

On the Growth of Mistakes in Differentially Private Online Learning: A Lower Bound Perspective

Daniil Dmitriev

Department of Mathematics, ETH Zurich and ETH AI Center, Zurich, Switzerland

DANIIL.DMITRIEV@AI.ETHZ.CH

Kristóf Szabó

Department of Mathematics, ETH Zurich, Zurich, Switzerland

KRIST.SZ13@GMAIL.COM

Amartya Sanyal

Max Planck Institute for Intelligent Systems, Tübingen, Germany

AMSA@DI.KU.DK

Editors: Shipra Agrawal and Aaron Roth

Abstract

In this paper, we provide lower bounds for Differentially Private (DP) Online Learning algorithms. Our result shows that, for a broad class of (ϵ, δ) -DP online algorithms, for number of rounds T such that $\log T \leq O(1/\delta)$, the expected number of mistakes incurred by the algorithm grows as $\Omega(\log T)$. This matches the upper bound obtained by [Golowich and Livni \(2021\)](#) and is in contrast to non-private online learning where the number of mistakes is independent of T . To the best of our knowledge, our work is the first result towards settling lower bounds for DP-Online learning and partially addresses the open question in [Sanyal and Ramponi \(2022\)](#).

1. Introduction

With the increasing need to protect the privacy of sensitive user data while conducting meaningful data analysis, Differential Privacy (DP) ([Dwork et al., 2006](#)) has become a popular solution. DP algorithms ensure that the impact of any single data sample on the output is limited, thus safeguarding individual privacy. Several works have obtained DP learning algorithms for various learning problems in both theory and practice.

However, privacy does not come for free and often leads to a statistical (and sometimes computational) cost. The classical solution for non-private Probably Approximately Correct (PAC) learning ([Valiant, 1984](#)) is via Empirical Risk Minimisation (ERM) that computes the best solution on the training data. Several works ([Bassily et al., 2014](#); [Chaudhuri et al., 2011](#)) have shown that incorporating DP into ERM incurs a compulsory statistical cost that depends on the dimension of the problem. In the well-known setting of PAC learning with DP, [Kasiviswanathan et al. \(2011\)](#) provided the first guarantees that all finite VC classes can be learned with a sample size that grows logarithmically in the size of the class. This line of research was advanced by subsequent works ([Beimel et al., 2013a, 2014](#); [Feldman and Xiao, 2014](#)), resulting in the findings of [Alon et al. \(2022\)](#) which established a surprising equivalence between non-private online learning and Approximate DP-PAC learning.

Unlike the setting of PAC learning, Online learning captures a sequential game between a learner and an adversary. The adversary knows everything about the learner’s algorithm except its random bits. In this work we consider a setting where, for a known hypothesis class \mathcal{H} , the adversary chooses a sequence of data points $\{x_1, \dots, x_t\}$ and the target hypothesis $f^* \in \mathcal{H}$ prior to engaging with the learner. Then, the adversary reveals these data points one by one to the learner, who must

offer a prediction for each. After each prediction, the adversary reveals the true label for that point. The learner’s performance is evaluated by comparing the incurred mistakes against the theoretical minimum that could have been achieved by an optimal hypothesis in hindsight. Known as the *realisable oblivious mistake bound* model, the seminal work of Littlestone (1988) showed that i) the number of mistakes incurred by any learner is lower-bounded by the Littlestone dimension (more precisely, $\text{Ldim}(\mathcal{H})/2$) of the target class \mathcal{H} and ii) there is an algorithm that makes at most $\text{Ldim}(\mathcal{H})$ mistakes. This algorithm is commonly referred to as the Standard Optimal Algorithm (SOA).

Recall that certain problem classes possess finite Vapnik-Chervonenkis (VC) dimensions but infinite Littlestone dimensions (such as the one-dimensional threshold problem). This, together with the equivalence between non-private online learning and DP-PAC learning (Alon et al., 2022) implies that there exists a fundamental separation between DP-PAC learning and non-private PAC learning. In other words, some learning problems can be solved with vanishing error, as the amount of data increases, in PAC learning but will suffer unbounded error in DP-PAC learning. This implication was first proven for pure DP by Feldman and Xiao (2014) and later for approximate DP by Alon et al. (2019). With the debate on the sample complexity of approximate DP-PAC learning resolved, we next ask whether a similar gap exists between online learning with DP and non-private online learning. Golowich and Livni (2021) addressed this by introducing the Differentially Private Standard Optimal Algorithm (DP-SOA), which suffers a mistake count, that increases logarithmically with the number of rounds T compared to a constant error rate in non-private online learning (Littlestone, 1988). This difference suggests a challenge in DP online learning, where errors increase indefinitely as the game continues. The question of whether this growing error rate is an unavoidable aspect of DP-online learning was posed as an open question by Sanyal and Ramponi (2022).

Main Result In this work, we provide evidence that this additional cost is inevitable. Consider any hypothesis class \mathcal{H} and for a learning algorithm \mathcal{A} . Let $\mathbb{E}[M_{\mathcal{A}}]$ be the expected number of mistakes incurred by \mathcal{A} and let T be the total number of rounds for which the game is played.

We obtain a lower bound on $\mathbb{E}[M_{\mathcal{A}}]$ under some assumptions on the learning algorithm \mathcal{A} . Informally, we say an algorithm \mathcal{A} is β -concentrated (see Definition 7 for a formal definition) if there is some output sequence that it outputs with probability at least $1 - \beta$ in response to a *non-distinguishing* input sequence. A *non-distinguishing* input sequence is a (possibly repeated) sequence of input data points such that there exists some $f_1, f_2 \in \mathcal{H}$ which cannot be distinguished just by observing their output on the non-distinguishing input sequence. We prove a general statement for any hypothesis class in Theorem 1 but show an informal corollary below.

Corollary 1 (Informal Corollary of Theorem 1) *There exists a hypothesis class \mathcal{H} with $\text{Ldim}(\mathcal{H}) = 1$ (see Definition 2), such that for any $\varepsilon, \delta > 0, T \leq \exp(1/(32\delta))$, and any online learner \mathcal{A} that is (ε, δ) -DP and 0.1-concentrated, there is an adversary, such that*

$$\mathbb{E}[M_{\mathcal{A}}] = \tilde{\Omega}\left(\frac{\log T}{\varepsilon}\right), \quad (1.1)$$

where $\tilde{\Omega}$ hides logarithmic factors in ε . For $T > \exp(1/(32\delta))$, $\mathbb{E}[M_{\mathcal{A}}] = \tilde{\Omega}(1/\delta)$.

While the above result uses a hypothesis class of Littlestone dimension one, our main result in Section 4.1 also holds for any hypothesis class, even with Littlestone dimension greater than one.

Utilising the Point_N hypothesis class (see Definition 4) in Corollary 1, we demonstrate that the minimum number of mistakes a DP online learner must make is bounded below by a term that increases logarithmically with the time horizon T . This holds if the learning algorithm is concentrated at least when T is less than or equal to $\exp(1/(32\delta))$. This contrasts with non-private online learning, where the number of mistakes does not increase with T in hypothesis classes with bounded Littlestone dimension, even if the learner is concentrated.¹ Our result also shows that the analysis of the algorithm of Golowich and Livni (2021), which shows an upper bound of $\Omega(\log(T/\delta))$ for DP-SOA is tight as long as $T \leq \exp(1/(32\delta))$. However, as illustrated in Figure 1, this is not a limitation as for larger T , since a simple *Name & Shame* algorithm incurs lesser mistake than DP-SOA albeit at vacuous privacy levels (see discussion after Theorem 1).

In fact, the assumption of concentrated learners is not overly restrictive given that known DP-online learning algorithms exhibit this property, as detailed in Section 4.2. Notably, the DP-SOA presented by Golowich and Livni (2021) which is the sole DP online learning method known to achieve a mistake bound $O(\log(T))$, is concentrated as shown in Corollary 2. This suggests that the lower bound holds for all potential DP online learning algorithms.

Additionally, we extend our result to another class of DP online algorithms, which we refer to as *uniform firing* algorithms, that are in essence juxtaposed to concentrated algorithms. These algorithms initially select predictors at random until a certain confidence criterion is met, prompting a switch to a consistent predictor—this transition, or ‘firing’, is determined by the flip of a biased coin (with bias p_t), where the likelihood of firing increases with each mistake. However, the choice of how p_t increases and the selection of the predictor upon firing depend on the algorithm’s design. For this specific type of algorithms, particularly in the context of learning the Point_3 hypothesis class, Proposition 4 establishes a lower bound on mistakes that also grows logarithmically with T .

Section 5.1 discusses Continual Observation (Chan et al., 2011; Dwork et al., 2010b), another popular task within sequential DP. We show that results on DP Continual Counters can be used to derive upper bounds in the online learning setting. Nonetheless, it is not clear whether lower bounds for that setting can be transferred to DP-Online learning. In addition, these upper bounds suffer a dependence on the hypothesis class size.

Finally, we point out that, to the best of our knowledge we are unaware of any algorithms in the literature for pure DP online learning. Our lower bound in Corollary 1 immediately provides a lower bound for pure DP. Similarly, DP Continual Counters provide a method for achieving upper bounds, specifically for the Point_N classes, albeit with a linear dependency on N . Obtaining tight upper and lower bounds remains an interesting direction for future research.

2. Preliminaries

We provide the necessary definitions for both Online Learning and Differential Privacy.

2.1. Online Learning

We begin by defining the online learning game between a learner \mathcal{A} and an adversary \mathcal{B} . Let $T \in \mathbb{N}_+$ denote number of rounds and let \mathcal{X} be some domain.

Definition 1 (General game) *Let $\mathcal{H} \subseteq \{0, 1\}^{\mathcal{X}}$ be a hypothesis class of functions from \mathcal{X} to $\{0, 1\}$. The game between a learner \mathcal{A} and adversary \mathcal{B} is played as follows:*

1. All deterministic algorithms are 0-concentrated by definition.

- adversary \mathcal{B} picks $f^* \in \mathcal{H}$ and a sequence $x_1, \dots, x_T \in \mathcal{X}$
- for $t = 1, \dots, T$:
 - learner \mathcal{A} outputs a current prediction $\hat{f}_t \in \{0, 1\}^{\mathcal{X}}$ (see Definition 3)
 - \mathcal{A} receives $(x_t, f^*(x_t))$
- let $M_{\mathcal{A}} = \sum_{t=1}^T \mathbb{I} \{ \hat{f}_t(x_t) \neq f^*(x_t) \}$

In this work, we study the following min-max problem:

$$\min_{\mathcal{A}} \max_{\mathcal{B}} \mathbb{E} [M_{\mathcal{A}}] = \min_{\mathcal{A}} \max_{\mathcal{B}} \sum_{t=1}^T \mathbb{P} \left[\hat{f}_t(x_t) \neq f^*(x_t) \right], \quad (2.1)$$

where probability is taken over the randomness in \mathcal{A} . We refer to the random variable $M_{\mathcal{A}}$ as the mistake count of \mathcal{A} . The optimal mistake count of this game can be characterised by a combinatorial property of the hypothesis class \mathcal{H} , called *Littlestone dimension*, first shown by Littlestone (1988).

Littlestone dimension To define *Littlestone dimension*, we need the concept of *mistake tree*. A mistake tree \mathcal{T} is a complete binary tree, where each internal node v corresponds to some $x_v \in \mathcal{X}$. Each root-to-leaf path of the tree – denoted as $x_1, x_2, \dots, x_d, x_{\text{leaf}}$ – is associated with a label sequence y_1, \dots, y_d , where $y_i = \mathbb{I} \{ x_{i+1} \text{ is the right child of } x_i \}$.

We say that \mathcal{T} is *shattered* by \mathcal{H} , if for every possible root-to-leaf labeled path $((x_1, y_1), \dots, (x_d, y_d))$, there exists $f \in \mathcal{H}$, such that $f(x_i) = y_i$ for all $i \in [d]$. This concept leads us to the formal definition of the Littlestone dimension as follows:

Definition 2 *Littlestone dimension* ($\text{Ldim}(\mathcal{H})$) of a hypothesis class \mathcal{H} is defined as the maximum depth of any mistake tree that can be shattered by \mathcal{H} .

Littlestone (1988) proved that for any hypothesis class \mathcal{H} , there exists a deterministic learner, called Standard Optimal Algorithm (SOA) such that $M_{\text{SOA}} \leq \text{Ldim}(\mathcal{H})$. Furthermore, for any learner \mathcal{A} there exists a (possibly random) adversary \mathcal{B} , such that $\mathbb{E} [M_{\mathcal{A}}] \geq \frac{\text{Ldim}(\mathcal{H})}{2}$, where expectation is taken with respect to the randomness in \mathcal{B} . However, the SOA learner is not restricted to output a hypothesis in \mathcal{H} while making its predictions. Such learning algorithms are classified as *improper learners*, which we define formally below:

Definition 3 (Proper and Improper learner) A learner \mathcal{A} for a hypothesis class \mathcal{H} is called *proper* if its output is restricted to belong to \mathcal{H} . Any learner that is not proper is called *improper*.

It is worth noting that most learners in online learning are improper learners though their output may only be simple mixtures of hypothesis in \mathcal{H} (Hanneke et al., 2021). We illustrate the importance of improper learners with a simple hypothesis class that we heavily use in the rest of this paper.

Definition 4 (Point class) For $N \in \mathbb{N}_+$, for domain $\mathcal{X} = [N] := \{1, \dots, N\}$, define

$$\text{Point}_N := \left\{ f^{(i)}, i \in [N] \right\}, \quad \text{where} \quad f^{(i)}(x) = \mathbb{I} \{ i = x \}. \quad (2.2)$$

Note that $\text{Ldim}(\text{Point}_N) = 1$ for any $N > 1$. A simple algorithm to learn Point_N predicts 0 for every input until it makes a mistake. The input i where it incurred the mistake must correspond to the true target concept $f^{(i)}$. This algorithm is improper as no hypothesis f in Point_N predicts 0 universally over the whole domain.

2.2. Differential Privacy

In this work, our goal is to study learners that satisfy the DP guarantee (Dwork et al., 2006) defined formally below.

Definition 5 (Approximate differential privacy) *An algorithm \mathcal{A} is said to be (ϵ, δ) -DP, if for any two input sequences $\tau = ((x_1, y_1), \dots, (x_T, y_T))$ and $\tau' = ((x'_1, y'_1), \dots, (x'_T, y'_T))$, such that there exists **only one** t with $(x_t, y_t) \neq (x'_t, y'_t)$, it holds that*

$$\mathbb{P}(\mathcal{A}(\tau) \in S) \leq \exp(\epsilon) \mathbb{P}(\mathcal{A}(\tau') \in S) + \delta, \quad (2.3)$$

where S is any set of possible outcomes.

When $\delta = 0$ we recover the definition of *pure differential privacy*, denoted by ϵ -DP. Note that for online learner \mathcal{A} , output at step t depends only on first $t - 1$ elements of the input sequence. In the setting of offline learning, the inputs τ, τ' can be thought of as two datasets of length T and \mathcal{A} as the learning algorithm that outputs one hypothesis f (not necessarily in \mathcal{H}). If \mathcal{A} simultaneously satisfies DP and is a PAC learner (Valiant, 1984), it defines the setting of DP-PAC learning (Kasiviswanathan et al., 2011). However, the setting of DP-online learning is more nuanced due to two reasons.

Privacy of Prediction or Privacy of Predictor The first complexity arises from what the privacy adversary observes when altering an input, termed as its *view*. Since \mathcal{A} provides an output hypothesis $\hat{f}_t \in \{0, 1\}^{\mathcal{X}}$ at every time step $t \in [T]$ as shown in Definition 1, the adversary’s view could encompass the entire list of predictors. Our work, like Golowich and Livni (2021), focuses on this scenario, where the output set is $S \subseteq \{0, 1\}^{\mathcal{X} \times T}$. Nevertheless, certain studies restrict the adversary’s view to only the predictions, excluding the predictors themselves (Beimel et al., 2013b; Dwork and Feldman, 2018; Naor et al., 2023). In this setting, it is also important to assume that the adversary only observes the predictions on the inputs that it did not change (Kaplan et al., 2023; Kearns et al., 2015); thus, they have $S \subseteq \{0, 1\}^{(T-1)}$.

Oblivious and Adaptive adversary The second complexity is about whether the online adversary \mathcal{B} pre-selects all input points or adaptively chooses the next point based on the learner \mathcal{A} ’s previous response. Although the former, known as an *oblivious* adversary, seems less potent, this difference does not manifest itself in non-private learning (Cesa-Bianchi and Lugosi, 2006). However, this distinction becomes significant in the context of DP online learning. Adaptive adversaries, by design, leverage historical data in their decision-making process. While works like Kaplan et al. (2023) focus on adaptive adversaries, others like Kearns et al. (2015) concentrate on oblivious ones, and Golowich and Livni (2021) examine both. Our contribution lies in setting lower bounds against the simpler scenario of oblivious adversaries.

3. Related work

Understanding which hypothesis classes can be privately learned is an area of vibrant research and was started in the context of Valiant’s PAC learning model (Valiant, 1984). A hypothesis class \mathcal{H} is considered PAC-learnable if there exists an algorithm \mathcal{A} , which can utilize a polynomial-sized²,

2. The term ‘polynomial-sized’ refers to a sample size that is polynomial in the PAC parameters, including the error rate, confidence level, size of the hypothesis class, and the dimensionality of the input space.

independent, and identically distributed (i.i.d.) sample D from any data distribution to produce a hypothesis $h \in \mathcal{H}$ that achieves a low classification error with high probability on that distribution. In the context of DP-PAC learning, as defined by [Kasiviswanathan et al. \(2011\)](#), the learner \mathcal{A} must also satisfy DP constraint with respect to the sample D . The overarching objective in this research domain is to find a clear criterion for the private learnability of hypothesis classes, analogous to the way learnability has been characterized in non-private settings—through the Vapnik-Chervonenkis (VC) dimension for offline learning ([Blumer et al., 1989](#)) and the Littlestone dimension for online learning ([Ben-David et al., 2009](#); [Littlestone, 1988](#)).

[Kasiviswanathan et al. \(2011\)](#) started this line of research by showing that the sample complexity of DP-PAC learning a hypothesis class \mathcal{H} is $O(\log(|\mathcal{H}|))$. [Beimel et al. \(2014\)](#) showed that the VC dimension does not dictate the sample complexity for proper pure DP-PAC learning of the Point_N class. However, they showed that if the setting is relaxed to improper learning then this sample complexity can be improved, thus showing a separation between proper and improper learning, something that is absent in the non-private PAC model. [Beimel et al. \(2013a\)](#) sharpened this result by constructing a new complexity measure called *probabilistic representation dimension* and proving that this measure characterises improper pure DP exactly.

By leveraging advanced tools from communication complexity theory, they refined the understanding of the probabilistic representation dimension and demonstrated that the sample complexity for learning a notably simple hypothesis class, denoted as Line_p , under approximate improper DP-PAC conditions, is significantly lower than the corresponding lower bound established for pure contexts.

Relaxing the notion of pure DP to approximate DP, [Beimel et al. \(2013b\)](#) showed that the sample complexity for proper approximate DP-PAC learning can be significantly lower than proper pure DP-PAC learning, thereby showing a separation between pure and approximate DP in the context of proper DP-PAC learning. The inquiry into whether a similar discrepancy exists in improper DP-PAC learning was resolved by [Feldman and Xiao \(2014\)](#) who proved a separation between pure and approximate DP in the improper DP-PAC learning model. To do this, they first proved a sharper characterisation of the probabilistic representation dimension using concepts from communication complexity. Then, they showed that the sample complexity for learning a notably simple hypothesis class, denoted as Line_p , under approximate improper DP-PAC conditions, is significantly lower than the corresponding lower bound established for pure DP.

[Feldman and Xiao \(2014\)](#) were also the first to obtain lower bounds for DP-PAC learning that grows as $\Omega(\text{Ldim}(\mathcal{H}))$, albeit limited to the pure DP setting. [Alon et al. \(2019\)](#) showed that it is possible to obtain a lower bound for approximate DP that grows as $\Omega(\log^*(\text{Ldim}(\mathcal{H})))$ thus marking a clear distinction between non-private and approximate DP-PAC learning. This finding illustrated that DP-PAC learning’s complexity could align with that of online learning, which is similarly governed by the Littlestone dimension. In a series of subsequent works, see [Alon et al. \(2022\)](#); [Ghazi et al. \(2021\)](#), a surprising connection was established between private offline learning and non-private online learning. In particular, classes that are privately offline learnable are precisely those with finite Littlestone dimension.

This naturally highlights a similar question of private online learning, in particular whether DP further limits which classes are learnable in the DP-Online learning model. [Golowich and Livni \(2021\)](#) provided an algorithm, called DP-SOA, which has expected number of mistakes growing as $O(2^{2^{\text{Ldim}}} \log T)$. Interestingly, unlike SOA in the non-private online setting, DP-SOA’s mistake count increases with number of steps T in the online game. When considering adaptive adversaries,

the upper bound on mistakes escalates to $O(\sqrt{T})$. Under a slightly weaker definition of DP, known as Challenge-DP, where the privacy adversary only sees the predictions and not the whole predictor function, Kaplan et al. (2023) obtained an upper bound of $O(\log^2(T))$ for both adaptive and oblivious adversaries. However, it is not clear from these works, whether the dependence on T is unavoidable. A related setting is that of *continual observation under DP* where such a dependence is indeed unavoidable under the pure DP model. However, the results from continual observation do not immediately transfer to online learning as discussed in Section 5.1.

4. Lower Bound for Private Online learning under Concentration assumption

In this section, we provide the main result of our work along with their proof. Before stating the main result in Theorem 1, we need to define the concept of *distinguishing tuple* and β -concentrated learners.

Definition 6 Given $f_0, f_1 \in \mathcal{H}$ and $x_{\text{eq}}, x_{\text{dif}} \in \mathcal{X}$, we call the tuple $(f_0, f_1, x_{\text{eq}}, x_{\text{dif}})$ distinguishing, if it satisfies both $f_0(x_{\text{eq}}) = f_1(x_{\text{eq}})$ and $f_0(x_{\text{eq}}) = f_0(x_{\text{dif}}) \neq f_1(x_{\text{dif}})$.

A distinguishing tuple means that there are two functions (f_0, f_1) , and two input points $(x_{\text{eq}}, x_{\text{dif}})$, such that only one of these points can effectively differentiate between the two functions. The absence of a distinguishing tuple implies a restricted hypothesis class: either \mathcal{H} is a singleton ($|\mathcal{H}| = 1$), or it contains precisely two inversely related functions ($|\mathcal{H}| = 2$ with $f_1 = 1 - f_0$). In the latter case, every input point contains information distinguishing f_1 and f_2 . This implies that there is no difference between input sequences, and the mistake bound will not depend on T . For the purposes of our analysis, we proceed under the assumption that a distinguishing tuple always exists.

Let $(f_0, f_1, x_{\text{eq}}, x_{\text{dif}})$ be a distinguishing tuple and suppose that adversary chooses $f^* \in \{f_0, f_1\}$. Knowing only information on $f^*(x_{\text{eq}})$ does not help to tell apart f_0 from f_1 . Furthermore, if an algorithm is ‘too confident’, meaning that it strongly prefers output of f_0 on x_{dif} over output of f_1 , it will necessarily make a mistake on x_{dif} if $f^* = f_1$. We will use this basic intuition to obtain our main lower bound and we will call such learners ‘concentrated’, as defined below.

For input τ , index $t \in [T]$ and an input point $x \in \mathcal{X}$, we denote $\mathcal{A}(\tau)_t[x]$ to be the value of the t -th output function of $\mathcal{A}(\tau)$ evaluated at point x .

Definition 7 An algorithm \mathcal{A} is called β -concentrated, if there exists a distinguishing tuple $(f_0, f_1, x_{\text{eq}}, x_{\text{dif}})$, such that

$$\mathbb{P}[\forall t \in [T], \mathcal{A}(\tau_0)_t[x_{\text{dif}}] = f_0(x_{\text{dif}})] \geq 1 - \beta, \quad (4.1)$$

where $\tau_0 = ((x_{\text{eq}}, f_0(x_{\text{eq}})), \dots, (x_{\text{eq}}, f_0(x_{\text{eq}})))$.

Note that τ_0 from Definition 7 is a ‘dummy’, *non-distinguishing* input, as it does not contain any information to distinguish f_0 from f_1 .

4.1. Main Result

We now show that if the learner is both differentially private and concentrated, it will necessarily suffer a large (logarithmic in T) number of mistakes in the game of Definition 1.

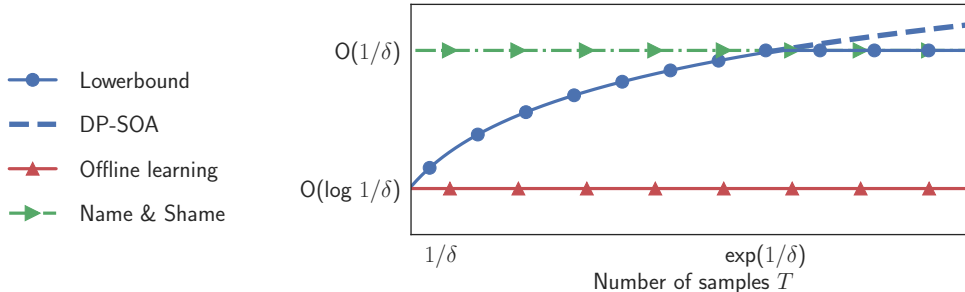


Figure 1: Lower bound from Theorem 1 vs. existing upper bounds. We assume that ε, δ are fixed and for simplicity ignore dependence on ε . X-axis corresponds to the number of samples, growing from $T \sim 1/\delta$ to $T \sim \exp(1/\delta)$ and larger. Y-axis shows the expected number of mistakes, $\mathbb{E}[M_{\mathcal{A}}]$.

Theorem 1 *Let \mathcal{H} be an arbitrary hypothesis class. Let $\varepsilon > 0$, $\delta \leq \varepsilon^2$ and $T \leq \exp(1/(32\delta))$. If, for some $\delta \leq \beta \leq 1/10$, \mathcal{A} is a β -concentrated (ε, δ) -DP online learner of \mathcal{H} , then there exists an adversary \mathcal{B} , such that*

$$\mathbb{E}[M_{\mathcal{A}}] = \tilde{\Omega}\left(\frac{\log T/\beta}{\varepsilon}\right), \quad (4.2)$$

where $\tilde{\Omega}$ contains logarithmic in ε factors. For $T > \exp(1/(32\delta))$, $\mathbb{E}[M_{\mathcal{A}}] = \tilde{\Omega}(1/\delta)$.

Before comparing our lower bound with known upper bounds, we first discuss the condition $T \leq \exp(1/(32\delta))$. Our bound suggests that for sufficiently large T , specifically when T exceeds $\exp(1/(32\delta))$, the dependency on T is in fact not needed. This can be seen from the following simple *Name & Shame* algorithm: initialize an empty set S ; at each step, apply SOA and output $\text{SOA}(S)$; upon receiving a new entry $(x_t, f^*(x_t))$, add it to S with probability δ . Clearly, this algorithm is $(0, \delta)$ -DP, since for any fixed input point, it only depends on this point with probability δ . Furthermore, each time the algorithm incurs a mistake, it adds this mistake to S with probability δ . Since the algorithm runs SOA on S , it will make on expectation at most $\text{Ldim}(\mathcal{H})/\delta$ mistakes. This algorithm is, however, not ‘conventionally’ private, since it potentially discloses a δ -fraction of the data, namely the set S .

Now, we compare result of Theorem 1 with known upper bounds in Figure 1. Recall that DP-SOA in Golowich and Livni (2021) obtained an upper bound which increases logarithmically with the time horizon T . Figure 1 shows that for $T \leq \exp(1/(32\delta))$, the dependency on T is necessary, thereby showing the tightness of DP-SOA (Golowich and Livni, 2021). For larger T , the aforementioned *Name & Shame* actually outperforms DP-SOA, indicating that, for fixed ε, δ , *online algorithm always compromises privacy at very large T*. However, even for smaller T , the plot shows that the number of mistake must grow with increasing T until it matches the *Name & Shame* algorithm.

The main idea behind the construction of the lower bound is the following: assume that there is a distinguishing tuple $(f_0, f_1, x_{\text{eq}}, x_{\text{dif}})$, such that \mathcal{A} on $(x_{\text{eq}}, \dots, x_{\text{eq}})$ with high probability only outputs functions that are equal to f_0 at x_{dif} . Then, the adversary picks $f^* = f_1$ and its goal is to select time steps t to insert x_{dif} , such that (i) \mathcal{A} with high probability will make a mistake at t , and (ii) \mathcal{A} will not be able to ‘extract a lot of information’ from this mistake. As we show, both of these conditions can be guaranteed using the DP property and concentration assumption of \mathcal{A} . Note that if adversary inserts x_{dif} serially starting from $t = 1$, \mathcal{A} can have regret much smaller than $\log T$. For example, if \mathcal{A} first incurs a constant number of mistakes, the group privacy property allows a

considerable change in the output distribution to only output $f_1 = f^*$. But this is possible only if \mathcal{A} can predict in advance, where it will make the mistake in the future.

Therefore, the adversary needs to ‘disperse’ the points x_{dif} across T time steps, such that \mathcal{A} cannot predict, where it might make the next mistake. We begin our construction by first inserting a point at the beginning. Then, depending on whether the learner anticipates more points x_{dif} in the first half or in the second half of the sequence, we insert x_{dif} in the half where it least expects it. Note that because of the concentration and DP assumption, the learner cannot anticipate points x_{dif} in both halves simultaneously (recall that on input sequence consisting of only x_{eq} , learner does not output a correct function for x_{dif}). By continuing this construction recursively, we are able to insert $\Omega(\log T)$ points x_{dif} , such that with constant probability, on each of them \mathcal{A} will make a mistake.

Proof [of Theorem 1] Since \mathcal{A} is 0.1-concentrated, there exists a distinguishing tuple $(f_0, f_1, x_{\text{eq}}, x_{\text{dif}})$, such that

$$\mathbb{P}[\forall t \in [T], \mathcal{A}(\tau_0)_t[x_{\text{dif}}] = f_0(x_{\text{dif}})] \geq 0.9, \quad (4.3)$$

where $\tau_0 = ((x_{\text{eq}}, f_0(x_{\text{eq}})), \dots, (x_{\text{eq}}, f_0(x_{\text{eq}})))$. WLOG assume that $f_0(x_{\text{eq}}) = f_0(x_{\text{dif}}) = 0$. When $T > \exp(1/(32\delta))$, by simply only using the first $\exp(1/(32\delta))$ rounds, we obtain the required lower bound $\Omega(1/\delta)$. In the remaining, we assume that $T \leq \exp(1/(32\delta))$. Furthermore, let k be the largest integer, such that $2^k - 1 \leq T$. Note that $\frac{1}{2} \log T \leq k \leq 2 \log T \leq 1/(16\delta)$. For simplicity, we also assume that $T = 2^k - 1$. We start with the case $\varepsilon = \varepsilon_0 = \log(3/2)$ and pick $f^* = f_1$ for the adversary.

Note that to show $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log T/\beta)$, we can construct two adversaries, first achieving $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log 1/\beta)$ (**Case I**) and the other $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log T)$ (**Case II**). We now prove each of them.

Case I: We start with showing the adversary for the former bound, i.e., such that $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log 1/\beta)$. The concentration assumption implies that for any $t \in [T]$, $\mathbb{P}(\mathcal{A}(\tau_0)_t[x_{\text{dif}}] = 1) \leq \beta$. Therefore, if τ_k contains k copies of the point $(x_{\text{dif}}, 1)$, by applying DP property of \mathcal{A} k times, we obtain that for any $t \in [T]$, $\mathbb{P}(\mathcal{A}(\tau_k)_t[x_{\text{dif}}] = 1) \leq \beta \exp(k\varepsilon_0) + \delta \left(\frac{\exp(k\varepsilon_0) - 1}{\exp(\varepsilon_0) - 1} \right) \leq (k\delta + \beta) \exp(k\varepsilon_0)$. If $k = \frac{1}{8\varepsilon_0} \log(1/\beta)$, we derive that for any $t \in [T]$,

$$\mathbb{P}(\mathcal{A}(\tau_k)_t[x_{\text{dif}}] = 1) \leq (k\delta + \beta) \exp(k\varepsilon_0) \leq \frac{\delta \log(1/\delta)}{8\delta^{1/8}\varepsilon_0} \leq \frac{1}{3} \delta^{1/2} \log(1/\delta) \leq \frac{1}{2}. \quad (4.4)$$

This implies that for the sequence τ_k , expected number of mistakes is $\mathbb{E}[M_{\mathcal{A}}] \geq \frac{1}{16\varepsilon_0} \log(1/\beta)$.

Case II: In the previous case, it did not matter where exactly the points $(x_{\text{dif}}, 1)$ are inserted. However, if we want to prove the lower bound $\Omega(\log T)$, this is no longer true and one needs to be careful with the placement of the inserted points.

In the following, we construct a sequence $\mathbf{c} = (x_1, \dots, x_T)$, such that expected number of mistakes of \mathcal{A} will be large. We proceed iteratively, maintaining scalar sequences (l_i) , (r_i) , and sequences $\mathbf{c}^{(i)}$ such that

1. $\mathbf{c}^{(i)}$ contains exactly i points x_{dif} on the prefix $[1, l_i - 1]$,
2. $p_i := \mathbb{P}[\forall t \in [1, l_i - 1], \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 0] \geq 1/2 - 4i\delta$,
3. $q_i := \mathbb{P}[\forall t \in [1, l_i - 1], \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 0 \text{ and } \exists t \in [l_i, r_i], \text{ s.t. } \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 1] \leq 2\delta$.

Assume that we obtain these sequences until $i = k \geq \frac{1}{2} \log T$. Then on the event $\{\forall t \in [1, l_k - 1], \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 0\}$, \mathcal{A} will make k mistakes. Since $k \leq 1/(16\delta)$, from the second property above $p_k \geq 1/4$. This implies that for the sequence $\mathbf{c}^{(k)}$ we have $\mathbb{E}[M_{\mathcal{A}}] \geq p_k k \geq \frac{1}{8} \log T$.

We construct the sequences by induction, starting with $\mathbf{c}^{(0)} = (x_{\text{eq}}, \dots, x_{\text{eq}})$. For each i , $\mathbf{c}^{(i)}$ will differ from $\mathbf{c}^{(i+1)}$ at exactly one point. This allows us to use DP property of \mathcal{A} in order to compare outputs on $\mathbf{c}^{(i)}$ and $\mathbf{c}^{(i+1)}$. From δ -concentrated assumption, we can pick $l_0 = 1, r_0 = T$ which gives $p_0 = 1$ and $q_0 = 0.1$ (we interpret $\mathbb{P}[\forall t \in \emptyset \dots] = 1$).

Given $\mathbf{c}^{(i)}$ we construct $\mathbf{c}^{(i+1)}$ by substituting l_i -th input point with x_{dif} :

1. Let $(\mathbf{c}^{(i+1)})_j = (\mathbf{c}^{(i)})_j$ for all $j \neq l_i$,
2. Set $(\mathbf{c}^{(i+1)})_{l_i} = x_{\text{dif}}$.

Now we compute l_{i+1}, r_{i+1} and bound p_{i+1}, q_{i+1} . To do this, we first introduce

$$\begin{aligned} p'_i &:= \mathbb{P} \left[\forall t \in [1, l_i], \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 0 \right] \geq p_i - q_i, \\ q'_i &:= \mathbb{P} \left[\forall t \in [1, l_i], \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 0 \text{ and } \exists t \in [l_i + 1, r_i], \text{ s.t. } \mathcal{A}(\mathbf{c}^{(i)})_t[x_{\text{dif}}] = 1 \right] \leq q_i, \end{aligned} \quad (4.5)$$

which just account for the shift $l_i \rightarrow l_i + 1$. Next, for $m_i = (l_i + r_i)/2$ and for any $\mathbf{x} \in \{x_{\text{eq}}, x_{\text{dif}}\}^T$, define

$$\begin{aligned} Q(\mathbf{x}) &:= \{ \forall t \in [1, l_i], \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 0 \text{ and } \exists t \in [l_i + 1, r_i], \text{ s.t. } \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 1 \}, \\ Q_1(\mathbf{x}) &:= \{ \forall t \in [1, l_i], \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 0 \text{ and } \exists t \in [l_i + 1, m_i], \text{ s.t. } \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 1 \}, \\ Q_2(\mathbf{x}) &:= \{ \forall t \in [1, m_i], \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 0 \text{ and } \exists t \in [m_i + 1, r_i], \text{ s.t. } \mathcal{A}(\mathbf{x})_t[x_{\text{dif}}] = 1 \}. \end{aligned} \quad (4.6)$$

Clearly, for any \mathbf{x} , $Q(\mathbf{x}) = Q_1(\mathbf{x}) \cup Q_2(\mathbf{x})$ with $Q_1(\mathbf{x}) \cap Q_2(\mathbf{x}) = \emptyset$. Therefore,

$$q'_i = \mathbb{P} \left[Q(\mathbf{c}^{(i)}) \right] = \mathbb{P} \left[Q_1(\mathbf{c}^{(i)}) \right] + \mathbb{P} \left[Q_2(\mathbf{c}^{(i)}) \right]. \quad (4.7)$$

We can use DP property of \mathcal{A} when comparing outputs on $\mathbf{c}^{(i)}$ and $\mathbf{c}^{(i+1)}$ to get

$$\begin{aligned} \min \left(\mathbb{P} \left[Q_1(\mathbf{c}^{(i+1)}) \right], \mathbb{P} \left[Q_2(\mathbf{c}^{(i+1)}) \right] \right) &\leq \frac{1}{2} \left(\mathbb{P} \left[Q_1(\mathbf{c}^{(i+1)}) \right] + \mathbb{P} \left[Q_2(\mathbf{c}^{(i+1)}) \right] \right) \\ &= \frac{1}{2} \mathbb{P} \left[Q(\mathbf{c}^{(i+1)}) \right] \leq \frac{1}{2} \left(\exp(\varepsilon_0) \mathbb{P} \left[Q(\mathbf{c}^{(i)}) \right] + \delta \right) \\ &= \exp(\varepsilon_0) q'_i / 2 + \delta / 2 \leq \frac{3}{4} q_i + \delta / 2. \end{aligned} \quad (4.8)$$

If $\mathbb{P} \left(Q_1(\mathbf{c}^{(i+1)}) \right) \leq \mathbb{P} \left(Q_2(\mathbf{c}^{(i+1)}) \right)$ we set $l_{i+1} := l_i + 1, r_{i+1} := m_i$, which gives $q_{i+1} = \mathbb{P} \left(Q_1(\mathbf{c}^{(i+1)}) \right)$ and $p_{i+1} = p'_i \geq p_i - q_i$. Otherwise, we set $l_{i+1} = m_i + 1, r_{i+1} = r_i$, with $q_{i+1} = \mathbb{P} \left(Q_2(\mathbf{c}^{(i+1)}) \right)$ and $p_{i+1} \geq p'_i - q'_i \geq p_i - 2q_i$. Overall, with our choice of l_{i+1}, r_{i+1} :

$$q_{i+1} \leq \frac{3}{4} q_i + \delta / 2 \quad \text{and} \quad p_{i+1} \geq p_i - 2q_i. \quad (4.9)$$

By geometric sum properties, we have that $q_i \leq \frac{1}{10} (3/4)^i + 2\delta$, and therefore, $p_i \geq 1/2 - 4i\delta$. Finally, note that $r_{i+1} - l_{i+1} = (r_i - l_i)/2 - 1$. In the beginning we have $r_0 - l_0 = 2^k - 2$, therefore

$r_i - l_i = 2^{k-i} - 2$. Thus, we can repeat this process $\frac{1}{2} \log T \leq k \leq 2 \log T \leq 1/(16\delta)$ times. By construction, all inserted points x_{dif} for $c^{(i)}$ lie on the prefix $[1, l_i - 1]$. In order to extend to other values of ε , we either insert more points per step (if $\varepsilon < \varepsilon_0$), or divide the segment into more parts (if $\varepsilon > \varepsilon_0$). We refer to Appendix C for the argument in these cases. This concludes the proof. ■

4.2. Examples of β -concentrated online learners for Point_N

Next, we show that several learners that could be used for online learning Point_N , satisfy Definition 7. In particular, in Lemma 1 we prove that any improper learner that learns Point_N using a finite union of hypothesis in Point_N (defined as Multi-point in Definition 8) is concentrated. Corollary 2 implies that the only existing private online learning algorithm DP-SOA is concentrated.

Definition 8 (Multi-Point class) For $1 \leq K \leq N \in \mathbb{N}_+$, we define Point_N^K to be the class of functions that are equal to 1 on at most K points.

Lemma 1 Let $\beta > 0$ and $K, T \in \mathbb{N}_+$. For any $N \geq 3KT^2/\beta$, any learner \mathcal{A} of Point_N that only uses Point_N^K as its output set is β -concentrated.

Proof Set $x_{\text{eq}} = 1$ and $f_1 = \mathbb{I}\{\cdot = 2\}$. Let $\tau_0 = ((x_{\text{eq}}, f_1(x_{\text{eq}})), \dots, (x_{\text{eq}}, f_1(x_{\text{eq}})))$ and $\mathcal{A}(\tau_0) = (\hat{f}_1, \dots, \hat{f}_T)$. For $t \in [T]$, let $Q_t = \{x \in [N], \text{ such that } \mathbb{P}(\hat{f}_t(x) = 1) \geq \beta/T\}$. We can write

$$\begin{aligned} \sum_{x \in [N]} \mathbb{P}(\hat{f}_t(x) = 1) &= \sum_{x \in [N]} \sum_{\substack{f \in \text{Point}_N^K \\ f(x)=1}} \mathbb{P}(\hat{f}_t = f) \\ &= \sum_{f \in \text{Point}_N^K} \sum_{\substack{x \in [N] \\ f(x)=1}} \mathbb{P}(\hat{f}_t = f) \leq K \sum_{f \in \text{Point}_N^K} \mathbb{P}(\hat{f}_t = f) = K, \end{aligned} \tag{4.10}$$

which implies that $|Q_t| \leq K[T/\beta]$. Therefore, by union bound, as long as $N \geq 3KT^2/\beta$, there exists $x_{\text{dif}} \in ([N] \setminus \bigcup_{t=1}^T Q_t) \setminus \{1, 2\}$. Applying union bound again, we obtain that

$$\mathbb{P}(\exists t, \text{ such that } \hat{f}_t(x_{\text{dif}}) = 1) \leq \sum_{t=1}^T \mathbb{P}(\hat{f}_t(x_{\text{dif}}) = 1) \leq \beta. \tag{4.11}$$

Thus, \mathcal{A} is β -concentrated and the distinguishing tuple is $(f_1, \mathbb{I}\{\cdot = x_{\text{dif}}\}, x_{\text{eq}}, x_{\text{dif}})$. ■

In Lemma 1 we assumed that K is a constant. Note that as long as $K = o(N)$, by taking N large enough one can obtain the same result. We also remark that the SOA for Point_N , which is an improper learner, uses Point_N^1 as the output hypothesis class. [Beimel et al. \(2014\)](#) also uses improper learners consisting of union of points for privately learning Point_N in the pure DP-PAC model; their hypothesis class is a subset of the Multi-Point class.³

While their motivation was to explicitly construct a smaller hypothesis class that *approximates* Point_N well, on a more general note, it is also common to use larger hypothesis classes in computational learning theory for improper learning. This is particularly useful for benefits like computational efficiency (e.g., 3-DNF vs 3-CNF) and robustness ([Diakonikolas et al., 2019](#)). The purpose

3. See Section 4.1 in [Beimel et al. \(2014\)](#) to see the conditions that the hypothesis class needs to satisfy.

of Lemma 1 is to show that natural improper learners using larger hypothesis classes for Point_N , like Definition 8, are also concentrated. Furthermore, as we show next, the only generic private online learner that we are aware of— DP-SOA, is also concentrated.

Corollary 2 DP-SOA for Point_N is β -concentrated for any $N \geq 3T^2/\beta$.

Proof Output of DP-SOA is equal to the output of SOA algorithm on some input sequence. Note that SOA for Point_N is always equal to either (i) an all-zero function, or (ii) a target function f^* , which implies that it lies in Point_N^1 . Therefore, by Lemma 1, we obtain that DP-SOA is also β -concentrated. ■

Independent work In a concurrent and independent work, Cohen et al. (2024) established that learners that *are not* concentrated must nevertheless suffer large expected regret for the specific case of Point_T . They construct an adversary, strongly relying on the large size of the function class. In particular, generalizing for a smaller function classes, e.g. Point_3 remains an interesting open question. We obtain a partial progress in this direction, see Section 5.2.

5. Discussion

5.1. Connection to DP under continual observation

Continual observation under DP (Dwork et al., 2010b) is the process of releasing statistics continuously on a stream of data while preserving DP over the stream. One of simplest problems in this setting is DP-Continual counting where a counter $\mathcal{C} : \{0, 1\}^T \rightarrow \mathbb{N}_+^T$ is used. We say \mathcal{C} is deemed a $(T, \alpha, \beta, \varepsilon)$ -DP continual counter if \mathcal{C} is ε -DP with respect to its input and with probability at least $1 - \beta$ satisfies

$$\max_{t \leq T} |\mathcal{C}(\tau)_t - \sum_{i \leq t} \tau_i| \leq \alpha. \tag{5.1}$$

The proof of Proposition 3 illustrates a straightforward method to convert DP continual counters to DP online learners for the Point_N hypothesis class, thereby transferring upper bounds from DP continual counting to DP online learning. More precisely, the reduction results in a DP online learning algorithm for Point_N with an additional \sqrt{N} factor to the privacy parameter and number of mistakes bounded by α . We believe this argument can also be extended to other finite classes by adjusting the mistake bound with an additional factor that depends on the size of the class.

Proposition 3 For sufficiently small $\varepsilon, \beta \geq 0$ and any $\alpha \geq 0$, let \mathcal{C} be a $(T, \alpha, \beta, \varepsilon)$ -DP continual counter. Then, for any $\delta > 0$, an (ε', δ) -DP online learner \mathcal{A} for Point_N exists ensuring $\mathbb{E}[\mathcal{M}_{\mathcal{A}}] \leq \alpha$ with $\varepsilon' = \varepsilon \sqrt{3N \log(1/\delta)}$.

Several works (see Chan et al. (2011)) have identified counters with $\alpha = O\left((\log T)^{1.5}/\varepsilon\right)$ which immediately implies a mistake bound that also scales as $\text{poly} \log(T)/\varepsilon$ using Proposition 3. Ignoring⁴ the dependence on N , this almost matches the lower bound proposed in Theorem 1. On the other hand, Dwork et al. (2010a) have shown a lower bound of $\Omega\left(\frac{\log T}{\varepsilon}\right)$ for α for any pure DP-continual counter. However it is not clear how to convert the lower bound for DP-continual learning

4. We do not optimise this as this dependence is not the focus of this work, for this argument consider $N = O(1)$. However, we believe it can be reduced using a DP continual algorithm for MAXSUM; see Jain et al. (2023)

to DP online learning. This is because 1) Equation (5.1) asks for a uniformly (across t) accurate counter whereas the mistake bound in Equation (2.1) is a global measure and 2) this lower bound is for pure DP whereas our setting is approximate DP. We leave this question for the future work.

5.2. Beyond concentration assumption

In Section 4.1, we analyzed a class of algorithms, namely concentrated algorithms, containing the only existing DP online learning algorithm, DP-SOA, and provided a matching logarithmic in T lower bound. A natural question is whether this lower bound can be extended to *any* learner. While we do not provide a general answer, we show that our construction can be made more general. Here, we analyze a different class of algorithms, called *firing algorithms*.

For simplicity of exposition, consider the hypothesis class Point_3 , and let $f^* \in \{f^{(1)}, f^{(2)}\}$. Let x^* be such that $f^* = f^{(x^*)}$ and assume that adversary only considers sequences consisting of $(x^*, 1)$ and $(3, 0)$. Let \mathcal{D} be a distribution on $\{\text{all-zero function}, f^{(1)}, f^{(2)}\}$. At each point t , a firing algorithm \mathcal{A} computes $p_t = p_t(\# \text{mistakes up until } t) \in [0, 1]$ with $p_t(0) = 0$, and samples a $\xi_t \sim \text{Bern}(p_t)$. If $\xi_t = 1$, the algorithm commits to the correct hypothesis henceforth; otherwise, it outputs $f_t \sim \mathcal{D}$. Note that if the adversary introduces $(3, 0)$ at step t , \mathcal{A} is guaranteed not to err, and a single mistake suffices to identify (non-privately) the correct hypothesis.

When \mathcal{D} has support on only one of $\{\text{all-zero function}, f^{(1)}, f^{(2)}\}$, it yields a 0-concentrated algorithm. Furthermore, the continual observation algorithm can also be viewed as a firing algorithm with a proper choice of \mathcal{D} . We analyze the opposite to the concentrated algorithms, in particular when $\mathcal{D} = \text{Unif}\{f^{(1)}, f^{(2)}\}$. We call such learners *uniform firing algorithms* and we also obtain a logarithmic in T lower bound for them.

Proposition 4 *Let \mathcal{A} be an (ε, δ) -DP uniform firing algorithm for class Point_3 . Then, if $\log T = O(1/\delta)$, there exists an adversary, such that $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log T)$.*

Proof is provided in Appendix B and is similar to the proof of Theorem 1, but requires a more delicate construction.

5.3. Pure differentially private online learners

While we have so far focused on Approximate DP with $\delta > 0$, in this section we briefly discuss Online learning under pure DP. Note that Lemma 1 and Theorem 1 immediately imply the following lower bound on pure differentially private online learners for Point_N .

Corollary 5 *Let $\varepsilon > 0$, $\beta > 0$, $K, T \in \mathbb{N}_+$ and $N \geq 3KT^2/\beta$. For any ε -DP learner \mathcal{A} which uses only Point_N^K as its output set, there exists an adversary \mathcal{B} , such that*

$$\mathbb{E}[M_{\mathcal{A}}] = \Omega\left(\min\left(\frac{\log T/\beta}{\varepsilon}, T\right)\right). \quad (5.2)$$

Proof Lemma 1 implies that \mathcal{A} must be β -concentrated. Furthermore, since \mathcal{A} is ε -DP, it is also (ε, β) -DP, and, thus, the existence of adversary with large regret follows from Theorem 1. ■

For (ε, δ) -DP online algorithms, there exists an upper bound provided by Golowich and Livni (2021). However, to the best of our knowledge, not much is known about ε -DP online algorithms.

One way to obtain such algorithm is by leveraging existing results from the continual observation literature as done in Proposition 3. Under the same assumptions as in Proposition 3, using basic composition instead of advanced composition in the last step results in an ε' scaling as $N\varepsilon$. However, this only works for Point_N and not for more general classes. For arbitrary finite hypothesis classes, it is possible to use a DP continual counter, similar to above to obtain a mistake bound that also scales with the size of the class. However, the question of whether generic learners can be designed remains open. Another interesting open question is whether the dependence on the size of the class e.g., N for Point_N , is necessary.

5.4. Open problems

This work relies on assumptions about properties (Concentrated Assumption in Definition 7 or uniform firing Assumption in Section 5.2) of the learning algorithms to show a lower bound for any hypothesis class. On the other hand, Cohen et al. (2024) do not need any assumption on their algorithm but their result only holds for large Point classes (and it cannot immediately be transferred to small Point classes, e.g. Point_3). To fully settle the problem of lower bounds for online DP learners, the limitations of both this work and that of Cohen et al. (2024) need to be addressed. We propose the following conjecture which we believe is true.

Conjecture 6 (Lower bound for Approximate DP) *For any $\varepsilon, \delta > 0$ and any (ε, δ) -DP learner \mathcal{A} of Point_3 , there exists an adversary \mathcal{B} , such that $\mathbb{E}[M_{\mathcal{A}}] = \Omega\left(\min\left(\frac{\log T/\delta}{\varepsilon}, 1/\delta\right)\right)$.*

A straightforward implication of our main result is that the $\log T$ lower bound also holds for pure DP online learners under the same assumptions on the algorithm. In Section 5.1, we showed a generic reduction from DP continual counters to DP online learners for Point_N with regret $\text{poly} \log(T)$. This reduction can be extended to pure DP online learner by using basic composition in Proposition 3, with an additional cost of \sqrt{N} . In Conjecture 7, we raise the question whether for small N , the dependence on T can be lowered to $\log T$ and showed to be tight.

Conjecture 7 (Upper and lower bounds for Pure DP) *For any $\varepsilon > 0$,*

1. Upper bound: *there exists ε -DP learner of Point_3 , s.t. for any adversary, $\mathbb{E}[M] = O\left(\frac{\log T}{\varepsilon}\right)$.*
2. Lower bound: *for any ε -DP learner of Point_3 , there exists an adversary, s.t. $\mathbb{E}[M] = \Omega\left(\frac{\log T}{\varepsilon}\right)$.*

Note that the lower bound part of Conjecture 7 follows from Conjecture 6, but can also be viewed through connection to the continual observation model. In the latter regime, an $\Omega(\log T)$ lower bound was shown in Dwork et al. (2014).

Acknowledgments

DD is supported by ETH AI Center doctoral fellowship and ETH Foundations of Data Science initiative. AS and DD thank Yevgeny Seldin for hosting them in University of Copenhagen which helped in the discussion of this work and Amir Yehudayoff for helpful discussions. DD also thanks Afonso Bandeira and Petar Nizic-Nikolac for fruitful discussions during the course of the project.

References

- Noga Alon, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private pac learning implies finite littlestone dimension. In *Symposium on Theory of Computing (STOC)*, 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.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Symposium on foundations of computer science (FOCS)*, 2014.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Characterizing the sample complexity of private learners. In *Conference on Innovations in Theoretical Computer Science (ITCS)*, 2013a.
- Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, 2013b.
- Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine learning*, 2014.
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *Conference On Learning Theory (COLT)*, 2009.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. Learnability and the vapnik-chervonenkis dimension. *Journal of the ACM (JACM)*, 1989.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- T-H Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *Transactions on Information and System Security (TISSEC)*, 2011.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research (JMLR)*, 2011.
- Edith Cohen, Xin Lyu, Jelani Nelson, Tamás Sarlós, and Uri Stemmer. Lower bounds for differential privacy under continual observation and online threshold queries. *arXiv preprint arXiv:2403.00028*, 2024.
- Ilias Diakonikolas, Themis Gouleakis, and Christos Tzamos. Distribution-independent pac learning of halfspaces with massart noise. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- Cynthia Dwork and Vitaly Feldman. Privacy-preserving prediction. In *Conference On Learning Theory (COLT)*, 2018.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography (TCC)*, 2006.

- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Symposium on Theory of computing (STOC)*, 2010a.
- Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Symposium on Foundations of Computer Science (FOCS)*, 2010b.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Vitaly Feldman and David Xiao. Sample complexity bounds on differentially private learning via communication complexity. In *Conference on Learning Theory (COLT)*, 2014.
- Badih Ghazi, Noah Golowich, Ravi Kumar, and Pasin Manurangsi. Sample-efficient proper pac learning with approximate differential privacy. In *Symposium on Theory of Computing (STOC)*, 2021.
- Noah Golowich and Roi Livni. Littlestone classes are privately online learnable. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Steve Hanneke, Roi Livni, and Shay Moran. Online learning with simple predictors and a combinatorial characterization of minimax in 0/1 games. In *Conference on Learning Theory (COLT)*, 2021.
- Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam Smith. The price of differential privacy under continual observation. In *International Conference on Machine Learning (ICML)*, 2023.
- Haim Kaplan, Yishay Mansour, Shay Moran, Kobbi Nissim, and Uri Stemmer. Black-box differential privacy for interactive ml. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 2011.
- Michael Kearns, Malleesh M Pai, Ryan Rogers, Aaron Roth, and Jonathan Ullman. Robust mediators in large games. *arXiv:1512.02698*, 2015.
- Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 1988.
- Moni Naor, Kobbi Nissim, Uri Stemmer, and Chao Yan. Private everlasting prediction. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Amartya Sanyal and Giorgia Ramponi. Open problem: Do you pay for privacy in online learning? In *Conference on Learning Theory (COLT)*, 2022.
- Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 1984.

Appendix A. Proof of Proposition 3

Proof The proof for Proposition 3 simply follows from first instantiating $|\mathcal{H}|$ separate $(T, \alpha, \beta, \varepsilon')$ -DP continual counters, one for each $h \in \mathcal{H}$.

Then we construct the learning algorithm as follows: at each step t , the learner outputs the hypothesis that predicts 0 uniformly. This is possible as the learning algorithm is improper. Then the adversary provides the learner with $(x_t, f^*(x_t))$. The algorithm then privately updates the counter \mathcal{C}_h with $b_h = \mathbb{I}\{f^*(x_t) = h(x_t)\} \wedge f^*(x_t)$. Then, as soon as any counter's value (say \mathcal{C}_{h^*}) surpasses α , only predict using h^* .

Note that this algorithm only incurs a mistake at step t if $b_{h^*} = 1$ at step t . However, every time $b_{h^*} = 1$, the counter for f^* gets updated by one. Thus, after α updates it can be guaranteed with probability $1 - \beta$ that $h^* = f^*$. To get the final expected mistake bound, set $\beta = \alpha$ and this completes the proof.

The privacy proof follows from applying strong composition to the different counters. ■

Appendix B. Lower bound for a non-concentrated algorithm

Consider the problem of learning function class Point_3 , such that $f^* \in \{f^{(1)}, f^{(2)}\}$. Let x^* be such that $f^* = f^{(x^*)}$ and assume that the adversary only considers sequences consisting of $(x^*, 1)$ and $(3, 0)$. In this section, we show a lower bound for one special type of non-concentrated algorithm: *firing algorithms*.

At each round t , the algorithm computes $p_t = p_t(\# \text{mistakes up until } t) \in [0, 1]$ with $p_t(0) = 0$, and samples $\xi_t \sim \text{Bern}(p_t)$. If $\xi_t = 1$, from this point on, the algorithm always outputs the correct hypothesis. In this case, we say that \mathcal{A} has ‘fired’ at step t . Otherwise, it outputs $f_t \sim \text{Unif}\{f^{(1)}, f^{(2)}\}$. We call such learners *uniform firing algorithms*. Note that \mathcal{A} never errs on points $(3, 0)$, since both $f^{(1)}$ and $f^{(2)}$ are equal to zero at $x = 3$. Essentially, we analyze a distinguishing tuple $(f^{(1)}, f^{(2)}, 3, x^*)$.

Proof [of Proposition 4] WLOG, we assume $\delta > 1/T^2$. For the elements of the input sequence, we denote $0 := (3, 0)$ and $1 := (x^*, 1)$. Similar to the proof of Theorem 1, we proceed by creating sequences τ_i , with $\tau_0 = (0, \dots, 0)$ and τ_i containing exactly i copies of 1.

Assume that there is a scalar sequence (r_i) and a sequence of subset of events (G_i) , such that

1. $\mathbb{P}(G_i) \geq 1 - O(i\delta) \geq 1/2$,
2. $\mathbb{P}(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i \mid G_i) = O(\delta)$.

This implies that we can upper bound probability that \mathcal{A} outputs f^* for any $t \leq r_i$ as follows (for a subset of events A , we denote its complement by \bar{A} and $\mathbb{P}_i(A) := \mathbb{P}(A \mid G_i)$):

$$\begin{aligned}
 \mathbb{P}(\mathcal{A}(\tau_i)_t = f^*) &= \mathbb{P}_i(\mathcal{A}(\tau_i)_t = f^* \mid G_i) \mathbb{P}(G_i) + \mathbb{P}_i(\mathcal{A}(\tau_i)_t = f^* \mid \bar{G}_i) \mathbb{P}(\bar{G}_i) \\
 &\leq \mathbb{P}_i(\mathcal{A}(\tau_i)_t = f^*) + O(i\delta) \\
 &= O(i\delta) + \mathbb{P}_i(\mathcal{A}(\tau_i)_t = f^* \mid \mathcal{A} \text{ 'fired' before } r_i) \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ 'fired' before } r_i) \\
 &\quad + \mathbb{P}_i(\mathcal{A}(\tau_i)_t = f^* \mid \mathcal{A}(\tau_i) \text{ did not 'fire' before } r_i) \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ did not 'fire' before } r_i) \\
 &\leq \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ 'fired' before } r_i) + \frac{1}{2} + O(i\delta) \\
 &\leq \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + \frac{1}{2} + O(i\delta) \\
 &= \frac{1}{2} + O(i\delta).
 \end{aligned} \tag{B.1}$$

For $i = O(1/\delta)$, we obtain that $\mathbb{P}(\mathcal{A}(\tau_i)_t \neq f^*) = \Omega(1)$. If τ_i contains exactly i points $(x^*, 1)$ on the prefix $[1, r_i]$, then, by picking $i = \Omega(\log T)$ (recall that $\log T = O(1/\delta)$) and summing the latter probabilities, we obtain that $\mathbb{E}[M_{\mathcal{A}}] = \Omega(\log T)$.

We now construct sequences $(r_i), (G_i), (\tau_i)$, also including a scalar sequence (l_i) , such that $[l_i, r_i]$ are nested segments and $(r_i), (G_i)$ satisfy the aforementioned requirements. We proceed by induction. First, define $l_0 = 1, r_0 = T/2, \tau_0 = (0, \dots, 0)$, and $G_0 = \Omega$, i.e., so that $\mathbb{P}(G_0) = 1$. Note that $\mathbb{P}(\mathcal{A}(\tau_0) \text{ outputs only } f^* \text{ starting from } r_0) = 2^{-T/2} \leq \delta$. Next, assume that we already defined l_i, r_i, G_i . We divide a segment $[l_i, r_i]$ into three parts $[l_i, lm_i], [lm_i, rm_i], [rm_i, r_i]$, with $lm_i = l_i + (r_i - l_i)/3$ and $rm_i = l_i + 2(r_i - l_i)/3$. We have the following decomposition:

$$\begin{aligned}
 &\mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) \\
 &= \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i \text{ and outputs only } f^* \text{ starting from } [lm_i, rm_i]) \\
 &\quad + \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i \text{ and does not only output } f^* \text{ on } [lm_i, rm_i]).
 \end{aligned} \tag{B.2}$$

As long as $rm_i - lm_i \geq \sqrt{T}$, we obtain

$$\begin{aligned}
 &\mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ on } [lm_i, rm_i]) \\
 &= \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ fires before } lm_i) + \sum_{t=1}^{rm_i - lm_i} 2^{-t} \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ fires on } lm_i + t) + 2^{-(lm_i - rm_i)} \\
 &\leq 2\mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + 2^{-\sqrt{T}}.
 \end{aligned} \tag{B.3}$$

We let $\tau_{i+1} = \tau_i$, except that $(\tau_{i+1})_{l_i} = 1$. Then, using DP property we have that

$$\begin{aligned}
 &\mathbb{P}_i(\mathcal{A}(\tau_{i+1}) \text{ outputs only } f^* \text{ starting from } r_i) \\
 &\leq \exp(\varepsilon) \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + \delta / \mathbb{P}(G_i) \\
 &\leq \exp(\varepsilon) \mathbb{P}_i(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + 2\delta.
 \end{aligned} \tag{B.4}$$

Therefore, either

$$\begin{aligned}
 &\mathbb{P}(\mathcal{A}(\tau_{i+1}) \text{ outputs only } f^* \text{ starting from } r_i \text{ and outputs only } f^* \text{ on } [lm_i, rm_i]) \\
 &\leq \frac{1}{2} (\exp(\varepsilon) \mathbb{P}(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + \delta),
 \end{aligned} \tag{B.5}$$

or

$$\begin{aligned} & \mathbb{P}(\mathcal{A}(\tau_{i+1}) \text{ outputs only } f^* \text{ starting from } r_i \text{ and does not only output } f^* \text{ on } [lm_i, rm_i]) \\ & \leq \frac{1}{2}(\exp(\varepsilon) \mathbb{P}(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i) + \delta). \end{aligned} \quad (\text{B.6})$$

In the former case, we can reiterate with $l_{i+1} = l_i + 1$, $r_{i+1} = lm_i$, and $G_{i+1} = G_i$, using that

$$\mathbb{P}(\mathcal{A}(\tau_{i+1}) \text{ outputs only } f^* \text{ starting from } lm_i) \leq \mathbb{P}(\mathcal{A}(\tau_{i+1}) \text{ outputs all ones on } [lm_i, rm_i]). \quad (\text{B.7})$$

In the latter case, we set $G_{i+1} = G_i \cap \{\mathcal{A}(\tau_{i+1}) \text{ does not only output } f^* \text{ on } [lm_i, rm_i]\}$. Note that $\mathbb{P}(G_{i+1}) = \mathbb{P}(G_i) \mathbb{P}(\mathcal{A}(\tau_{i+1}) \text{ does not only output } f^* \text{ on } [lm_i, rm_i]) \geq 1 - O((i+1)\delta)$. We reiterate with $l_{i+1} = rm_i$ and $r_{i+1} = r_i$.

Altogether, as long as $rm_i - lm_i \geq \sqrt{T}$, we can always maintain that

1. $\mathbb{P}(G_{i+1}) \geq 1 - O(i\delta)$,
2. $\mathbb{P}(\mathcal{A}(\tau_i) \text{ outputs only } f^* \text{ starting from } r_i \mid G_i) = O(\delta)$.

We can therefore continue this process at least $\Omega(\log T)$ times, which implies a large number of mistakes as discussed before. \blacksquare

Proof of Proposition 4 resembles the proof of Theorem 1 with several important differences. The main difficulty comes from the fact that by just looking at the output of the algorithm, it is impossible to say with certainty, whether it ‘fired’ or not. And we need to do this, in order to bound the marginal probability of outputting f^* (which is 1/2 if \mathcal{A} did not fire and 1 if it fired), see Equation (B.1).

To overcome this issue, instead of splitting the segment into two parts and reiterating in either left or right, we split into three parts and reiterate in the first or the third. Second segment plays a special role, to ‘detect’ whether algorithm fired or not. From the properties of the algorithm, if the full segment is equal to all f^* , we can confidently say that algorithm fired before this segment. Event G_i corresponds to a ‘good’ event, meaning that \mathcal{A} did not fire and *did not look like it fired*, i.e., in the middle segment there exist one output not equal to f^* .

Finally, we would like to highlight that there is nothing specific about the choice $\mathcal{D} = \text{Unif}(f^{(1)}, f^{(2)})$. The main property we needed for the proof is that

$$\mathbb{P}(\mathcal{A} \text{ outputs all } f^* \text{ on a specific subsequence of large enough length} \mid \mathcal{A} \text{ did not fire}) \leq \delta.$$

Therefore, we expect that similar technique works for other choices of \mathcal{D} , beyond a uniform distribution.

Appendix C. Extending finite lower bound to other values of ε

Smaller ε Let $d \in \mathbb{N}_+$ and $\varepsilon = \varepsilon_0/d$, where $\varepsilon_0 = \log(3/2)$. WLOG we assume that $T = d(2^k - 1)$. We repeat the same construction as in the proof of Theorem 1 for the case ε_0 , but now instead of inserting only one point at each step, we insert d points. Note that this way sequences $\mathbf{c}^{(i+1)}$ and $\mathbf{c}^{(i)}$ are d points away from each other, therefore, by applying DP property of \mathcal{A} d times, we obtain the following equivalent version of Equation (4.8):

$$\min \left(\mathbb{P} \left[Q_1(\mathbf{c}^{(i+1)}) \right], \mathbb{P} \left[Q_2(\mathbf{c}^{(i+1)}) \right] \right) \leq \frac{1}{2} \left(\exp(\varepsilon_0) \mathbb{P} \left[Q(\mathbf{c}^{(i)}) \right] + A\delta \right), \quad (\text{C.1})$$

where $A = 1 + \exp(\varepsilon_0/d) + \exp(2\varepsilon_0/d) + \dots + \exp((d-1)\varepsilon_0/d) \leq d \exp(\varepsilon_0)$. Therefore, multiplying δ by a factor $d \exp(\varepsilon_0)$, we recover the previous setting. Since only $\log 1/\delta$ appears in the final bound, and since ε_0 is a constant, we only suffer an extra $\log d \sim \log 1/\varepsilon$ term, which is subleading.

Larger ε When $\varepsilon > \varepsilon_0$, instead of dividing the sequence in two parts, we will divide it into s parts, for some integer $s \geq 3$, obtaining

$$\min \left(\mathbb{P} \left[Q_1(\mathbf{c}^{(i+1)}) \right], \dots, \mathbb{P} \left[Q_s(\mathbf{c}^{(i+1)}) \right] \right) \leq \frac{1}{s} \left(\exp(\varepsilon) \mathbb{P} \left[Q(\mathbf{c}^{(i)}) \right] + \delta \right). \quad (\text{C.2})$$

We need to ensure that $\exp(\varepsilon)/s \leq 3/4$, in order for the proof to go through, which follows if $\log s = O(\varepsilon)$. Finally, note that instead of repeating process $\Omega(\log T)$ times, we repeat $\Omega\left(\frac{\log T}{\log s}\right) = \Omega\left(\frac{\log T}{\varepsilon}\right)$ times, which matches the required mistake bound.