

Protecting the Undeleted in Machine Unlearning

Aloni Cohen   

University of Chicago, IL, USA

Refael Kohen   

Tel Aviv University, Israel

Kobbi Nissim   

Dept. of Computer Science, Georgetown University, Washington, DC, USA

Uri Stemmer   

Tel Aviv University, Israel

Abstract

Machine unlearning aims to remove specific data points from a trained model, often striving to emulate “perfect retraining”, i.e., producing the model that would have been obtained had the deleted data never been included. We demonstrate that this approach, and security definitions that enable it, carry significant privacy risks for the *remaining* (undeleted) data points. We present a reconstruction attack showing that for certain tasks, which can be computed securely without deletions, a mechanism adhering to perfect retraining allows an adversary controlling merely $\omega(1)$ data points to reconstruct almost the entire dataset simply by issuing deletion requests. We survey existing definitions for machine unlearning, showing they are either susceptible to such attacks or too restrictive to support basic functionalities like exact summation. To address this problem, we propose a new security definition that specifically safeguards undeleted data against leakage caused by the deletion of other points. We show that our definition permits several essential functionalities, such as bulletin boards, summations, and statistical learning.

2012 ACM Subject Classification Security and privacy

Keywords and phrases Unlearning, data deletion, privacy

Digital Object Identifier 10.4230/LIPIcs.FORC.2026.17

Related Version *Full Version:* <https://arxiv.org/abs/2602.16697>

Funding *Refael Kohen:* Work partially supported by the Israel Science Foundation (grant 1419/24) and the Blavatnik Family Foundation.

Kobbi Nissim: Work partly supported by National Science Foundation award CCF-2217678 “Co-design of law and computer science for privacy in sociotechnical software systems” and by a gift to Georgetown University.

Uri Stemmer: Work partially supported by the Israel Science Foundation (grant 1419/24) and the Blavatnik Family Foundation.

1 Introduction

Sensitive personal data underlies the training of machine learning models across many domains, necessitating compliance with legal data protection standards such as the EU General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). These legal standards provide individuals with a number of protections and rights, including the right to request that their specific information be deleted, a.k.a. the Right to be Forgotten (RTBF). This legal provision gave rise to the field of *Machine Unlearning*, a branch of machine learning focused on enabling the removal of elements from the training data without needing to rebuild models from scratch.



© Aloni Cohen, Refael Kohen, Kobbi Nissim, and Uri Stemmer;
licensed under Creative Commons License CC-BY 4.0

7th Symposium on Foundations of Responsible Computing (FORC 2026).

Editor: Huijia (Rachel) Lin; Article No. 17; pp. 17:1–17:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We now make our discussion more formal and survey several influential definitions for machine unlearning. Consider a data curator $\mathcal{C}$ that executes an algorithm $\mathcal{A}$ on an initial input dataset $D_0 = (x_1, \dots, x_n)$ to obtain an outcome h_0 (say a hypothesis or a statistic). Throughout the execution, at every time i , the curator $\mathcal{C}$ receives a *deletion request* of the form “delete point x_j ” and needs to respond with an updated outcome h_i that should reflect the state of the dataset after this i th deletion request. Starting from the work of Cao and Yang [3], *machine unlearning* was defined as a computational problem, where the ideal solution is for the curator to retrain the learning algorithm $\mathcal{A}$ from scratch in every step on the remaining dataset. The goal of the machine unlearning algorithm is to emulate this ideal solution (perfectly or approximately) faster, i.e., without having to spend the work needed to re-run $\mathcal{A}$ for every deletion request. For example, if $\mathcal{A}$ returns the exact sum $h_0 = x_1 + \dots + x_n$, deleting x_j could be done via a single subtraction $h_1 = h_0 - x_j$ (instead of summing the n numbers again). This way, deletion is performed in constant rather than linear time.

Following Cao and Yang [3], several other definitions were suggested, including by [14, 17, 18, 22, 23], with the same goal of mimicking/approximating the behavior of retraining from scratch in every iteration (“the ideal solution”) while saving on runtime. These definitions differ in the approximation metrics they consider and in their environmental settings (such as whether or not the deletion requests are adaptive to past outputs). The focus of our work is not on runtime; therefore, for our purposes, all these definitions can be grouped together as *definitions that are satisfied by perfect retraining*. See Section 1.2 for a survey of prior definitions.

Perfect retraining and privacy. While the right to delete one’s information is rooted in privacy standards such as the GDPR, the relationship between data deletion and privacy is not straightforward. In particular, it is well-known that definitions satisfied by perfect retraining fail under a simple “differencing attack” where an adversary learns information about the deleted item by examining the output hypothesis before and after a deletion (e.g., in the exact sum example, if an attacker sees the outcome before and after Alice deletes her input, then the attacker learns Alice’s input exactly).

Nevertheless, one may argue that such compromise of Alice’s personal data is acceptable as long as Alice is aware of her privacy risks and still wants to exercise control over her data and delete it from the model. However, as was initially observed by [4], there are cases where the deletion of Alice’s data can compromise the privacy of *other* users who did not ask for their data to be deleted. We believe that enabling such behavior is unacceptable; thus, we argue that definitions that aim to mimic or approximate perfect retraining should be used with caution.

A simple example where the deletion of one party’s data causes another’s to be revealed is in median computation. Given n data points $x_1 \leq x_2 \leq \dots \leq x_n$ their median is $x_{\frac{n+1}{2}}$ if n is odd and $x_{\frac{n}{2}}$ otherwise. Observe that the initial median computation reveals one data point. Assume n is odd and that Alice’s entry is greater than the median. When Alice asks to delete her data, a value smaller than the median is exposed – the input provided by some other user. Variants of this example were identified by [4, 16, 5]. See Appendix A for another example in the context of k -means clustering.

1.1 Our contributions

Certain tasks are inherently incompatible with perfect-retraining. The median example outlined above is perhaps somewhat controversial because the exact median itself is a highly non-private computation, even without any deletions. Thus, one might argue that the privacy

violation generated in this context is more because of the functionality itself rather than because of the deletions. Taking this argument to the limit, consider the functionality f where $f(D) = \emptyset$ if $|D| \geq n$, and $f(D) = D$ otherwise. Now, if the initial dataset D is of size n , then by deleting her data Alice causes the function f to leak all of the (remaining) dataset. This motivates the following question.

► **Question 1.** Is there a task that can be solved privately on static data, yet fundamentally compromises the privacy of undeleted data the moment deletion requests are supported?

We remark that it is trivial to construct specific, poorly designed algorithms whose privacy guarantees are broken by deletions.¹ In contrast, Question 1 asks whether there exists a task for which *any* algorithm would provably be vulnerable. Indeed, in Section 2, we present such a task: while it can be computed with differential privacy in the static setting, an attacker controlling a small number of parties can force *any* perfect-retraining algorithm to leak almost the entire dataset, simply by issuing sequential deletion requests.

► **Theorem 2 (informal).** *There exists a task T such that*

- (1) *There exists an $(o(1), o(1))$ -differentially private algorithm $\mathcal{A}$ such that for any dataset D we have that $\mathcal{A}(D)$ solves T w.r.t. D with high probability.*
- (2) *Perfect retraining w.r.t. T allows an attacker controlling merely $\omega(1)$ points to reconstruct almost all of the dataset. Specifically, there is an attacker $\mathcal{B}$ and a set of points $B = \{b_1, \dots, b_r\}$ of size $|B| = \omega(1)$ such that the following holds. (B is the set of points controlled by $\mathcal{B}$.) For any dataset D containing B , and any algorithm $\mathcal{A}$ for T , given $(\mathcal{A}(D), \mathcal{A}(D \setminus \{b_1\}), \mathcal{A}(D \setminus \{b_1, b_2\}), \dots, \mathcal{A}(D \setminus B))$, the attacker $\mathcal{B}$ can reconstruct almost all of D .*

By item (1), without deletions, the task T can be solved while guaranteeing strong privacy protections for the input dataset. Nevertheless, by item (2), an attacker that controls merely $\omega(1)$ points and can request their deletion from the dataset can reconstruct almost all of the dataset (reconstruct $\approx n$ points by deleting $\omega(1)$ points). As we mentioned, many existing definitions for machine unlearning are satisfied by perfect retraining. Theorem 2 shows that these definitions do not rule out mechanisms that leak their entire dataset after just $\omega(1)$ deletion requests.

► **Remark 3.** Item (1) of Theorem 2 is stated with (ε, δ) -DP for $\delta = o(1)$. This is not standard, but still prevents the attacker from reconstructing, say, 1% of the rows in D . The formal details are given in Appendix B. In particular, this can be restated with arbitrary (ε, δ) parameters, at the cost requiring the attacker described in Item (2) to control $\omega\left(\frac{\log(1/\delta)}{\varepsilon^2}\right)$ points.

► **Remark 4.** There are several privacy-oriented definitions in the literature that are *not* satisfied by perfect retraining (and thus might not be susceptible to our attacks). Perhaps the first example of such a (folklore) definition would be one that requires the use of differential privacy (DP) with small enough privacy parameters to capture group privacy for all the deleted users.² However, to the best of our knowledge, all existing definitions that are not satisfied by perfect retraining are too restrictive in the sense that they disallow natural and important functionalities, such as a simple sum or a bulletin board. Again, we refer the reader to Section 1.2 for a survey of prior definitions.

¹ For example, consider an algorithm that, upon every deletion, recomputes and releases a different DP computation on its (remaining) dataset, where all of these DP computations are done using the same random seed. While every single output is indeed private, the vector of all outputs is not.

² That is, at the outset of the execution, the curator releases some DP computation on the initial dataset (with very small privacy parameters) and then simply ignores all deletion requests.

A new definition for deletion safety. Motivated by these vulnerabilities, we propose a new framework for machine unlearning. Unlike prior definitions that focus on removing the influence of the deleted point to mimic a world where it never existed, our definition focuses on protecting the privacy of the remaining, *undeleted* elements. Informally, we require that an attacker who observes an initial model and submits a sequence of deletion requests gains no additional information about the remaining dataset beyond what is inherent in the initial model and the deleted values themselves. We formalize this via a simulation-based definition, and provide variants handling non-adaptive and adaptive adversaries, as well as a relaxation allowing for bounded leakage.

Examples satisfying our new definition. We show that our new definition is satisfied by stateless algorithms. This allows us to show that the exact sum and the bulletin board functionalities, which were disallowed by previous rigorous definitions, are compatible with our definition. Furthermore, we provide examples for mechanisms, which were studied by prior work under the umbrella of perfect retraining, that are nevertheless compatible with our new definition. In other words, even though our work highlights security vulnerabilities of perfect retraining as a security notion, this does not necessarily mean that existing algorithms in the literature are bad; indeed, some of them could be justified using our new security notion.

1.2 A survey of prior definitions

Definitions that are satisfied by perfect retraining. Our attacks apply to all definitions which are satisfied by perfect retraining. This includes all the definitions in the literature in which perfect retraining is the “stated goal”/“optimum” [3, 14, 17, 18, 22, 23]. It also includes definitions where retraining from scratch is sufficient even if not the stated goal.

One such definition, called *deletion-as-control*, was introduced by Cohen et al. [7]. Recall that we consider a data curator $\mathcal{C}$ that holds an initial dataset D and sequentially receives deletion requests for elements from D (and updates its responses accordingly). Informally, deletion-as-control requires that the internal state of the curator $\mathcal{C}$ in the real-world (where, say, Alice contributes her information and then asks to delete it) would be distributed similarly to its state in a “hypothetical world” where Alice does not contribute her data in the first place, but all other parties (except the curator) behave as if she did. Note that this definition is satisfied by perfect retraining because in perfect retraining the curator need not have a state besides the dataset, which is identical in the two executions (after the deletion).

Another notable definition by Godin and Lamontagne [15] requires that the internal state of the curator $\mathcal{C}$ after Alice deletes her data should be simulatable from the view of all other parties (except Alice and the curator). Note that perfect retraining satisfies this definition, because the curator does not have a state besides the dataset, which consists of information jointly held by the parties excluding Alice.

Definitions that are too restrictive. Not all definitions are satisfied by perfect retraining; against these our attacks do not apply. As we argue next, these definitions can be too restrictive in some settings. In particular, all the definitions we discuss next rule out simple and important functionalities including exact sums and bulletin board [7]. As a concrete example to have in mind, a bulletin board should allow users to publicly post messages and to later remove them.

Perhaps the most influential definition in this context, called *deletion-as-compliance*, was introduced by Garg et al. [13]. Suppose that Alice contributes and later deletes her data by interacting with the curator $\mathcal{C}$. Informally, the requirement in the definition of Garg et al. [13] is that both the internal state of the curator $\mathcal{C}$ as well as the *view* of all other parties (except the curator and Alice) should be distributed similarly to the way they would have been distributed in an ideal world in which Alice never contributed her data in the first place. This means, in particular, that the view of these other parties (capturing everything they see throughout the execution) cannot be affected by Alice. Bulletin boards and exact sums are impossible: the view of other parties is inherently affected by Alice's contribution.

This issue was addressed by Gao et al. [12] who presented a weaker variant of the definition of Garg et al. [13] with the goal of being more permissive. Informally, the definition of Gao et al. [12] is as follows. Suppose that Alice contributes *two* records and later deletes *one* of them (by interacting with the curator $\mathcal{C}$). The requirement is that all other parties together should not be able to distinguish (based on their view of the execution) whether Alice deleted the first or the second record. This definition is indeed more permissive and, in particular, permits the *fact that a deletion has occurred* to be visible (but the identity of the deleted record needs to be kept hidden). However, this definition still rules out the exact sum and the bulletin board functionalities for the same reasons as above.

Finally, Chourasia and Shah [6] studied privacy risks that may arise when the adversary that selects the sequence of deletion requests is *adaptive* to the prior outputs of the curator.³ Their work was motivated by observation that if the sequence of deletions is adaptive, then this sequence itself might continue to reveal information about a record even after this record was deleted from the system. Informally, their definition requires that for every adaptive adversary Q (that adaptively specifies deletion requests), every initial dataset D , every record $y \in D$, and every time step i , the output of the curator after receiving $i - 1$ deletion requests from Q and then a request to delete y is simulatable from $D \setminus \{y\}$.⁴ Notice that the simulator of Chourasia and Shah [6] is given all the data except for a single record, including every undeleted record. This is very much at odds with the ethos of our goal of privacy for those who remain.

Still, as Chourasia and Shah [6] prove, the fact that this must hold for any *adaptive* adversary Q forces the curator to provide some form of privacy protections to the records in D . (Roughly, the functionality can be non-private for at most a single time step.) Thus it disallows simple functionalities such as the exact sum or the bulletin board. To see this for the bulletin board functionality, consider an adaptive Q that behaves as follows: If Alice's public message mentions Bob, then Q will request to delete Alice then Bob. Otherwise, Q will request to delete Alice then Charlie. Since the simulator is executed without Alice's record, it cannot simulate which record to delete next.

Note that the definition of Chourasia and Shah [6] argues about every single time step separately. In Section 2.2, we show that such definitions are vulnerable to a type of attack we call a *generalized differencing attack*.

³ This work extends the work of Gupta et al. [18] that we mentioned above. Gupta et al. [18] also studied adaptivity in machine unlearning, tackling the concern that algorithms that aim to mimic perfect retraining efficiently might lose their guarantees when the sequence of deletion is selected adaptively.

⁴ More specifically, their definition requires that for every adaptive adversary Q and every time step i there exists a simulator π_i^Q such that the following holds. For every initial dataset D and every record $y \in D$, the output distribution of the curator at time i , after receiving $i - 1$ deletion requests (chosen adaptively by Q) and then receiving a request to delete y , is similar to the outcome distribution of the simulator $\pi_i^Q(D \setminus \{y\})$.

2 Attacks

In this section, we present several attacks in which an attacker controlling some of the records in a dataset can use the ability to delete these records in order to learn information about *other* records in the dataset.

2.1 Reconstruction attacks

Let us reconsider the median example, with the goal of releasing an *approximate* median rather than the exact median. Specifically, consider an algorithm that given a dataset D returns a differentially private (DP) approximation for the median. This breaks the attack described in Section 1.1, as it prevents the attacker from directly learning information about the remaining points from their DP median. Still, by deleting more and more points the attacker gets to see more and more releases (approximate medians), which together might allow it to infer information about the remaining points. If we were to ensure that *all releases together* satisfy DP, as in the *continual observation model* of Dwork et al. [10], then this would provably prevent such an attack. However, this can come at a high price in terms of utility. Specifically, Jain et al. [19] presented problems that can be solved efficiently with differential privacy in the *single shot* setting (where only one output is released), but not in the *continual* setting, where the input dataset evolves over time and DP must be preserved across all outputs together.

We now briefly describe a variant of one of the attacks of Jain et al. [19], tailored to our setting. Essentially, the only conceptual modification from the attacks presented by Jain et al. [19] is that we highlight the fact that it is not just that DP cannot be maintained in this continual setting, but it is in fact the privacy of “undeleted” points that get compromised.⁵ This will not yet result in the attack described in Theorem 2 as the number of users that the attacker needs to control will be relatively large. Consider the following problem: Given a multiset $T \subseteq [n]$, the goal is to return an approximation to

$$\text{CountMod}(T) \triangleq \sum_{i \in [n]} \mathbf{1}\{\text{count}_T(i) \bmod 3 = 2\},$$

where $\text{count}_T(i)$ is the number of occurrences of i in T . Note that CountMod is a function of sensitivity 1, and hence could be approximated with DP (in the single shot setting) with error roughly $1/\epsilon$.

Now consider a dataset (multiset) $T = D \cup S \subseteq [n]$, where D is a random subset of $[n]$ of size $n/2$, and where S is the multiset of size $3n^2$ containing $3n$ copies of every $i \in [n]$. We consider an attacker that “controls” S (in the sense that it can request the deletion of points from S) and aims to learn the dataset D . Now observe that by controlling S , the attacker can force any mechanism for CountMod to provide it with answers to *counting queries* w.r.t. D . In more detail, a *counting query* is defined via a predicate $q : [n] \rightarrow \{0, 1\}$, and the value of such a query on D is $q(D) \triangleq \sum_{i \in D} q(i)$. Given such a counting query q , the attacker acts as follows:

⁵ This also holds in the attacks of Jain et al. [19], but not explicit as this was not their focus; they aimed to show that the error must be large as a function of the size of the evolving dataset (which includes both the deleted and undeleted elements in our context).

1. For every $i \in [n]$ such that $q(i) = 1$, request the deletion of 2 copies of i from S (and thus from T).
2. Obtain a release $a_q \approx \text{CountMod}(T)$. Note that $\text{CountMod}(T) = q(D)$, because, by Step 1, for every $i \in [n]$ such that $q(i) = 0$ we have that $\text{count}_T(i) \bmod 3 \neq 2$, and for every $i \in [n]$ such that $q(i) = 1$ we have that $\text{count}_T(i) \bmod 3 = 2$ if and only if $i \in D$.
3. For every $i \in [n]$ such that $q(i) = 1$, request the deletion of one copy of i from S (and thus from T). Note that this “resets” the modifications we did in Step 1 in the sense that for every i , its number of occurrences in S is a multiple of 3.

This allows the attacker to obtain answers to (at least) n arbitrary counting queries w.r.t. D . This is known to allow the attacker to reconstruct D unless the approximation error is $\Omega(\sqrt{n})$ by the results of [11].

Reducing the number of deletions. In the last attack, the attacker controlled $O(n^2)$ points and were able to reconstruct $\approx n$ “undeleted” points. By slightly modifying the definition of the CountMod problem, it can be shown that controlling $O(n)$ points also suffices. The idea is that in the construction of [11], the n counting queries are fixed in advance (Fourier queries) and so we could “hard-code” them into the definition of the problem. Let these queries be denoted as $q_1, \dots, q_n$, and let k be a parameter (think about $k \approx \frac{1}{\epsilon}$). Now let T be a dataset (multiset) containing elements from $[n] \cup \{\star\}$, and consider the following task. Let $\text{Star}(T)$ denote the number of $\star$'s in T . The goal is to approximate the value of q_i on $D = \{x \in T : x \neq \star\}$ for an integer i satisfying

$$\left\lfloor \frac{\text{Star}(T) - k}{3k} \right\rfloor \leq i \leq \left\lceil \frac{\text{Star}(T) + k}{3k} \right\rceil. \quad (1)$$

Note that this definition of the problem does not require the algorithm to compute $\text{Star}(T)$ exactly; it suffices to estimate it up to an error of k , which can be done privately provided that $k \approx \frac{1}{\epsilon}$. Also note that when $\text{Star}(T)$ is a multiple of $3k$, say $\text{Star}(T) = 3kr$, then there is exactly one “allowed” index i , namely $i = r$.

Now consider a dataset (multiset) $T = D \cup S \subseteq [n] \cup \{\star\}$, where D is a random subset of $[n]$ of size $n/2$ and where S is the multiset containing $3kn$ copies of $\star$. Now observe that by controlling S and deleting its elements one by one, the attacker gets (approximate) answers to all of the queries $q_n, q_{n-1}, \dots, q_1$ on the dataset D , and could hence reconstruct D unless the error is $\Omega(\sqrt{n})$. So now the attacker controls $O(n)$ points and reconstructs $O(n)$ other (“undeleted”) points.

Further reducing the number of deletions. We can further reduce the number of points that the attacker controls (and deletes) by modifying the task such that the algorithm is required to release approximate answers to *multiple* queries together, rather than just one. More specifically, as before, let $q_1, \dots, q_n$ denote a (fixed) set of n counting queries, and let $k, t \in \mathbb{N}$ be parameters where t divides n . Given a dataset (multiset) $T \subseteq [n] \cup \{\star\}$, the new goal is to approximate the value of each of $q_{(i-1)t+1}, \dots, q_{it}$ on $D = \{x \in T : x \neq \star\}$ for an integer i satisfying Inequality (1). As before, a valid value for i can be computed in a differentially private manner, provided that $k \gtrsim \frac{1}{\epsilon}$. In addition, as long as $t = o(n)$, a (single shot) DP algorithm can release answers to t queries with error $o(\sqrt{n})$. Thus, the problem can be solved with privacy in the single shot setting (without deletions). However, with deletions, it now suffices for the attacker to control merely $\omega(1)$ points. Specifically, consider a dataset (multiset) $T = D \cup S \subseteq [n] \cup \{\star\}$, where D is a random subset of $[n]$ of size $n/2$ and where S is the multiset containing $|S| = 3kn/t = \omega(1)$ copies of $\star$. Now observe that by

controlling S and deleting its elements one by one, the attacker gets (approximate) answers to all of the queries $q_n, q_{n-1}, \dots, q_1$ on the dataset D , and could hence reconstruct D unless the error is $\Omega(\sqrt{n})$. The upside here is that the attacker needs to control very few points ($\omega(1)$ points) and by deleting them the attacker can reconstruct a dataset of size n . This results in Theorem 2. The formal details are given in Appendix B.

2.2 Generalized differencing attacks

As we mentioned, definitions that aim to mimic/approximate the behavior of perfect retraining are susceptible to “differencing attacks”, such as the simple counting example, where seeing two consecutive releases (before and after a deletion) allows the attacker to learn the deleted point. We now present simple generalizations of this attack that extend to other types of definitions:

1. Consider a learning algorithm $\mathcal{A}$ that takes a dataset D containing n numbers, represented in binary (say using d bits). The algorithm ignores the data and returns a random binary vector of length d , denoted as z . Next, upon receiving a request to delete a point x , the algorithm deletes x from the data and returns $z \oplus x$. This, of course, allows the same “differencing attack”. However, note that now the marginal distribution of every single release given by the algorithm is simply uniform over d -bit vectors, independent of the data. This shows that definitions that only argue about every single time step individually, like the definition of Chourasia and Shah [6] mentioned above, can be vulnerable to differencing attacks even if they require that the distribution of every release is independent of the data.⁶
2. Somewhat artificially, a similar attack could be leveraged to reveal information about “undeleted” points. The attack begins the same, with the initial release being a random d -bit vector z . Next, upon receiving a request to delete a point x , the algorithm deletes x from the data, samples a random point y from the remaining data points, and returns $z \oplus y$.

3 Our new definition(s)

Consider a setting where a curator wishes to compute a functionality f on an initial dataset D while supporting ongoing deletion requests. We present security definitions that specifically safeguard undeleted data against leakage caused by the deletion of other points. We conceptually separate our focus from the following two types of privacy risks:

1. Our security definition does not address privacy risks that are inherent to the functionality f itself. This mirrors the paradigm of secure Multi-Party Computation (MPC) in cryptography: an MPC protocol guarantees that nothing is leaked beyond the output of the chosen function f , but whether evaluating f is acceptable from a privacy standpoint is left as an external, system-level choice. We adopt a similar philosophy, accepting that the initial release of $f(D)$, before any deletion is made, may inherently reveal information about D . We instead wish to limit the information subsequently revealed by deleting elements, not taking for granted the leakage that would come from recomputing f after each request.

⁶ For this we need to assume that either the curator is allowed to hold a state (the vector z), or that the deletion algorithm receives an input both the point to be deleted (the point x) and the previous release (the vector z in our case).

2. We separate the risks incurred to *deleted* points from the risks incurred to *undeleted* points. As we mentioned, the machine unlearning community is already well-aware that there are cases where, by deleting her data, Alice exposes herself to privacy risks (such as in a differencing attack). Nevertheless, Alice should have the right to make a calculated, informed decision to delete her data even in the face of such risks.

Our point is not that leakage via f itself or w.r.t. deleted points is unimportant. Rather, we argue that even if we accept these two types of leakage, there are still other critical risks that need to be addressed. Specifically, our work focuses on the scenario where some users' deletions compromise the privacy of the remaining users who never opted to delete their data. Our goal is to formulate a security definition guaranteeing that Bob's privacy is not harmed by other people deleting their data, beyond what is inherently revealed about Bob's data from the initial release and the value of the deleted points (which the adversary may learn).

3.1 Non-adaptive attacker

We start with the simplest variant of our definition that handles *non-adaptive* attackers who decide on the sequence of deletions ahead of time. To simplify the presentation, we do not distinguish between the “algorithm” and the “curator” and consider one stateful algorithm $\mathcal{A}$ with the following syntax. Initially $\mathcal{A}$ is executed on a dataset D and outputs an initial output $z_0 \leftarrow \mathcal{A}(D)$. Then, for k rounds $i = 1, 2, \dots, k$, algorithm $\mathcal{A}$ obtains a point $y_i \in D$ and returns a modified output $z_i \leftarrow \mathcal{A}(D, y_1, \dots, y_i)$. The mindset is that z_i represents the result of “deleting” y_i from D , but this is not part of the formal definition. We also write $(\mathcal{A}(D), \mathcal{A}(D, y_1), \mathcal{A}(D, y_1, y_2), \dots, \mathcal{A}(D, y_1, y_2, \dots, y_k))$ to denote the vector of outcomes obtained by running $\mathcal{A}$ on D and then “deleting” one by one the points $y_1, \dots, y_k$.⁷

► **Definition 5** (Non-adaptive variant). *Algorithm $\mathcal{A}$ is k -undeleted-safe if there is a simulator Sim such that for every dataset D of size n , and every sequence $(y_1, \dots, y_{k'})$ containing $k' \leq k$ distinct points from D it holds that*

$$\text{Sim}(\mathcal{A}(D), y_1, \dots, y_{k'}) \approx (\mathcal{A}(D), \mathcal{A}(D, y_1), \mathcal{A}(D, y_1, y_2), \dots, \mathcal{A}(D, y_1, y_2, \dots, y_{k'})).$$

In words, there is a simulator that takes only the initial output $\mathcal{A}(D)$ and the values of the deleted points $(y_1, \dots, y_{k'})$, and can simulate the entire sequence of updated outputs generated by $\mathcal{A}$. Intuitively, this formalizes the goal stated above: an attacker observing the entire sequence of unlearning updates learns nothing more about the remaining dataset than what they could already compute themselves from the initial release and the deleted values.

► **Remark 6.** Definition 5 can be instantiated with any notion of similarity between the distributions, such as statistical indistinguishability, computational indistinguishability, or the similarity notion used by differential privacy.

3.2 Static adaptive attacker

We now extend the definition to the case where an attacker $\mathcal{B}$ can adaptively decide on the *order* of deletions. The attacker is static in the sense that the points it controls (and is allowed to delete) are predetermined. We assume that there are no other deletions in the game. In addition, the attacker might possess some side-information about the other points

⁷ We stress that algorithm $\mathcal{A}$ can have an internal state, which is omitted from this vector notation for simplicity.

17:10 Protecting the Undeleted in Machine Unlearning

in the dataset D , which we denote as $\text{SideInfo}(D)$. Formally, for a dataset $D = (x_1, \dots, x_n)$, a set $Y \subseteq [n]$ of size $k' \leq k$ denoting the indices of the points in D controlled by $\mathcal{B}$, and a side-information function SideInfo , we consider the following game between algorithm $\mathcal{A}$ and the attacker $\mathcal{B}$:

StaticGame($\mathcal{A}, D=(x_1, \dots, x_n), \mathcal{B}, Y \subseteq [n], \text{SideInfo}$):

1. Let $z_0 \leftarrow \mathcal{A}(D)$.
2. The attacker $\mathcal{B}$ obtains the controlled points $D|_Y \triangleq \{(i, x_i) : i \in Y\}$, the side-information $\text{SideInfo}(D)$, and the initial output z_0 .
3. For $\ell = 1, 2, \dots, |Y|$:
 - a. The attacker $\mathcal{B}$ chooses an index $j_\ell \in Y \setminus \{j_1, j_2, \dots, j_{\ell-1}\}$
 - b. Algorithm $\mathcal{A}$ receives a deletion request for the point x_{j_ℓ} and returns an updated output z_ℓ , which is given to $\mathcal{B}$.
4. The attacker $\mathcal{B}$ outputs its view of the interaction, which includes its internal randomness, the indices of the controlled points Y , its side-information $\text{SideInfo}(D)$, and all the outputs it received from $\mathcal{A}$ during the interaction.

► **Definition 7** (Static adaptive variant). *Algorithm $\mathcal{A}$ is k -undeleted-safe if for every attacker $\mathcal{B}$ there is a simulator Sim such that for every side-information function SideInfo , every dataset D of size n , and every set of indices $Y \subseteq [n]$ of size $|Y| \leq k$ it holds that*

$$\text{Sim}(\mathcal{A}(D), D|_Y, \text{SideInfo}) \approx \text{StaticGame}(\mathcal{A}, D, \mathcal{B}, Y, \text{SideInfo}).$$

Note that any algorithm $\mathcal{A}$ that satisfies Definition 7 also satisfies Definition 5, i.e., Definition 7 is (perhaps) harder to obtain. This holds by taking $\mathcal{B}$ to be the attacker that fixes the ordering of the deletions from Y before the interaction begins, and setting SideInfo to be empty.

3.3 Dynamic adaptive attacker

In the previous two definitions, the set of indices/points that the attacker controlled was fixed ahead of time. We now handle the case where the attacker can dynamically choose the indices to control and delete. So now Y , the set of indices controlled by the attacker, is not fixed in advance. It is initially empty, and throughout the execution, the attacker dynamically decides on indices from $[n]$ to “corrupt”. Once corrupted, these indices are added to Y . Formally, given a dataset $D = (x_1, \dots, x_n)$ we consider the following game between a stateful algorithm $\mathcal{A}$ and an attacker $\hat{\mathcal{B}}$ with side-information function SideInfo . This game has two modes, parameterized by the bit $b \in \{0, 1\}$, where $b = 0$ corresponds to a “real world” execution in which the attacker learns the sequence of modified outputs of $\mathcal{A}$ after every deletion, and where $b = 1$ corresponds to an “ideal world” execution in which this information is hidden from the attacker.

DynamicGame_b($\mathcal{A}, \hat{\mathcal{B}}, D=(x_1, \dots, x_n), \text{SideInfo}$):

1. Let $\ell = 0$ and $z_0 \leftarrow \mathcal{A}(D)$.
2. The attacker $\hat{\mathcal{B}}$ obtains $\text{SideInfo}(D)$ and the initial output z_0 .
3. Initiate $Y = R = \emptyset$. % Here Y denotes the set of indices controlled by the attacker, and R denotes the indices it had already deleted.
4. While $|R| < k$ and the attacker $\hat{\mathcal{B}}$ has not ended the loop:

- a. Set $\ell \leftarrow \ell + 1$.
 - b. The attacker $\hat{\mathcal{B}}$ chooses a set of indices $I_\ell \subseteq [n]$, possibly empty, under the restriction that $|Y \cup I_\ell| \leq k$, and learns x_i for every $i \in I_\ell$. Set $Y \leftarrow Y \cup I_\ell$.
 - c. The attacker $\hat{\mathcal{B}}$ chooses a set of indices $J_\ell \subseteq Y \setminus R$, possibly empty. Set $R \leftarrow R \cup J_\ell$.
 - d. If $b = 0$ and $J_\ell \neq \emptyset$ then $\mathcal{A}$ receives a deletion request for the points in J_ℓ and returns an updated output z_ℓ , which is given to the attacker. If $b = 1$, do nothing.
 - e. The attacker $\hat{\mathcal{B}}$ chooses whether to continue or to end the loop.
5. The attacker $\hat{\mathcal{B}}$ outputs v . % In a “real world” execution, where $b = 0$, this v can be thought of as the view of the attacker $\hat{\mathcal{B}}$ at the end of the execution. However, in an “ideal world” execution, where $b = 1$, we think of $\hat{\mathcal{B}}$ as a *simulator* that aims to simulate some other attacker, and so the value v is not necessarily the view of $\hat{\mathcal{B}}$ itself, but rather a “simulated view” that should correspond to the view of the simulated attacker at the end of a real world execution.
 6. The output of the game is (Y, v) . % We include Y in the output of the game to force the simulator (when $b = 1$) to choose its corruptions similarly to the real world attacker (when $b = 0$). Otherwise, for example, the simulator could corrupt k indices and use them to simulate a real-world attacker that only corrupts $k/2$ indices.

► **Definition 8** (Dynamic adaptive variant). *Algorithm $\mathcal{A}$ is k -undeleted-safe if for every attacker $\mathcal{B}$ there is a simulator Sim such that for every side-information function SideInfo and every dataset D of size n it holds that*

$$\text{DynamicGame}_1(\mathcal{A}, \text{Sim}, D, \text{SideInfo}) \approx \text{DynamicGame}_0(\mathcal{A}, \mathcal{B}, D, \text{SideInfo}).$$

Note that any algorithm $\mathcal{A}$ that satisfies Definition 8 also satisfies Definition 7, i.e., Definition 8 is (perhaps) harder to obtain. This holds by taking $\mathcal{B}$ to be the attacker that fixes the set Y before the interaction begins.

► **Remark 9.** As we mentioned, in our definitions we accept that the *initial* release of $f(D)$, before any deletion is made, may inherently reveal information about D . We could strengthen our definitions by allowing the adversary to choose the timing of the first release it observes. In this scenario, the first release provided to the attacker would not necessarily be the global initial release. Instead, the attacker might “wake up” at a later point and request the current release, potentially after several deletions have already occurred. The requirement would then dictate that the attacker learns nothing beyond what can be deduced from this first observed release and the values of the points deleted subsequently.

4 Algorithms satisfying our definition and extensions

It is immediate from Definition 8 that it is satisfied by any *stateless* algorithm $\mathcal{A}$, where by *stateless* we mean that for any D and deleted points $y_1, \dots, y_k$ it holds that $(\mathcal{A}(D), \mathcal{A}(D, y_1), \dots, \mathcal{A}(D, y_1, y_2, \dots, y_k))$ is strictly determined by $(\mathcal{A}(D), y_1, \dots, y_k)$. This includes interesting functionalities like a bulletin board and exact summations. Note that $\mathcal{A}(D)$ does not have to be deterministic for this argument to go through.

► **Example 10.** Consider the algorithm $\mathcal{A}$ that takes a dataset containing n numbers from $[0, 1]$ and returns a DP estimation for their summation. Then, upon receiving a deletion request for a point x_i , the algorithm simply returns $(z - x_i)$ where z is the previously released sum. This algorithm is stateless, so it satisfies our new definition, and it provides DP w.r.t. undeleted points.

Note that this task can also be solved in the continual observation model using the classical “tree algorithm” of Dwork et al. [10], which would guarantee DP for both the deleted and the undeleted points. However, this provably requires error that scales with $\Omega(\log k)$ [9]. The simple algorithm described in Example 10 shows that if we only want to provide privacy for undeleted points, then a single DP calculation (the initial noisy summation) suffices, resulting in error independent of k .

It is also immediate from our definition that some algorithms/functionalities do *not* satisfy it. For example, as we mentioned, the exact median allows the attacker to learn undeleted points, and thus does not satisfy our new definition.⁸ So suppose that we are interested in computing such function f that is not directly compatible with our framework. The question then becomes: “*how much additional information is needed, beyond f , in order for the simulation to become possible?*”. For example, observe that if SideInfo contains all of D (the side-information available to the simulator and the attacker), then simulating the exact median, or any other functionality, becomes trivial. We aim to quantify this “additional leakage” that suffices for simulation, and minimize it. In particular, in some cases, it might be reasonable to allow for additional leakage that satisfies DP. We can argue about the additional leakage via two primary approaches:

- (1) We could explicitly give the simulator some additional leakage function g of the dataset D , in order for the simulation to become possible. This has the benefit of making the leakage explicit; hopefully allowing us to argue about it more easily and try to minimize it.
- (2) Instead (or maybe in addition) to considering such a leakage function g , it might also be useful to replace f with a *different* function $\hat{f}$ that approximates it in some sense while being more suited to our definitions.⁹ Computing $\hat{f}$ instead of f might require less, or no, additional leakage (in the form of g). However, this might introduce another type of leakage because in our definition we do not protect against leakage caused by the initial computation, which would now be according to $\hat{f}$ instead of f .

We formulate option (1) in the following definition.¹⁰

► **Definition 11** (Dynamic adaptive variant, with leakage). *Let g be a leakage function. Algorithm $\mathcal{A}$ is (k, g) -undeleted-safe with leakage if for every attacker $\mathcal{B}$ there is a simulator $\text{Sim}(\cdot)$ that takes an additional input and then plays in DynamicGame_1 such that the following holds. For every side-information function SideInfo and every dataset D of size n it holds that*

$$\text{DynamicGame}_1(\mathcal{A}, \text{Sim}(g(D)), D, \text{SideInfo}) \approx \text{DynamicGame}_0(\mathcal{A}, \mathcal{B}, D, \text{SideInfo}).$$

Note that any algorithm $\mathcal{A}$ that is k -undeleted-safe according to Definition 8 is also (k, g) -undeleted-safe according to Definition 11 with $g \equiv \perp$. On the other extreme, for $g(D) = D$ Definition 11 becomes trivial in the sense that *any* algorithm $\mathcal{A}$ would be (k, g) -undeleted-safe (this holds because if the simulator knows D then it can simulate all of the execution of DynamicGame_0). When g is “bounded”, say a DP function, then the definition might be considered reasonable.

⁸ To see this, consider a dataset D containing 10 copies of the point 0, one “middle” point w which is either $w = 1/4$ or $w = 3/4$, and 12 copies of the point 1. The exact median is 1. Suppose that the attacker controls one of the copies of 1 (so $k = 1$ in this example). By deleting this controlled point, the exact median becomes w , which is then revealed to the attacker. But without deletions the simulator does not get any information about w and hence could not conduct the simulation.

⁹ This is common practice in the literature on differential privacy: When designing private analogous to non-private computations, we typically compute noisy/approximate variants of the original functionality.

¹⁰ For concreteness, we state this variant of the definition only w.r.t. Definition 8. The same can be said w.r.t. Definitions 7 and 5.

4.1 Approximate median

Suppose that we are interested in an algorithm that initially takes a dataset D containing n elements from a domain X and releases an approximation z_0 to the median of D . Then, for k rounds, the algorithm receives the i -th deletion request and outputs a modified answer z_i . We evaluate the approximation of the median for the (remaining) data at any given moment using its *rank error*. That is, if the current dataset is D_i and the algorithm returns an answer z_i , then the error of z_i w.r.t. D_i is defined as

$$\text{error}_{D_i}(z_i) := |\{x \in D_i : x \text{ is strictly between } z_i \text{ and } \text{median}(D_i)\}|.$$

We aim to design an algorithm for approximating the median, while satisfying our security definition, that guarantee small worst-case error across all their $(k+1)$ releases. Naively, one could solve this task by first computing and releasing the exact median, and then ignoring all deletion requests. This stateless algorithm satisfies our security definition and guarantees error at most k deterministically. Alternatively, we could extend Example 10 to this setting as follows.

► **Example 12.** Consider the algorithm $\mathcal{A}$ that takes a dataset D containing n points from some ordered domain X and operates as follows. Initially, $\mathcal{A}$ computes a differentially private data structure P that provides approximate answers to range queries w.r.t. D . Specifically, for every $a \leq b \in X$ it ensures that $P(a, b) \approx |\{x \in D : a \leq x \leq b\}|$.¹¹ Note that this, in particular, allows us to compute an approximate median of the dataset D . Now, throughout the execution, we simply subtract deleted elements from the approximate counts given by P . Specifically, at any moment throughout the execution, we estimate the number of (remaining) elements in a range $[a, b]$ as $P(a, b) - \#_{\text{del}}(a, b)$, where $\#_{\text{del}}(a, b)$ denotes the (exact) number of deleted elements between a and b . As in Example 10, this provides DP w.r.t. undeleted points, without increasing the error introduced by P .

4.2 A recipe for undeleted-safe algorithms

Example 12 suggests the following informal recipe for fitting a functionality f to our framework:

- (1) Identify a set of sufficient statistics for the function f .
- (2) Estimate these statistics with differential privacy.
- (3) Throughout the execution, maintain these statistics as a function of the deleted points.

Of course, this recipe is only useful when the set of sufficient statistics we identify is compatible with differential privacy. It turns out that this is indeed the case for some functionalities that were studied by prior work on machine unlearning.

In particular, Cao and Yang [3] presented the following “SQ framework” for machine unlearning: Let f be any function which can be computed in the non-adaptive *statistical query* (SQ) model. That is, there exists a collection of predicates C such that for any dataset D , the value of $f(D)$ can be computed by post-processing $\{\sum_{x \in D} c(x)\}_{c \in C}$, i.e., the sum of each of the predicates in C on the elements in D . As Cao and Yang [3] showed, tracking $\{\sum_{x \in D} c(x)\}_{c \in C}$ throughout the execution allows us to maintain the current value of f without recomputing it from scratch. When $|C|$ is relatively small, this results in an efficient algorithm for perfect retraining. This SQ framework is particularly suited for our

¹¹There are several such constructions in the literature, e.g., by [1, 2, 20, 8].

definition, as computing the initial values of $\{\sum_{x \in D} c(x)\}_{c \in C}$ can be done in a DP manner (again, assuming that $|C|$ is small). Afterwards, as in Examples 10 and 12, we could update these values as a function of the deleted points alone.

References

- 1 Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *APPROX-RANDOM*, volume 8096 of *Lecture Notes in Computer Science*, pages 363–378. Springer, 2013. doi:10.1007/978-3-642-40328-6_26.
- 2 Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil P. Vadhan. Differentially private release and learning of threshold functions. In *FOCS*, pages 634–649. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.45.
- 3 Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pages 463–480. IEEE, 2015. doi:10.1109/SP.2015.35.
- 4 Nicholas Carlini, Matthew Jagielski, Chiyuan Zhang, Nicolas Papernot, Andreas Terzis, and Florian Tramèr. The privacy onion effect: Memorization is relative. *Advances in Neural Information Processing Systems*, 35:13263–13276, 2022.
- 5 Yeshwanth Cherapanamjeri, Sumegha Garg, Nived Rajaraman, Ayush Sekhari, and Abhishek Shetty. The space complexity of learning-unlearning algorithms. *CoRR*, abs/2506.13048, 2025. doi:10.48550/arXiv.2506.13048.
- 6 Rishav Chourasia and Neil Shah. Forget unlearning: Towards true data-deletion in machine learning. In *International conference on machine learning*, pages 6028–6073. PMLR, 2023. URL: <https://proceedings.mlr.press/v202/chourasia23a.html>.
- 7 Aloni Cohen, Adam Smith, Marika Swanberg, and Prashant Nalini Vasudevan. Control, confidentiality, and the right to be forgotten. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 3358–3372, 2023. doi:10.1145/3576915.3616585.
- 8 Edith Cohen, Xin Lyu, Jelani Nelson, Tamás Sarlós, and Uri Stemmer. Optimal differentially private learning of thresholds and quasi-concave optimization. In *STOC*, pages 472–482. ACM, 2023. doi:10.1145/3564246.3585148.
- 9 Edith Cohen, Xin Lyu, Jelani Nelson, Tamás Sarlós, and Uri Stemmer. Lower bounds for differential privacy under continual observation and online threshold queries. In *COLT*, volume 247 of *Proceedings of Machine Learning Research*, pages 1200–1222. PMLR, 2024. URL: <https://proceedings.mlr.press/v247/cohen24b.html>.
- 10 Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In *STOC*, pages 715–724. ACM, 2010. doi:10.1145/1806689.1806787.
- 11 Cynthia Dwork and Sergey Yekhanin. New efficient attacks on statistical disclosure control mechanisms. In *Annual International Cryptology Conference*, pages 469–480. Springer, 2008. doi:10.1007/978-3-540-85174-5_26.
- 12 Ji Gao, Sanjam Garg, Mohammad Mahmoody, and Prashant Nalini Vasudevan. Deletion inference, reconstruction, and compliance in machine (un) learning. *arXiv preprint*, 2022. arXiv:2202.03460.
- 13 Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. Formalizing data deletion in the context of the right to be forgotten. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 373–402. Springer, 2020. doi:10.1007/978-3-030-45724-2_13.
- 14 Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data deletion in machine learning. *Advances in neural information processing systems*, 32, 2019.

- 15 Jonathan Godin and Philippe Lamontagne. Deletion-compliance in the absence of privacy. In *2021 18th International Conference on Privacy, Security and Trust (PST)*, pages 1–10. IEEE, 2021. doi:10.1109/PST52912.2021.9647774.
- 16 Yuechun Gu, Jiajie He, and Keke Chen. Auditing approximate machine unlearning for differentially private models. In *ICDM*, pages 1253–1262. IEEE, 2025.
- 17 Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. *arXiv preprint*, 2019. arXiv:1911.03030.
- 18 Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. Adaptive machine unlearning. *Advances in Neural Information Processing Systems*, 34:16319–16330, 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/87f7ee4fdb57bdfd52179947211b7ebb-Abstract.html>.
- 19 Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam D. Smith. The price of differential privacy under continual observation. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 14654–14678. PMLR, 2023. URL: <https://proceedings.mlr.press/v202/jain23b.html>.
- 20 Haim Kaplan, Katrina Ligett, Yishay Mansour, Moni Naor, and Uri Stemmer. Privately learning thresholds: Closing the exponential gap. In *COLT*, volume 125 of *Proceedings of Machine Learning Research*, pages 2263–2285. PMLR, 2020. URL: <http://proceedings.mlr.press/v125/kaplan20a.html>.
- 21 Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982. doi:10.1109/TIT.1982.1056489.
- 22 Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pages 931–962. PMLR, 2021. URL: <http://proceedings.mlr.press/v132/neel21a.html>.
- 23 Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning. *Advances in Neural Information Processing Systems*, 34:18075–18086, 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/9627c45df543c816a3ddf2d8ea686a99-Abstract.html>.

A k -means clustering on the line

Consider a perfect retraining algorithm for k -means that, upon deletion, simply (re)runs Lloyd’s algorithm [21] on the remaining data points.¹² Furthermore, for the sake of this example, let us consider the simple case where $k = 2$ and the data points are on the real line.

We consider the following experiment. Initially, a dataset D containing n points is sampled from some fixed (known) distribution $\mathcal{D}$. Then, αn random points from the dataset D are given to an attacker $\mathcal{A}$. We refer to these points as the points that the attacker *controls*, and denote them as $A \subseteq D$. The game now continues in iteration where in every iteration:

1. The attacker $\mathcal{A}$ picks a point $a \in A \cap D$ and requests it to be deleted from D .
2. We remove a from D , rerun Lloyd’s algorithm on D to obtain a new tuple (c'_1, c'_2) and give this tuple to $\mathcal{A}$.

The goal of the attacker is to identify points which it does not control, i.e., discover points from $D \setminus A$. We show empirically that this can be done successfully for several choices for the underlying distribution $\mathcal{D}$, including the uniform distribution over an interval, and a mixture of two normal distributions.

¹²Recall that Lloyd’s algorithm for k -means is an iterative refinement procedure that alternates between assigning data points to the nearest cluster center and updating the cluster centers to be the mean of their assigned points. These iterative refinement procedure continues until convergence (when assignments no longer change).

The intuition behind the attack is as follows. Suppose that the attacker knows, in addition to the current cluster centers (c_1, c_2) , the *sizes* of the clusters (m_1, m_2) , where $m_1 + m_2 = |D|$. By deleting a point from A which is larger than c_2 , the attacker causes the right cluster to shift slightly to the left, thereby potentially causing points from the left cluster to “jump” to the right cluster. If it happens that exactly one point jumps between the clusters, then the attacker learns this point exactly, as it knows the previous and new averages (i.e., centers) of the left cluster (and we assumed that the attacker knows the cluster sizes). We show empirically that this indeed happens several times throughout the execution.¹³

It remains, therefore, to explain how the attacker can learn the sizes of the clusters from the cluster centers. Suppose that by deleting a point a^* from the right side of the line, the attacker caused both centers to change, from $(c_1^{\text{old}}, c_2^{\text{old}})$ to $(c_1^{\text{new}}, c_2^{\text{new}})$. Let $m_1^{\text{old}}, m_2^{\text{old}}, m_1^{\text{new}}, m_2^{\text{new}}$ denote the sizes of the corresponding clusters before and after the change. Let P denote the sum of the points that jumped from the left to the right cluster and let $p = m_1^{\text{old}} - m_1^{\text{new}} = m_2^{\text{new}} - m_2^{\text{old}} + 1$ denote the number of jumping points, where the $(+1)$ is because of the deleted point a^* . Also let $n = m_1^{\text{old}} + m_2^{\text{old}} = m_1^{\text{new}} + m_2^{\text{new}} + 1$ denote the size of the dataset before the deletion, which is known to the attacker. We have that

$$P = m_1^{\text{old}} \cdot c_1^{\text{old}} - m_1^{\text{new}} \cdot c_1^{\text{new}} = m_1^{\text{old}} \cdot c_1^{\text{old}} - (m_1^{\text{old}} - p) \cdot c_1^{\text{new}}.$$

Now, the sum of the points in the right cluster, after the change, can be written in two ways: Either

$$m_2^{\text{new}} \cdot c_2^{\text{new}} = (m_2^{\text{old}} - 1 + p) \cdot c_2^{\text{new}} = (n - m_1^{\text{old}} - 1 + p) \cdot c_2^{\text{new}}, \quad (2)$$

or it can be written as

$$m_2^{\text{old}} \cdot c_2^{\text{old}} + P - a^* = (n - m_1^{\text{old}}) \cdot c_2^{\text{old}} + m_1^{\text{old}} \cdot c_1^{\text{old}} - (m_1^{\text{old}} - p) \cdot c_1^{\text{new}} - a^*. \quad (3)$$

Equating (2) and (3) we get one equation with two unknown variables, m_1^{old} and p . Isolating p from this equation allows us to express p as a function of m_1^{old} as follows:

$$p = \frac{m_1^{\text{old}} \cdot (c_1^{\text{old}} - c_1^{\text{new}} - c_2^{\text{old}} + c_2^{\text{new}}) + n \cdot (c_2^{\text{old}} - c_2^{\text{new}}) + c_2^{\text{new}} - a^*}{c_2^{\text{new}} - c_1^{\text{new}}}$$

Now recall that both p and m_1^{old} must be integers. Heuristically, by trying different values for m_1^{old} until the resulting p is an integer, we succeed in finding the correct values for p and m_1^{old} , and thus learn the sizes of the clusters exactly. The code for this simulation is available on GitHub.

B An attacker controlling $\omega(1)$ points

In this section we give the formal details for an attack scenario in which the attacker controls (and deletes) merely $\omega(1)$ points and is still able to reconstruct $\approx n$ “uncontrolled” (or “undeleted”) points.

¹³ If more than one point jumps between the clusters, say r points, then the attacker could heuristically delete a point from the other side of the line, thereby causing j points to potentially jump in the other direction. The attack now works if $j \in \{1, r-1, r+1\}$ since by looking at two consecutive jumps of $r, r+1$ points the attacker could isolate a single point.

► **Definition 13** (The batch queries (BQ) problem). Let $n, k, t \in \mathbb{N}$ and $\alpha \geq 0, \beta \in [0, 1]$ be parameters, and let $Q = (q_1, \dots, q_n)$ be a collection of n counting queries over $[n]$. That is, every q is a predicate $q : [n] \rightarrow \{0, 1\}$. Let $T \subseteq [n] \cup \{\star\}$ be a multiset and let $\vec{a} \in \mathbb{R}^t$ be a real valued vector. We say that $\vec{a}$ is α -accurate for T if there exists an integer i such that

$$\left\lfloor \frac{\text{Star}(T) - k}{3k} \right\rfloor \leq i \leq \left\lfloor \frac{\text{Star}(T) + k}{3k} \right\rfloor,$$

and

$$\frac{1}{t} \|\vec{a} - (q_{(i-1)t+1}(T \setminus \{\star\}), \dots, q_{it}(T \setminus \{\star\}))\|_2^2 \leq \alpha.$$

We consider two types of algorithms:

- **The one shot setting.** Let $\mathcal{A}$ be an algorithm that takes a multiset $T \subseteq [n] \cup \{\star\}$ and returns a vector in $\mathbb{R}^t$. We say that this algorithm (α, β) -solves the one shot variant of the $\text{BQ}_{Q,n,k,t}$ problem if for every T , with probability at least $1 - \beta$ it holds that $\mathcal{A}(T)$ is α -accurate w.r.t. T .
- **The continual setting.** Let $r \in \mathbb{N}$ be a parameter. Consider an algorithm $\mathcal{A}$ that is initially instantiated with a multiset $T \subseteq [n] \cup \{\star\}$ and returns a vector $\vec{a}_0 \in \mathbb{R}^t$. Then, in every round $\ell = 1, 2, \dots, r$, algorithm $\mathcal{A}$ obtains a “deletion request” $z_\ell \in [n] \cup \{\star\}$ and returns an updated vector $\vec{a}_\ell \in \mathbb{R}^t$. We say that this algorithm (α, β) -solves the $\text{BQ}_{Q,n,k,t}$ problem over r rounds in the continual setting if for every T and every adaptive strategy for selecting the z_ℓ ’s, with probability at least $1 - \beta$, for every ℓ it holds that $\vec{a}_\ell$ is α -accurate w.r.t. $T \setminus \{z_1, \dots, z_\ell\}$.

► **Theorem 14.** Let $n, k, t \in \mathbb{N}$ and $\beta, \varepsilon, \delta \in [0, 1]$ be such that

$$k \geq \frac{2}{\varepsilon} \log \frac{2}{\beta} \quad \text{and} \quad t = o\left(\frac{\varepsilon^2 n}{\log(1/\delta)}\right) \quad \text{and} \quad n = \omega\left(\frac{\log(1/\delta) \log(1/\beta)}{\varepsilon^2}\right).$$

Then, for any set Q of counting queries, there exists an (ε, δ) -DP algorithm for (α, β) -solving the $\text{BQ}_{Q,n,k,t}$ problem in the one-shot setting for $\alpha = o(n)$.

Proof. The construction is straightforward via the Laplace or Gaussian mechanism. Specifically, given the multiset T , we first use the Laplace mechanism to find an appropriate index i . That is, we compute:

$$\hat{i} \leftarrow \frac{\text{Star}(T) + \text{Lap}(\frac{2}{\varepsilon})}{3k}.$$

This step satisfies $(\frac{\varepsilon}{2}, 0)$ -DP. Furthermore, the chosen index $\hat{i}$ is valid (according to Definition 13) whenever $|\text{Lap}(\frac{2}{\varepsilon})| \leq k$, which happens with probability at least $1 - \beta/2$ provided that $k \geq \frac{2}{\varepsilon} \log \frac{2}{\beta}$.

Given $\hat{i}$, all that remains is to answer to corresponding set of t queries from Q , which we again do via the Gaussian mechanism. Specifically, we answer each of these t queries with independent noise sampled from $\mathcal{N}(0, \sigma^2)$ for $\sigma = \frac{\sqrt{16t \log(1/\delta)}}{\varepsilon}$. By composition theorems for DP, this satisfies $(\frac{\varepsilon}{2}, \delta)$ -DP, and thus (again by composition), the overall algorithm satisfies (ε, δ) -DP as required.

Furthermore, by standard tail bounds for sums of independent squared normal random variables (aka the chi squared distribution), with probability at least $1 - \beta/2$, our average squared errors (denoted as α) satisfies

$$\alpha \leq \frac{\sigma^2 \left(t + 2\sqrt{t \ln(2/\beta)} + 2 \ln(2/\beta) \right)}{t} = \frac{16 \log(1/\delta) \left(t + 2\sqrt{t \ln(2/\beta)} + 2 \ln(2/\beta) \right)}{\varepsilon^2},$$

17:18 Protecting the Undeleted in Machine Unlearning

which is $o(n)$ provided that

$$t = o\left(\frac{\varepsilon^2 n}{\log(1/\delta)}\right) \quad \text{and} \quad n = \omega\left(\frac{\log(1/\delta) \log(1/\beta)}{\varepsilon^2}\right),$$

as required. $\blacktriangleleft$

► **Theorem 15.** *Let $n, k, t \in \mathbb{N}$ and $\alpha \geq 0$, $\beta \in [0, 1]$ be such that $\alpha \leq o(n)$. There exists a set of n counting queries Q and there exists an attacker $\mathcal{B}$ such that the following holds. For any algorithm $\mathcal{A}$ for (α, β) -solving the $\text{BQ}_{Q,n,k,t}$ problem in the continual setting over $r = \lceil n/t \rceil$ rounds, and for any set $D \subseteq [n]$, if $\mathcal{A}$ is instantiated on the multiset T containing D as well as $3kr$ copies of $\star$, and then interacts with $\mathcal{B}$ over r rounds (as in Definition 13), then with probability at least $1 - \beta$, at the end of the interaction the attacker $\mathcal{B}$ outputs a set $\hat{D}$ such that $|D \Delta \hat{D}| \leq \alpha$.*

► **Corollary 16.** *Combining Theorems 14 and 15, for $n, k, t \in \mathbb{N}$ satisfying $k \geq \omega(1)$, and $n = \omega(1)$, and $t = o(n)$, we get that*

1. *The $\text{BQ}_{Q,n,k,t}$ problem can be $(o(n), \beta)$ -solved in the one shot setting while satisfying $(o(1), o(1))$ -DP for any constant β .*
2. *Any algorithm for $(o(n), \beta)$ -solving the $\text{BQ}_{Q,n,k,t}$ problem in the continual setting allows an attacker controlling merely $\omega(1)$ points to reconstruct almost all of the other points from the dataset with probability at least $1 - \beta$.*

Proof of Theorem 15. The construction follows from the results of [11], who presented a set of n counting queries Q and an attacker $\mathcal{B}_{\text{DY}}$ such that for every set $D \subseteq [n]$, given approximate answers $a_q \approx q(D)$ for every q in Q satisfying

$$\frac{1}{n} \sum_{q \in Q} (a_q - q(D))^2 \leq o(n),$$

the attacker $\mathcal{B}_{\text{DY}}$ can efficiently reconstruct a set $\hat{D}$ with $|D \Delta \hat{D}| \leq o(n)$.

We implement this attacker in our setting as follows. Let T be the multiset containing D as well as $3kr$ copies of $\star$. Now consider an attacker $\mathcal{B}$ that plays against an algorithm $\mathcal{A}$ in the continual batch queries problem. Initially, $\mathcal{A}$ is instantiated with T . Then, in every round, the attacker $\mathcal{B}$ deletes $3k$ copies of $\star$. By the definition of the BQ problem, this allows $\mathcal{B}$ to obtain approximate answers a_q to every $q \in Q$ such that with probability at least $1 - \beta$ we have

$$\frac{1}{n} \sum_{q \in Q} (a_q - q(D))^2 \leq o(n).$$

This allows $\mathcal{B}$ to run $\mathcal{B}_{\text{DY}}$ to obtain a suitable reconstruction $\hat{D}$. $\blacktriangleleft$