu^4\cdot 4\delta^2}}\big)
$$

Therefore, we have the following confidence ellipsoid (regularized with parameter λ):

$$
\{\theta : \sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\big[\mathbf{x}_s^\top(\theta_{p-1}-\theta_\star)\big]^2 + \lambda\|\theta\|_2^2 \leq \lambda S^2 + 1 + \frac{4}{c_\mu}B_P + \frac{8R_{max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{4+\frac{8}{c_\mu}B_P+\frac{64R_{max}^2}{c_\mu^4\cdot 4\delta^2}}\big)\}
$$

And this can be rewritten as a ellipsoid centered at ridge regression estimator $\hat{\theta}_{t,i} = V_{t,i}^{-1}b_{t,i}$, where $V_{t,i} = \lambda I + \sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\mathbf{x}_s\mathbf{x}_s^\top$ and $b_{t,i} = \sum_{p=1}^{P}\sum_{s=n_{p-1}+1}^{n_p}\mathbf{x}_s z_s$ (recall that ONS's prediction at time s is denoted as $z_s = \mathbf{x}_s^\top\theta_{p-1}$), i.e., $\forall t \in [T]$

$$
\|\theta_\star - \hat{\theta}_{t,i}\|_{V_{t,i}}^2 \leq \lambda S^2 + 1 + \frac{4}{c_\mu}B_P + \frac{8R_{max}^2}{c_\mu^2}\log\big(\frac{1}{\delta}\sqrt{4+\frac{8}{c_\mu}B_P+\frac{64R_{max}^2}{c_\mu^4\cdot 4\delta^2}}\big)\} + \hat{\theta}_{t,i}^\top b_{t,i} - \sum_{s=1}^{n_P} z_s^2
$$

with probability at least $1-\delta$. Then taking union bound over all N clients, we have, $\forall t \in [T], i \in [N]$

$$\|\theta_\star - \hat{\theta}_{t,i}\|^2_{V_{t,i}} \leq \lambda S^2 + 1 + \frac{4}{c_\mu} B_P + \frac{8R^2_{max}}{c^2_\mu} \log\big(\frac{N}{\delta}\sqrt{4 + \frac{8}{c_\mu} B_P + \frac{64R^2_{max}}{c^4_\mu \cdot 4\delta^2}}\big)\} + \hat{\theta}^\top_{t,i} b_{t,i} - \sum_{s=1}^{n_P} z^2_s$$

with probability at least $1-\delta$. □

• Regret and Communication Upper Bounds for FedGLB-UCB$_3$

The regret and communication cost of FedGLB-UCB$_3$ is given in the following theorem.

Theorem F.5 (Regret and Communication Cost Upper Bound of FedGLB-UCB$_3$)**.** *Under the condition that the learning rate of ONS* $\gamma = \frac{1}{2}\min(\frac{1}{4S\sqrt{k^2_\mu S^2 + R^2_{\max}}}, \frac{c_\mu}{(k^2_\mu S^2 + R^2_{\max})\sqrt{NT}})$*, and the total number of global synchronizations* $B = \sqrt{NT}$*, the cumulative regret* R_T *has upper bound*

$$R_T = O\left(\frac{k_\mu(k_\mu + R_{max})}{c_\mu} d N^{3/4} T^{3/4} \log(NT/\delta)\right)$$

with probability at least $1-\delta$*. The cumulative communication cost has upper bound*

$$C_T = O(N^{1.5}\sqrt{T})$$

Proof. Similar to the proof for the previous two variants of FedGLB-UCB, we divide the epochs into 'good' and 'bad' ones according to the determinant ratio, and then bound their cumulative regret separately.

Recall that the instantaneous regret $r_{t,i}$ incurred by client $i \in [N]$ at time step $t \in [T]$ has upper bound

$$\begin{aligned}
\frac{r_{t,i}}{k_\mu} &\leq \mathbf{x}^\top_{t,\star}\theta_\star - \mathbf{x}^\top_{t,i}\theta_\star \leq \mathbf{x}^\top_{t,i}\tilde{\theta}_{i,t} - \mathbf{x}^\top_{t,i}\theta_\star \\
&= \mathbf{x}^\top_{t,i}(\tilde{\theta}_{i,t} - \hat{\theta}_{t,i}) + \mathbf{x}^\top_{t,i}(\hat{\theta}_{t,i} - \theta_\star) \\
&\leq \|\mathbf{x}_{t,i}\|_{V^{-1}_{t,i}} \|\tilde{\theta}_{i,t} - \hat{\theta}_{t,i}\|_{V_{t,i}} + \|\mathbf{x}_{t,i}\|_{V^{-1}_{t,i}} \|\hat{\theta}_{t,i} - \theta_\star\|_{V_{t,i}} \\
&\leq 2\alpha_{t,i}\|\mathbf{x}_{t,i}\|_{V^{-1}_{t,i}}
\end{aligned}$$

Note that due to the update schedule $\mathcal{S}$, we have $\max_{p\in[P]}(n_p - n_{p-1}) = \frac{NT}{B}$. Then based on Corollary F.3.1, $\alpha_{t,i} = O(\frac{k_\mu + R_{max}}{c_\mu}\sqrt{d\log(NT)}\sqrt{\frac{NT}{B}})$, so we have, $\forall t \in [T], i \in [N]$,

$$r_{t,i} = O(\frac{k_\mu(k_\mu + R_{max})}{c_\mu}\sqrt{d\log(NT)}\sqrt{\frac{NT}{B}})\|\mathbf{x}_{t,i}\|_{A^{-1}_{t,i}}$$

with probability at least $1-\delta$.

Therefore, the cumulative regret for the 'good epochs' is $REG_{good} = O(\frac{k_\mu(k_\mu + R_{max})}{c_\mu} d \frac{NT}{\sqrt{B}} \log(NT))$.

Using the same argument as in the proof for FedGLB-UCB$_1$, the cumulative regret for each 'bad ' epoch is upper bounded by $2k_\mu S \frac{NT}{B}$. Since there can be at most $O(d\log NT)$ 'bad epochs', the cumulative regret for all the 'bad epochs' is upper bounded by

$$REG_{bad} = O(dNT\log(NT) \cdot \frac{k_\mu S}{B})$$

Combining the regret incurred in both 'good' and 'bad' epochs, we have

$$R_T = O\big(\frac{k_\mu(k_\mu + R_{max})}{c_\mu} d \frac{NT}{\sqrt{B}} \log(NT) + dNT\log(NT) \cdot \frac{k_\mu S}{B}\big)$$

To recover the regret in centralized setting, we can $B = NT$, which leads to $R_T = O(\frac{k_\mu(k_\mu + R_{max})}{c_\mu} d\sqrt{NT}\log(NT))$. However, this incurs communication cost $C_T = N^2 T$. Alternatively, if we set $B = \sqrt{NT}$, we have $R_T = O(\frac{k_\mu(k_\mu + R_{max})}{c_\mu} d N^{3/4} T^{3/4} \log(NT))$, and $C_T = O(N^{1.5}\sqrt{T})$. □

G Additional Explanation about Figure 2

In Section 5, we used the scatter plots to present the experiment results. Here we provide more explanation about how to interpret these figures. As mentioned earlier, each dot in Figure 2 and Figure 3 denotes the cumulative communication cost (x-axis) and regret (y-axis) that an algorithm (FedGLB-UCB, its variants, or DisLinUCB) with certain threshold value of D or B (labeled next to the dot) has obtained at iteration T.

Here, Figure 4 shows how the cumulative regret/reward and communication cost of five algorithms change over the course of federated bandit learning in our evaluations on synthetic dataset, (their final results at iteration T are used to plot five dots in Figure 2). By carefully examining the relationship between their regret and communication cost, we can see that in Figure 4, FedGLB-UCB ($D = 5.0$), FedGLB-UCB$_1$ ($B = 10.0$), FedGLB-UCB$_2$ ($B = 10.0$), and FedGLB-UCB$_3$ ($B = 5000.0$) incur similar total communication cost, but FedGLB-UCB ($D = 5.0$) attains much smaller regret than the others. Meanwhile, FedGLB-UCB ($D = 5000.0$) attains almost the same regret as FedGLB-UCB$_2$ ($B = 10.0$), but its communication cost is much lower.

Figure 4 also depicts how the communication was controlled in FedGLB-UCB under its event triggered protocol (e.g., generally a decreasing frequency of communication comparing to the scheduled updated in its variants). This shows that FedGLB-UCB strikes the best regret/reward-communication trade-off among the algorithm instances in comparison. However, this line chart can only accommodate a limited range of trade-off settings for these algorithms, to attain a reasonable visibility. In comparison, the scatter plots in Figure 2 provide a much more thorough view of how well the algorithms balance regret/reward and communication cost, by covering a large range of trade-off settings.

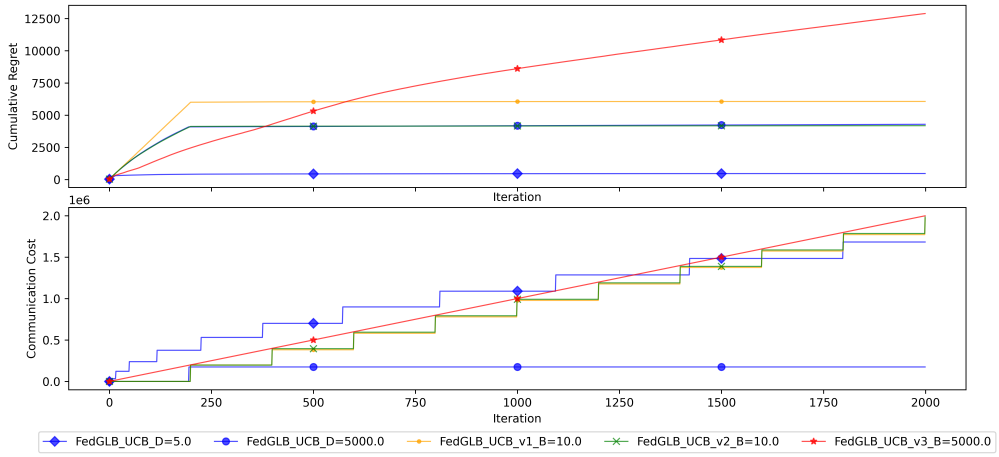

Figure 4: Experiment results showing regret and communication cost over time.