Optimal query complexity for private sequential learning against eavesdropping

Jiaming Xu
Duke University

Kuang Xu
Stanford University

Dana Yang
Duke University

Abstract

We study a learner-private sequential learning problem, motivated by the privacy and security concerns due to eavesdropping attacks. A learner tries to estimate an unknown scalar value, by sequentially querying an external database and receiving binary responses; meanwhile, a third-party adversary observes the learner's queries but not the responses. The learner's goal is to design a querying strategy with the minimum number of queries (optimal query complexity) so that she can accurately estimate the true value, while the eavesdropping adversary even with the complete knowledge of her querying strategy cannot. We develop new querying strategies and analytical techniques and use them to prove almost-matching upper and lower bounds on the optimal query complexity, obtaining a complete characterization of the optimal query complexity as a function of the estimation accuracy and the desired levels of privacy.

1 Introduction

Rapid developments in machine learning and data science have compelled organizations and individuals to increasingly rely on data to solve inference and decision problems. It quickly became clear, however, that collecting and disseminating data in bulk can expose data owners to serious privacy breaches (Dwork, 2008). To address the privacy concerns of data owners, researchers and practitioners have been advocating a new learning framework, known as *learning with external workers* (Konečnỳ et al., 2015). Under this framework, instead of allowing a learner to possess the entire data set and conduct analysis in an offline manner, data sets are kept secure by their owners, and the learner must interact with data owners by submitting queries and receiving responses.

While substantial progress has been achieved in protecting data owners' privacy in such systems (Dwork, 2008; Geyer et al., 2017; Song et al., 2013; Agarwal et al., 2018), the *learner's* privacy has largely been overlooked. Because a learner has to communicate frequently with data owners in order to perform analysis, their queries can be subject to eavesdropping by a third-party adversary. That adversary, in turn, could use the observed queries to reconstruct the learned model, thus allowing them to free-ride at the learner's expense, or worse, leverage such information in future sabotages.

In this paper, we focus on understanding how to protect the learner's privacy against eavesdropping attacks, and precisely quantifying the fundamental privacy-complexity trade-offs in such an interactive learning system. We base our analysis on the Private Sequential Learning model proposed by Tsitsiklis et al. (2020). Suppose that a learner is trying to estimate an unknown target value $X^* \in [0,1]$, by submitting n queries sequentially, $(q_1, \ldots, q_n) \in [0,1]^n$, for some $n \in \mathbb{N}$. For each query q_i, the learner receives a binary response $r_i = \mathbb{1}\{X^* \geq q_i\}$, indicating the position of X^* relative to the query, where $\mathbb{1}\{\cdot\}$ denotes the indicator function. Meanwhile, there is an adversary who observes all of the learner's queries $(q_1, \ldots, q_n)$, but not the responses $(r_1, \ldots, r_n)$. The adversary then tries to estimate X^*. The learner's goal is to design a querying strategy with a minimal n (optimal query complexity) so that she can estimate X^* up to an additive error of $\epsilon/2$ with probability 1 (accuracy), while *no adversary* can estimate X^* up to an additive error of $\delta/2$ with probability larger than $1/L$ for some integer $L \geq 2$ (privacy), even if they are equipped with the complete knowledge of the learner's querying strategy. The parameter L thus captures the learner's privacy level. In the special case of

Proceedings of the 24[th] International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130.

$L = 1$ (corresponding to having no privacy constraint) this learning model reduces to the classical problem of sequential search with binary feedback, with numerous applications such as data transmission with feedback (Horstein, 1963), finding the roots of a continuous function (Waeber et al., 2011), and even the game of "twenty questions" (Jedynak et al., 2012).

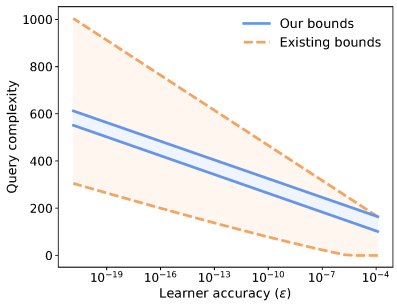

Figure 1: Our results in Theorem 1 (solid) versus the best known bounds (dashed, upper bound in Tsitsiklis et al. (2020), lower bound in Xu (2018)) under the noiseless Bayesian setting, with $L = 15$ and $\delta = 4\epsilon^{0.5}$. The figure is cut off to the right at $\delta = 1/L$, beyond which no strategy can be private.

Our contributions. The primary contributions of our paper are two fold:

1. We settle the optimal query complexity of the Private Sequential Learning problem in both the Bayesian setting (where X^* is random) and deterministic setting (where X^* is fixed but arbitrary). We do so by establishing query complexity upper and lower bounds that almost match in the entire parameter range, thus obtaining a complete picture of the optimal query complexity as a function of the estimation accuracy and the learner's privacy level. Our results substantially improve upon the best known upper and lower bounds (Tsitsiklis et al., 2020; Xu, 2018), and the improvements are most drastic over an important range of parameters, where both the adversary and the learner aim to locate X^* within small errors; see Figure 1 for an illustrative example.

2. We propose and analyze a new variant of the private learning model with noisy responses, an important feature that is especially salient in real-world operations and machine learning applications where the functional evaluations are often stochastic. In this setting, we prove upper and lower bounds on the optimal query complexity which match up to multiplicative constants that only depend on the level of noise in the responses. This mirrors the best known characterizations available in the non-private version of the problem, which also has a dependency on the noise level (Waeber et al., 2013).

Methodological contribution. Our results are rooted in new insights into the nature of learner-private sequential learning problems, which could have broader implications. For instance, in the Bayesian setting, one driving insight is that the portion of the learning process that demands the most privacy protection is *after* the learner has already obtained a reasonably accurate estimate of X^*. We further demonstrate that, as an implication of this observation, the querying trajectory of an optimal learner can be roughly divided into two phases:

1. A *pure-learning* phase, where the primary objective is to narrow the search down to a smaller interval that contains X^* and privacy is not a top priority.

2. A *private-refinement* phase, where the learner refines her estimate of X^* within the said interval, while allocating significantly more querying budget towards obfuscation.

We develop new learning strategies and analytical techniques to make this intuition precise. The algorithmic implications are significant. On the one hand, it suggests more efficient learner strategies that allocate more obfuscation budget towards the latter stages of learning. On the other hand, it can be used to design more powerful adversary strategies that focus on latter stages of a learner's query trajectory, and obtain a more accurate estimator. In analogous manner, our analysis of the deterministic setting also leverages a two-phase approach, although the obfuscation budget is now skewed towards the earlier stages of learning.

Related Literature. The Private Sequential Learning model differs significantly from the existing literature on private iterative learning, such as private stochastic gradient descent (Song et al., 2013; Abadi et al., 2016; Agarwal et al., 2018), private online learning (Jain et al., 2012), and private distributed learning (Geyer et al., 2017; McMahan et al., 2018; Nasr et al., 2019; Melis et al., 2019). The focus therein is to protect *data owners' privacy* by preventing the adversary inferring about a data owner from the outputs of learning algorithms, often using the notion of differential privacy (Dwork, 2008). In that setting, a common privacy-preserving mechanism is to inject calibrated noise at each iteration of the learning algorithms. In contrast, our work aims to protect *the learner's privacy* by preventing the adversary inferring the learned model from the learner's queries. As a result, our problem setup, privacy-preserving mechanisms, and main results are significantly different from those in this literature. Our focus on a decision maker's obfuscation task is related, at a high level, to recent studies of information-theoretically sound obfuscation in vari-

ous sequential decision-making problems (Fanti et al., 2015; Luo et al., 2016; Tsitsiklis and Xu, 2018; Erturk and Xu, 2019; Tang et al., 2020). However, most of these models focus on protecting data and information already in the position of the decision-maker. As such, they do not address the unique privacy challenges arising in learning, where the learner has to protect a piece of information that they themselves are just in the process of discovering.

1.1 Motivating Examples

Learning the optimal price: As discussed in Tsitsiklis et al. (2020), dynamics similar to those in the Private Sequential Learning model arise in the domain of dynamic pricing. Suppose a company is conducting market experiments to determine the release price of a product. The goal is to learn a global parameter about the entire consumer base, e.g., X^* equals the highest price to charge so that at least 50% of the consumers would purchase the product. At each epoch of the experiment, the company samples a subset of the consumers and experiment on a test price (query). Under the sequential learning model, the response r_i corresponds to the indicator function of whether at least 50% of the sampled consumers would purchase the product at price q_i. Note that, due to individual differences and the sampling process, the response is a noisy version of its population variant, which can be captured by the noisy variant of the model we study in this paper. In this example, a competitor (adversary) can easily access the sequence of test prices by participating in the experiments, but does not observe the responses. The optimal query complexity refers to the minimum number of epochs the company takes to estimate X^* accurately, while making sure the eavesdropping adversary cannot infer the final release price. Notice the distinction between our privacy incentive and the incentive to protect the data owners' privacy. The latter aims to ensure that the query sequence does not reveal the price each individual participant is willing to pay, which varies from person to person and can be far from X^*.

Federated Learning: Federated Learning is an emerging machine learning model training paradigm that has been gaining traction (Konečnỳ et al., 2015, 2016). In Federated Learning, a central learner trains a global model by aggregating local model updates across a large number of users. Specifically, the learner aims to estimate the optimal model parameter that minimizes the population risk, i.e., $\theta^* \in \arg\min_\theta L(\theta) \triangleq \mathbb{E}[\ell(Z,\theta)]$, where ℓ denotes the loss function and the average is taken over the underlying data distribution of Z. Each user has access to a local empirical risk function $\ell_u(\theta) = \frac{1}{|\mathcal{S}_u|}\sum_{j\in\mathcal{S}_u}\ell(Z_j,\theta)$ defined over the local data sample $\{Z_j\}_{j\in S_u}$. At each iteration i, the learner broadcasts the current model parameter estimator θ_i to the users. Using their local data, each user u runs multiple steps of gradient descent $\theta \leftarrow \theta - \eta\nabla\ell_u(\theta)$ starting from θ_i, and sends the model update θ_i^u back to the learner. The learner aggregates all the model updates to produce the model parameter for the next iteration θ_{i+1}.

When training with thousands of users, as the learner lacks enough administrative power over those external workers, the Federated Learning system is highly vulnerable to eavesdropping attacks (Kairouz et al., 2019). An honest-but-curious adversary can participate in the training stage by pretending to be an user, and eavesdrop on the sequence of broadcasted model parameters $\{\theta_i\}$ and steal the learned model. Our private sequential learning problem can be viewed as an abstraction of such eavesdropping attack faced by the central learner in Federated Learning. In particular, the true model parameter θ_* is assumed to be in one dimension and the model parameter estimator θ_i is viewed as a *query*. Then assume that instead of the local model update θ_i^u, each user u sends back to the learner only the directional information $\text{sign}(\theta_i^u - \theta_i)$ of the update. Suppose the direction of the local updates indicates the direction of the optimal model parameter under the local loss function, that is, $\text{sign}(\theta_i^u - \theta_i) = \text{sign}(\theta_*^u - \theta_i)$, where $\theta_*^u \in \arg\min_\theta \ell_u(\theta)$. Under this assumption, the majority vote of $\mathbb{1}\{\theta_i^u \geq \theta_i\}$ can be viewed as a noisy version of $\mathbb{1}\{\theta_* \geq \theta_i\}$.[1] The majority vote corresponds to the response r_i under the binary search model.

We remark that in Federated Learning, communication bandwidth is a scarce resource. Thus, efficient use of the queries is of fundamental importance. Although the binary search model is only an abstraction of the general Federated Learning framework, studying the trade-off between accuracy, privacy and query complexity under the binary search model can provide valuable insights on the algorithm design in Federated Learning. We discuss this example in more detail in the supplementary material.

2 Problem formulation

Consider the problem of learning some unknown true value $X^* \in [0,1]$. Let $\widehat{X}$ be the learner's estimator of X^* and $\widehat{\widehat{X}}$ be the adversary's. The learner submits queries $q_1, q_2, ... \in [0,1]$ sequentially. Each time a query q_i is submitted, the learner receives a response

[1] For example, if ℓ is the ℓ_1 loss, then θ_* is the (population) median of Z, and θ_*^u is the sample median for user u. Therefore, $\mathbb{1}\{\theta_*^u \geq \theta_i\}$ is distributed as Bernoulli with mean above 1/2 if $\theta_* \geq \theta_i$ and below 1/2 otherwise.

r_i. When the responses are noiseless, $r_i = \mathbb{1}\{X^* \geq q_i\}$. Under the noisy response setting, we assume that

$$\begin{aligned} r_i &\sim \text{Bernoulli}(p) && \text{if } X^* \geq q_i, \\ r_i &\sim \text{Bernoulli}(1-p) && \text{if } X^* < q_i \end{aligned}$$

for some $p \in (1/2, 1)$. That is, each observed response can be erroneous with probability $1 - p$.

The learner's query q_i can depend on all the past queries and responses, and is allowed to incorporate outside randomness. Since all random variables and all random vectors with finite alphabets can be simulated from a random variable uniformly distributed on $[0, 1]$, without loss of generality, let $Y \sim \text{Unif}[0, 1]$ be the random seed that the learner may use to generate queries. Then q_i can be written as $f_{i-1}(q_1, ..., q_{i-1}, r_1, ..., r_{i-1}, Y)$ for some function f_{i-1}. Note that the first query q_1 is submitted without any information and is only a function of Y. Thus we have $q_2 = f_1(q_1, r_1, Y) = f_1(f_0(Y), r_1, Y) := \phi_1(r_1, Y)$. It is easy to see that all q_i can be written iteratively as a function of only the past responses and Y, i.e., $q_i = \phi_{i-1}(r_1, ..., r_{i-1}, Y)$.

Then a querying strategy ϕ is defined by an initial mapping $f_0 : [0, 1] \to [0, 1]$ used to generate q_1 from Y, a sequence of mappings $(\phi_i)_i$ with $\phi_i : \{0, 1\}^i \times [0, 1] \to [0, 1]$ used to generate the rest of the query sequence, and a final estimator $\widehat{X}$, which can depend on Y and all the queries and responses. The adversary's estimator $\widetilde{X}$, on the contrary, is formed with only access to the queries and the querying strategy ϕ but not the responses or the random seed Y.

The goal of the learner is to design a querying strategy to ensure that she can accurately estimate X^*, but the adversary cannot. Different ways to quantify the estimators' performance arise naturally when the responses are noisy versus noiseless. We discuss the two settings separately.

2.1 Noiseless responses

Following Tsitsiklis et al. (2020), we consider both the Bayesian setting where $X^* \in [0, 1]$ is uniformly distributed on $[0, 1]$ and the setting where X^* is deterministic. The two settings call for different definitions for accuracy and privacy.

Bayesian setting We assume X^* is uniformly distributed on $[0, 1]$, which is independent from the random seed Y, as the learner does not know the true value X^* a priori. We say a strategy ϕ is

- ϵ-accurate for $\epsilon > 0$, if $\mathbb{P}\{|\widehat{X} - X^*| \leq \epsilon/2\} = 1$;
- (δ, L)-private for $\delta > 0$ and an integer $L \geq 2$, if there is *no* adversary $\widetilde{X}$ such that $\mathbb{P}\{|\widetilde{X} - X^*| \leq \delta/2\} > \frac{1}{L}$.

Deterministic setting Suppose X^* is a deterministic but arbitrary number on $[0, 1]$. Then the only source of randomness in the querying strategy is from Y. We say a strategy ϕ is

- ϵ-accurate for $\epsilon > 0$, if $\mathbb{P}\{|\widehat{X} - X^*| \leq \epsilon/2\} = 1$, $\forall X^* \in [0, 1]$;
- (δ, L)-private for $\delta > 0$ and an integer $L \geq 2$, if for each query sequence $\bar{q}$, the δ-covering number[2] of the information set $\mathcal{I}(\bar{q})$ is at least L. The information set is defined as the set of all true values that could lead to the query sequence $\bar{q}$ under strategy ϕ with non-negligible probability. Note that the query sequence q is a random vector that depends on X^* and Y, i.e., $q = q(X^*, Y)$. Formally we define

$$\mathcal{I}(\bar{q}) = \{X^* \in [0, 1] : \mathbb{P}\{q(X^*, Y) = \bar{q}\} > 0\}.$$

For both Bayesian and deterministic settings, we define the optimal query complexity as

$$N(\epsilon, \delta, L) = \min\{n : \exists \phi \text{ that is both } \epsilon\text{-accurate and } (\delta, L)\text{-private and submits at most } n \text{ queries}\}.$$

Note that for a larger δ or a larger L, the (δ, L)-private constraint is a stronger requirement. Therefore $N(\epsilon, \delta, L)$ is monotone nondecreasing in δ and L.

Without loss of generality, we focus on the regime of parameters $2\epsilon \leq \delta \leq 1/L$. To see why, note that on one end of the spectrum, if $\delta > 1/L$, then the adversary can make an arbitrary guess to break the privacy constraint: simply choosing $\widetilde{X} = 1/2$ yields $\mathbb{P}\{|\widetilde{X} - X^*| \leq \delta/2\} = \delta > 1/L$. In this regime the (δ, L)-privacy constraint is too strong to be satisfied by any querying strategy. On the other end of the spectrum, the regime $\delta \leq 2\epsilon$ is arguably not that interesting, as it is unnatural to require an adversary, who only have access to queries but not responses, to estimate X^* almost as accurately as the learner does.

2.2 Noisy responses under the Bayesian setting

The noisy response setting is a new variant of the binary sequential search model that we focus on. We consider the Bayesian formulation where $X^* \sim \text{Unif}[0, 1]$. Since the responses contain noise, no learner that submits a finite number of queries can estimate accurately with probability one. Hence the definition for the learner's accuracy needs to be modified. We consider the following two natural definitions.

[2] The δ-covering number of $A \subseteq \mathbb{R}$ is defined as the size of the smallest set $\mathcal{N}$, such that $\cup_{r \in \mathcal{N}}[r - \delta/2, r + \delta/2] \supseteq A$.

(a) (accurate on average) We say a querying strategy is ϵ-accurate for $\epsilon > 0$ if $\mathbb{E}|\widehat{X} - X^*| \le \epsilon/2$;

(b) (accurate with high probability) We say a querying strategy is (ϵ, M)-accurate for $\epsilon > 0$ and $M \ge 2$ (M is not necessarily an integer) if $\mathbb{P}\{|\widehat{X} - X^*| > \epsilon/2\} \le 1/M$.

The definition of privacy is the same as that in the noiseless case. A querying strategy is called (δ, L)-private if no adversary's estimator $\widetilde{X}$ can achieve $\mathbb{P}\{|\widetilde{X} - X^*| \le \delta/2\} > 1/L$.

Define the optimal query complexity as the minimum number of queries needed for a querying strategy to be both ϵ-accurate (resp. (ϵ, M)-accurate) and (δ, L)-private, denoted as $N_{\mathsf{avg}}(\epsilon, \delta, L)$ (resp. $N_{\mathsf{whp}}(\epsilon, M, \delta, L)$).

3 Main results

3.1 Noiseless responses

When the responses are noiseless, we prove almost matching upper and lower bounds on the optimal query complexity in both the Bayesian and deterministic settings.

Bayesian setting The following is our first main result.[3]

Theorem 1 (Bayesian setting)**.** *If $2\epsilon \le \delta \le 1/L$, then*

$$N(\epsilon, \delta, L) \ge L \log \frac{\delta}{\epsilon} + \log \frac{1}{L\delta} - 1 - 2L,$$
$$N(\epsilon, \delta, L) \le L \log \frac{\delta}{\epsilon} + \log \frac{1}{L\delta} - 1 + 2L.$$

Across the entire parameter regime, our result captures the impact of privacy requirement up to an additive gap of $4L$. For example, in the regime where δ,L stay as fixed constants while $\epsilon \to 0$, Theorem 1 indicates that

$$N(\epsilon, \delta, L) \approx L \log \frac{1}{\epsilon} + L \log \delta + \log \frac{1}{L\delta}.$$

The dominating term is $L \log(1/\epsilon)$. Since the cost of non-private search is $\log(1/\epsilon)$, attained by the bisection search, the cost of privacy is roughly a *multiplicative* factor of L.

On the other end of the spectrum where δ, ϵ both go to 0 proportionally, i.e. $\delta = c\epsilon$, we have

$$N(\epsilon, \delta, L) \approx \log(1/\epsilon) + (L \log c + \log(1/L) - \log c).$$

[3]Here we ignore all non-integrality issues. See the supplementary material for the general statements.

The cost of privacy is characterized by the factor in the parentheses, with a leading term $L \log(\delta/\epsilon)$ that only scales linearly in L.

Therefore according to Theorem 1, the optimal query complexity is highly sensitive to δ. As δ grows, we observe a phase transition where the cost of privacy grows from additive to multiplicative. In contract to our result, the best known upper bound (Tsitsiklis et al., 2020) and lower bound (Xu, 2018) are as follows.

$$L \log \frac{\delta}{\epsilon} - 3L \log\log \frac{\delta}{\epsilon} \le N(\epsilon, \delta, L) \le L \log \frac{1}{L\epsilon} + L.$$

There is a large gap between these bounds. For example, the previous upper bound suggests a multiplicative cost of privacy across the entire parameter range, and fails to uncover the additive regime.

Deterministic setting Under the deterministic setting, we establish upper and lower bounds that match up to only 8 queries.

Theorem 2 (Deterministic setting)**.** *If $2\epsilon \le \delta \le 1/L$, then*

$$N(\epsilon, \delta, L) \ge \max\left\{\log \frac{1}{\epsilon} + L, \log \frac{\delta}{\epsilon} + 2L\right\} - 8,$$
$$N(\epsilon, \delta, L) \le \max\left\{\log \frac{1}{\epsilon} + L, \log \frac{\delta}{\epsilon} + 2L\right\}.$$

Similar to the Bayesian setting, as δ grows, we observe a phase transition in the behavior of the optimal query complexity, with a sharp threshold at roughly $\delta = 2^{-L}$. When δ is below the threshold, we have $N(\epsilon, \delta, L) \approx \log(1/\epsilon) + L$. That is, the overhead due to privacy, in terms of L, is an additive factor of L. For δ above the threshold, we have $N(\epsilon, \delta, L) \approx \log(1/\epsilon) + 2L + \log\delta$. The cost of privacy is an additive factor of $2L + \log\delta$, which is always between L and $2L$ for $\delta > 2^{-L}$.

In comparison, the best known upper and lower bounds in prior work (Tsitsiklis et al., 2020) are

$$\log \frac{\delta}{\epsilon} + 2L \le N(\epsilon, \delta, L) \le \log \frac{1}{L\epsilon} + 2L,$$

which are not tight except for the lower bound in the large δ regime, and fail to capture the behavior of the optimal query complexity when δ is small. In particular, the previous upper bound is completely independent of δ, and the previous lower bound reduces to a trivial bound of $\log(1/\epsilon)$ when $\delta < 2^{-2L}$. Examples of the comparison between our results and the existing bounds are illustrated in Figures 1 and 2, under the Bayesian and deterministic settings, respectively.

The sharp bounds in Theorems 1 and 2 can be extend to multiple dimensions. We show that for

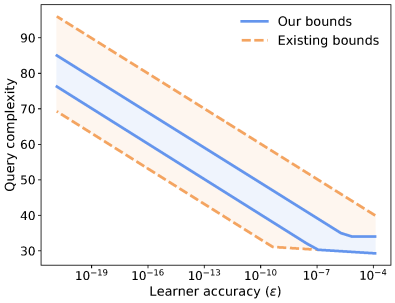

Figure 2: Our results in Theorem 2 (solid) versus the best known bounds (dashed) for the noiseless deterministic setting, $L = 15$ and $\delta = 4\epsilon^{0.9}$.

$X^* \in \mathbb{R}^d$, the optimal query complexity $N_d(\epsilon, \delta, L)$ in d dimensions is approximately $d(\log(1/(L^{1/d}\delta)) + L^{1/d}\log(\delta/\epsilon))$ under the Bayesian setting, and $d\max\{\log(1/\epsilon)+L^{1/d}, \log(\delta/\epsilon)+2L^{1/d}\}$ under the deterministic setting. See the supplementary material for the precise statements and proofs of the multidimensional results.

3.2 Noisy responses

Under the Bayesian setting, we extend our results to allow for noisy responses, which is a model variant that has not been studied before. Recall that each response is corrupted with probability $1-p$. Below is our main result in the noisy response model. In short, we are able to characterize the optimal query complexities up to constants that only depend on p.

Theorem 3. *Assume that $4\epsilon \le \delta \le 1/L$. Then*

$$N_{\mathsf{avg}}(\epsilon, \delta, L) \asymp_p L\log\frac{\delta}{\epsilon} + \log\frac{1}{\epsilon},$$
$$N_{\mathsf{whp}}(\epsilon, M, \delta, L) \asymp_p L\log\frac{M\delta}{\epsilon} + \log\frac{1}{\epsilon},$$

where $\asymp_p$ denotes matching upper and lower bounds up to multiplicative constants that depend on p.

The exact form for the p-dependent constants is contained in the supplementary material. In particular, all the constants go to infinity at the same rate of $(p-1/2)^{-2}$ as $p \to 1/2$.

Note that by manipulating the constants, Theorem 3 implies that $N_{\mathsf{avg}}(\epsilon, \delta, L) \asymp_p L\log(\delta/\epsilon) + \log(1/(L\delta))$, which matches with the query complexity in the noiseless setting, stated in Theorem 1. In other words, as noise is introduced to the model, the optimal query complexity is only impacted by the noise level through the multiplicative constants.

When M is large, definition (b) in Section 2.2 is a stronger condition on the learner's accuracy than (a). As a result an extra additive factor of order $L\log M$ shows up in $N_{\mathsf{whp}}(\epsilon, M, \delta, L)$. The intuition behind the $L\log M$ factor is that there must be at least $\Omega(\log M)$ queries near X^* to achieve an error probability of $1/M$; these queries then need to be duplicated L times to disguise the location of X^* from the adversary.

Compared with the noiseless response setting where we obtained tight control on the optimal query complexity, the results for the noisy response case are only up to multiplicative constants that depend on p. We remark that even for the vanilla noisy binary search problem without privacy consideration, the precise p-dependent constant remains an open problem.

4 New insights and algorithmic ideas

In this section we highlight the new algorithmic ideas and key analysis techniques that allowed us to obtain sharp bounds. Full proofs of the results are deferred to the supplementary material.

A fundamental difficulty in the design of private learning strategies is that the learner's queries simultaneously serve two, sometimes competing, goals: (1) to gather information about the target X^*, and (2) to deceive the adversary as to the target's location. More specifically, our approach begins by recognizing that a single learner query can be used to accomplish one, or multiple, of the following tasks:

(i) to obtain information in order to identify a small interval that contains X^* (diameter of this interval depends on the privacy level and adversary accuracy);

(ii) to obtain information in order pin-point X^* within the said small interval, down to an ϵ-accuracy;

(iii) to serve as a "decoy" to throw off the adversary.

The key to our analysis is to put emphasis on understanding the interaction of the queries serving these different types of tasks. We design more efficient querying strategies where multiple tasks are accomplished simultaneously. We also provide sharp analysis on the maximum number of queries that can be used for more than one purpose. In what follows, we demonstrate how these principles manifest in our algorithms and analyses.

4.1 The Bayesian setting with noiseless responses

Recall that privacy is breached if an adversary can estimate X^* up to an additive error of $\delta/2$ with probability at least $1/L$. That makes δL an appropriate choice

for the diameter of the small interval. Queries leading to the identification of such an interval (type (i)) do not significantly compromise the learner's privacy, because the interval's size is too large for the adversary to extract useful information. Beyond this point, however, the learner must submit further queries to narrow the range of X^* down to ϵ, and these queries must be carefully obfuscated. In other words, effective queries at this point should serve to accomplish tasks (ii) and (iii) simultaneously. As such, in the design of the optimal strategy, we divide learning into two phases:

1. a *pure-learning phase*, corresponding to task (i), where the sole focus of the learner is to identify a small interval containing the target X^*, and

2. a *private-refinement phase*, corresponding to tasks (ii) and (iii), where queries serve to simultaneously refine and obfuscate a fine-grained estimate of X^*.

Compared with the single-phase strategy in Tsitsiklis et al. (2020), our strategy shifts all obfuscation efforts into the later phase. In the pure learning phase, we first runs a bisection search to locate X^* within a length $L\delta$ interval J. In the private-refinement phase, borrowing the idea of replication (Xu, 2018), we divide J into L length δ subintervals. In the true subinterval containing X^*, the queries are submitted via the bisection search. Meanwhile, cloned queries are submitted in the other $L-1$ subintervals in parallel, to ensure that the adversary cannot infer the true subinterval with probability higher than $1/L$.

The bisection search in the first phase takes $\log(1/(L\delta))$ queries. The second phase consists of a $\log(\delta/\epsilon)$-cost bisection search replicated L times. Thus the query complexity of this two-phase querying strategy is roughly $\log(1/(L\delta)) + L\log(\delta/\epsilon)$.

Although this two-phase strategy is intuitively motivated, it is far from evident why it is information-theoretically optimal. To establish its optimality, we design an intelligent strategy of the adversary, and use that to show an almost matching lower bound on the optimal query complexity.

The previous lower bound considers an adversary who adopts the *proportional-sampling* strategy, where the adversary proportionally samples from all queries to produce an estimator $\widetilde{X}$. To improve upon the previous bound, we consider a more intelligent adversary's strategy named *truncated proportional-sampling*, where the adversary disregards the first K queries and proportionally samples from the rest of the queries. We show that under this adversary's strategy, any learner's strategy with less than roughly $\log(1/(L\delta)) + L\log(\delta/\epsilon)$ queries is either non-private or non-accurate. Deliberately discarding the first K queries is seemingly counter-intuitive, and is crucial for obtaining the tight lower bound. To show why ignoring information leads to information-theoretically optimal estimation is highly non-trivial. The heuristic reasoning is that the first few queries could negatively impact the adversary's estimator since they are very unlikely to be close to X^*. Another difficulty is to determine the number of queries to discard. Through our analysis, we discover a subtle but interesting duality between the learner and the adversary: the number of queries disregarded by the adversary K is exactly the number of queries submitted in the *pure-learning phase* under the learner's optimal strategy. However, in the proof of our lower bound, this argument with discarded queries is effective against any learner strategy, regardless of what the learner tries to achieve with these early queries.

Below is an outline of the lower bound proof. Under a truncated proportional-sampling adversary, for any querying strategy that is (δ, L)-private, we have

$$\begin{aligned}\frac{1}{L} \geq& \mathbb{P}\left\{\left|\widetilde{X} - X^*\right| \leq \frac{\delta}{2}\right\} \\ =& \frac{\sum_{i=K+1}^{n} \mathbb{P}\{q_i \in [X^* - \delta/2, X^* + \delta/2]\}}{n-K},\end{aligned}$$

where n is the total number of queries. We then show that by choosing $K \approx \log(1/(L\delta))$, for any querying strategy that is ϵ-accurate, the expected number of queries of $q_{K+1}, ..., q_n$ in $[X^* - \delta/2, X^* + \delta/2]$ is at least about $\log(\delta/\epsilon)$. Therefore $n \geq K + L\log(\delta/\epsilon)$ is a rough lower bound for the optimal query complexity.

4.2 The deterministic setting with noiseless responses

The story is slightly different under the deterministic setting. Unlike the Bayesian setting where the adversary only needs to perform well when averaging over a random X^*, here the adversary can no longer make guesses. Knowing that the adversary cannot guess, the learner only needs to worry about privacy breaches in a δ-width interval containing X^*. Moreover, before reaching this interval, the learner can ensure privacy by injecting possible alternative locations of X^* along each query sequence. That, recalling the queries' three tasks from the start of this section, corresponds to reusing queries for tasks (i) and (iii). In particular, we will design a query strategy that mirrors the two-phase architecture described for the Bayesian setting, with a coarse learning phase followed by that of a refinement. However, obfuscation efforts (iii) are now implemented during the first phase.

Recall that under the deterministic setting, a querying strategy is called (δ, L)-private if for each query sequence $\bar{q}$, the δ-covering number of the information set $\mathcal{I}(\bar{q})$ is at least L. Similar to the idea in Tsitsiklis et al. (2020), we achieve (δ, L)-privacy by planting L possible locations of X^* into the query sequence. Each possibility is planted via a "guess", which refers to a pair of queries ϵ-apart. These guesses are used to throw the adversary off track, since the adversary cannot tell whether a guess is correct (i.e. whether X^* is between the pair of queries), without observing the responses.

Compared to the strategy in Tsitsiklis et al. (2020) where the guesses are planted on an even grid on $[0, 1]$, we propose a scheme that is much more efficient. The high-level idea of our algorithm is to maximize the number of queries reused for tasks (i) and (iii). In other words, we want the guesses to not only help conceal the location of X^* from the adversary, but also help the learner location X^* at the same time.

When $\delta \leq 2^{-L}$, we submit the guesses along the trajectory of a bisection search in order to optimize efficiency. That is, the first guess is at 1/2 (i.e. the learner submits 2 queries at $1/2$,$1/2 + \epsilon$); the second guess is at 1/4 or 3/4, depending on whether X^* is above or below 1/2; so on and so forth. However, once one of the guesses is tested to be correct, to keep this finding from the adversary, the learner continues to submit the remaining guesses via a "fake" bisection. By fake bisection we mean a simulated bisection search where the binary responses are generated *i.i.d.* Bernoulli(1/2).

When $\delta > 2^{-L}$, the guesses submitted via bisection are not all δ-apart from each other, hence (δ, L)-privacy is not guaranteed. To resolve this problem, we design a more sophisticated strategy that searches for X^* in a less aggressive manner, so as to not approach the true value too rapidly. See the supplementary material for a detailed construction of the querying strategy.

Once all the guesses are submitted, we run a final bisection search to narrow down the range of X^* to an ϵ-length interval. As before, if a guess has been tested to be correct, the learner switches to a fake bisection instead.

Under the strategy described above, the first L guesses consist of $2L$ queries, and they help narrow down the range of X^* into a length $\max\{2^{-L}, \delta\}$ intervals. In total the learner submits $2L + \log(\max\{2^{-L}, \delta\}/\epsilon) = \max\{\log(1/\epsilon) + L, \log(\delta/\epsilon) + 2L\}$ queries.

When $\delta > 2^{-L}$, the lower bound shown in Xu (2018) is tight. Our analysis to obtain the sharp lower bound when $\delta \leq 2^{-L}$ is as follows. We dissect the query sequence, and separately investigates those queries that can be used both for protecting privacy and searching for X^*, and those queries that only fulfill one purpose. More specifically, for each fixed querying strategy that is both ϵ-accurate and (δ, L)-private, we show the existence of an interval I roughly of length δ and $X^* \in I$ such that when X^* is the true value, there are

(a) at least $\log(1/\delta) - 3$ queries outside of I that are at least δ away from each other.

(b) at least $\log(\delta/\epsilon)$ queries in I.

(c) at least $L - 5$ pairs of queries outside of I that are no more than ϵ apart.

The three sets of queries above roughly correspond to the tasks (i),(ii),(iii). Under the assumption $\delta \geq 2\epsilon$, since the queries in (a) are all δ apart, at most $L - 5$ of them can overlap with queries in (c). As such, we conclude that when X^* is the true value, at least $(\log(1/\delta) - 3) + \log(\delta/\epsilon) + (L - 5) = \log(1/\epsilon) + L - 8$ queries are needed. The full proof is contained in the supplementary material. The key challenge lies in showing the existence of an I that satisfies the above.

4.3 The Bayesian setting with noisy responses

The intuitions we gained from the noiseless response model continue to apply. However when the responses are noisy, we encounter some fundamental challenges as a result of the noise in the responses. The main difficulty is in tracking the posterior distributions of X^* for both the learner and the adversary simultaneously. On the one hand, we want the posterior distribution of X^* given the responses to concentrate fast, so that the learner can accurately estimate X^* despite the noises. On the other hand, we need the posterior distribution given the queries to not concentrate too rapidly, ensuring that the adversary cannot accurately learn. This greatly complicates the design and the analysis of the optimal querying strategy, as a closed-form expression for the posterior distributions is out of reach in this case.

For proof of the upper bound in the noisy response model, an intractable posterior distribution of X^* makes it challenging to show that a strategy is private. Even queries that are far from X^* can change the shape of the posterior distribution and potentially leak the location of X^* to the adversary. To overcome this difficulty, our analysis involves the design of a querying strategy that forces certain conditional independence structures between the query sequence and some local neighborhood of X^*. We then use the

conditional independence to carefully control the privacy leakage across all phases of learning.

Our querying strategies relies on an existing search algorithm known as the Burnashev-Zigangirov(BZ) algorithm (Burnashev and Zigangirov, 1974). Suppose $[0, 1]$ is divided into $1/\Delta$ (assumed to be an integer) equal length subintervals. The BZ algorithm estimates the subinterval that contains X^* by recursively querying the endpoints of the subintervals and updating the belief distribution of X^*. See the supplementary material for a description of the BZ algorithm and its statistical properties.

The idea behind the construction of the querying strategies inherits from the construction under the noiseless response setting. Recall that under the querying strategy described in Section 4.1, the learner first runs bisection search to locate X^* within a length $L\delta$ interval. She then runs replicated bisection on the L length δ subintervals, submitting queries via the bisection search in the true subinterval containing X^* and cloning those queries in the other $L-1$ subintervals. When the responses are noisy, firstly we replace the bisection searches with the BZ algorithm. Moreover, the learner can no longer discern the true interval by querying the endpoints of the subinterval only once. Instead we need to query each endpoint enough times, so that via a *maximum-likelihood* type procedure, the learner can estimate the true subinterval with high enough certainty.

For the lower bound proof, we get around the intractable posterior distributions using two sets of tools. Part of the proof utilizes information-theoretic arguments. The key step is to establish an upper bound on the *rate of information transfer*, which governs the speed at which the learner can gather information from the responses. That allowed us to lower bound the expected number of queries in a small interval that contains X^*. For the second part of the proof, we reduce the learner's estimation problem to a family of binary hypothesis testing problems between pairs of hypotheses on a small interval containing X^*. We then bound the testing errors from below using the Bhattacharyya coefficient (Kailath, 1967).

5 Conclusion and future work

Motivated by privacy and security concerns in applications such as Federated Learning and online price learning, we study a sequential learning problem that focuses on protecting the learner's privacy against eavesdropping. A learner aims to estimate a value by sequentially submitting queries and receiving binary responses, while ensuring an adversary who observes queries but not responses cannot estimate well. We design new querying strategies and prove upper bounds on the optimal query complexity. We also derive almost-matching lower bounds, showing that our querying strategies are nearly optimal. The results are further extended to when the unknown value is in high dimensions, and when the binary responses are noisy. An important future direction is to investigate how to protect the learner's privacy in more general online convex optimization problems, such as stochastic gradient descent algorithms, where the adversary observes the query x_t but not the stochastic gradient $g(x_t)$ at each iteration t.

Acknowledgements

J. Xu is supported by the NSF Grants IIS-1838124, CCF-1850743, and CCF-1856424. D. Yang is supported by the NSF Grant CCF-1850743. The authors want to thank all anonymous reviewers for the valuable feedback.

References

Abadi, M., A. Chu, I. Goodfellow, B. McMahan, I. Mironov, K. Talwar, and L. Zhang (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp. 308–318. ACM.

Agarwal, N., A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan (2018). cpsgd: Communication-efficient and differentially-private distributed SGD. In *Advances in Neural Information Processing Systems*, pp. 7564–7575.

Burnashev, M. V. and K. Zigangirov (1974). An interval estimation problem for controlled observations. *Problemy Peredachi Informatsii 10*(3), 51–61.

Dwork, C. (2008). Differential privacy: A survey of results. In M. Agrawal, D. Du, Z. Duan, and A. Li (Eds.), *Theory and Applications of Models of Computation*, Berlin, Heidelberg, pp. 1–19. Springer Berlin Heidelberg.

Erturk, M. S. and K. Xu (2019). Dynamically protecting privacy, under uncertainty. *arXiv preprint arXiv:1911.08875*.

Fanti, G., P. Kairouz, S. Oh, and P. Viswanath (2015). Spy vs. spy: Rumor source obfuscation. In *ACM SIGMETRICS Performance Evaluation Review*, Volume 43, pp. 271–284. ACM.

Geyer, R. C., T. Klein, and M. Nabi (2017). Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*.

Horstein, M. (1963). Sequential transmission using noiseless feedback. *IEEE Transactions on Information Theory 9*(3), 136–143.

Jain, P., P. Kothari, and A. Thakurta (2012). Differentially private online learning. In *Conference on Learning Theory*, pp. 24–1.

Jedynak, B., P. I. Frazier, and R. Sznitman (2012). Twenty questions with noise: Bayes optimal policies for entropy loss. *Journal of Applied Probability 49*(1), 114–136.

Kailath, T. (1967). The divergence and bhattacharyya distance measures in signal selection. *IEEE transactions on communication technology 15*(1), 52–60.

Kairouz, P., H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. (2019). Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*.

Konečnỳ, J., B. McMahan, and D. Ramage (2015). Federated optimization: Distributed optimization beyond the datacenter. *arXiv preprint arXiv:1511.03575*.

Konečnỳ, J., B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon (2016). Federated learning: Strategies for improving communication efficiency. In *NIPS Workshop on Private Multi-Party Machine Learning*.

Luo, W., W. P. Tay, and M. Leng (2016). Infection spreading and source identification: A hide and seek game. *IEEE Transactions on Signal Processing 64*(16), 4228–4243.

McMahan, B., D. Ramage, K. Talwar, and L. Zhang (2018). Learning differentially private recurrent language models. In *International Conference on Learning Representations (ICLR)*.

Melis, L., C. Song, E. De Cristofaro, and V. Shmatikov (2019). Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 691–706. IEEE.

Nasr, M., R. Shokri, and A. Houmansadr (2019). Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 739–753. IEEE.

Song, S., K. Chaudhuri, and A. D. Sarwate (2013). Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, pp. 245–248. IEEE.

Tang, W., W. Wang, G. Fanti, and S. Oh (2020). Privacy-utility tradeoffs in routing cryptocurrency over payment channel networks. In *Abstracts of the 2020 SIGMETRICS/Performance Joint International Conference on Measurement and Modeling of Computer Systems*, pp. 81–82.

Tsitsiklis, J. N. and K. Xu (2018). Delay-predictability trade-offs in reaching a secret goal. *Operations Research 66*(2), 587–596.

Tsitsiklis, J. N., K. Xu, and Z. Xu (2020). Private sequential learning. *Forthcoming in Operations Research*. arXiv:1805.02136.

Waeber, R., P. I. Frazier, and S. G. Henderson (2011). A Bayesian approach to stochastic root finding. In *Proceedings of the 2011 Winter Simulation Conference (WSC)*, pp. 4033–4045. IEEE.

Waeber, R., P. I. Frazier, and S. G. Henderson (2013). Bisection search with noisy responses. *SIAM Journal on Control and Optimization 51*(3), 2261–2279.

Xu, K. (2018). Query complexity of Bayesian private learning. In *Advances in Neural Information Processing Systems*, pp. 2431–2440.