Proceedings of Machine Learning Research vol xxx:1–5, 2022

Do you pay for Privacy in Online learning?

Giorgia Ramponi GIORGIA.RAMPONI@AI.ETHZ.CH
ETH AI Center, Zurich, Switzerland
Amartya Sanyal AMARTYA.SANYAL@AI.ETHZ.CH
ETH AI Center, Zurich, Switzerland

Online learning, in the mistake bound model, is one of the most fundamental concepts in learning theory. Differential privacy, instead, is the most widely used statistical concept of privacy in the machine learning community. It is then clear that defining problems which are online differential privacy learnable is of great interest. In this paper, we pose the question on if the two problems are the same from a learning perspective, i.e., is privacy for free in the online learning framework?

Keywords: Online Learning, Differential Privacy, Mistake Bound Model

1. Introduction

Online learning, in the mistake bound model, is one of the most fundamental concepts in learning theory. Let $X = \bigcup X_n$ be the instance space. The learner, in this model, receives at each timestep t an unlabelled example $x_t \in X$, predicts a label $\widehat{y}_t$ corresponding to x_t, and then receives the true label y_t for x_t. During this interaction, the learner maintains a working hypothesis h_t, which it uses to predict $\widehat{y}_t = h_t(x_t)$, and then uses the true label y_t to update the working hypothesis to h_{t+1}. The performance of the learner is measured by the number of *mistakes* the algorithm makes, i.e.,:

$$\text{Mistakes}\left(T, (x_t, y_t)_{t=1}^{\infty}\right) := \sum_{t=1}^{T} \left(h_t(x_t) \neq y_t\right).$$

Given this definition of performance, a hypothesis class $\mathcal{C}$ on the instance space $X = \bigcup X_n$ is said to be *online learnable* in the mistake bound model if there exists a learner L that makes at most $\text{poly}\left(n, size\left(c\right)\right)$ mistakes on any sequence of samples consistent with a concept $c \in \mathcal{C}$, where p is some polynomial. This is also known as the *realisable setting*.

Another relevant concept in learning theory is privacy. The most widely used statistical notion of privacy in the machine learning literature is *differential privacy*. An (ϵ, δ)-differentially private (randomised) algorithm is guaranteed to output *similar* distributions over the output space of the algorithm when presented with inputs that only differ in one element. More formally, in the offline setting, a learning algorithm $\mathcal{A} : \mathcal{X} \to \mathcal{Y}$ is said to be (ϵ, δ)-differentially private if, for any two datasets S_1, S_2 that differ in just one element, we have that $\mathbb{P}\left[\mathcal{A}\left(S_1\right) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[\mathcal{A}\left(S_2\right) \in Q\right] + \delta$ where $Q \subseteq \mathcal{Y}$ is any subset of the output space of the algorithm. We define differential privacy in the online setting in Definition 3.

Some previous works (Jain et al., 2012; Agarwal and Singh, 2017; Abernethy et al., 2019) treat the problem of constructing online learning algorithms (mostly in the regret minimization setting (Shalev-Shwartz and Singer, 2007)) maintaining the differential-privacy properties. However, it is still not clear how these two problems (non-private mistake bound and private mistake-bound) are connected and if there exists some problem which is online learnable in the mistake bound model but not private online learnable in the mistake bound model. In other words, the open problem presented in this paper concerns a fundamental question about learning:

"Is privacy for free in the online learning framework?"

© 2022 G. Ramponi & A. Sanyal.

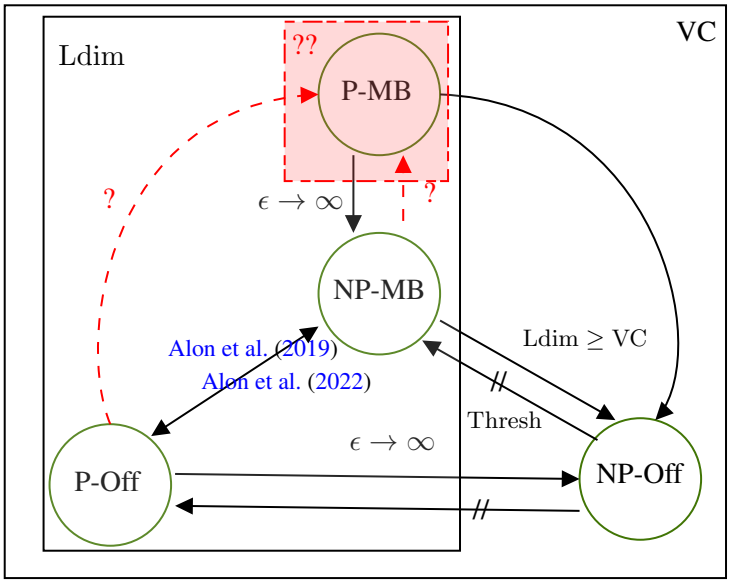

Figure 1: The figure summerizes the relation between the mentioned four online learning problems: Non-Private Offline Learning (NP-Off), Non-Private Mistake Bound Learning (NP-MB), Private-Offline Learning (P-Off) and Private Mistake-Bound model (P-MB).

2. Related works on learnability

In this section we discuss relevant literature on characterising the different learning problems introduced in the previous section and establishing the connections between them. We summarise these relations in Figure 1.

Non-Private Offline Learning The (non-private) offline learning (NP-off) is the most classical learning problem in learning theory. This was formalised by the seminal paper of Valiant (1984) as *Probably approximately correct* (PAC) learnability. A hypothesis class is said to be (α, β)-PAC learnable if there is an algorithm that when given access to a number of samples polynomial in $\frac{1}{\alpha}, \frac{1}{\beta}$, and the problem size returns a hypothesis that achieves error less than α with probability greater than $1 - \beta$. Here, the problem size simply refers to the minimal size of a representation of a hypothesis from the hypothesis class. It is now well known that the Vapnik–Chervonenkis dimension (VC) dimension (Vapnik, 1999) exactly characterises non-private offline learnability in that any hypothesis class with finite VC dimension is learnable in the PAC model (and vice versa).

Non-Private Online Learning As discussed before, a hypothesis class $\mathcal{H}$ is said to be learnable in the online mistake bound model if there is a finite M and an online algorithm $\mathcal{A}$ such that $\mathcal{A}$ makes at most M mistakes on any sequence of data labelled with some $h \in \mathcal{H}$. Interestingly, it is also possible to characterise online learnability using a different combinatorial measure of the hypothesis class called the Littlestone dimension, which we define in Definition 1. Littlestone (1988) proved that for any hypothesis class $\mathcal{H}$, there exists an online learning algorithm that makes at most $\mathrm{Ldim}(\mathcal{H})$ mistakes on any sequence labelled by some $h \in \mathcal{H}$, thereby characterising online learnability.

Definition 1 (Littlestone dimension (Littlestone, 1988)) *The littlestone dimension of a hypothesis class $\mathcal{H}$, denoted as* $\mathrm{Ldim}(\mathcal{H})$ *is the depth of the largest tree that can be shattered by $\mathcal{H}$, where we define* "shattering a tree" *in Definition 2.*

Definition 2 (Shattering a tree) *Consider a full binary tree of depth d such that each node is labelled by some $x \in \mathcal{X}$. For a set of labels $\{y_i\}_{i=1}^{d}$, define its corresponding path as starting from the root and taking the left child when $y = -1$ and the right child when $y = +1$. The tree is said to be shattered by some $h \in \mathcal{H}$ if for every set of labels in $\{-1, 1\}^d$, its corresponding path can be shattered by some $h \in \mathcal{H}$ i.e. for all x_i in the path, $h(x_i) = y_1$.*

Private Offline Learning The non-private offline PAC learnability definition was extended to the case of differentially private learnability by Raskhodnikova et al. (2008). A hypothesis class is $(\epsilon, \delta, \alpha, \beta)$-differentially private PAC learnable if there exists an (α, β)-PAC learning algorithm that is also (ϵ, δ)-differentially private. Raskhodnikova et al. (2008) showed that any problem that is PAC learnable is also learnable by a differentially private learning algorithm but the required number of samples is dependant on the size of the input space in addition to the VC dimension, which can be arbitrarily larger than the VC dimension. This left open the question of whether the sample complexity can be characterised exactly by a combinatorial measure of the complexity of the hypothesis class.

Alon et al. (2019) resolved a part of the question by answering that the required number of samples is at least $\Omega(\log^*(\mathrm{Ldim}(\mathcal{H})))$ where $\log^*$ is the iterated logarithm. Alon et al. (2022) proved the reverse side and concluded that any class with finite littlestone dimension can be learned offline privately with a finite number of samples. More specifically they showed that any hypothesis class with a finite Littlestone dimension d is private learnable with number of samples double exponential in d. This concludes that private offline learnability is exactly characterised by the littlestone dimension (Littlestone, 1988), which in turn exactly characterises online learnability in the mistake bound model thereby showing an equivalence between the two learning regimes. However, this leaves open the question of whether private learnability, with a suitable definition, is harder than non-private offline learnability.

3. Open Problem

In this section we expose the research question introduced in this paper. Before it, we introduce the concept of $\{\epsilon, \delta\}$-differentially private online learning algorithm.

Definition 3 ($\{\epsilon, \delta\}$-differential online privacy) *Let $\mathcal{H}$ be a set of hypotheses $\mathcal{H} = \bigcup_{n=1}^{\infty} \mathcal{H}_n$ over the input space $\mathcal{X} = \bigcup_{n=1}^{\infty} \mathcal{X}_n$. Then an online algorithm $\mathcal{A}$ is $\{\epsilon, \delta\}$-online differentially private if for all $T \in \mathbb{N}$, for any two sequences of points S_T and S'_T that differs in at most one entry the following holds:*

$$Pr(\mathcal{A}(S_T) \in \mathcal{S}) \leq e^{\epsilon} Pr(\mathcal{A}(S'_T) \in \mathcal{S}) + \delta$$

The question that we pose is if every problem that is online learnable it is also online privately learnable, in other words, if the set of problems solvable in these two learnability classes are the same. This question can be solved proving one of the two following theorems, where theorem 1 implies that there exists a problem which is online learnable but non-online private learnable and theorem 2 implies, instead, the opposite.

Theorem 1 *There exists a set of hypotheses $\mathcal{H} = \bigcup_{n=1}^{\infty} \mathcal{H}_n$ over the input space $\mathcal{X} = \bigcup_{n=1}^{\infty} \mathcal{X}_n$ such that for all $T \in \mathbb{N}$, for any sequence of points $S_T = \{(x_1, h^*(x_1)), \ldots, (x_T, h^*(x_T))\}$, such that $h^* \in \mathcal{H}$,*

1. *(Online learnable) there exists an online algorithm $\mathcal{A}$ that does not make more than M mistakes on the sequence S_T for some $M < \infty$.*

2. *(Not privately online learnable) any (ϵ, δ)-differentially private online algorithm makes at least $M' \geq M + \alpha(\epsilon, \delta, T)$ mistakes,*

where $\alpha : \mathbb{R} \times [0,1] \to \mathbb{N}$ is such that $\alpha(\epsilon, \delta, T) \gtrsim_{\delta} \frac{\sqrt{T}}{\epsilon}$.

Theorem 1 claims that there exists some hypothesis class that is non-privately online learnable but any private online algorithm makes infinite mistakes when $\epsilon \lesssim \sqrt{T}$. Here, the symbols $\gtrsim$ and $\lesssim$ mean greater than or less than upto a multiplicative constant and $\gtrsim_{\delta}$ ignores the dependance on δ. As we know that a non-private algorithm can solve the problem with small number of mistakes, it is natural to expect that any hardness result would only hold for a sufficiently small ϵ. We next state another hypothesis which states that any non-privately online learnable hypothesis class is also privately online learnable as long as the privacy parameter ϵ is bounded away from zero. We state this in Theorem 2 below.

Theorem 2 *Let $\mathcal{H}$ be any hypothesis class that is online learnable. Then, there exists a positive monotonically decreasing function $\gamma : \mathbb{R}_+ \to \mathbb{R}_+$ such that for all $h \in \mathcal{H}, T \in \mathbb{N}$, and $\gamma(T) \gtrsim \epsilon \gtrsim \sqrt{T}$ there exists an (ϵ, δ)-differentially private online algorithm that makes a finite number of mistakes for any sequence of points $S_T = \{(x_1, h^*(x_1)), \ldots, (x_T, h^*(x_T))\}$ of length T labelled by $h \in \mathcal{H}$, where $\gamma(z) < \sqrt{z}$ for all $z \in \mathbb{R}_+$.*

Theorem 2 states, instead, that for every hypothesis class that is online learnable, there exists an (ϵ, δ)-differentially private online learning algorithm as long as ϵ is not too small. The lower bound on ϵ is natural to consider as imposing the condition that ϵ goes to zero essentially necessitates a trivial algorithm where we cannot hope to have any reasonable mistake bound.

By definition, all privately online learnable problems are also non-privately online learnable (take $\epsilon \to 0$). Therefore, one possible implication of a proof for Theorem 1 would be the definition of a combinatorial measure that is even more restrictive than littlestone definition (Ldim is more restrictive than VC), which we are not aware of and is perhaps of even wider interest to the learning theory community.

However, recent results from Bousquet et al. (2021) in the context of *universal learning* (which is another definition of learnability in the same spirit as PAC learning) suggests that a combinatorial measure that is more restrictive than the Littlestone dimension is unlikely. In particular, they show that there are only three possible rates in universal learning with the fastest being characterised by the littlestone dimension and the slowest by VC dimension. This makes Theorem 2 more likely. Some initiali progress towards this has been made by Golowich and Livni (2021). An interesting outcome of a proof for Theorem 2 is a general algorithm to convert an online learner to a private online learner. Another interesting implication of this could be a characterisation of γ which would establish a lower bound for privacy in online learning algorithms. We promise a wheel of parmigiano reggiano to whoever proves Theorem 1 or a tub of biriyani for solving Theorem 2.

References

Jacob D Abernethy, Young Hun Jung, Chansoo Lee, Audra McMillan, and Ambuj Tewari. Online learning via the differential privacy lens. *Advances in Neural Information Processing Systems*, 32, 2019.

Naman Agarwal and Karan Singh. The price of differential privacy for online learning. In *International Conference on Machine Learning*, pages 32–40. PMLR, 2017.

Noga Alon, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private pac learning implies finite littlestone dimension. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 852–860, 2019.

Noga Alon, Mark Bun, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private and online learnability are equivalent. *ACM Journal of the ACM (JACM)*, 2022.

Olivier Bousquet, Steve Hanneke, Shay Moran, Ramon Van Handel, and Amir Yehudayoff. A theory of universal learning. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 532–541, 2021.

Noah Golowich and Roi Livni. Littlestone classes are privately online learnable. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=4bKbEP9b65v.

Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Conference on Learning Theory*, pages 24–1. JMLR Workshop and Conference Proceedings, 2012.

Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 2(4):285–318, 1988.

Sofya Raskhodnikova, Adam Smith, Homin K Lee, Kobbi Nissim, and Shiva Prasad Kasiviswanathan. What can we learn privately. In *Proceedings of the 54th Annual Symposium on Foundations of Computer Science*, pages 531–540, 2008.

Shai Shalev-Shwartz and Yoram Singer. Online learning: Theory, algorithms, and applications. 2007.

Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 1999.