
TIGHT BOUNDS FOR MACHINE UNLEARNING VIA DIFFERENTIAL PRIVACY

YIYANG HUANG AND CLÉMENT L. CANONNE

University of New South Wales

e-mail address: hi.yiyangh@gmail.com

University of Sydney

e-mail address: clement.canonne@sydney.edu.au

ABSTRACT. We consider the formulation of “machine unlearning” of Sekhari, Acharya, Kamath, and Suresh (NeurIPS 2021), which formalizes the so-called “right to be forgotten” by requiring that a trained model, upon request, should be able to ‘unlearn’ a number of points from the training data, as if they had never been included in the first place. Sekhari et al. established some positive and negative results about the number of data points that can be successfully unlearned by a trained model without impacting the model’s accuracy (the “deletion capacity”), showing that machine unlearning could be achieved by using differentially private (DP) algorithms. However, their results left open a gap between upper and lower bounds on the deletion capacity of these algorithms: our work fully closes this gap, obtaining tight bounds on the deletion capacity achievable by DP-based machine unlearning algorithms.

1. INTRODUCTION

Machine learning models trained on user data are now routinely used virtually everywhere, from recommendation systems to predictive models. In many cases, this user data itself includes some sensitive information (e.g., healthcare or race) or private aspects (customer habits, geographic data), sometimes even protected by law. To address this issue – that the models trained on sensitive datasets must not leak personal or private information – in a principled fashion, one of the leading frameworks is that of *differential privacy* (DP) [DMNS06], which has *de facto* become the standard for privacy-preserving machine learning over the past decade.

Key words and phrases: Differential privacy, Unlearning, Machine unlearning, Optimization.

This work was performed while Y.H. was an undergraduate student at the University of New South Wales, and a part-time Research Assistant at The University of Sydney. C.C. is supported by an ARC DECRA (DE230101329) and an unrestricted gift from Google Research.

At its core, DP requires that the output of a randomized algorithm M not change drastically if one to modify one of the datapoints: that is, if X, X' are two datasets only differing in *one* user’s data, then for all possible outputs S of the algorithm one should have roughly the same probability of observing S under both inputs:

$$\Pr[M(X) \in S] \leq e^\varepsilon \Pr[M(X') \in S] + \delta$$

where $\varepsilon > 0$ and $\delta \in (0, 1]$ quantify the privacy guarantee (the smaller values, the better the privacy; see Section 2 for formal definitions). Intuitively, an algorithm M being (ε, δ) -DP means that its output does not reveal much about any particular user’s data, since the output would be nearly identical had this user’s data been completely different.

While the use of differential privacy can mitigate many privacy concerns, it does come with some limitations. The first is the overhead it brings: that is, ensuring differential privacy for a learning task typically incurs an overhead in the number of data points needed to achieve the same accuracy guarantee. Perhaps more importantly, DP does not address all possible privacy threats: even if a ML model is trained on a sensitive dataset in a differentially private way, the dataset may still be subject to some attacks – e.g., if the server where the training data is stored is itself compromised. Somewhat tautologically: DP is not a silver bullet, and only provides meaningful guarantees against the threat models it was meant to address.

Another type of concern focuses on the individual *right to maintain control on one’s own data*: broadly speaking, this is asking that each user can (under some reasonable circumstances) require that their personal data and information be removed from a company’s collected data and trained models. This so-called “right to be forgotten,” which allows people to request that their data be deleted entirely from an ML system, has been passed into legislation or is considered in some form or another by various countries or entities, prominently the European Union’s General Data Protection Regulation (GDPR), the California Privacy Rights Act (CCRA), Canada’s proposed Consumer Privacy Protection Act (CPPA), and most recently in Australia [Kar23].

However, translating this “right to be forgotten” into practice comes with a host of challenges, starting with the work of [CY15], which initiated a new area of research in ML and computer science, that of *machine unlearning*. Following this original paper, [GGV20] then provided a formal definitional framework for machine unlearning, drawing upon cryptographic concepts. A naive technical solution would be for a given company to keep the original training set at all times, and, upon a deletion request by a user, remove this user’s data from the set before retraining the whole model on the result. This, of course, comes up with two major drawbacks: first, the cost to the company, in terms of time and computational resources, of retraining a large model on a regular basis. Second, the *privacy cost*, as keeping the training set for an indefinite time in order to be able to handle the deletion requests leaves the door open to potential attacks and data breaches. Fortunately, there have been, over the past few years, a flurry of better (and more involved) approaches to machine unlearning, to handle deletion requests much more efficiently, and requiring to maintain much less of the training set (see, e.g., [BCC⁺21, NHR⁺25], and related work below).

The above discussion, still, brings to light an important question: *is machine unlearning, paradoxically, at odds with (differential) privacy? What is the connection between the two notions: are they complementary, or is there a trade-off between them?*

This is the main question this work sets out to address. Our starting point is the probabilistic definition of machine unlearning set forth by Sekhari, Acharya, Kamath, and Suresh [SAKS21], itself reminiscent of the definition of Differential Privacy (see Definition 2.5 for the formal statement): a pair of algorithms (A, \bar{A}) is an (ε, δ) -*unlearning algorithm* if (1) $A: \mathcal{X}^* \rightarrow \mathcal{W}$ is a (randomized) learning algorithm which, given a dataset $X \subseteq \mathcal{X}^*$, outputs model parameters $A(X) \in \mathcal{W}$; and (2) $\bar{A}: \mathcal{X}^* \times \mathcal{W} \times \mathcal{T} \rightarrow \mathcal{W}$ which, on input a set of *deletion requests* $U \subseteq X$, previous model parameters w , and some succinct additional “side information” $T(X) \in \mathcal{T}$ about the original dataset, output updated model parameters $w' \in \mathcal{W}$ from which the data from U has been unlearned, that is, such that

$$\Pr[\bar{A}(U, A(X), T(X)) \in W] \leq e^\varepsilon \Pr[\bar{A}(\emptyset, A(X \setminus U), T(X \setminus U)) \in W] + \delta$$

and

$$\Pr[\bar{A}(\emptyset, A(X \setminus U), T(X \setminus U)) \in W] \leq e^\varepsilon \Pr[\bar{A}(U, A(X), T(X)) \in W] + \delta$$

for every possible set $W \subseteq \mathcal{W}$ of model parameters. Loosely speaking, this requires that the outcomes of (a) training a model M via A on the dataset X then unlearning some of the original training data $U \subseteq X$ from M using \bar{A} , and (b) training a model M' via A directly on the dataset $X \setminus U$ then unlearning nothing via \bar{A} , be nearly indistinguishable.

In their paper, Sekhari et al. [SAKS21] focus on generalization guarantees of unlearning algorithm, i.e., what can be achieved by unlearning algorithms when focusing on population loss, namely, when aiming to minimize

$$F(w) := \mathbb{E}_{x \sim \mathcal{D}}[f(w, x)]$$

given a prespecified loss function $f: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$, where the expectation is over the draw of a new datapoint from the underlying distribution p on the sample space. The quality of a learning algorithm A is then measured by the expected excess risk

$$R(f, A) := \mathbb{E} \left[F(A(X)) - \inf_{w \in \mathcal{W}} F(w) \right]$$

where the expectation is taking over the random choice of a dataset $X \sim \mathcal{D}^n$ of size n , and the randomness of A itself. The focus of [SAKS21], as is ours, is then to quantify the *deletion capacity* achievable for (ε, δ) -unlearning given a prespecified loss function, that is, the maximum number of data points one can ask to be forgotten (maximum size of the subset U) before the excess risk increases by more than some threshold (see Definition 2.6).

In their paper, [SAKS21] draw a connection between DP learning algorithms and unlearning ones, showing that DP learning algorithms do provide *some* unlearning guarantees out-of-the-box, and that one can achieve non-trivial unlearning guarantees for convex loss functions by leveraging the literature on differentially private optimization and learning. A key observation they make is that any unlearning algorithm with deletion capacity m of the form $\bar{A}(U, A(S), T(S)) = \bar{A}(A(S))$ is, in fact, a differentially private algorithm *with the neighboring relation being taken with respect to edit distance at most m* . Due to this correspondence, throughout we will refer to as “DP-based” this restricted class of unlearning algorithms of the form $\bar{A}(U, A(S), T(S)) = \bar{A}(A(S))$.

One of their main results is showing that these DP-based unlearning algorithms, which crucially *do not rely on any side information* (the additional input $T(X) \in \mathcal{T}$ provided to the unlearning algorithm \bar{A}), *nor on the set of requests itself*, can handle strictly fewer deletion requests than general unlearning algorithms which *do* rely on such side information.

Their results, however, do not fully characterize the deletion capacity of these “DP-based” machine unlearning algorithms, leaving a significant gap between their upper and lower bounds. We argue that fully understanding this quantity is crucial, as DP-based unlearning algorithms are *exactly* those for which there is no conflict between the two notions of DP and unlearning – *instead, this class of algorithms is the one for which they work hand in hand*. This is in contrast to the more general unlearning algorithms relying on maintaining and storing side information about the training set, as this side information can make their deployment susceptible to privacy breaches.

1.1. Our contributions. The main contribution of our paper is a tight bound on the “amount of unlearning” achievable by *any* such “DP-based” machine unlearning algorithm. For the sake of exposition, we state in this section informal versions of our results.

Theorem 1.1 (Unlearning For Convex Loss Functions (Informal; see Theorems 3.1 and 3.3)). Let $f: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$ be a 1-Lipschitz convex loss function, where $\mathcal{W} \subseteq \mathbb{R}^d$ is the parameter space. There exists an (ε, δ) -machine unlearning algorithm which, trained on a dataset $S \subseteq \mathcal{X}^n$, does not store *any* side information about the training set besides the learned model, and does not rely on the set of deletion requests, yet can unlearn up to

$$m = O\left(\min\left(\frac{n\varepsilon\alpha}{\sqrt{d}\log(1/\delta)}, \frac{n\sqrt{\varepsilon}\alpha}{\sqrt{d}}\right)\right)$$

datapoints without incurring excess population risk greater than α . Moreover, this is tight: there exists a 1-Lipschitz *linear* loss function such that no DP-based-machine unlearning algorithm can unlearn $\Omega\left(\frac{n\varepsilon\alpha}{\sqrt{d}\log(1/\delta)}\right)$ data points without excess population risk α .

Our techniques also allow us to easily establish the analogue for *strongly* convex optimization:

Theorem 1.2 (Unlearning For Strongly Convex Loss Functions (Informal)). Let $f: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$ be a 1-Lipschitz *strongly* convex loss function. There exists an (ε, δ) -machine unlearning algorithm which, trained on a dataset $S \subseteq \mathcal{X}^n$, does not store *any* side information about the training set besides the learned model, and does not rely on the set of deletion requests, yet can unlearn up to

$$m = O\left(\min\left(\frac{n\varepsilon\sqrt{\alpha}}{\sqrt{d}\log(1/\delta)}, \frac{n\sqrt{\varepsilon}\alpha}{d}\right)\right)$$

datapoints without incurring excess population risk greater than α . Moreover, this is nearly tight: there exists a 1-Lipschitz strongly convex loss function such that no such DP-based machine unlearning algorithm can unlearn $\Omega\left(\frac{n\varepsilon\sqrt{\alpha}}{\sqrt{d}}\right)$ data points without excess population risk α .

We note that, prior to our work, only bounds for the convex loss function case were known, with an upper bound of $m = \tilde{O}(n\varepsilon\alpha/\sqrt{d}\log(e\varepsilon/\delta))$ (loose by polylogarithmic factors for $\varepsilon = O(1)$, as well as an $1/\sqrt{\varepsilon}$ factor for $\varepsilon \gg 1$) and a limited lower bound stating that $m \geq 1$ is only possible if $n\varepsilon/\sqrt{d} = \Omega(1)$. We point out that while the high-privacy ($\varepsilon \ll 1$) is sought as desirable, it is common in real-life, deployed systems to use much larger values of

ε , typically $\varepsilon \gg 1$. We also provide analogous bounds for the case of *pure* unlearning, i.e., $\delta = 0$ (Theorems 3.10 and 3.11).

Our next contribution, motivated by the similarity of the formalisation of machine unlearning (without side information) and that of differential privacy, is to establish the analogue of key properties of DP for machine unlearning, namely, *post-processing* and *composition* of machine unlearning algorithms. To do so, we first identify a natural property of machine unlearning algorithms, which, when satisfied, will allow for the composition properties:

Assumption 1.3 (Unlearning Laziness). *An unlearning algorithm (\bar{A}, A) is said to be lazy if, when provided with an empty set of deletion requests, the unlearning algorithm \bar{A} does not update the model. That is, $\bar{A}(\emptyset, A(X), T(X)) = A(X)$ for all datasets X .*

We again emphasize that this laziness property is not only intuitive, it is also satisfied by several existing unlearning algorithms, and in particular those proposed in [SAKS21]. (For instance, one can check by inspection that when the input delete requests set U is empty, [SAKS21, Algorithm 1] just adds zero-variance, zero-mean noise to its input $\hat{w} = A(S)$, ignoring the side information $T(S)$ entirely: or, in other terms, it does not perform any computation and returns $A(S)$ unchanged.)

Theorem 1.4 (Post-processing of unlearning). Let (\bar{A}, A) be an (ε, δ) -unlearning algorithm taking no side information. Let $f: \mathcal{W} \rightarrow \mathcal{W}$ be an arbitrary (possibly randomized) function. Then $(f \circ \bar{A}, A)$ is also an (ε, δ) -unlearning algorithm.

Under our laziness assumption, we also establish the following:

Theorem 1.5 (Chaining of unlearning). Let (\bar{A}, A) be a *lazy* (ε, δ) -unlearning algorithm taking no side information, and able to handle up to m deletion requests. Then, the algorithm which uses (\bar{A}, A) to sequentially unlearn k disjoint deletion requests $U_1, \dots, U_k \subseteq X$ such that $|\cup_i U_i| \leq m$, outputting

$$\bar{A}(U_k, \dots, \bar{A}(U_1, A(X)) \dots)$$

is an (ε', δ') -unlearning algorithm, with $\varepsilon' = k\varepsilon$ and $\delta' = \delta \cdot \frac{e^{k\varepsilon} - 1}{e^\varepsilon - 1} = O(k\delta)$ (for $k = O(1/\varepsilon)$).

and, finally,

Theorem 1.6 (Advanced grouposition of unlearning). Let $(\bar{A}_1, A), \dots, (\bar{A}_k, A)$ be *lazy* (ε, δ) -unlearning (with common learning algorithm A) taking no side information, and define the grouposition of those unlearning algorithms, \tilde{A} as

$$\tilde{A}(U, A(X)) = f(\bar{A}_1(U, A(X)), \dots, \bar{A}_k(U, A(X))).$$

where $f: \mathcal{W}^k \rightarrow \mathcal{W}$ is any (possibly randomized) function. Then, for every $\delta' > 0$, (\tilde{A}, A) is an (ε', δ') -unlearning taking no side information, where $\varepsilon' = \frac{1}{2}k\varepsilon^2 + \varepsilon\sqrt{2k \ln(1/\delta')}$.

1.2. Related work. Albeit recent, the field of machine unlearning has already received considerable attention from the ML community since introduced in [CY15]. There emerged an array of studies and papers focusing on practical solutions and their empirical performance. We focus in this section on the works most closely related to ours, mostly theoretical.

Literature in machine unlearning related to differential privacy concentrates in two branches: (1) adversarial: models are prone to attacks when attackers have access to both before and after version when the deletion requests are processed, and (2) conceptual: the conceptual similarity of machine unlearning and differential privacy.

The original, stringent definition of unlearning requires $\varepsilon = 0$ (full deletion of the user’s data, as if it had never been included in the training set in the first place) in contrast to differential privacy that allows $\varepsilon > 0$, leaving a possibility for “memorization.” To relax this definition, [GGVZ19] proposed the probabilistic version of unlearning. The formalization of the machine unlearning problem was proposed by [GGV20], adapted from [GGVZ19]’s definition, as a *deletion compliance framework*. In their work, they claimed the free unlearning effects of differential privacy. However, the main difference with our work is that [GGV20] relied on the statistical distance in quantifying error (privacy) whereas we adapt the usual (ε, δ) -DP notion. In addition, their work does not investigate the actual learning performance and quantifies the size of samples (deletion capacity) that we can unlearn.

After the formalization of the unlearning problem, a major line of work that involves both unlearning with differential privacy emerged to investigate adversarial scenarios. For instance, [SBB⁺20] discussed differential privacy as a defense mechanism to combat inference and reconstruction attacks, similarly for [CZW⁺21] and [GGMV22].

There has also been considerable past theoretical work on machine unlearning and DP that are related to our study, yet most works center upon *empirical risk minimization* of unlearning algorithms (e.g. [GGHvdM20, ISCZ20], [NRS21], [UMR⁺21] etc.), which seeks to find an approximate minimizer on the remaining dataset after deletion. Closest to our work is the recent paper of [SAKS21], which formulated the notion of machine unlearning used in our paper and focused on *population loss minimization* of approximating unlearning algorithm (i.e., allowing $\varepsilon > 0$). Their objectives, however, were somewhat orthogonal to ours, as they focused for a large part on minimizing the space requirements for the side information $T(X)$ provided to the unlearning algorithm (while our paper focuses on algorithms which do *not* rely on any such side information, prone to potential privacy leaks). While their work, to motivate this focus, established partial bounds on the deletion capacity of unlearning algorithm that do not take in extra statistics, these bounds were not tight, and one of our main contributions is closing this gap.

Following [SAKS21], the notion of *online* unlearning algorithm – which receive the deletion requests sequentially – was put forward and studied in [SW22], again with memory efficiency with respect to the side information in mind; however, their primary focus is on the empirical performance of unlearning algorithm. In a different direction, [GJN⁺21], drawing upon techniques from differential privacy, provides a generic transformation showing data deletion guarantees against *oblivious* requests can be converted into analogous guarantees for the more challenging setting of *adaptively chosen* sequences of data deletion requests.

Another work related closely to ours is the notion of *certified data removal* proposed by [GGHvdM20]. The main difference between (ε, δ) -certified removal and the definition from

[SAKS21] is that, in the former, the unlearning mechanism requires access not only to the samples to be deleted (the set $U \subseteq X$), but also to the full original training set X : this is exactly the type of constraints our work seeks to avoid, due to the risk of data breach this entails.

Finally, the work of [CSSV23] provides a different formalization of the right to be forgotten, through a new notion the authors introduce, that of *deletion-as-control*. They then systematically explore the connections between this notion (incomparable to that used in our work) and variants of the standard threat model for differential privacy; specifically, pan-privacy, and privacy under continual observation.

1.3. Organization of the paper. We first provide in Section 2 necessary background and notion for differential privacy, learning, and the formulation of machine unlearning used throughout the paper. We then provide a detailed outline of the proof of our main result, Theorem 1.1, in Section 3, before providing the post-processing and composition theorems (Theorems 1.4, 1.5, and 1.6) in Section 4. We conclude with a discussion.

2. PRELIMINARIES

In this section, we recall some notions and results that we will extensively rely on in our proofs and theorems, starting with differential privacy.

2.1. Differential Privacy.

Definition 2.1 ((Central) Differential Privacy (DP)). *Fix $\varepsilon > 0$ and $\delta \in [0, 1]$. An algorithm $M: \mathcal{X}^n \rightarrow \mathcal{Y}$ satisfies (ε, δ) -differential privacy (DP) if for every pair of neighboring datasets X, X' , and every (measurable) subset $S \subseteq \mathcal{Y}$:*

$$\Pr[M(X) \in S] \leq e^\varepsilon \Pr[M(X') \in S] + \delta.$$

We further say that M satisfies pure differential privacy (ε -DP) if $\delta = 0$, otherwise it is approximate differential privacy.

We now recall another notion of differential privacy in terms of Renyi Divergence, from [BS16].

Definition 2.2 (Zero-Concentrated Differential Privacy (zCDP)). *A randomized algorithm $M: \mathcal{X}^n \rightarrow \mathcal{Y}$ satisfies ρ -zCDP if for every neighboring datasets (differing on a single entry) $X, X' \in \mathcal{X}^n$, and for all $a \in (1, \infty)$:*

$$D_a(M(X) \| M(X')) \leq \rho a$$

where $D_a(M(X) \| M(X'))$ is the a -Rényi divergence between the distributions of $M(X)$ and $M(X')$.

We will rely in our proofs on the following group privacy property of zCDP:

Proposition 2.3 (k -distance group privacy of ρ -zCDP [BS16, Proposition 1.9]). *Let $M: \mathcal{X}^n \rightarrow \mathcal{Y}$ satisfy ρ -zCDP. Then, M is $(k^2\rho)$ -zCDP for every $X, X' \in \mathcal{X}^n$ that differs in at most k entries.*

2.2. Learning. We also will require some definitions on learning, specifically with respect to minimizing population loss. Fix any loss function $f: \mathcal{W} \times \mathcal{X}$, where \mathcal{W} is the (model) parameter space and \mathcal{X} is the sample space. Then, the *population loss* is defined as

$$F(w) := \mathbb{E}_{x \sim D}[f(w, x)]$$

in which the expectation is over the distribution of x (one sample) and w is the learning output.

Then, let $F^* = \inf_{w \in \mathcal{W}} F(w)$ be the minimized population risk and w^* be the corresponding minimizer. Consider a learning algorithm $A: \mathcal{X}^n \rightarrow \mathcal{W}$ that takes in dataset $S \in \mathcal{X}^n$ and returns hypothesis $w := A(S) \in \mathcal{W}$. Then we can define *excess risk* as:

$$R(f, A) := \mathbb{E}[F(A(S)) - F^*]$$

where the expectation is over the randomness of A and S .

Finally, we define the sample complexity of learning as follows (see, e.g., [SAKS21, Definition 1]).

Definition 2.4 (Sample complexity of learning). *The α -sample complexity of a problem is defined as:*

$$n(\alpha) := \min\{n \mid \exists A \text{ s.t. } \mathbb{E}[F(A(S))] - F^* \leq \alpha, \forall \mathcal{D}\}$$

2.3. Unlearning. As previously discussed, we rely on the definition of unlearning proposed in by [SAKS21], and maintain same notation. Note that $T(S)$ denotes the data statistics (which could be the entire dataset S or any form of statistic) available to \bar{A} .

Definition 2.5 ((ε, δ) -unlearning). *For all S of size n and delete requests $U \subseteq S$ such that $|U| \leq m$, and $W \subseteq \mathcal{W}$, a learning algorithm A and an unlearning algorithm \bar{A} is (ε, δ) -unlearning if:*

$$\Pr[\bar{A}(U, A(S), T(S)) \in W] \leq e^\varepsilon \Pr[\bar{A}(\emptyset, A(S \setminus U), T(S \setminus U)) \in W] + \delta$$

and

$$\Pr[\bar{A}(\emptyset, A(S \setminus U), T(S \setminus U)) \in W] \leq e^\varepsilon \Pr[\bar{A}(U, A(S), T(S)) \in W] + \delta,$$

Our results will be phrased in terms of the deletion capacity, which captures the number of deletion requests an unlearning algorithm can handle before seeing a noticeable drop in its output accuracy:

Definition 2.6 (Deletion Capacity). *Let $\varepsilon, \delta > 0$, S be a dataset of size n drawn i.i.d. from \mathcal{D} and let $\ell(w, z)$ be a loss function. For a pair of learning and unlearning algorithm A, \bar{A} that are (ε, δ) -unlearning, the deletion capacity $m_{\varepsilon, \delta}^{A, \bar{A}}$ is defined as the maximum size of deletions requests set $|U|$ that we can unlearn without doing worse in excess population risk than α :*

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) := \max\{m \mid \mathbb{E} \left[\max_{U \subseteq S: |U| \leq m} F(\bar{A}(U, A(S), T(S))) - F^* \right] \leq \alpha\}$$

where $F^* = \min_{w \in \mathcal{W}} F(w)$.

3. MAIN RESULT

In this section, we provide a detailed outline and prove of our main result on unlearning for convex loss functions, Theorem 1.1, for which we prove the upper and lower bounds separately.

Theorem 3.1 (Deletion capacity from unlearning via DP, Lower Bound). Suppose $\mathcal{W} \subseteq \mathbb{R}^d$, and fix any Lipschitz convex loss function. Then there exists a lazy (ε, δ) -unlearning algorithm (\bar{A}, A) , where \bar{A} has the form $\bar{A}(U, A(S), T(S)) := A(S)$ (and thus, in particular, is DP-based) with deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) = \Omega\left(\min\left(\frac{n\varepsilon\alpha}{\sqrt{d \log(1/\delta)}}, \frac{n\sqrt{\varepsilon}\alpha}{\sqrt{d}}\right)\right)$$

where the constant in the $\Omega(\cdot)$ only depends on the properties of the loss function.

Proof of Theorem 3.1. The proof follows the same outline as in [SAKS21]. The main change is that instead of the standard DP grouposition (after using the learning guarantees provided by [BFTT19, Theorem 3.2]), we rely on a more recent learning algorithm providing zCDP guarantees, which allows us to leverage the better grouposition properties of concentrated differential privacy:

Lemma 3.2 (zCDP mini-batch noisy SGD [FKT20]). *Fix any L -Lipschitz convex loss function over a convex subset \mathcal{B} of \mathbb{R}^d of diameter D . Then there exists an algorithm A which satisfies $(\rho^2/2)$ -zCDP with excess population loss:*

$$\mathbb{E}\left[F(\theta) - \min_{\theta \in \mathcal{B}} F(\theta)\right] = O\left(DL \cdot \left(\frac{1}{\sqrt{n}} + \frac{\sqrt{d}}{\rho n}\right)\right)$$

where the expectation is taken over the randomness of A .

In more detail, we first establish a tighter bound for algorithm that achieves m -entries group privacy via Lemma 3.2. [FKT20] provides a zCDP version of [BFTT19, Theorem 3.2] with $\rho^2/2$ -zCDP; by group privacy, we get $\frac{m^2\rho^2}{2}$ -zCDP by Proposition 2.3 for neighboring datasets differing in m entries. Then, translating $\frac{1}{2}m^2\rho^2$ -zCDP to (ε, δ) -DP yields

$$\varepsilon = O\left(m^2\rho^2 + m\rho\sqrt{\log(1/\delta)}\right)$$

(see, e.g., [BS16, Proposition 3]). By the above discussion, using this zCDP-private learning algorithm with

$$\rho := \Theta\left(\min\left(\frac{\varepsilon}{m\sqrt{\log(1/\delta)}}, \frac{\sqrt{\varepsilon}}{m}\right)\right)$$

we get an excess population loss bounded by

$$O\left(DL\left(\frac{1}{\sqrt{n}} + \frac{m\sqrt{d \log(1/\delta)}}{\varepsilon n} + \frac{m\sqrt{d}}{\sqrt{\varepsilon} n}\right)\right) \quad (3.1)$$

It only remains to show how the claimed deletion capacity bound follows from this excess population risk guarantee. Construct, as discussed earlier, an unlearning algorithm \bar{A} that returns the input without making any changes (and in particular does not require any

additional statistics $T(S)$, and satisfies the laziness assumption). Since A is (ε, δ) -DP, for any set $U \subseteq S, |U| = m$, and $W \subseteq \mathcal{W}$,

$$\begin{aligned}\Pr[A(S) \in W] &\leq e^\varepsilon \Pr[A(S') \in W] + \delta \\ \Pr[A(S') \in W] &\leq e^\varepsilon \Pr[A(S) \in W] + \delta\end{aligned}$$

But since $\bar{A}(U, A(S)) = A(S)$, this readily yields, letting $S' := S \setminus U$:

$$\begin{aligned}\Pr[\bar{A}(U, A(S)) \in W] &\leq e^\varepsilon \Pr[\bar{A}(\emptyset, A(S')) \in W] + \delta \\ \Pr[\bar{A}(\emptyset, A(S')) \in W] &\leq e^\varepsilon \Pr[\bar{A}(U, A(S)) \in W] + \delta\end{aligned}$$

which implies that (A, \bar{A}) is indeed (ε, δ) -unlearning for U of size (up to) m .

Recalling the definition of deletion capacity, we finally deduce from (3.1) the deletion capacity with excess population risk less than α :

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) \geq m = \Omega\left(\min\left(\frac{\varepsilon n \alpha}{\sqrt{d \log(1/\delta)}}, \frac{\sqrt{\varepsilon n \alpha}}{\sqrt{d}}\right)\right)$$

where the $O(\cdot)$ hides constant factors depending only on the loss function (namely, the Lipschitz constant L , and the diameter D). \square

Theorem 3.3 (Deletion capacity from unlearning via DP, Upper Bound). There exists a Lipschitz convex loss function (indeed, *linear*) for which any DP-based (ε, δ) -unlearning algorithm (\bar{A}, A) must have deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) = O\left(\frac{\varepsilon n \alpha}{\sqrt{d \log(1/\delta)}}\right).$$

Proof of Theorem 3.3. The proof relies on the key observation (already made in [SAKS21, Section 3.1]) that the class of (ε, δ) -machine unlearning algorithms we refer to as DP-based algorithms which can tolerate up to m deletions are, in fact, (ε, δ) -differentially private algorithms *with the neighboring relation being taken with respect to edit distance at most m* . This simple observation allows us to lift DP lower bounds for *this more stringent neighboring relation* to the machine unlearning setting.

We will consider the following linear (and therefore convex and Lipschitz) loss function $\mathcal{L}(\theta, S)$:

$$\mathcal{L}(\theta, S) := -\langle \theta, \sum_{i=1}^n x_i \rangle \tag{3.2}$$

for dataset S of n points $x_1, \dots, x_n \in \{-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d$. We also define the 1-way marginal query, i.e. average, as:

$$q(S) := \frac{1}{n} \sum_{i=1}^n x_i. \tag{3.3}$$

To establish our deletion capacity lower bound with respect to this loss function, we will proceed in three stages: the first is to relate population loss (what we are interested in) to *empirical* loss – which allows us to focus on the existence of a “hard dataset.” The second step is then to establish a sample complexity lower bound on the empirical risk for this loss function of any (ε, δ) -DP algorithm, via a reduction to differentially private computing of 1-marginals. This step is similar to the one underlying the weaker lower bound of [SAKS21]

(itself relying on an argument of [BFTT19]), although a more careful choice of building blocks for the reduction already allows us to obtain an improvement by logarithmic factors.

Third, lift this DP lower bound to a stronger lower bound for DP with respect to edit distance m . This step, while elementary, is novel, as it morally corresponds to establishing the converse of the grouposition property of differential privacy for our specific setting, a converse which does *not* hold in general. Our argument, relatively simple, will crucially rely on the linearity of our loss function.

The first step, which relates population to empirical loss, is required for the use of the specific lower bound we will rely on, from [SU16]. This lower bound relies on a reduction to “good” fingerprinting codes, using a construction of such codes from [SU15]. Importantly, the fingerprinting codes in question (which correspond, in the reduction, to the dataset $S \in \{\pm 1\}^{n \times d}$) have i.i.d. rows, and so the lower bound does indeed lead to a statistical lower bound over a randomly sampling dataset.

For the second step, our starting point is the following lower bound of Steinke and Ullman:

Theorem 3.4 (Lower bound for one-way marginals [SU16, Main Theorem]). For every $\varepsilon \in (0, 1]$ and $2^{-o(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$, and for every $\alpha \leq 1/10$, if $A: \{-1, +1\}^{n \times d} \rightarrow [-1, +1]^d$ is (ε, δ) -differentially private and $\mathbb{E}_{\mathcal{A}}[\|\mathcal{A}(S) - q(S)\|_1] \leq \alpha d$ (i.e., with average-case accuracy α) for all $S \in \{\pm 1\}^{n \times d}$, then we must have

$$n = \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{\varepsilon \alpha}\right).$$

Using this lower bound as a blackbox, we then can adapt the argument of [BST14, Lemma 5.1, Part 2] to obtain the following slightly stronger statement:

Lemma 3.5 (Lower bound for Privately Computing 1-way Marginals). *Let $n, d \in \mathbb{N}$, $\varepsilon \in (0, 1]$, and $2^{-o(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$. For all $\alpha \leq 1/10$, if \mathcal{A} is (ε, δ) -differentially private. Then, for $S \subseteq \{\pm \frac{1}{\sqrt{d}}\}^{n \times d}$, one must have*

$$\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] = \min\left(\alpha, \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{n \varepsilon}\right)\right),$$

where $q(S) = \frac{1}{n} \sum_{i=1}^n x_i$ as before. Moreover, this still holds under the assumption that $\|q(S)\|_2 \in [\frac{M-1}{n}, \frac{M+1}{n}]$, where $M = \Omega(\min(n\alpha, \frac{\sqrt{d \log(1/\delta)}}{\varepsilon}))$.

Proof of Lemma 3.5. Our proof follows the same outline as in [BST14], but using the result of Theorem 3.4 as a black-box instead of the packing argument of [BST14]. Before doing so, we have to translate the result from Theorem 3.4 into our setting, and handle the slightly different choice of parameterization ($\{\pm 1\}^d$ instead of $\{\pm 1/\sqrt{d}\}^d$).

Let $n_\alpha := C \cdot \frac{\sqrt{d \ln(1/\delta)}}{\varepsilon \alpha}$, where $C > 0$ is strictly smaller than the constant hidden in the $\Omega(\cdot)$ of Theorem 3.4. By contradiction, suppose that, for some $n \leq n_\alpha$, we have an (ε, δ) -differentially private algorithm \mathcal{A} that takes in a dataset $S \subseteq \{\pm \frac{1}{\sqrt{d}}\}^{n \times d}$ and outputs an estimate $\mathcal{A}(S)$ of $q(S)$ with expected L_2 error α . Rescaling, we get that the algorithm \mathcal{A}'

which, on input $S' \subseteq \{\pm 1\}^{n \times d}$, computes $S := S'/\sqrt{d} \subseteq \{\pm \frac{1}{\sqrt{d}}\}^{n \times d}$ and outputs $\sqrt{d} \cdot \mathcal{A}(S)$ (1) is (ε, δ) -DP by post-processing, and (2) since q is linear, has error related to that of \mathcal{A} by

$$\mathbb{E}[\|\mathcal{A}'(S') - q(S')\|_2] = \sqrt{d} \cdot \mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] \leq \sqrt{d} \cdot \alpha. \quad (3.4)$$

However, by Theorem 3.4, \mathcal{A}' must have expected L_1 error at least αd since $n \leq n_\alpha$. By Cauchy–Schwarz,

$$\alpha d < \mathbb{E}[\|\mathcal{A}'(S') - q(S')\|_1] \stackrel{\text{CS}}{\leq} \sqrt{d} \cdot \mathbb{E}[\|\mathcal{A}'(S') - q(S')\|_2] \stackrel{(3.4)}{\leq} \sqrt{d} \cdot (\alpha \sqrt{d}) = \alpha d$$

leading to a contradiction. So for $n \leq n_\alpha$, any (ε, δ) -DP algorithm to estimate q must have expected L_2 error at least α , i.e., $\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] \geq \alpha$. Further, one can see by inspection of the proof of Theorem 3.4 that $\|q(S)\|_2$ satisfies the assumption in the ‘‘Moreover.’’

Now, for $n > n_\alpha$, assuming, for simplicity and without loss of generality that $n - n_\alpha$ is even, we use a padding argument to establish the other part of the bound. Let \mathcal{A} be any (ε, δ) -differentially private algorithm for answering q on datasets of size n . Suppose for the sake of contradiction, that \mathcal{A} satisfies

$$\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] < \frac{n_\alpha}{n} \cdot \alpha \quad (3.5)$$

for every dataset S of size n .

Fix an arbitrary point $\mathbf{c} \in \{\pm 1/\sqrt{d}\}^d$. Given any dataset $S = (x^{(1)}, \dots, x^{(n_\alpha)}) \in \{\pm 1\}^{d \times n_\alpha}$ of size n_α , we construct \hat{S} of size n as follows. Its first n_α entries are $x^{(1)}, \dots, x^{(n_\alpha)}$; then for the remaining $n - n_\alpha$, we have (1) the first $\lceil \frac{n - n_\alpha}{2} \rceil$ (i.e. the first half) of those entries are all copies of \mathbf{c} , and (2) the remaining $\lfloor \frac{n - n_\alpha}{2} \rfloor$ are copies of $-\mathbf{c}$.

Note that we have

$$q(\hat{S}) = \frac{n_\alpha}{n} q(S)$$

for every S , and in particular $\|q(\hat{S})\|_2$ satisfies the assumption in the ‘‘Moreover.’’

Now, we define an algorithm $\hat{\mathcal{A}}$ for approximating q on datasets of size n_α as follows. On input $S \in \{\pm 1\}^{d \times n_\alpha}$, $\hat{\mathcal{A}}$:

- (1) Computes $\hat{S} \in \{\pm 1\}^{d \times n}$ as above
- (2) Outputs $\frac{n}{n_\alpha} \mathcal{A}(\hat{S})$

Since \mathcal{A} is already differentially private, $\hat{\mathcal{A}}$ is also (ε, δ) -DP due to the post-processing property of differential privacy. Moreover,

$$\mathbb{E}[\|\hat{\mathcal{A}}(S) - q(S)\|_2] = \mathbb{E}\left[\left\|\frac{n}{n_\alpha} \mathcal{A}(\hat{S}) - \frac{n}{n_\alpha} q(\hat{S})\right\|_2\right] = \frac{n}{n_\alpha} \mathbb{E}\left[\left\|\mathcal{A}(\hat{S}) - q(\hat{S})\right\|_2\right] \stackrel{(3.5)}{<} \frac{n}{n_\alpha} \cdot \frac{n_\alpha}{n} \alpha = \alpha$$

and so $\hat{\mathcal{A}}$ achieves expected error strictly smaller than α on datasets of size n_α ; which contradicts the first part of the lower bound we already established. So for $n > n_\alpha$, any (ε, δ) -DP algorithm to estimate q must have expected L_2 error at least $\frac{n_\alpha}{n} \cdot \alpha = C \cdot \frac{\sqrt{d \ln(1/\delta)}}{n\varepsilon}$.

Finally, we we have shown that for every n and every $\varepsilon > 0$, there is a constant $C > 0$ such that every (ε, δ) -differentially private algorithm \mathcal{A} answering the linear query q must have,

on some dataset S of size n , expected L_2 error at least

$$\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] = \min\left(\alpha, C \cdot \frac{\sqrt{d \ln(1/\delta)}}{n\varepsilon}\right)$$

proving the lemma. \square

Combining the above with the argument strategy of [BST14, Theorem 5.3] finally yields the main lemma for the second step of our proof for Theorem 3.1:

Lemma 3.6 (Lower bound on empirical loss of (ε, δ) -DP algorithms). *Let $n, d \in \mathbb{N}$, $\varepsilon \in (0, 1]$, and $2^{-o(n)} \leq \delta \leq 1/n^{1+\Omega(1)}$. For every (ε, δ) -differentially private algorithm with output θ^{priv} , there is a dataset $S = \{x_1, \dots, x_n\} \subseteq \{-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d$ such that*

$$\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] = \min\left(\alpha, \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{n\varepsilon}\right)\right)$$

where $\theta^* := \frac{\sum_{i=1}^n x_i}{\|\sum_{i=1}^n x_i\|_2}$ is the minimizer of $\mathcal{L}(\theta, S) := -\langle \theta, \frac{1}{n} \sum_{i=1}^n x_i \rangle$ which is linear and, as such, Lipschitz and convex.

Proof of Lemma 3.6. This proof follows the same structure as that of [BST14, Theorem 5.3] but adapt the bound in terms of expectation.

First, observe that for any $\theta \in \mathbb{B}$ and dataset S we have:

$$\mathcal{L}(\theta, S) - \mathcal{L}(\theta^*, S) = \frac{1}{2} \|q(S)\|_2 \|\theta - \theta^*\|_2^2,$$

since $\|\theta - \theta^*\|_2^2 = \|\theta^*\|_2^2 + \|\theta\|_2^2 - 2\langle \theta, \theta^* \rangle = 2(1 - \langle \theta, \theta^* \rangle)$ using the fact that $\theta^*, \theta \in \mathbb{B}$ have $\|\theta\|_2, \|\theta^*\|_2 = 1$.

Suppose that there is an (ε, δ) -differentially private algorithm \mathcal{A} that outputs θ^{priv} such that, for every dataset $S \subseteq \{-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d$, we have:

$$\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] \leq \Delta$$

for a sufficiently small constant $C > 0$, and some $\Delta \geq 0$. We will prove a lower bound on Δ . To do so, recall $q(S) = \theta^* \cdot \|q(S)\|_2$; and that the lower bound from Lemma 3.5 still holds when the dataset S is promised to be such that $q(S) \in [(M \pm 1)/n]$, for $M = \Theta(\min(n\alpha, \sqrt{d \log(1/\delta)})/\varepsilon)$.

Consider the algorithm (private by post-processing) \mathcal{A} which outputs $\mathcal{A}(S) = \frac{M}{n} \theta^{\text{priv}}$. Then, for any dataset S such that $\|\sum_{i=1}^n x_i\|_2 \in [M - 1, M + 1]$,

$$\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] \leq \mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2^2]^{1/2} = \mathbb{E}\left[\left\|\frac{M}{n} \theta^{\text{priv}} - q(S)\right\|_2^2\right]^{1/2}.$$

On the other hand,

$$\begin{aligned}
\mathbb{E} \left[\left\| \frac{M}{n} \theta^{\text{priv}} - q(S) \right\|_2^2 \right] &\leq 2 \left(\mathbb{E} [\|q(S)\|_2^2 \|\theta^{\text{priv}} - \theta^*\|_2^2] + \mathbb{E} \left[\left\| \frac{M}{n} \theta^{\text{priv}} - \|q(S)\|_2 \theta^{\text{priv}} \right\|_2^2 \right] \right) \\
&= 4 \|q(S)\|_2 \mathbb{E} [\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] + 2 \left(\frac{M}{n} - \|q(S)\|_2 \right)^2 \\
&\leq \frac{4(M+1)}{n} \mathbb{E} [\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] + \frac{2}{n^2} \\
&\hspace{15em} (\text{as } n \|q(S)\|_2 \in [M-1, M+1]) \\
&\leq \frac{4(M+1)\Delta}{n} + \frac{2}{n^2}
\end{aligned}$$

By Lemma 3.5, we know that $\mathbb{E}[\|\mathcal{A}(S) - q(S)\|_2] \geq \min\left(\alpha, C \cdot \frac{\sqrt{d \ln(1/\delta)}}{n\varepsilon}\right)$, for some absolute constant $C > 0$, in the worst case. Hence, we must have

$$\frac{\Delta \cdot M}{n} \geq \min\left(\alpha^2, C^2 \cdot \frac{d \ln(1/\delta)}{n^2 \varepsilon^2}\right);$$

recalling the setting of M , we get $\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] = \min\left(\alpha, \Omega\left(\sqrt{\frac{d \log(1/\delta)}{n\varepsilon}}\right)\right)$. \square

The above lemma establishes a lower bound on the empirical loss of any (ε, δ) -differentially private algorithm. To derive from this our claimed lower bound on unlearning algorithms, we need to introduce a dependence on m , the deletion capacity (i.e., number of points to unlearn). This is done in the last (third) step of our argument, via a reduction which establishes a (restricted) converse to the grouposition property of DP.

Recall that an algorithm $M: \mathcal{X}^n \rightarrow \mathcal{Y}$ satisfies (ε, δ) -DP for edit distance m if for every pair of neighboring datasets X, X' that differ in up to m entries, and every $S \subseteq \mathcal{Y}$:

$$\Pr[M(X) \in S] \leq e^\varepsilon \Pr[M(X') \in S] + \delta.$$

We apply this m -edit distance (ε, δ) -DP on Lemma 3.6 by a reduction that shows: for any differentially private algorithm with respect to edit distance at most m must incur an empirical loss given by Lemma 3.6.

Lemma 3.7. *Suppose there exists an m -edit distance (ε, δ) -DP algorithm \mathcal{M} that takes in a dataset S of size n to approximate $q(S)$ as defined in (3.3), with empirical loss γ . Then, we can construct a 1-edit distance (i.e., standard) (ε, δ) -DP algorithm \mathcal{M}' that, on input a dataset S' of $N = n/m$ data points, approximates $q(S')$ to error γ .*

Proof of Lemma 3.7. The reduction is quite simple: given \mathcal{M} , construct \mathcal{M}' as follows for $N = \frac{n}{m}$ inputs:

$$\mathcal{M}'(x_1, \dots, x_N) = \mathcal{M}(\underbrace{x_1, \dots, x_1}_m, \underbrace{x_2, \dots, x_2}_m, \dots, \underbrace{x_N, \dots, x_N}_m).$$

We immediately have that \mathcal{M}' is (ε, δ) -DP for the usual 1-edit distance between datasets, since \mathcal{M} is DP with respect to edit distance m . The sample complexity and error bound then follows directly from $n = N \times m$, where $n \geq N, N \in \mathbb{N}, m \geq 1$, and the fact that $q(x_1, \dots, x_N) = q(x_1, \dots, x_1, x_2, \dots, x_2, \dots, x_N, \dots, x_N)$ due to the definition of q . \square

Combining Lemma 3.7 with Lemma 3.6, we get that any m -edit distance (ε, δ) -DP algorithm \mathcal{M} approximating q on datasets of size $n = N \times m$ must have error γ at least

$$\min\left(\alpha, \Omega\left(\frac{\sqrt{d \log(1/\delta)}}{N\varepsilon}\right)\right) = \min\left(\alpha, \Omega\left(\frac{m\sqrt{d \log(1/\delta)}}{n\varepsilon}\right)\right)$$

which, reorganising the terms and recalling the definition of deletion capacity, yields the claimed bound on $m_{\varepsilon, \delta}^{A, \bar{A}}$, and hence completes the proof of Theorem 3.3. \square

3.1. Proof of Theorem 1.2. The proof of Theorem 1.2 (the strongly convex case), restated below, is analogous to those of Theorems 3.1 and 3.3, but using [FKT20, Theorem 5.1] for the upper bound (instead of Lemma 3.2) and [BST14, Theorem 5.5] for the lower bound (instead of Theorem 3.4).

Theorem 3.8 (Unlearning For Strongly Convex Loss Functions (Theorem 1.2, restated)). Let $f: \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$ be a 1-Lipschitz *strongly* convex loss function. There exists a DP-based (ε, δ) -machine unlearning algorithm which, trained on a dataset $S \subseteq \mathcal{X}^n$, can unlearn up to

$$m = \Omega\left(\min\left(\frac{n\varepsilon\sqrt{\alpha}}{\sqrt{d \log(1/\delta)}}, \frac{n\sqrt{\varepsilon\alpha}}{d}\right)\right)$$

datapoints without incurring excess population risk greater than α . Moreover, this is nearly tight: there exists a 1-Lipschitz strongly convex loss function such that no DP-based machine unlearning algorithm can unlearn $\Omega(\frac{n\varepsilon\sqrt{\alpha}}{\sqrt{d}})$ data points without excess population risk α .

Sketch. As mentioned above, for the achievability result we can invoke the following result of Feldman, Koren, and Talwar:

Lemma 3.9 (zCDP strongly convex optimization [FKT20, Theorem 5.1]). *Fix any L -Lipschitz, λ -strongly convex loss function over a convex subset \mathcal{B} of \mathbb{R}^d . Then there exists an algorithm A which satisfies $(\rho^2/2)$ -zCDP with excess population loss:*

$$\mathbb{E}\left[F(\theta) - \min_{\theta \in \mathcal{B}} F(\theta)\right] = O\left(\frac{L^2}{\lambda} \cdot \left(\frac{1}{\sqrt{n}} + \frac{d}{\rho^2 n^2}\right)\right)$$

where the expectation is taken over the randomness of A .

Proceeding as in the proof of Theorem 3.1, leveraging grouposition of concentrated DP and choosing

$$\rho := \Theta\left(\min\left(\frac{\varepsilon}{m\sqrt{\log(1/\delta)}}, \frac{\sqrt{\varepsilon}}{m}\right)\right)$$

leads to an excess loss

$$O\left(\frac{L^2}{\lambda} \left(\frac{1}{n} + \frac{m^2 d \log(1/\delta)}{\varepsilon^2 n^2} + \frac{m^2 d}{\varepsilon n^2}\right)\right)$$

and so a bound $m_{\varepsilon, \delta}^{A, \bar{A}} \geq m = \Omega\left(\frac{\sqrt{\lambda}}{L} \min\left(\frac{\varepsilon n \sqrt{\alpha}}{\sqrt{d \log(1/\delta)}}, \frac{n \sqrt{\varepsilon \alpha}}{d}\right)\right)$. The impossibility result, as mentioned above, follows from using the same argument as in Theorem 3.3, using the $\Omega(d/(\varepsilon\sqrt{\alpha}))$ sample complexity lower bound of [BST14, Theorem 5.5] for (ε, δ) -DP strongly convex optimization. \square

3.2. Pure unlearning. We also present analogous deletion capacity bounds for the ε -unlearning case:

Theorem 3.10 (Deletion capacity from unlearning via DP, Lower Bound). Suppose $\mathcal{W} \subseteq \mathbb{R}^d$, and fix any L -Lipschitz (strongly) convex loss function. Then there exists a lazy (ε, δ) -unlearning algorithm (\bar{A}, A) , where \bar{A} has the form $\bar{A}(U, A(S), T(S)) := A(S)$ (thus, in particular, is DP-based) with deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) = \Omega\left(\frac{\varepsilon n \alpha^2}{d}\right).$$

Theorem 3.11 (Deletion capacity from unlearning via DP, Upper Bound). There exists a Lipschitz convex loss function for which any DP-based ε -unlearning algorithm (\bar{A}, A) must have deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) = O\left(\frac{\varepsilon n \alpha}{d}\right).$$

3.2.1. *Proof of Theorems 3.3 and 3.4.* We first restate a lower bound of Bassily, Smith, and Thakurta [BST14, Theorem 5.5], which we will rely on.

Lemma 3.12 (Lower bound on empirical loss of ε -DP algorithms). *Let $n, d \in \mathbb{N}, \varepsilon > 0$. For every ε -differentially private algorithm with output θ^{priv} , there is a dataset $S = \{x_1, \dots, x_n\} \subseteq \{-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}}\}^d$ such that*

$$\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] = \min\left(1, \Omega\left(\frac{d}{n\varepsilon}\right)\right)$$

where $\theta^* := \frac{\sum_{i=1}^n x_i}{\|\sum_{i=1}^n x_i\|_2}$ is the minimizer of $\mathcal{L}(\theta, S) := -\langle \theta, \frac{1}{n} \sum_{i=1}^n x_i \rangle$ (which is linear and, as such, Lipschitz and convex).

Proof. From [BST14, Theorem 5.2], we have that, for some constant $c > 0$,¹

$$\Pr\left[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S) \geq c \cdot \min\left(1, \frac{d}{\varepsilon n}\right)\right] \geq 1/2$$

Then since $\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S) \geq 0$,

$$\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] \geq c \cdot \min(n, d/\varepsilon) \cdot \Pr[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S) \geq c \cdot \min(n, d/\varepsilon)]$$

and so $\mathbb{E}[\mathcal{L}(\theta^{\text{priv}}, S) - \mathcal{L}(\theta^*, S)] \geq \frac{c}{2} \cdot \min(1, \frac{d}{n\varepsilon})$. \square

¹Note that [BST14, Theorem 5.2] is stated for the total (not average) loss, hence the renormalization by a factor n here.

Theorem 3.13 (Upper bound on number of removable samples for ε -unlearning algorithm). Suppose $\mathcal{W} \subseteq \mathbb{R}^d$, and fix any Lipschitz convex loss function. Then any lazy ε -unlearning algorithm (\bar{A}, A) , where \bar{A} has the form $\bar{A}(U, A(S), T(S)) := A(S)$ (thus, in particular, takes no side information) must have deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) \leq O\left(\frac{n\alpha\varepsilon}{d}\right).$$

Proof. Using what we have from Lemma 3.12, we can apply the similar reductions as in the proof of Theorem 3.1 (from ERM to SCO, and from 1-distance DP to m -distance DP through Lemma 3.7).

Hence the lower bound on the expected excess risk of generalization error is given by, writing $N = \frac{n}{m}$, where N is the total number of samples in the reduction in Lemma 3.7,

$$\min\left(1, \Omega\left(\frac{d}{N\varepsilon}\right)\right) = \min\left(1, \Omega\left(\frac{md}{n\varepsilon}\right)\right)$$

and by rearranging with the expected excess risk bound α , we get the desired deletion capacity m . \square

Theorem 3.14 (Lower bound on number of removable samples for ε -unlearning algorithm). Suppose $\mathcal{W} \subseteq \mathbb{R}^d$, and fix any L -Lipschitz Δ -(strongly) convex loss function. Then there exists a lazy (ε, δ) -unlearning algorithm (\bar{A}, A) , where \bar{A} has the form $\bar{A}(U, A(S), T(S)) := A(S)$ (and thus, in particular, takes no side information) with deletion capacity

$$m_{\varepsilon, \delta}^{A, \bar{A}}(\alpha) = \Omega\left(\frac{\varepsilon n \alpha^2 \Delta}{dL}\right).$$

Proof Sketch. By adapting the proof of given in [SSSS09, Section 5.4], which relies on an expected guarantee before applying Markov's inequality, we get the following variant of [BST14, Theorem F.1]:

Theorem 3.15. Let ℓ be an L -Lipschitz, Δ -strongly convex loss function. Then, for any parameter θ (possibly depending on \mathcal{D}), the following holds over the randomness of sampling the data set \mathcal{D} :

$$\mathbb{E}[\text{ExcessRisk}(\theta)] \leq \sqrt{\frac{2L^2}{n\Delta}} \mathbb{E}\left[\sqrt{(\mathcal{L}(\theta, \mathcal{D}) - \mathcal{L}(\theta^*(\mathcal{D}), \mathcal{D}))}\right] + \frac{4L^2}{n\Delta}$$

In particular, applying this to the output $\hat{\theta} = \hat{\theta}(\mathcal{D})$ of any randomized algorithm, and using Jensen's inequality, we get

$$\mathbb{E}\left[\text{ExcessRisk}(\hat{\theta})\right] \leq \sqrt{\frac{2L^2}{n\Delta}} \sqrt{\mathbb{E}\left[\mathcal{L}(\hat{\theta}, \mathcal{D}) - \mathcal{L}(\theta^*(\mathcal{D}), \mathcal{D})\right]} + \frac{4L^2}{n\Delta} \quad (3.6)$$

where the expectation is over the draw of the dataset \mathcal{D} and the randomness of the algorithm.

Then, applying this theorem with the upper bound of the ε -DP algorithm from [BST14], we have

$$\mathbb{E}\left[\text{ExcessRisk}(\hat{\theta})\right] \leq C \cdot \left(\sqrt{\frac{L^2}{n\Delta}} \sqrt{\frac{d}{\varepsilon}} + \frac{L^2}{n\Delta}\right) \quad (3.7)$$

for some absolute constant $C > 0$.

By using Lemma 3.7 again, we have:

$$\mathbb{E}[\text{ExcessRisk}(\hat{\theta})] \leq C \cdot \left(\sqrt{\frac{mdL^2}{\varepsilon n \Delta}} + \frac{mL^2}{n\Delta} \right) \quad (3.8)$$

which gives the desired lower bound. \square

4. THE COMPOSITION PROPERTIES OF (ε, δ) -UNLEARNING

In this section, we establish the post-processing (Theorem 1.4), chaining (Theorem 1.5), and grouposition (Theorem 1.6) properties of (ε, δ) -unlearning, the last two under our “laziness” assumption (Assumption 1.3).

Theorem 4.1 (Post-processing of unlearning (Theorem 1.4, restated)). Let (\bar{A}, A) be an (ε, δ) -unlearning algorithm taking no side information. Let $f: \mathcal{W} \rightarrow \mathcal{W}$ be an arbitrary possibly randomized function. Then $(f \circ \bar{A}, A)$ is also an (ε, δ) -unlearning algorithm.

Proof. The proof follows exactly same as post-processing property of differential privacy. We consider the case that f is a deterministic function here without loss of generality.

Let $T = \{r \in \mathbb{R}^d \mid f(r) \in \mathcal{Y}\}$ and $\mathcal{Y} \subseteq \mathbb{R}^d$. Consider for any $\mathcal{Y} \subseteq \mathbb{R}^d$:

$$\begin{aligned} \Pr[f(\bar{A}(A(S), U)) \in \mathcal{Y}] &= \Pr[\bar{A}(A(S), U) \in T] \\ &\leq e^\varepsilon \Pr[\bar{A}(A(S), U) \in T] + \delta \\ &= e^\varepsilon \Pr[f(\bar{A}(A(S), U)) \in \mathcal{Y}] + \delta \end{aligned}$$

which concludes the proof. \square

Under our laziness assumption, we can establish bounds on applying an unlearning algorithm repeatedly when the overall deletion request is within the deletion capacity:

Theorem 4.2 (Chaining of unlearning (Theorem 1.5, restated)). Let (\bar{A}, A) be a *lazy* (ε, δ) -unlearning algorithm taking no side information, and able to handle up to m deletion requests. Then, for any k , the algorithm which uses (\bar{A}, A) to sequentially unlearn k disjoint deletion requests $U_1, \dots, U_k \subseteq X$ such that $|\cup_i U_i| \leq m$ by outputting

$$\bar{A}(U_k, \dots, \bar{A}(U_1, A(X)) \dots)$$

is an (ε', δ') -unlearning algorithm, with $\varepsilon' = k\varepsilon$ and $\delta' = \delta \cdot \frac{e^{k\varepsilon} - 1}{e^\varepsilon - 1} = O(k\delta)$ (for $k = O(1/\varepsilon)$).

Proof. We proceed by induction on $n \geq 1$. Given a pair of (ε, δ) -unlearning algorithm (\bar{A}, A) and deletion requests $U_1, \dots, U_n \subseteq S \in \mathbb{R}^{n \times d}$ such that $|\cup_i U_i| \leq m := m_{\varepsilon, \delta}^{\bar{A}, A}$ with $U_i \cap U_j = \emptyset, \forall i \neq j$ for $i, j \in [n]$.

(1) For $n = 1$: for every measurable set T , if $|U_1| \leq m$

$$\Pr[\bar{A}(A(S), U_1) \in T] \leq e^{n\varepsilon} \Pr[\bar{A}(A(S \setminus U_1), \emptyset)] + \delta$$

by the definition of (ε, δ) -unlearning. Hence the case $n = 1$ holds.

(2) Assume the induction hypothesis holds for some $n \geq 1$:

$$\Pr[\bar{A}(\dots \bar{A}(A(S), U_1), \dots, U_n) \in T] \leq e^{n\varepsilon} \Pr[\bar{A}(A(S \setminus \bar{U}_n), \emptyset)] + \sum_{j=0}^{n-1} e^{j\varepsilon} \cdot \delta, \quad \text{for all } T \quad (4.1)$$

whenever U_1, \dots, U_n are disjoint sets with $|\cup_i U_i| \leq m$.

(3) Then for $n + 1$: fixing an arbitrary set T ,

$$\begin{aligned} \Pr[\bar{A}(\dots \bar{A}(A(S), U_1), \dots, U_{n+1}) \in T] &\stackrel{(4.1)}{\leq} e^{n\varepsilon} \Pr[\bar{A}(\bar{A}(A(S \setminus \bar{U}_n), \emptyset), U_{n+1})] + \sum_{j=0}^{n-1} e^{j\varepsilon} \cdot \delta \\ &= e^{n\varepsilon} \Pr[\bar{A}(A(S \setminus \bar{U}_n), U_{n+1})] + \sum_{j=0}^{n-1} e^{j\varepsilon} \cdot \delta \\ &\leq e^{(n+1)\varepsilon} \Pr[\bar{A}(A(S \setminus \bar{U}_{n+1}), \emptyset) \in T] + \sum_{j=0}^{(n+1)-1} e^{j\varepsilon} \cdot \delta \end{aligned}$$

where the first and third inequalities result from the definition of (ε, δ) -unlearning (using that $|\cup_i U_i| \leq m$) and the second equality is due to the Laziness Assumption 1.3.

Hence, by induction, the claim holds for all $n \in \mathbb{N}$. \square

Theorem 4.3 (Advanced grouposition of unlearning (Theorem 1.6, restated)). Let $(\bar{A}_1, A), \dots, (\bar{A}_k, A)$ be *lazy* (ε, δ) -unlearning (with common learning algorithm A) taking no side information, and define the composition of those unlearning algorithms, \tilde{A} , as

$$\tilde{A}(U, A(X)) = f(\bar{A}_1(U, A(X)), \dots, \bar{A}_k(U, A(X))).$$

where $f: \mathcal{W}^k \rightarrow \mathcal{W}$ is any (possibly randomized) function. Then, for every $\delta' > 0$, (\tilde{A}, A) is an (ε', δ') -unlearning algorithm taking no side information, where $\varepsilon' = \frac{1}{2}k\varepsilon^2 + \varepsilon\sqrt{2k \ln(1/\delta')}$.

Proof. The proof follows the same argument as in [Vad17, Lemma 2.4]. We consider the case of $\delta > 0$ only as the $\delta = 0$ is the same.

Fix two datasets, S (original dataset) and $S' := S \setminus U$ (“forgotten dataset”) where U is the set of delete requests with $|U| \leq m_{\varepsilon, \delta}^{\bar{A}, A}$. Note that S, S' differs in m entries.

For an output $y = (y_1, \dots, y_k) \in \mathcal{Y}$, define “memory” loss (analogous to privacy loss in differential privacy) to be:

$$\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) = \ln \frac{\Pr[\mathcal{A}(A(S), U) = y]}{\Pr[\mathcal{A}(A(S'), \emptyset) = y]}$$

where $|\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y)| \leq \varepsilon$.

Then, by [Vad17, Lemma 1.5] we know that $\bar{A}_i(A(S), U), \bar{A}_i(A(S'), \emptyset)$ are (ε, δ) -indistinguishable, hence there are events $E = E_1 \wedge \dots \wedge E_k, E' = E'_1 \wedge \dots \wedge E'_k$ such that w.p. at least $1 - k\delta$,

for all $y_i, i \in [k]$,

$$\begin{aligned} \mathbb{E} \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) \right] &= \mathbb{E} \left[\ln \frac{\Pr[\mathcal{A}(A(S), U) = y \mid E]}{\Pr[\mathcal{A}(A(S'), \emptyset) = y \mid E']} \right] \\ &= \sum_{i=1}^k \mathbb{E} \left[\ln \left(\frac{\Pr[\bar{A}_i(A(S), U) = y \mid E_i]}{\Pr[\bar{A}_i(A(S'), \emptyset) = y \mid E'_i]} \right) \right] \\ &= \sum_{i=1}^k \mathbb{E} \left[\mathcal{L}_{\bar{A}_i}^{S \rightarrow S'}(y) \right] \end{aligned}$$

where we observe that the expectation of the loss is just KL-divergence between the distributions of $\bar{A}_i(A(S), U)$ and $\bar{A}_i(A(S'), \emptyset)$ conditioned on E and E' . Hence:

$$\mathbb{E} \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) \right] = \sum_{i=1}^k \text{D}_{\text{KL}}(\bar{A}_i(A(S), U) \parallel \bar{A}_i(A(S'), \emptyset)) \leq \frac{1}{2} k \varepsilon^2$$

where the inequality is a result from [BS16, Proposition 3.3] when $\alpha = 1$. This proposition is applicable because the conditional distributions of \bar{A}_i are (ε, δ) -indistinguishable, which shares the max-divergence definition.

Then by Hoeffding's inequality, as the loss is bounded by $[-\varepsilon, \varepsilon]$, we have

$$\begin{aligned} \Pr \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) > \varepsilon' \right] &= \Pr \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) > \frac{1}{2} k \varepsilon^2 + t \right] \\ &\leq \Pr \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) > \mathbb{E} \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) \right] + t \right] \\ &\leq \exp \left(-\frac{t^2}{2k\varepsilon^2} \right) \end{aligned}$$

We set $t = \varepsilon \sqrt{2k \ln(1/\delta')}$, so that the RHS above is δ' , and $\varepsilon' := \frac{1}{2} k \varepsilon^2 + \varepsilon \sqrt{2k \ln(1/\delta')}$.

Hence, for any set $T \subseteq \mathcal{Y}$:

$$\begin{aligned} \Pr[\mathcal{A}(A(S), U) \in T] &\leq \Pr \left[\mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) > \varepsilon' \right] + \sum_{y \in T: \mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) \leq \varepsilon'} \Pr[\mathcal{A}(A(S), U) = y] \\ &\leq \delta' + \sum_{y \in T: \mathcal{L}_{\mathcal{A}}^{S \rightarrow S'}(y) \leq \varepsilon'} e^{\varepsilon'} \Pr[\mathcal{A}(A(S'), \emptyset) = y] \\ &\leq \delta' + e^{\varepsilon'} \Pr[\mathcal{A}(A(S'), \emptyset) \in T] \end{aligned}$$

where the second inequality is from the definition of unlearning. Thus, along with an application of [Vad17, Lemma 1.5], this proves that $\mathcal{A} = (\bar{A}_1, \dots, \bar{A}_k)$ is indeed $(\varepsilon', \delta' + k\delta)$ -unlearning w.r.t. learning algorithm A . \square

5. DISCUSSION AND FUTURE WORK

Our work fully characterizes the deletion capacity of any unlearning algorithm (\bar{A}, A) minimizing population risk under both convex and strongly convex loss functions, when only given the model parameters (output of the learning algorithm) and the set of deletion requests. This restriction, namely that the unlearning algorithm does not rely on any additional side

information, is motivated by the potential privacy risks storing non-private side information can entail.

We hope our work will lead to further study of the interplay between differential privacy and machine unlearning, and to additional study of the “DP-like” properties of machine unlearning, such as the postprocessing and composition properties our present work has identified. In view of the myriad applications these properties have had in privacy-preserving algorithm design, we believe that their analogue for machine unlearning will prove very useful.

One final remark before concluding, we believe it is important to reiterate a key point about our work. One could view our main results as focusing on what can be achieved by unlearning algorithms “which, upon a deletion request, do not do anything,” since they leverage the guarantees already provided by differential privacy. However, this does *not* mean that the algorithm “does nothing” *overall*. Instead, the point here is that the algorithms considered already satisfy the very stringent notion of DP, and as such already paid some utility cost to provide this guarantee; *having paid that price*, they benefit from some “unlearning bonus” for free.

Put differently, the aim of this paper is not to promote or justify when it comes to unlearning, but instead to characterize exactly how much unlearning guarantees come “for free” if one decides or already has to offer the strong guarantee of differential privacy. Thus, our aim is not to discourage unlearning-only solutions when DP is not required; but instead, by understanding the interplay between DP and unlearning, to show that the joint (and seemingly daunting) requirement of both differential privacy and right to be forgotten requirement is more affordable than it seems.

REFERENCES

- [BCC⁺21] Lucas Bourtole, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 141–159. IEEE, 2021.
- [BFTT19] Raef Bassily, Vitaly Feldman, Kunal Talwar, and Abhradeep Guha Thakurta. Private stochastic convex optimization with optimal rates. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems*, pages 11279–11288, 2019.
- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658, 2016.
- [BST14] Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 464–473. IEEE Computer Society, 2014.

- [CSSV23] Aloni Cohen, Adam D. Smith, Marika Swanberg, and Prashant Nalini Vasudevan. Control, confidentiality, and the right to be forgotten. In *CCS*, pages 3358–3372. ACM, 2023.
- [CY15] Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *IEEE Symposium on Security and Privacy*, pages 463–480. IEEE Computer Society, 2015.
- [CZW⁺21] Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang. When machine unlearning jeopardizes privacy. In Yongdae Kim, Jong Kim, Giovanni Vigna, and Elaine Shi, editors, *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, November 15 - 19, 2021*, pages 896–911. ACM, 2021.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography, TCC '06*, pages 265–284, Berlin, Heidelberg, 2006. Springer.
- [FKT20] Vitaly Feldman, Tomer Koren, and Kunal Talwar. Private stochastic convex optimization: Optimal rates in linear time. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, page 439–449, New York, NY, USA, 2020. Association for Computing Machinery.
- [GGHvdM20] Chuan Guo, Tom Goldstein, Awni Y. Hannun, and Laurens van der Maaten. Certified data removal from machine learning models. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 3832–3842. PMLR, 2020.
- [GGMV22] Ji Gao, Sanjam Garg, Mohammad Mahmoody, and Prashant Nalini Vasudevan. Deletion inference, reconstruction, and compliance in machine (un)learning. *Proc. Priv. Enhancing Technol.*, 2022(3):415–436, 2022.
- [GGV20] Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. Formalizing data deletion in the context of the right to be forgotten. In *EUROCRYPT (2)*, volume 12106 of *Lecture Notes in Computer Science*, pages 373–402. Springer, 2020.
- [GGVZ19] Antonio A. Ginart, Melody Y. Guan, Gregory Valiant, and James Y. Zou. Making ai forget you: Data deletion in machine learning. In *Neural Information Processing Systems*, 2019.
- [GJN⁺21] Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. Adaptive machine unlearning. In *NeurIPS*, pages 16319–16330, 2021.
- [ISCZ20] Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Y. Zou. Approximate data deletion from machine learning models: Algorithms and evaluations. 2020.
- [Kar23] Paul Karp. Australia to consider european-style right to be forgotten privacy laws. *The Guardian*, 2023.
- [NHR⁺25] Thanh Tam Nguyen, Thanh Trung Huynh, Zhao Ren, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning. *ACM Trans. Intell. Syst. Technol.*, July 2025.
- [NRS21] Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In Vitaly Feldman, Katrina

- Ligett, and Sivan Sabato, editors, *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, volume 132 of *Proceedings of Machine Learning Research*, pages 931–962. PMLR, 2021.
- [SAKS21] Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. In *Advances in Neural Information Processing Systems*, volume 34, pages 18075–18086. Curran Associates, Inc., 2021.
- [SBB⁺20] Ahmed Salem, Apratim Bhattacharya, Michael Backes, Mario Fritz, and Yang Zhang. Updates-Leak: Data set inference and reconstruction attacks in online learning. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1291–1308. USENIX Association, August 2020.
- [SSSSS09] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada, June 18-21, 2009*, 2009.
- [SU15] Thomas Steinke and Jonathan R. Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *COLT*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 1588–1628. JMLR.org, 2015.
- [SU16] Thomas Steinke and Jonathan R. Ullman. Between pure and approximate differential privacy. *J. Priv. Confidentiality*, 7(2), 2016.
- [SW22] Vinith M. Suriyakumar and Ashia C. Wilson. Algorithms that approximate data removal: New results and limitations. In *NeurIPS*, 2022.
- [UMR⁺21] Enayat Ullah, Tung Mai, Anup Rao, Ryan A. Rossi, and Raman Arora. Machine unlearning via algorithmic stability. In Mikhail Belkin and Samory Kpotufe, editors, *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 4126–4142. PMLR, 2021.
- [Vad17] Salil P. Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography*, pages 347–450. Springer International Publishing, 2017.