Private Sequential Learning (Extended Abstract)

John N. Tsitsiklis JNT@MIT.EDU
LIDS, Massachusetts Institute of Technology, Cambridge, MA 02139

Kuang Xu KUANGXU@STANFORD.EDU
Graduate School of Business, Stanford University, Stanford, CA 94305

Zhi Xu ZHIXU@MIT.EDU
LIDS, Massachusetts Institute of Technology, Cambridge, MA 02139

Editors: Sébastien Bubeck, Vianney Perchet and Philippe Rigollet

Abstract

We formulate a private learning model to study an intrinsic tradeoff between privacy and query complexity in sequential learning. Our model involves a learner who aims to determine a scalar value, v^*, by sequentially querying an external database and receiving binary responses. In the meantime, an adversary observes the learner's queries, though not the responses, and tries to infer from them the value of v^*. The objective of the learner is to obtain an accurate estimate of v^* using only a small number of queries, while simultaneously protecting her privacy by making v^* provably difficult to learn for the adversary. Our main results provide tight upper and lower bounds on the learner's query complexity as a function of desired levels of privacy and estimation accuracy. We also construct explicit query strategies whose complexity is optimal up to an additive constant.[1]

Keywords: sequential learning, privacy, bisection algorithm.

1. Introduction

Organizations and individuals often rely on relevant data to solve decision problems. Sometimes, such data are beyond the immediate reach of a decision maker and must be acquired by interacting with an external entity or environment. However, these interactions may be monitored by a third-party adversary and subject the decision maker to potential privacy breaches, a possibility that has become increasingly prominent as information technologies and tools for data analytics advance.

The present paper studies a decision maker, henceforth referred to as the learner, who acquires data from an external entity in an interactive fashion by submitting sequential queries. The *interactivity* benefits the learner by enabling her to tailor future queries based on past responses and thus reduce the number of queries needed, while, at the same time, exposes the learner to substantial privacy risk: the more her queries depend on past responses, the easier it might be for an adversary to use the observed queries to infer those past responses. Our main objective is to articulate and understand an intrinsic *privacy versus query complexity tradeoff* in the context of such a Private Sequential Learning model.

1. Extended abstract. Full version appears as [arXiv:1805.02136, v2].

2. The Private Sequential Learning Model

The Private Sequential Learning model involves a *learner* who aims to determine a particular *true value*, v^*. The true value is a scalar in some bounded subset of $\mathbb{R}$. Without loss of generality, we assume that v^* belongs to the interval[2] $[0, 1)$ and that the learner knows that this is the case. The true value is stored in an external database. In order to learn the true value, the learner interacts with the database by submitting queries as follows. At each step k, the learner submits a *query* $q_k \in [0, 1)$, and receives from the database a *response*, r_k, indicating whether v^* is greater than or equal to the query value, i.e.,

$$r_k = \mathbb{I}(v^* \geq q_k),$$

where $\mathbb{I}(\cdot)$ stands for the indicator function. Furthermore, each query is allowed to depend on the responses to previous queries, through a learner strategy, to be defined shortly.

Denote by N the total number of learner queries, and by $\epsilon > 0$ the learner's desired accuracy. After receiving the responses to N queries, the learner aims to produce an estimate $\hat{x}$, for v^*, that satisfies

$$|\hat{x} - v^*| \leq \frac{\epsilon}{2}.$$

In the meantime, there is an *adversary* who is also interested in learning the true value, v^*. The adversary has no access to the database, and hence seeks to estimate v^* by free-riding on observations of the learner queries. Let $\delta > 0$ be an accuracy parameter for the adversary. We assume that the adversary can observe the values of the queries but not the responses, and knows the learner's query strategy. Based on this information, and after observing all of the queries submitted by the learner, the adversary aims to generate an estimate, $\hat{x}^a$, for v^*, that satisfies

$$|\hat{x}^a - v^*| \leq \frac{\delta}{2}.$$

2.1. Learner Strategy

The queries that the learner submits to the database are generated by a (possibly randomized) *learner strategy*, in a sequential manner: the query at step k depends on the queries and their responses up until step $k-1$, as well as on a discrete random variable Y. In particular, the random variable Y allows the learner to randomize if needed, and we will refer to Y as the *random seed*. Without loss of generality, we assume that Y is uniformly distributed over $\{1, 2, \ldots, \mathcal{Y}\}$, where $\mathcal{Y}$ is a large integer. Formally, fixing $N \in \mathbb{N}$, a learner strategy ϕ of length N is comprised of two parts:

1. A finite sequence of N query functions, $(\phi_1, \ldots, \phi_N)$, where each ϕ_k is a mapping that takes as input the values of the first $k-1$ queries submitted, the corresponding responses, as well as the realized value of Y, and outputs the kth query q_k.

2. An estimation function ϕ^E, which takes as input the N queries submitted, the corresponding responses, and the realized value of Y, and outputs the final estimate $\hat{x}$ for the true value v^*.

We will denote by Φ_N the set of all learner strategies of length N, defined as above.

2. We consider a half-open interval here, which allows for a cleaner presentation, but the essence is not changed if the interval is closed.

2.2. Information Available to the Adversary

We summarize in this subsection the information available to the adversary. First, the adversary is aware that the true value v^* belongs to $[0, 1)$. Second, we assume that the adversary can observe the values of the queries but not the corresponding responses, and that the learner strategy ϕ is known to the adversary. In particular, the adversary observes the value of each query q_k, for $k = 1, \ldots, N$, and knows the N mappings, $\phi_1, \phi_2, \ldots, \phi_N$. This means that if the adversary had access to the values $r_1, r_2, \ldots, r_{k-1}$ and the realized value of Y, she would know exactly what q_k is for step k. While it may seem that an adversary who sees both the learner strategy and her actions is too powerful to defend against, we will see in the sequel that the learner will still be able to implement effective and efficient obfuscation by exploiting the randomness of Y.

3. Private Learner Strategies

In this section, we introduce and formally define private learning strategies, the central concept of this paper. A private learner strategy must always make sure that its estimate is close to the true value v^*, while keeping the adversary's probability of correct detection of v^* sufficiently small. Our goal in this section is to formalize those ideas.

3.1. Information Set

Recall from Section 2.2 that the adversary knows the values of the queries and the learner strategy. We will now convert this knowledge into a succinct representation: the *information set* of the adversary. Fix a learner strategy, ϕ. Denote by $\mathcal{Q}(x)$ the set of query sequences that have a positive probability of appearing under ϕ, when the true value v^* is equal to x:

$$\mathcal{Q}(x) = \{\overline{q} \in [0, 1)^N : \mathbb{P}_\phi(Q = \overline{q}) > 0\}, \tag{1}$$

where Q is a vector-valued random variable representing the sequence of learner queries, whereas $\overline{q}$ stands for a typical realization; the probability is measured with respect to the randomness in the learner's random seed, Y.

Definition 1 *Fix $\phi \in \Phi_N$. The information set for the adversary, $\mathcal{I}(\overline{q})$, is defined by:*

$$\mathcal{I}(\overline{q}) = \Big\{x \in [0, 1) : \overline{q} \in \mathcal{Q}(x)\Big\}, \quad \overline{q} \in [0, 1)^N. \tag{2}$$

From the viewpoint of the adversary, the information set represents all possible true values that are consistent with the queries observed. As such, it captures the amount of information that the learner reveals to the adversary.

3.2. $(\epsilon, \delta, L)-$Private Strategies

A private learner strategy should achieve two aims: accuracy and privacy. Accuracy can be captured in a relatively straightforward manner, by measuring the absolute distance between the learner's estimate and the true value. An effective measure of the learner's privacy, on the other hand, is more subtle, as it depends on what the adversary is able to infer. To this end, we develop in this subsection a privacy metric by quantifying the "effective size" of the information set $\mathcal{I}(\overline{q})$ described in Definition 1. Intuitively, since the information set contains all possible realizations of the true

value, v^*, the larger the information set, the more difficult it is for the adversary to pin down the true value.

Definition 2 *Fix $\delta > 0$, $L \in \mathbb{N}$, and a set $\mathcal{E} \subset \mathbb{R}$. We say that a collection of L closed intervals $[a_1, b_1]$, $[a_2, b_2]$, $\ldots$, $[a_L, b_L]$, is a (δ, L) cover for $\mathcal{E}$ if $\mathcal{E} \subset \bigcup_{1 \le j \le L} [a_j, b_j]$, and $b_j - a_j \le \delta$ for all j.*

*We say that a set $\mathcal{E}$ is (δ, L)-**coverable** if it admits a (δ, L) cover. In addition, we define the δ-**cover number** of a set $\mathcal{E}$, $C_\delta(\mathcal{E})$, as*

$$C_\delta(\mathcal{E}) \triangleq \min \{L \in \mathbb{N} : \ \mathcal{E} \text{ is } (\delta, L)\text{-coverable}\}. \tag{3}$$

We are now ready to define (ϵ, δ, L)-private learner strategies.

Definition 3 (Private Learner Strategy) *Fix $\epsilon > 0$, $\delta > 0$, $L \ge 2$, with $L \in \mathbb{N}$. A learner strategy $\phi \in \Phi_N$ is (ϵ, δ, L)-private if it satisfies the following:*

1. *Accuracy constraint: the learner estimate accurately recovers the true value, with probability one:*
$$\mathbb{P}\Big(\big|\hat{x}(x, Y) - x\big| \le \epsilon/2\Big) = 1, \quad \forall\, x \in [0, 1),$$
where the probability is measured with respect to the randomness in Y.

2. *Privacy constraint: for every $x \in [0, 1)$ and every possible sequence of queries $\overline{q} \in \mathcal{Q}(x)$, the δ-cover number of the information set for the adversary, $C_\delta\big(\mathcal{I}(\overline{q})\big)$, is at least L, i.e.,*
$$C_\delta\big(\mathcal{I}(\overline{q})\big) \ge L, \quad \forall\, \overline{q} \in \mathcal{Q}(x). \tag{4}$$

The accuracy constraint requires that a private learner strategy always produce an accurate estimate within the error tolerance ϵ, for any possible true value in $[0, 1)$. The privacy constraint controls the size of the information set induced by the sequence of queries generated, and the parameter L can be interpreted as the learner's privacy level: since the intervals used to cover the information set are of length at most δ, each interval can be thought of as representing a plausible guess for the adversary. Therefore, the probability of the adversary successfully estimating the location of v^* is essentially inversely proportional to the number of intervals needed to cover the information set, which is at most $1/L$.

3.3. Example Applications

We examine two illustrative example applications of our model.

Example 1 - learning an optimal price. A firm is to release a new product and would like to identify a revenue maximizing price, p^*, prior to the product launch. The firm believes that the revenue function, $f(p)$, is strictly concave and differentiable as a function of the price, p, but has otherwise little additional information. A sequential learning process is employed to identify p^* over a series of epochs: in epoch k, the firm assesses how the market responds to a test price, p_k, and receives a binary feedback as to whether $f'(p_k) \ge 0$ or $f'(p_k) < 0$. This may be achieved, for instance, by contracting a consulting firm to conduct market surveys on the price sensitivity around p_k. The firm would like to estimate p^* with reasonable accuracy over a small number of epochs, but is wary that a competitor might be able to observe the surveys and deduce from them the value

of p^* ahead of the product launch. In the context of Private Sequential Learning, the firm is the learner, the competitor is the adversary, the revenue-maximizing price is the true value, and the test prices are the queries. The binary response on the revenue's price sensitivity indicates whether the revenue-maximizing price is less than the current test price.

Example 2 - online optimization with private weights. In the previous example, the adversary is a third-party entity who does not observe the responses to the queries. We now illustrate in this example that the Private Sequential Learning model can also describe a situation where the adversary is the database to which queries are submitted, and thus has partial knowledge of the responses; the connection is made precise in Xu (2017).

Consider a learner who wishes to identify the maximizer, x^*, of a function $f(x) = \sum_{i=1}^{m} \alpha_i f_i(x)$ over some bounded interval $\mathcal{X} \subset \mathbb{R}$, where $\{f_i(\cdot)\}_{1\leq i\leq m}$ is a collection of strictly concave differentiable constituent functions, and $\{\alpha_i\}_{1\leq i\leq m}$ are positive (private) weights representing the importance that the learner associates with each constituent function. The learner knows the weights but does not have information about the constituent functions; such knowledge is to be acquired by querying an external database. During epoch k, the learner submits a test value, x_k, and receives from the database the derivatives of all constituent functions at x_k, $\{f_i'(x_k)\}_{1\leq i\leq m}$. Using the weights, the learner can then compute the derivative $f'(x_k)$, whose sign serves as a binary indicator of the position of the maximizer x^* relative to the current test value. The database, which possesses complete information about the constituent functions but does not know the weights, would like to infer from the learner's querying pattern the maximizing value x^* or possibly the weights themselves. The query strategies we develop for Private Sequential Learning can also be applied in this setting.

3.4. Related Work

Our work is related in spirit to differential privacy (Dwork, 2008; Dwork and Roth, 2014) and the private information retrieval problem in cryptography (Kushilevitz and Ostrovsky, 1997; Chor et al., 1998; Gasarch, 2004). However, in contrast to differential privacy, our definition of privacy measures the adversary's ability to perform a *specific* inference task. It is also substantially weaker than the ones studied in private information retrieval: the adversary may still obtain *some* information on the value the learner is searching for. This relaxation of the privacy requirement allows the learner to deploy richer and more sample-efficient query strategies. In a different model, Tsitsiklis and Xu (2018) study the issue of privacy in a sequential decision problem, where an agent attempts to reach a particular node in a graph, traversing it in a way that obfuscates her intended destination against an adversary who observes her past trajectories. However, a major new element in our model is that the learner strives to *learn* a piece of information of which she herself has no prior knowledge. The central conflict of trying to learn something while preventing others from learning the same information sets our work apart from the extant literature.

4. Main Result

The learner's overall objective is to employ the minimum number of queries while satisfying the accuracy and privacy requirements. We state our main theorem in this section, which establishes lower and upper bounds for the query complexity of a private learner strategy, as a function of the adversary accuracy δ, learner accuracy ϵ, and learner privacy level, L. Recall that Φ_N is the set of

learner strategies of length N. Define $N^*(\epsilon, \delta, L)$ to be the minimum number of queries needed across all (ϵ, δ, L)-private learner strategies,

$$N^*(\epsilon, \delta, L) = \min\big\{N \in \mathbb{N} : \Phi_N \text{ contains at least one } (\epsilon, \delta, L)\text{-private strategy}\big\}. \tag{5}$$

Our result will focus on the regime of parameters where

$$0 < 2\epsilon < \delta \leq 1/L. \tag{6}$$

Having $2\epsilon < \delta$ corresponds to a scenario where the learner would like to identify the true value with high accuracy, while the adversary is aiming for a coarse estimate. Note that the regime where $\delta < \epsilon$ is arguably much less interesting, because it is not natural to expect the adversary, who is not engaged in the querying process, to have a higher accuracy requirement than the learner. The requirement that $\delta \leq 1/L$ stems from the following argument. If $\delta \geq 1/(L-1)$, then the entire interval $[0, 1)$ is trivially $(\delta, L-1)$-coverable, and $C_\delta\big(\mathcal{I}(\overline{q})\big) \leq C_\delta\big([0,1)\big) \leq L - 1 < L$. Thus, the privacy constraint is automatically violated, and no private learner strategy exists. To obtain a nontrivial problem, we therefore only need to consider the case where $\delta < 1/(L-1)$, which is only sightly broader than the regime $\delta \leq 1/L$ that we consider. The following theorem is the main result of this paper[3].

Theorem 4 (Query Complexity of Private Sequential Learning) *Fix $\epsilon > 0$, $\delta > 0$, and a positive integer $L \geq 2$, such that $2\epsilon < \delta \leq 1/L$. Then,*

$$\max\left\{\log\frac{1}{\epsilon},\ \log\frac{\delta}{\epsilon} + 2L - 4\right\} \leq N^*(\epsilon, \delta, L) \leq \log\frac{1}{L\epsilon} + 2L. \tag{7}$$

The proof of the upper bound in Theorem 4 is constructive, providing a specific learner strategy that satisfies the bound. If we set $\delta = 1/L$, which corresponds to the worst case where the adversary's accuracy requirement is essentially as loose as possible, then Theorem 4 leads to the following corollary. It yields upper and lower bounds that are tight up to an additive constant of 4. In other words, the private learner strategy that we construct achieves essentially the optimal query-complexity in this scenario.

Corollary 5 *Fix $\epsilon > 0$ and a positive integer $L \geq 2$ such that $2\epsilon < 1/L$. The following holds.*

1. *If $L = 2$, we have*

$$\log\frac{1}{\epsilon} \leq N^*\Big(\epsilon, \frac{1}{L}, L\Big) \leq \log\frac{1}{\epsilon} + 4. \tag{8}$$

2. *If $L \geq 3$, we have*

$$\log\frac{1}{L\epsilon} + 2L - 4 \leq N^*\Big(\epsilon, \frac{1}{L}, L\Big) \leq \log\frac{1}{L\epsilon} + 2L. \tag{9}$$

3. All logarithms are taken with respect to base 2. To reduce clutter, non-integer numbers are to be understood as rounded upwards.

A main take-away from the above results is about the price of privacy: it is not difficult to see that in the absence of a privacy constraint, the most efficient strategy, using a bisection search, can locate the true value with $\log(1/\epsilon)$ queries. Our results thus demonstrate that the price of privacy is at most an *additive* factor of $2L$.

The proof of Theorem 4 is given in Tsitsiklis et al. (2018). For the upper bound, we construct a certain Opportunistic Bisection (OB) query strategy, where the learner augments a bisection search with $2L$ additional "opportunistic" queries in a randomized and symmetric manner, such the adversary cannot be certain whether the true value is discovered by the bisection search, or the opportunistic queries. For the lower bound, one may ask whether the additional $2L$ queries need to be *distinct* from the $\log(1/\epsilon)$ queries used by the bisection search, or essentially, whether the query complexity could be further reduced by "blending" the queries for obfuscation with those for identifying the true value in a more effective manner. However, we show that such "blending" is not possible.

References

Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–981, 1998.

Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.

William Gasarch. A survey on private information retrieval. In *Bulletin of the EATCS*. Citeseer, 2004.

Eyal Kushilevitz and Rafail Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *Foundations of Computer Science (FOCS)*, volume 97, pages 364–373, 1997.

John Tsitsiklis and Kuang Xu. Delay-predictability tradeoffs in reaching a secret goal. *Operations Research*, 66(2):587–596, 2018.

John N Tsitsiklis, Kuang Xu, and Zhi Xu. Private sequential learning. *arXiv preprint arXiv:1805.02136*, 2018. URL https://arxiv.org/abs/1805.02136.

Zhi Xu. Private sequential search and optimization. Master's thesis, Massachusetts Institute of Technology, 2017. URL http://hdl.handle.net/1721.1/112054.