MACHINE UNLEARNING OF FEDERATED CLUSTERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Federated clustering is an unsupervised learning problem that arises in a number of practical applications, including personalized recommender and healthcare systems. With the adoption of recent laws ensuring the "right to be forgotten", the problem of machine unlearning for federated clustering methods has become of significant importance. This work proposes the first known unlearning mechanism for federated clustering with privacy criteria that support simple, provable, and efficient data removal at the client and server level. The gist of our approach is to combine special initialization procedures with quantization methods that allow for secure aggregation of estimated local cluster counts at the server unit. As part of our platform, we introduce *secure compressed multiset aggregation (SCMA)*, which is of independent interest for secure sparse model aggregation. In order to simultaneously facilitate low communication complexity and secret sharing protocols, we integrate Reed-Solomon encoding with special evaluation points into the new SCMA pipeline and derive bounds on the time and communication complexity of different components of the scheme. Compared to completely retraining K-means++ locally and globally for each removal request, we obtain an average speed-up of roughly 84x across seven datasets, two of which contain biological and medical information that is subject to frequent unlearning requests.

1 INTRODUCTION

The availability of large volumes of user training data has contributed to the success of modern machine learning models. For example, most state-of-the-art computer vision models are trained on large-scale image datasets including Flickr (Thomee et al., 2016) and ImageNet (Deng et al., 2009). Organizations and repositories that collect and store user data must comply with privacy regulations, such as the recent European Union General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and the Canadian Consumer Privacy Protection Act (CPPA), all of which guarantee the right of users to remove their data from the datasets (*Right to be Forgotten*). Data removal requests frequently arise in practice, especially for sensitive datasets pertaining to medical records (numerous machine learning models in computational biology are trained using UK Biobank (Sudlow et al., 2015) which hosts a collection of genetic and medical records of roughly half a million patients (Ginart et al., 2019)). Removing user data from a dataset is insufficient to ensure sufficient privacy, since training data can often be reconstructed from trained models (Fredrikson et al., 2015; Veale et al., 2018). This motivates the study of *machine unlearning* (Cao & Yang, 2015) which aims to efficiently eliminate the influence of certain data points on a model. Naively, one can retrain the model from scratch to ensure complete removal, yet retraining comes at a high computational cost and is thus not practical when accommodating frequent removal requests. To avoid complete retraining, specialized approaches have to be developed for each unlearning application (Ginart et al., 2019; Guo et al., 2020; Bourtoule et al., 2021; Sekhari et al., 2021).

At the same time, federated learning (FL) has emerged as a promising approach to enable distributed training over a large number of users while protecting their privacy (McMahan et al., 2017; Chen et al., 2020; Kairouz et al., 2021; Wang et al., 2021; Bonawitz et al., 2021). The key idea of FL is to keep user data on their devices and train global models by aggregating local models in a communication-efficient and secure manner. Due to model inversion attacks (Zhu et al., 2019; Geiping et al., 2020), secure local model aggregation at the server is a critical consideration in FL, as it guarantees that the server cannot get specific information about client data based on their local models (Bonawitz et al., 2017; Bell et al., 2020; So et al., 2022; Chen et al., 2022). Since data privacy is the main goal in FL, it should be natural for a FL framework to allow for frequent data removal of a subset of

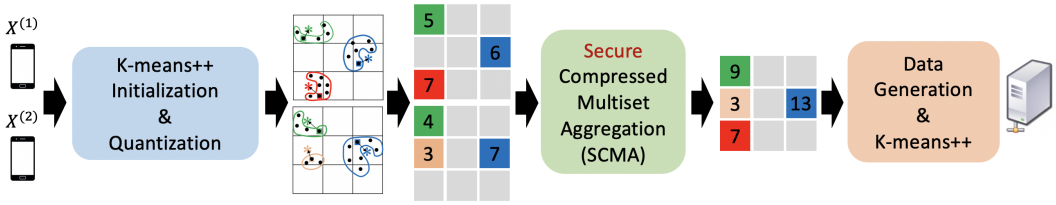

Figure 1: Overview of our proposed federated clustering scheme. K-means++ initialization and quantization are performed at each client in parallel. The SCMA procedure ensures that only the server knows the aggregated statistics of clients, without revealing who contributed the points in each individual cluster. The server generates points from the quantization bins with prescribed weights and performs full K-means++ clustering to infer the global model.

client data in a cross-silo setting (e.g., when several patients request their data to be removed in the hospital database), or the entire local dataset for clients in a cross-device setting (e.g., when users request apps not to track their data on their phones). This leads to the largely unstudied problem termed *federated unlearning* (Liu et al., 2021; Wu et al., 2022; Wang et al., 2022). However, existing federated unlearning methods do not guarantee exact or approximate unlearning (Ginart et al., 2019; Guo et al., 2020), they do not come with theoretical performance guarantees after model updates, and often, they are vulnerable to adversarial attacks.

Our contributions. We describe the first unlearning framework for a newly developed federated clustering approach (Fig. 1) which addresses the following problems:

- Unlearning of clusters at the client level via an intuitive application of K-means++ initialization without subsequent Lloyd's iterations. The advantage of using K-means++ is that the models do not need to be updated unless removed points were selected as initial centroids. At the server level, full K-means++ can be performed to improve the overall federated clustering performance.
- Enabling secure aggregation of sparse vectors via a novel approach termed *secure compressed multiset aggregation* (SCMA). The gist of SCMA is to encode quantized local data into information pertaining to the identifiers of quantization bins and their occupancies using specialized Reed-Solomon code evaluations, and followed by secure model aggregations.
- Ensuring privacy-accuracy-efficiency trade-offs through the introduction of a privacy criterion suitable for federated clustering and unlearning with simple communication protocols. In this context, we derive theoretical performance guarantees for the computational and communication complexity of model training and unlearning procedures.
- Compiling a collection of datasets for benchmarking unlearning of federated clusters, including two new datasets which contain anonymized cancer genomics and microbiome information that is subject to frequent unlearning requests. Experimental results reveal that our one-shot algorithm offers an average speed-up of roughly 84x compared to complete retraining across seven datasets.

2 Related Works

Machine unlearning. For *centralized* machine unlearning problems, two types of unlearning requirements were proposed in previous works: exact unlearning and approximate unlearning. For exact unlearning, the unlearned model is required to perform identically as a completely retrained model. To achieve this, Cao & Yang (2015) introduced distributed learners, Bourtoule et al. (2021) proposed sharding-based methods, Ginart et al. (2019) used quantization to eliminate the effect of removed data in clustering problems, and Chen et al. (2021) applied sharding-based methods to Graph Neural Networks. For approximate unlearning, the "differences" in behavior between the unlearned model and the completely retrained model should be appropriately bounded, similarly to what is done in the context of differential privacy. Following this latter direction, Guo et al. (2020) introduced the inverse Newton update for linear models, Sekhari et al. (2021) studied the generalization performance of approximately unlearned models, while Chien et al. (2022) extended the analysis to linearized Graph Neural Networks. A limited number of recent works also investigated data removal in the FL settings: Liu et al. (2021) proposed to use fewer iterations during retraining for federated unlearning, Wu et al. (2022) introduced Knowledge Distillation into the unlearning procedure to eliminate the effect of data points requested for removal, and Wang et al. (2022)

considered removing all data from one particular class via inspection of the internal influence of each channel in Convolutional Neural Networks. These federated unlearning methods are (mostly) empirical and do not come with theoretical guarantees for model performance after removal and/or for the unlearning efficiency. In contrast, our proposed federated clustering framework not only enables efficient data removal in practice, but also provides theoretical guarantees for the unlearned model performance and for the expected time complexity of the unlearning procedure.

Federated clustering. The idea of federated clustering is to perform clustering using data that resides at different edge devices. It is closely related to clustered FL (Sattler et al., 2020), whose goal is to learn several global models simultaneously, based on the cluster structure of the dataset, as well as personalization according to the cluster assignments of client data in FL (Mansour et al., 2020). One difference between federated clustering and distributed clustering (Guha et al., 2003; Ailon et al., 2009) is the assumption of data heterogeneity across different clients. Recent works (Ghosh et al., 2020; Dennis et al., 2021; Chung et al., 2022) exploit the non-i.i.d nature of client data to improve the performance of some learners. Another difference pertains to data privacy. Most previous methods were centered around the idea of sending exact (Dennis et al., 2021) or quantized client (local) centroids (Ginart et al., 2019) to the server, which may not be considered private as it leaks the data statistics or cluster information of all the clients. To avoid sending exact centroids, Li et al. (2022) proposes sending distances between data points and centroids to the server without revealing the membership of data points to any of the parties involved. Note that there is currently no formal definition of computational or information-theoretic secrecy/privacy for federated clustering problems, making it hard to compare methods addressing different aspects of FL. Our method introduces a simple-to-unlearn clustering process and new privacy mechanism that is intuitively appealing as it involves communicating obfuscated point counts of the clients to the server.

3 PRELIMINARIES

We start with a formal definition of the centralized K-means problem. Given a set of n points $\mathcal{X}$ arranged into a matrix $X \in \mathbb{R}^{n \times d}$, and the number of clusters K, the K-means problem asks for finding a set of points $\mathbf{C} = \{c_1, ..., c_K\}, c_k \in \mathbb{R}^d, \forall k \in [K]$ that minimizes the objective

$$\phi_c(X; \mathbf{C}) = \|X - C\|_F^2, \tag{1}$$

where $||\cdot||_F$ denotes the Frobenius norm of a matrix, $\|\cdot\|$ denotes the ℓ_2 norm of a vector, and $C \in \mathbb{R}^{n \times d}$ records the closest centroid in $\mathbf{C}$ to each data point x_i (i.e., $c_i = \arg\min_{c_j \in \mathbf{C}} \|x_i - c_j\|$). Without loss of generality, we make the assumption that the optimal solution is unique in order to facilitate simpler analysis and discussion, and denote the optimum by $\mathbf{C}^* = \{c_1^*, ..., c_K^*\}$. The set of centroids $\mathbf{C}^*$ induces an optimal partition $\bigcup_{k=1}^K \mathcal{C}_k^*$ over $\mathcal{X}$, where $\forall k \in [K], \mathcal{C}_k^* = \{x_i : ||x_i - c_k^*|| \leq ||x_i - c_j^*|| \; \forall i \in [n], j \in [K]\}$. We use $\phi_c^*(X)$ to denote the optimal value of the objective function for the centralized K-means problem. With a slight abuse of notation, we also use $\phi_c^*(\mathcal{C}_k^*)$ to denote the objective value computed over the optimal cluster $\mathcal{C}_k^*$. A detailed description of a commonly used approach for solving the K-means problem, K-means++, is available in Appendix A.

In federated clustering, the dataset $\mathcal{X}$ is no longer available at the centralized server. Instead, data is stored on L edge devices (clients) and the goal of federated clustering is to learn a global set of K centroids $\mathbf{C}_s$ over the union of client data. For simplicity, we assume that there exists no identical data points across clients, and that the overall dataset $\mathcal{X}$ is the union of the datasets $\mathcal{X}^{(l)}$ arranged as $X^{(l)} \in \mathbb{R}^{n^{(l)} \times d}$ on devices $l, \forall l \in [L]$. To make sure the global model can be learned securely, we allow all clients to communicate with each other as well as the server only to perform secure model aggregation (Bonawitz et al., 2017; Bell et al., 2020; So et al., 2022). The server will receive the aggregated cluster statistics of all clients, while each client will only know the cluster assignment of its own data $\mathcal{X}^{(l)}$ and not that of the global server clusters. In this case, the federated K-means problem asks for finding K global centroids $\mathbf{C}_s$ that minimize the objective

$$\phi_f(X; \mathbf{C}_s) = \sum_{l=1}^{L} \|X^{(l)} - C_s^{(l)}\|_F^2, \tag{2}$$

where $C_s^{(l)} \in \mathbb{R}^{n^{(l)} \times d}$ records the centroids of the global clusters that data points $\{x_i^{(l)}\}_{i=1}^{n^{(l)}}$ on client l belong to. A formal definition of global clusters is given in Definition 5.2. Here, ϕ_f stands for the objective of federated clustering, not to be confused with the centralized ϕ_c from Eq. (1).

Exact unlearning. For clustering problems, the *exact unlearning* criterion may be formulated as follows. Let $\mathcal{X}$ be a given dataset and $\mathcal{A}$ a (randomized) clustering algorithm that trains on $\mathcal{X}$ and outputs a set of centroids $\mathbf{C} \in \mathcal{M}$, where $\mathcal{M}$ is the chosen space of models. Let $\mathcal{U}$ be an unlearning algorithm that is applied to $\mathcal{A}(\mathcal{X})$ to remove the effects of one data point $x \in \mathcal{X}$. Then $\mathcal{U}$ is an exact unlearning algorithm if $\forall \mathbf{C} \in \mathcal{M}, x \in \mathcal{X}, \mathbb{P}(\mathcal{U}(\mathcal{A}(\mathcal{X}), \mathcal{X}, x) = \mathbf{C}) = \mathbb{P}(\mathcal{A}(\mathcal{X} \backslash x) = \mathbf{C})$. To avoid confusion, in certain cases, this criterion is referred to as *probabilistic (model) equivalence.*

Privacy-accuracy-efficiency trilemma. How to trade-off data privacy, model performance, communication and computational efficiency is a long-standing problem in distributed learning (Acharya & Sun, 2019; Chen et al., 2020; Gandikota et al., 2021) that also carries over to FL and federated clustering. Solutions that simultaneously address all these challenges in the latter context are still lacking. For example, Dennis et al. (2021) proposed a one-shot algorithm that takes model performance and communication efficiency into consideration by sending the *exact* centroids of each client to the server in a *nonanonymous* fashion. This approach may not be desirable under stringent privacy constraints as the server can gain information about individual client data. On the other hand, privacy considerations were addressed in Li et al. (2022) by performing K-means Lloyd's iterations anonymously via distribution of computations across different clients. Since the method relies on obfuscating pairwise distances for each client, it incurs computational overheads to hide the identity of contributing clients at the server and communication overheads due to interactive computations. None of the above methods is suitable for unlearning applications. To simultaneously enable unlearning and address the trilemma in the unlearning context, our privacy criterion involves transmitting *the number of client data points within local client clusters* in such a manner that the server cannot learn the data statistics of any specific client, but only the overall statistics of the union of client datasets. In this case, computations are limited and the clients on their end can perform efficient unlearning, unlike the case when presented with data point/centroid distances.

Random and adversarial removal. Most unlearning literature focuses on the case when all data points are equally likely to be removed, a setting known as *random removal*. However, adversarial data removal requests may arise when users are malicious in unlearning certain points that are critical for model training (i.e., boundary points in optimal clusters). We refer to such a removal request as *adversarial removal*. In Section 4, we provide theoretical analysis for both types of removal.

4 Machine Unlearning via Specialized Seeding

We exhibit next a useful connection between exact unlearning and the K-means++ initialization procedure. Throughout this section we consider the problem in the centralized setting so as to set the stage for the discussion of the federated clustering setting in Section 5. The first step is to introduce the unlearning mechanism based on the initial centroid set, as described in Alg. 1.

4.1 Performance Analysis

Since the centroids chosen through K-means++ initialization are true data points, the results returned by Alg. 1 are guaranteed to contain no information about the data points that have been removed. To verify that Alg. 1 is an exact unlearning method, we also need to check that $\mathbf{C}'$ is probabilistically equivalent to the models generated by rerunning the K-means++ initialization process on $\mathcal{X}'$, the set of point remaining after removal. This observation is best explained through an example, while a formal proof is provided in Appendix C.

Suppose that one would like to unlearn the data point x_i from the clustering model represented by an initial centroid set $\mathbf{C}$. For the first chosen centroid c_1 in $\mathbf{C}$, there are only two possibilities: $c_1 = x_i$ and $c_1 \neq x_i$. In the first case, we will have to re-run the initialization over $\mathcal{X}'$, which is retraining the model. In the second case, since we know $c_1 \neq x_i$, the probability of choosing c_1 from $\mathcal{X}$ as the first centroid becomes the conditional probability $\mathbb{P}(\text{choose } c_1 \text{ from } \mathcal{X} \text{ as the first centroid} | c_1 \neq x_i) = \frac{1}{n-1} = \mathbb{P}(\text{choose } c_1 \text{ from } \mathcal{X}' \text{ as the first centroid})$, where n is the number of data points in $\mathcal{X}$. Therefore, the first centroid c_1' returned by Alg. 1 can always be viewed as the first centroid obtained by re-running the initialization over $\mathcal{X}'$. This can be generalized for c_j $\forall j \in [K]$ and arbitrary $\mathcal{X}_R$. Since Alg. 1 is an exact unlearning method, we can have the performance guarantees in Lemma 4.1 based on Theorem 1.1 of Vassilvitskii & Arthur (2006).

Lemma 4.1. *For any set of data points $\mathcal{X}$ and removal set $\mathcal{X}_R$, assuming that the remaining dataset is $\mathcal{X}'$ and the centroid set returned by Alg. 1 is $\mathbf{C}'$, we have $\mathbb{E}(\phi_c(X'; \mathbf{C}')) \leq 8(\ln K + 2)\phi_c^*(X')$.*

4.2 COMPLEXITY ANALYSIS

Algorithm 1 Unlearning via K-means++ Init.

input: Dataset $\mathcal{X}$, centroid set $\mathbf{C}$ obtained by K-means++ initialization on $\mathcal{X}$, removal request set $\mathcal{X}_R = \{x_{r_1}, \ldots, x_{r_R}\}$.

if $c_j \notin \mathcal{X}_R \; \forall c_j \in \mathbf{C}$ **then**
  $\mathbf{C}' \leftarrow \mathbf{C}$
else
  $i \leftarrow (\arg\min_j c_j \in \mathcal{X}_R) - 1$
  if $i = 0$ **then**
    $\mathbf{C}' \leftarrow \varnothing, \mathcal{X}' \leftarrow \mathcal{X} \backslash \mathcal{X}_R$.
  else
    $\mathbf{C}' \leftarrow \{c_1, \ldots, c_i\}, \mathcal{X}' \leftarrow \mathcal{X} \backslash \mathcal{X}_R$.
  end if
  for $j = i+1, \ldots, K$ **do**
    Sample x from $\mathcal{X}'$ with prob $\frac{d^2(x, \mathbf{C}')}{\phi_c(X'; \mathbf{C}')}$.
    $\mathbf{C}' \leftarrow \mathbf{C}' \cup \{x\}$.
  end for
end if
return $\mathbf{C}'$

We show next that Alg. 1 supports removing R random points in expected time $O(RK^2d)$ and R adversarial points in expected time $O(RK^3\epsilon_1\epsilon_2 d)$. In contrast, the corresponding time complexity of retraining is $O(nKRd)$. To analyze the removal time complexity of Alg. 1, we first state two assumptions pertaining to optimal cluster sizes and outliers.

Assumption 4.2. *Let $\epsilon_1 = \frac{n}{K s_{\min}}$ be a constant denoting what we refer to as* cluster size imbalance, *where $s_{\min}$ equals the size of the smallest cluster in the optimal clustering; when $\epsilon_1 = 1$, all clusters are of the same size $\frac{n}{K}$.*

Assumption 4.3. *Assume that $\epsilon_2 \geq 1$ is a fixed constant. An outlier x_i in $\mathcal{X}$ satisfies*

$$\|x_i - c_j^*\| \leq \|x_i - c_k^*\|, \forall k \in [K]$$
$$\textit{and} \quad \|x_i - c_j^*\| > \sqrt{\frac{\epsilon_2 \phi_c^*(\mathcal{C}_j^*)}{|\mathcal{C}_j^*|}}. \tag{3}$$

The assumption is not overly restrictive, as outliers are commonly defined as points that are well-separated from all "regular" clusters: the distance between an outlier and its closest centroid is larger than a scaled proxy for the empirical standard deviation of the average distance between the cluster points and that centroid (Chien et al., 2018; Dennis et al., 2021). There also exist many different approaches for removing outliers for clustering purposes (Chawla & Gionis, 2013; Gan & Ng, 2017), which effectively reduce the probability of outliers to negligible values.

Assumptions 4.2 and 4.3 can be used to derive an expression for the probability of $x_i \in \mathbf{C}$ when x_i needs to be unlearned, which equals the probability of retraining with $\mathcal{X}'$.

Lemma 4.4. *Assume that the number of data points in $\mathcal{X}$ is n and the probability of the dataset containing at least one outlier is upper bounded by $O(1/n)$. Let $\mathbf{C}$ be the centroid set obtained by running K-means++ on $\mathcal{X}$. For unlearning a data point $x_i \in \mathcal{X}$ we have that for random removal: $\mathbb{P}(x_i \in \mathbf{C}) < O(K/n)$; for adversarial removal: $\mathbb{P}(x_i \in \mathbf{C}) < O\left(K^2\epsilon_1\epsilon_2/n\right)$.*

Based on Alg. 1, we need to reinitialize the centroids only when the point that requests removal lie in the original centroid set. We need $O(K)$ time to check if this is the case, and we need extra $O(nd)$ time to sample every new centroid due to the distance computation. Therefore, we arrive at an estimate of the expected removal time as shown in Theorem 4.5, which is independent of the problem size n.

Theorem 4.5. *Assume that the number of data points in $\mathcal{X}$ is n and the probability of the dataset containing at least one outlier is upper bounded by $O(1/n)$. Algorithm 1 supports removing R points with time $O(RK^2d)$ in expectation for random removals, and with time $O(RK^3\epsilon_1\epsilon_2 d)$ in expectation for adversarial removals. The time complexity of complete retraining equals $O(nKRd)$.*

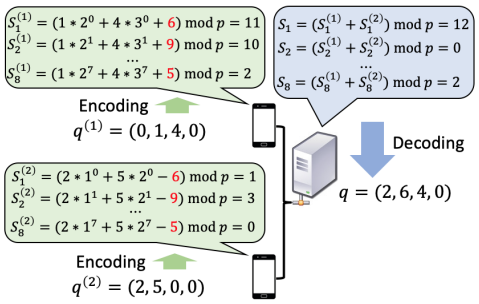

Figure 2: Example of the SCMA procedure for $K = 2, L = 2, B^d = 4, n = 12, p = 13$.

5 UNLEARNING FEDERATED CLUSTERS

We introduce next our federated clustering framework (Alg. 2) based on the discussion from Section 4. The removal procedure is described in Alg. 4. We introduce an oversampling coefficient α for clients,

following a common strategy for K-means clustering (Bahmani et al., 2012) to improve the final federated clustering performance. In practice, a constant value α suffices, and for simplicity, we set $\alpha = 1$ in our subsequent analysis. For Alg. 2 to satisfy the privacy criterion stated in Section 3, we need to design an efficient protocol for aggregating information at the server pertaining to the overall dataset without leaking individual client data statistics. This approach is discussed in Section 5.1. A simplified version of Alg. 2 is discussed in Appendix J, for applications where the privacy criterion is not an imperative concern.

Algorithm 2 Secure Federated Clustering

1: **input:** Dataset $\mathcal{X}$ distributed on L clients $(\mathcal{X}^{(1)}, \ldots, \mathcal{X}^{(L)})$, client oversampling coefficient α.
2: Run (αK)-means++ initialization on each client l in parallel, obtain the initial centroid sets $\mathbf{C}^{(l)}$, and record the corresponding cluster sizes $(|\mathcal{C}_1^{(l)}|, \ldots, |\mathcal{C}_{\alpha K}^{(l)}|)$, $\forall l \in [L]$.
3: Perform uniform quantization of $\mathbf{C}^{(l)}$ on each dimension, and flatten the quantization bins into a vector $q^{(l)}, \forall l \in [L]$.
4: Set $q_j^{(l)} = \left|\mathcal{C}_k^{(l)}\right|$ with j being the index of the quantization bin where $c_k^{(l)}$ lies in for $\forall k \in [\alpha K]$. Set $q_j^{(l)} = 0$ for all other indices.
5: Securely sum up $q^{(l)}$ at server by Algorithm 3, with the aggregated vector denoted as q.
6: For index $j \in \{t : q_t \neq 0\}$, sample q_j points based on pre-defined distribution and denote their union as new dataset $\mathcal{X}_s$ at server.
7: Run full K-means++ clustering at server with $\mathcal{X}_s$ to obtain the centroid set $\mathbf{C}_s$ at server.
8: **return** Each client retains its own centroid set $\mathbf{C}^{(l)}$, server retains $\mathcal{X}_s, q$ and $\mathbf{C}_s$.

Note that Step 6 of Alg. 2 involves generating q_j points for the j-th quantization bin. The simplest idea is to choose the center of the quantization bin as the representative point and assign weight q_j to it. Then, in Step 7, we can use the weighted K-means++ algorithm at the server to further reduce the computational complexity at the server side, by reducing the problem size at the server from n, which is the total number of data points, to KL. However, in practice we find that when the server computational power is not the bottleneck in the FL system, generating data points uniformly at random for the quantization bins with non-zero weights can often lead to better clustering performance. Thus, this is the default choice in our subsequent simulations.

5.1 SCMA at the Server

To utilize off-the-shelf secure aggregation methods developed for FL (Bonawitz et al., 2017; Bell et al., 2020; So et al., 2022), we first need to quantize the centroids of client data and convert the information into vectors. For simplicity, we use uniform quantization with step size γ for each dimension in Step 3 of Alg. 2. A detailed description of the quantization process is available in Appendix I. Once the vector representations $q^{(l)}$ of length B^d $(B = 1/\gamma)$, $l \in [L]$ of client data are generated, we can use standard secure model aggregation methods (Bonawitz et al., 2017; Bell et al., 2020; So et al., 2022) to sum up all $q^{(l)}$ securely and obtain the aggregated results q at the server. However, since the length of each $q^{(l)}$ is B^d, securely aggregating the whole vector would lead to very high communication complexity for each client if γ is small (i.e., if we need finer quantization) or d is large (i.e., if we have to cluster high dimensional data). Moreover, each $q^{(l)}$ is a sparse vector since the number of client centroids is much smaller than the number of quantization bins (i.e., $K \ll B^d$). It is inefficient and unnecessary to sum up $q^{(l)}$ for all dimensions. This motivates us to first compress the vectors and then perform the secure aggregation, and we refer to this process as SCMA.

By observing that there can be at most K nonzero entries in $q^{(l)}, \forall l \in [L]$ and at most KL nonzero entries in q, we invoke the Reed-Solomon code construction (Reed & Solomon, 1960) for designing SCMA. Let $\mathbb{F}_p = \{0, 1, \ldots, p-1\}$ be a finite field of prime order $p \geq \max\{n, B^d\}$. We treat the indices of the quantization bins as distinct elements from the underlying finite field, and use them as evaluation points of the encoder polynomial. In addition, we treat a nonzero entry $q_j^{(l)}$ in vector $q^{(l)}$ as a substitution error at the j-th entry in a codeword. Then, we use our SCMA scheme shown in Alg. 3, where the messages that the clients send to server can be treated as syndromes in Reed-Solomon decoding. The discussion about the differences between SCMA and a related problem termed private set union is available in Appendix M. Note that in our scheme, the server does not know $q^{(l)}, l \in [L]$ beyond the fact that $\sum_{l \in [L]} q^{(l)} = q$, which fits into our privacy criterion. This follows because $z_i^{(l)}$ is uniformly distributed over $\mathbb{F}_p$ and independently chosen for different $l \in [L]$ and $i \in [2KL]$. For details, please refer to Appendix K. One example of SCMA procedure is illustrated in Fig. 2.

5.2 PERFORMANCE ANALYSIS

Algorithm 3 SCMA

1: **input:** L different vectors $q^{(l)}$ of length B^d to be securely aggregated, a finite field $\mathbb{F}_p$.
2: Each client $l \in [L], i \in [2KL]$ communicates $\left(S_1^{(l)}, \ldots, S_{2KL}^{(l)}\right)$ to the server, where $S_i^{(l)} = (\sum_{j:q_j^{(l)} \neq 0} q_j^{(l)} \cdot j^{i-1} + z_i^{(l)}) \bmod p$, and $z_i^{(l)}$ is a random key uniformly distributed over $\mathbb{F}_p$ and hidden from the server. The keys $\{z_i^{(l)}\}_{l \in [L], i \in [2KL]}$ are generated offline using standard secure model aggregation so that $\sum_l z_i^{(l)} = 0$.
3: The server first computes the sum $S_i = \sum_{l \in [L]} S_i^{(l)}, \forall i \in [2KL]$. Given S_i, the server computes the coefficients of the polynomial $g(x) = \prod_{j:q_j \neq 0}(1 - j \cdot x)$ using the Berlekamp-Massey algorithm (Berlekamp, 1968; Massey, 1969). Then, the server factorizes $g(x)$ over the field $\mathbb{F}_p$ to determine the roots j^{-1}, $q_j \neq 0$, using the polynomial factorizing algorithm (Kedlaya & Umans, 2011). Finally, the server solves a set of $2KL$ linear equations $S_i = \sum_{l \in [L]} S_i^{(l)} = \sum_{j:q_j \neq 0} q_j \cdot j^{i-1}$ for $i \in [2KL]$, by considering q_j as unknowns and j^{i-1} as known coefficients for $q_j \neq 0$.
4: **return** q reconstructed at the server.

We describe next the performance guarantees of Alg. 2 w.r.t. the objective defined in Eq. (2), using the definition of global clusters from Definition 5.2 originally proposed in Dennis et al. (2021). Note that only the global clustering information can be sent to clients by the server, if requested (i.e., global centroids will not be interactively communicated between clients and server). We first consider the case where no quantization is used (see Alg. 5). The proof is relegated to Appendix G.

Lemma 5.1. *Suppose that the entire dataset across clients is denoted by $\mathcal{X}$, and the set of server centroids returned by Alg. 5 is $\mathbf{C}_s$. Then $\mathbb{E}\left(\phi_f(X; \mathbf{C}_s)\right) < O(\log^2 K) \cdot \phi_c^*(X)$.*

The result of Lemma 5.1 holds for any distribution of data across clients, and Alg. 5 may perform even better than for the case of homogeneous assignments. This is best explained through the following example. In the extreme case of data heterogeneity, each client stores a different cluster w.r.t. the global optimal clustering ($L = K$). Then Alg. 5 can be viewed as seeding each optimal cluster by a data point uniformly at random chosen from that cluster, through 1-means clustering at each client. This leads to a performance guarantee of $\mathbb{E}\left(\phi_f(X; \mathbf{C}_s)\right) = 2\phi_c^*(X)$, where the approximation factor is independent of K. A more detailed discussion how data heterogeneity can benefit our unlearning method is available in Appendix G.

Definition 5.2. *Suppose that the local clusters at client l are denoted by $\mathcal{C}_k^{(l)}, \forall k \in [K], l \in [L]$, and that the clusters at the server are denoted by $\mathcal{C}_k^s, \forall k \in [K]$. The global clustering equals $\mathcal{P}_k = \left\{x_i^{(l)} | x_i^{(l)} \in \mathcal{C}_j^{(l)}, c_j^{(l)} \in \mathcal{C}_k^s, \forall j \in [K], l \in [L]\right\}$, where $c_j^{(l)}$ is the centroid of $\mathcal{C}_j^{(l)}$. Note that $(\mathcal{P}_1, \ldots, \mathcal{P}_K)$ forms a partition of the entire dataset $\mathcal{X}$, and the representative centroid for $\mathcal{P}_k$ is defined as $c_{s,k}$ in $\mathbf{C}_s$ returned by Alg. 2.*

Next, we consider the distortion introduced by quantization. The performance guarantee is given in Theorem 5.3. Note that SCMA does not contribute to the distortion as it always returns the exact sum. Furthermore, following the same reasoning as in Lemma 4.1, one can show that the result of Theorem 5.3 also holds for the remaining dataset $\mathcal{X}'$ after running the unlearning procedure in Alg. 4.

Theorem 5.3. *Suppose that we performed uniform quantization with step size γ in Algorithm 2. Then we have $\mathbb{E}\left(\phi_f(X; \mathbf{C}_s)\right) < O(\log^2 K) \cdot \phi_c^*(X) + O(nd\gamma^2 \log K)$.*

Remark. *To make the second term in the upper bound a constant w.r.t. n, we can choose $\gamma = \Theta(1/\sqrt{n})$, which is a good choice in practice offering good performance across different datasets.*

5.3 COMPLEXITY ANALYSIS

We derived a cohort of in-depth analysis pertaining to the computational complexity of both types of data removals, SCMA at the client and server side, as well as the communication complexity at the clients. Due to space limitations, these results are summarized in Appendix B.

6 EXPERIMENTAL RESULTS

Algorithm 4 Unlearning of Federated Clusters

1: **input:** Dataset $\mathcal{X}$ distributed on L clients $(\mathcal{X}^{(1)}, \ldots, \mathcal{X}^{(L)})$, $(\mathbf{C}^{(l)}, \mathcal{X}_s, q, \mathbf{C}_s)$ obtained by Algorithm 2 on $\mathcal{X}$, removal request set $\mathcal{X}_R^{(l)}$ for single-client removal or $\mathcal{L}_R$ for multi-client removal.
2: **if** single-client removal **then**
3: Run Algorithm 1 on client l and update $q^{(l)}$ if client l has to perform retraining.
4: **else**
5: $q^{(l)} \leftarrow \mathbf{0}$ on client $l, \forall l \in \mathcal{L}_R$.
6: **end if**
7: Securely sum up $q^{(l)}$ at server by Algorithm 3, with the aggregated vector denoted as q'.
8: **if** $q' = q$ **then**
9: $\mathbf{C}'_s \leftarrow \mathbf{C}_s$.
10: **else**
11: Generate $\mathcal{X}'_s$ with q'.
12: Run full K-means++ at the server with $\mathcal{X}'_s$ to obtain $\mathbf{C}'_s$.
13: **end if**
14: **return** Client centroid sets $\mathbf{C}^{(l)\prime}$, server data $\mathcal{X}'_s, q'$ and centroids $\mathbf{C}'_s$.

To empirically characterize the trade-off between the efficiency of data removal and performance of our newly proposed federated clustering method, we compare it with baseline methods on both synthetic and real datasets. Due to space limitations, more in-depth experiments and discussions are delegated to Appendix N.

Datasets and baselines. We use one synthetic dataset generated by a Gaussian Mixture Model (Gaussian) and six real datasets (Celltype, Covtype, FEMNIST, Postures, TMI, TCGA) in our experiments. We preprocess the datasets such that the data distribution is non-i.i.d. across different clients. The symbol K' in Fig. 3 represents the maximum number of (true) clusters among clients, while K represents the number of true clusters in the global dataset. A detailed description of the data statistics and the preprocessing procedure is available in Appendix N. Since there is currently no off-the-shelf algorithm designed for unlearning federated clusters, we adapt DC-Kmeans (DC-KM) from Ginart et al. (2019) to apply to our problem setting, and use complete retraining as the baseline comparison method. To evaluate federated clustering performance on the complete dataset (before data removals), we also include the K-FED algorithm from Dennis et al. (2021) as the baseline method. In all plots, our Alg. 2 is referred to as MUFC. Note that in FL, clients are usually trained in parallel so that the estimated time complexity equals the sum of the longest processing time of a client and the processing time of the server.

Clustering performance. The clustering performance of all methods on the complete dataset is shown in the first row of Tab. 1. The loss ratio is defined as $\phi_f(X; \mathbf{C}_s)/\phi_c^*(X)$[1], which is the metric used to evaluate the quality of the obtained clusters. For the seven datasets, MUFC offered the best performance on TMI and Celltype, datasets for which the numbers of data points in different clusters are highly imbalanced. This can be explained by pointing out an important difference between MUFC and K-FED/DC-KM: the quantized centroids sent by the clients may have non-unit weights, and MUFC is essentially performing weighted K-means++ at the server. In contrast, both K-FED and DC-KM assign equal unit weights to the client's centroids. Note that assigning weights to the client's centroids based on local clusterings not only enables a simple analysis of the scheme but also improves the empirical performance, especially for datasets with highly imbalanced cluster distributions. For all other datasets except Gaussian, MUFC obtained competitive clustering performance compared to K-FED/DC-KM. The main reason why DC-KM outperforms MUFC on Gaussian data is that all clusters are of the same size in this case. Also note that DC-KM runs full K-means++ clustering for each client while MUFC only performs initialization. Although running full K-means++ clustering at the client side can improve the empirical performance on certain datasets, it also greatly increases the computational complexity during training and the retraining probability during unlearning, which is shown in Fig. 3. Nevertheless, we also compare the performance of MUFC with K-FED/DC-KM when running full K-means++ clustering on clients for MUFC in Appendix N.

We also investigated the influence of K' and γ on the clustering performance. Fig. 3(a) shows that MUFC can obtain a lower loss ratio when $K' < K$, indicating that data is non-i.i.d. distributed across clients. Fig. 3(b) shows that the choice of γ does not seem to have a strong influence on the clustering performance of Gaussian datasets, due to the fact that we use uniform sampling in Step 6 of Alg. 2 to generate the server dataset. Meanwhile, Fig. 3(c) shows that γ can have a significant influence on the clustering performance of real-world datasets, which agrees with our analysis in Theorem 5.3.

[1] $\phi_c^*(X)$ is approximated by running K-means++ multiple times and selecting the smallest objective value.

Table 1: Clustering performance of different federated clustering algorithms compared to centralized K-means++ clustering.

		TMI	Celltype	Gaussian	TCGA	Postures	FEMNIST	Covtype
Loss ratio	MUFC	1.24 ± 0.10	1.14 ± 0.03	1.25 ± 0.02	1.18 ± 0.05	1.10 ± 0.01	1.20 ± 0.00	1.03 ± 0.02
	K-FED	1.84 ± 0.07	1.72 ± 0.24	1.25 ± 0.01	1.56 ± 0.11	1.13 ± 0.01	1.21 ± 0.00	1.60 ± 0.01
	DC-KM	1.54 ± 0.13	1.46 ± 0.01	1.02 ± 0.00	1.15 ± 0.02	1.03 ± 0.00	1.18 ± 0.00	1.03 ± 0.02
Speed-up of MUFC (if no retraining is performed)		151x	1535x	2074x	483x	613x	53x	267x

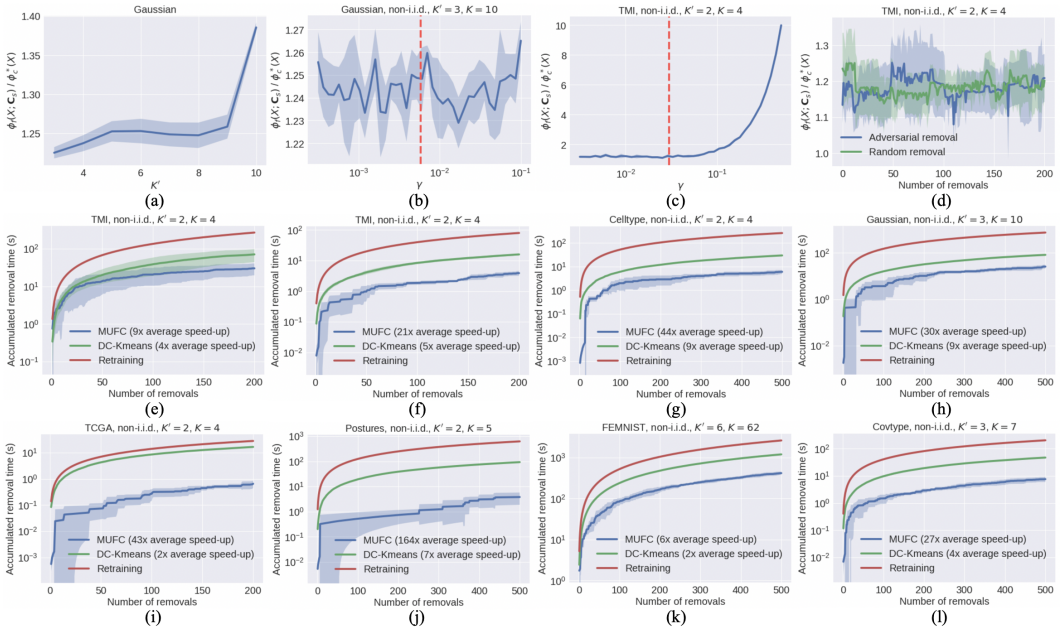

Figure 3: The shaded areas represent the standard deviation of results from different trails. (a) Influence of data heterogeneity on the clustering performance of MUFC: K' represents the maximum number of (global) clusters covered by the data at the clients, while $K' = 10$ indicates that the data points are i.i.d. distributed across clients. (b)(c) Influence of the quantization step size γ on the clustering performance of MUFC. The red vertical line indicates the default choice of $\gamma = 1/\sqrt{n}$, where n is the total number of data points across clients. (d) The change in the loss ratio after each round of unlearning. (e) The accumulated removal time for adversarial removals. (f)-(l) The accumulated removal time for random removals.

The red vertical line in both figures indicates the default choice of $\gamma = 1/\sqrt{n}$, where n stands for the number of total data points across clients.

Unlearning performance. Since K-FED does not support data removal, has high computational complexity, and its empirical clustering performance is worse than DC-KM (see Tab. 1), we only compare the unlearning performance of MUFC with that of DC-KM. For simplicity, we consider removing one data point from a uniformly at random chosen client l at each round of unlearning. The second row of Tab. 1 records the speed-up ratio w.r.t. complete retraining for one round of MUFC unlearning (Alg. 4) when the removed point does not lie in the centroid set selected at client l. Fig. 3(e) shows the accumulated removal time on the TMI dataset for adversarial removals, which are simulated by removing the data points with the highest contribution to the current value of the objective function at each round, while Fig. 3(f)-(l) shows the accumulated removal time on different datasets for random removals. The results show that MUFC maintains high unlearning efficiency compared to all other baseline approaches, and offers an average speed-up ratio of 84x when compared to complete retraining for random removals across seven datasets. We also report the change in the loss ratio of MUFC during unlearning in Fig. 3(d). The loss ratio remains nearly constant after each removal, indicating that our unlearning approach does not significantly degrade clustering performance. Similar conclusions hold for other tested datasets, as shown in Appendix N.

7 ETHICS STATEMENT

The seven datasets used in our simulations are all publicly available. Among these datasets, TCGA and TMI contain potentially sensitive biological data and are downloaded after logging into the database. We adhered to all regulations when handling this anonymized data and will only release the data processing pipeline and data that is unrestricted at TCGA and TMI. Datasets that do not contain sensitive information can be downloaded directly from their open-source repositories.

8 REPRODUCIBILITY STATEMENT

We upload our source code along with test datasets as supplementary materials to reproduce our results shown in Section 6. Detailed instructions are included as well in the source code.

REFERENCES

Jayadev Acharya and Ziteng Sun. Communication complexity in locally private distribution estimation and heavy hitters. In *International Conference on Machine Learning*, pp. 51–60. PMLR, 2019.

Nir Ailon, Ragesh Jaiswal, and Claire Monteleoni. Streaming k-means approximation. *Advances in neural information processing systems*, 22, 2009.

Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, and Sergei Vassilvitskii. Scalable k-means++. *arXiv preprint arXiv:1203.6402*, 2012.

Constance Beguier, Mathieu Andreux, and Eric W Tramel. Efficient sparse secure aggregation for federated learning. *arXiv preprint arXiv:2007.14861*, 2020.

James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrède Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1253–1269, 2020.

Elwyn R. Berlekamp. Algebraic coding theory. In *McGraw-Hill series in systems science*, 1968.

Jock A Blackard and Denis J Dean. Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables. *Computers and electronics in agriculture*, 24(3):131–151, 1999.

Kallista Bonawitz, Peter Kairouz, Brendan McMahan, and Daniel Ramage. Federated learning and privacy: Building privacy-preserving systems for machine learning and data science on decentralized data. *Queue*, 19(5):87–114, 2021.

Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.

Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.

Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečnỳ, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. LEAF: A Benchmark for Federated Settings. *arXiv preprint arXiv:1812.01097*, 2018.

Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE Symposium on Security and Privacy*, pp. 463–480. IEEE, 2015.

Sanjay Chawla and Aristides Gionis. k-means–: A unified approach to clustering and outlier detection. In *Proceedings of the 2013 SIAM international conference on data mining*, pp. 189–197. SIAM, 2013.

Min Chen, Zhikun Zhang, Tianhao Wang, Michael Backes, Mathias Humbert, and Yang Zhang. Graph unlearning. *arXiv preprint arXiv:2103.14991*, 2021.

Wei-Ning Chen, Peter Kairouz, and Ayfer Ozgur. Breaking the communication-privacy-accuracy trilemma. *Advances in Neural Information Processing Systems*, 33:3312–3324, 2020.

Wei-Ning Chen, Christopher A Choquette Choo, Peter Kairouz, and Ananda Theertha Suresh. The fundamental price of secure aggregation in differentially private federated learning. In *International Conference on Machine Learning*, pp. 3056–3089. PMLR, 2022.

Eli Chien, Chao Pan, and Olgica Milenkovic. Certified graph unlearning. *arXiv preprint arXiv:2206.09140*, 2022.

I Chien, Chao Pan, and Olgica Milenkovic. Query k-means clustering and the double dixie cup problem. *Advances in Neural Information Processing Systems*, 31, 2018.

Jichan Chung, Kangwook Lee, and Kannan Ramchandran. Federated unsupervised clustering with generative models. In *AAAI 2022 International Workshop on Trustable, Verifiable and Auditable Federated Learning*, 2022.

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. EMNIST: Extending MNIST to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pp. 2921–2926. IEEE, 2017.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Don Kurian Dennis, Tian Li, and Virginia Smith. Heterogeneity for the win: One-shot federated clustering. In *International Conference on Machine Learning*, pp. 2611–2620. PMLR, 2021.

Alfredo Eisinberg and Giuseppe Fedele. On the inversion of the vandermonde matrix. *Applied mathematics and computation*, 174(2):1384–1397, 2006.

Irem Ergun, Hasin Us Sami, and Basak Guler. Sparsified secure aggregation for privacy-preserving federated learning. *arXiv preprint arXiv:2112.12872*, 2021.

Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1322–1333, 2015.

Keith Frikken. Privacy-preserving set union. In *International Conference on Applied Cryptography and Network Security*, pp. 237–252. Springer, 2007.

Guojun Gan and Michael Kwok-Po Ng. K-means clustering with outlier removal. *Pattern Recognition Letters*, 90:8–14, 2017.

Venkata Gandikota, Daniel Kane, Raj Kumar Maity, and Arya Mazumdar. vqsgd: Vector quantized stochastic gradient descent. In *International Conference on Artificial Intelligence and Statistics*, pp. 2197–2205. PMLR, 2021.

Andrew Gardner, Christian A Duncan, Jinko Kanno, and Rastko Selmic. 3D hand posture recognition from small unlabeled point sets. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 164–169. IEEE, 2014a.

Andrew Gardner, Jinko Kanno, Christian A Duncan, and Rastko Selmic. Measuring distance between unordered sets of different sizes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 137–143, 2014b.

Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33:16937–16947, 2020.

Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020.

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.

Sudipto Guha, Adam Meyerson, Nina Mishra, Rajeev Motwani, and Liadan O'Callaghan. Clustering data streams: Theory and practice. *IEEE transactions on knowledge and data engineering*, 15(3): 515–528, 2003.

Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3832–3842. PMLR, 13–18 Jul 2020.

Xiaoping Han, Renying Wang, Yincong Zhou, Lijiang Fei, Huiyu Sun, Shujing Lai, Assieh Saadatpour, Ziming Zhou, Haide Chen, Fang Ye, et al. Mapping the mouse cell atlas by microwell-seq. *Cell*, 172(5):1091–1107, 2018.

John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979.

Carolyn Hutter and Jean Claude Zenklusen. The Cancer Genome Atlas: Creating Lasting Value beyond Its Data. *Cell*, 173(2):283–285, 2018.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.

Kiran S Kedlaya and Christopher Umans. Fast polynomial factorization and modular composition. *SIAM Journal on Computing*, 40(6):1767–1802, 2011.

Lea Kissner and Dawn Song. Privacy-preserving set operations. In *Annual International Cryptology Conference*, pp. 241–257. Springer, 2005.

Euiwoong Lee, Melanie Schmidt, and John Wright. Improved and simplified inapproximability for k-means. *Information Processing Letters*, 120:40–43, 2017.

Songze Li, Sizai Hou, Baturalp Buyukates, and Salman Avestimehr. Secure federated clustering. *arXiv preprint arXiv:2205.15564*, 2022.

Gaoyang Liu, Xiaoqiang Ma, Yang Yang, Chen Wang, and Jiangchuan Liu. Federaser: Enabling efficient client-level data removal from federated learning models. In *2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS)*, pp. 1–10. IEEE, 2021.

Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2): 129–137, 1982.

Meena Mahajan, Prajakta Nimbhorkar, and Kasturi Varadarajan. The planar k-means problem is np-hard. *Theoretical Computer Science*, 442:13–21, 2012.

Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020.

James Massey. Shift-register synthesis and bch decoding. *IEEE transactions on Information Theory*, 15(1):122–127, 1969.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.

Jane Peterson, Susan Garges, Maria Giovanni, Pamela McInnes, Lu Wang, Jeffery A Schloss, Vivien Bonazzi, Jean E McEwen, Kris A Wetterstrand, Carolyn Deal, et al. The NIH Human Microbiome Project. *Genome research*, 19(12):2317–2323, 2009.

Irving S Reed and Gustave Solomon. Polynomial codes over certain finite fields. *Journal of the society for industrial and applied mathematics*, 8(2):300–304, 1960.

Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE transactions on neural networks and learning systems*, 32(8):3710–3722, 2020.

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.

Jae Hong Seo, Jung Hee Cheon, and Jonathan Katz. Constant-round multi-party private set union using reversed laurent series. In *International Workshop on Public Key Cryptography*, pp. 398–412. Springer, 2012.

Jinhyun So, Corey J Nolet, Chien-Sheng Yang, Songze Li, Qian Yu, Ramy E Ali, Basak Guler, and Salman Avestimehr. Lightsecagg: a lightweight and versatile design for secure aggregation in federated learning. *Proceedings of Machine Learning and Systems*, 4:694–720, 2022.

Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine*, 12(3): e1001779, 2015.

Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.

Sergei Vassilvitskii and David Arthur. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035, 2006.

Michael Veale, Reuben Binns, and Lilian Edwards. Algorithms that remember: model inversion attacks and data protection law. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2133):20180083, 2018.

Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021.

Junxiao Wang, Song Guo, Xin Xie, and Heng Qi. Federated unlearning via class-discriminative pruning. In *Proceedings of the ACM Web Conference 2022*, pp. 622–632, 2022.

Chen Wu, Sencun Zhu, and Prasenjit Mitra. Federated unlearning with knowledge distillation. *arXiv preprint arXiv:2201.09441*, 2022.

Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019.