

# Secure Federated Clustering

Songze Li<sup>1,2</sup>, Sizai Hou<sup>1</sup>, Baturalp Buyukates<sup>3</sup> and Salman Avestimehr<sup>3</sup>

<sup>1</sup>IoT Thrust, The Hong Kong University of Science and Technology  
(Guangzhou)

<sup>2</sup>Department of Computer Science and Engineering, The Hong Kong  
University of Science and Technology

<sup>3</sup> Department of Electrical and Computer Engineering, University of  
Southern California

## Abstract

We consider a foundational unsupervised learning task of  $k$ -means data clustering, in a federated learning (FL) setting consisting of a central server and many distributed clients. We develop SecFC, which is a secure federated clustering algorithm that simultaneously achieves 1) *universal performance*: no performance loss compared with clustering over centralized data, regardless of data distribution across clients; 2) *data privacy*: each client's private data and the cluster centers are not leaked to other clients and the server. In SecFC, the clients perform Lagrange encoding on their local data and share the coded data in an information-theoretically private manner; then leveraging the algebraic structure of the coding, the FL network exactly executes the Lloyd's  $k$ -means heuristic over the coded data to obtain the final clustering. Experiment results on synthetic and real datasets demonstrate the universally superior performance of SecFC for different data distributions across clients, and its computational practicality for various combinations of system parameters. Finally, we propose an extension of SecFC to further provide membership privacy for all data points.

## 1 Introduction

Federated learning (FL) is an increasingly popular distributed learning paradigm, which enables decentralized clients to collaborate on training an AI model over their private data collectively, without surrendering the private data to a central cloud. When running the FedAvg algorithm introduced in [1], each client trains a local model with its private data, and sends the local model to a central server for further aggregation into a global model. Since its introduction, FL has demonstrated great potential in improving the performance and security of a

wide range of learning problems, including next-word prediction [2], recommender systems [3], and predictive models for healthcare [4].

In this work, we focus on a foundational unsupervised learning task - clustering. In an FL system, the clustering problem can be categorized into client clustering and data clustering. Motivated by heterogeneous data distribution across clients and the needs for training personalized models [5, 6, 7], client clustering algorithms are developed to partition the clients into different clusters to participate in training of different models [8, 9, 10, 11]. On the other hand, data clustering problem aims to obtain a partition of data points distributed on FL clients, such that some clustering loss is minimized. In [12], a model-based approach UIFCA was proposed to train a generative model for each cluster, such that each data point can be classified into one of the clusters using the models. Another recent work [13] considers the classical  $k$ -means clustering problem, and proposed an algorithm  $k$ -FED to approximately execute the Lloyd’s heuristic [14] in the FL setting.  $k$ -FED is communication-efficient as it only requires one-shot communication from the clients to the server.

While being the state-of-the-art federated data clustering algorithms, both UIFCA and  $k$ -FED are observed to be limited in the following two aspects: 1) The clustering performance heavily depends on the data distribution across clients. Specifically, the clustering performance of  $k$ -FED is guaranteed only when the data points on each client are from at most  $k' \leq \sqrt{k}$  clusters, and a non-empty cluster at each client has to contain a minimum number of points; 2) Leakage of data privacy. Just like other model-based FL algorithms, the data privacy of UIFCA is subject to model inversion attacks (see, e.g., [15, 16, 17, 18]). In  $k$ -FED, each client performs Lloyd’s algorithm on local data, and sends the local cluster centers that are averages of private data points directly to the server, which allows a curious server to infer about the private dataset (e.g., number of local clusters, and the distribution of local data points).

Motivated by these limitations, we develop a secure federated clustering protocol named SecFC, for solving a  $k$ -means clustering problem over an FL network. SecFC builds upon the classical Lloyd’s heuristic and has the following salient properties:

- **Universal performance:** The clustering performance of SecFC matches exactly that of running Lloyd on the collection of all clients’ data in a centralized manner, which is invariant to different data distributions at the clients;
- **Information-theoretical data privacy:** Each client’s private data is information-theoretically private from a certain number of colluding clients. The server does not know the values of the private data points and the cluster centers.

In the first phase of SecFC, each client locally generates some random noises and mixes them with its private data using Lagrange polynomial encoding [19], and then securely shares the coded data segments with the other clients. In the second phase, the clients and the server jointly execute the Lloyd’s iterations over coded data. As each client obtains a coded version of the entire dataset in the first phase, leveraging the algebraic structure of the coding scheme, each

client can compute a secret share of the distance between each pair of cluster center and data point. After receiving sufficient number of secret shares, the server is able to reveal the pair-wise distances between all cluster centers and all data points, so the center update can be carried out exactly as specified in the Lloyd’s algorithm. Importantly, the server can only recover the distance between a cluster center and a data point, without knowing their respective locations, which is the key enabler for SecFC to achieve performance and privacy simultaneously. We further extend SecFC to strengthen it with membership privacy. Specifically, clients privately align the IDs of their local data using private set union [20, 21, 22], such that data clustering is performed without revealing which data point belongs to which client.

We conduct extensive experiments to cluster synthetic and real datasets on FL systems, and compare SecFC with the centralized Lloyd’s algorithm and  $k$ -FED as our baselines. Under different data distributions across clients with different number of local clusters  $k'$ , SecFC consistently achieves almost identical performance as the centralized Lloyd, which outperforms  $k$ -FED. Also, unlike  $k$ -FED whose performance varies substantially with the value of  $k'$ , and deteriorates severely when  $k' = k$ , the performance of SecFC is almost invariant to the value of  $k'$ , indicating its universal superiority. We also evaluate the execution overheads of SecFC under various combinations of system parameters, which demonstrate the feasibility of applying SecFC to solve practical clustering tasks.

## 2 Background and Related Work

Clustering is an extensively studied learning task. We discuss existing works in three categories.

- **Centralized clustering.** In centralized clustering, a central entity, i.e., server, performs the data partition. A well-known centralized clustering problem is  $k$ -means clustering, where the aim is to find  $k$  centers such that the Euclidean distance from each data point to its closest center is minimized. Lloyd’s heuristic [14] is a popular iterative approach in performing  $k$ -means clustering due to its simplicity, even though its performance heavily depends on algorithm initialization. Unlike the Lloyd’s heuristic which does not provide any provable guarantees, Kanungo et al. [23] proposes an approximation algorithm for  $k$ -means clustering that is based on swapping cluster centers in and out, which yields a solution that is at most  $(9 + \epsilon)$  factor larger than the optimal solution. Awasthi and Sheffet [24] propose a variant of the Lloyd’s heuristic that is equipped with such an approximation algorithm to guarantee convergence in polynomial time.
- **Parallel and distributed clustering.** Distributed and parallel clustering implementations are particularly useful in the presence of large datasets that are prominent in many learning applications. The bottleneck in the centralized  $k$ -means clustering is the computation of distances among the data

points and the centers. In parallel implementations of  $k$ -means clustering, the idea is to split the dataset into different worker machines so that the distance computations are performed in parallel [25, 26]. In addition to  $k$ -means clustering, there have been parallel implementations of other clustering schemes such as the density-based clustering scheme DBSCAN [27]. When the dataset is distributed in nature, it may not be practical to send local datasets to the central entity. In such cases, distributed clustering is favorable albeit a performance loss compared to the centralized setting [28, 29, 30].

- **Federated clustering.** Clustering techniques have been used in FL systems to handle data heterogeneity [31, 9, 32, 13, 33], particularly focusing on personalized model training and client selection. References [31, 32, 9] focus on clustering the clients to jointly train multiple models, one for each cluster, whereas reference [13] clusters the client data points to reach more personalized models and perform a more representative client selection at each iteration. Unlike the  $k$ -means clustering which uses a heuristic to generate clusters, model-based clustering assumes that the data to be clustered comes from a probabilistic model and aims to recover that underlying distribution [34, 35]. In their recent work [12], using this approach, Chung et al. cluster heterogeneous client data in a federated setting by associating each data point with a cluster (model) with the highest likelihood. As also highlighted by [33], common to all these existing works on  $k$ -means clustering is the fact that none of them provides *data privacy*, in the sense that cluster center information as well as clients’ data are kept private. This motivates us to develop SecFC in this work, which is to the best of our knowledge, the first federated clustering algorithm with formal privacy guarantees.

Before we close this section, we briefly review some prior works on secure clustering.

**Privacy-preserving clustering.** When the clustering is performed on private client data as in the case of federated clustering, the issue of secure clustering emerges. The aim here is to cluster the entire dataset without revealing anything to the clients, including all intermediate values, but the output, i.e., final partition of data. Privacy-preserving clustering has been studied in the literature employing cryptographic techniques from secure multiparty computation, for the cases of horizontal, vertical, and arbitrary partitions of data [36, 37, 38, 39, 40, 41]. Among these works, we want to highlight [41], where authors implement a secret sharing-based secure  $k$ -means clustering scheme. Unlike the proposed SecFC scheme, however, [41] employs an additive secret sharing scheme and reveals the final cluster centers to the participants. In SecFC, clients use Lagrange coding to encode and share their private data. During the execution of SecFC, neither final nor any intermediate candidate center information is disclosed to the clients or the server, which is not the case in other earlier works [36, 37, 39] in the literature as well.

**Notation.** For a positive integer  $n$ , let  $[n]$  denote the set of integers  $\{1, \dots, n\}$ . Let  $|\mathcal{S}|$  denote the cardinality of a set  $\mathcal{S}$ .

## 3 Preliminaries and Problem Formulation

### 3.1 $k$ -means Clustering

Given a dataset  $\mathbf{X}$  consisting of  $m$  data points  $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{R}^d$ , each with features of dimension  $d$ , the  $k$ -means clustering problem aims to obtain a  $k$ -partition of  $\mathbf{X}$ , denoted by  $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_k)$ , such that the following clustering loss is minimized.

$$L(\mathcal{S}) = \sum_{h=1}^k \sum_{\mathbf{x}_i \in \mathcal{S}_h} \|\mathbf{x}_i - \boldsymbol{\mu}_h\|_2^2, \quad (1)$$

where  $\boldsymbol{\mu}_h$  denotes the center of cluster  $\mathcal{S}_h$ , which is computed as  $\boldsymbol{\mu}_h = \frac{1}{|\mathcal{S}_h|} \sum_{\mathbf{x}_i \in \mathcal{S}_h} \mathbf{x}_i$ .

A classical method for clustering is the Lloyd’s heuristic algorithm [14]. It is an iterative refinement method that alternates on two steps. The first step is to assign each data point to its nearest cluster center. The second step is to update each cluster center, as the mean of data points assigned to that cluster in the first step. The algorithm terminates when the data assignments no longer change. When the data points are well separated, Lloyd’s algorithm often converges in a few iterations and achieves good performance [42], but its complexity can grow superpolynomially in the worst case [43]. Despite of many variants of Lloyd’s algorithm to improve its clustering performance and computational complexity in general cases, we focus on solving the  $k$ -means clustering problem using the original Lloyd’s heuristic due to its simplicity.

### 3.2 Secure Federated $k$ -means

We consider solving a  $k$ -means clustering problem, over a federated learning system that consists of a central server and  $n$  distributed clients. The entire dataset consists of private data points arbitrarily distributed across the clients. For each  $j \in [n]$ , we denote the set of local data points at client  $j$  as  $\mathcal{D}_j$ ,  $j \in [n]$ . Specifically, we would like to run the Lloyd’s algorithm to obtain a  $k$ -clustering of the clients’ data, with the coordination of the server.

A similar problem was recently considered in [13], where a federated clustering protocol,  $k$ -FED, was proposed to approximately execute Lloyd’s heuristic with one-shot communication from the clients to the server. The basic idea of  $k$ -FED is to have each client  $j$  run Lloyd to generate a  $k_j$ -clustering of its local data  $\mathcal{D}_j$ , and send the cluster centers to the server, and then the server runs Lloyd again to cluster the centers into  $k$  groups, generating a  $k$ -clustering of the entire dataset. We note that  $k$ -FED is limited in the following two aspects: 1) to run  $k$ -FED, each client  $j$  needs to know the actual number of local clusters  $k_j$ , and the clustering performance is guaranteed only when the number of local clusters is less than  $\sqrt{k}$ , and the number of data points in each non-empty local cluster is large enough, which may hardly hold in practical scenarios; 2) the clients’ private data is leaked to the server through the client centers directly sent to the server, which are computed as the means of the data points in respective local clusters.

Motivated by the above limitations, we aim to design a federated  $k$ -means clustering algorithm that simultaneously achieves the following requirements.

- **Universal performance:** No loss of clustering performance compared with the centralized Lloyd’s algorithm on the entire dataset, regardless of the data distribution across the clients;
- **Data  $t$ -privacy:** We consider a threat model where the clients and the server are *honest-but-curious*. For some security parameter  $t < n$ , during the protocol execution, no information about a client’s private data should be leaked, even when any subset of up to  $t$  clients collude with each other. Further, the server should not know explicitly any private data of the clients, as well as the locations of the cluster centers.

We propose Secure Federated Clustering (or SecFC) that simultaneously satisfies the above performance and security requirements. In SecFC, private data is secret shared across clients such that information-theoretical privacy is guaranteed against up to  $t$  colluding clients. Using computations over coded data, the server decodes pair-wise distances between the data points and the centers, which are sufficient to advance Lloyd iterations. In this way, SecFC exactly recovers the center updates of centralized Lloyd’s algorithm, with the server knowing nothing about the data points and the centers.

## 4 The Proposed SecFC Protocol

As SecFC leverages cryptographic primitives to protect data privacy, which operates on finite fields, we assume that the elements of each data point are from a finite field  $\mathbb{F}_q$  of order  $q$ . In practice, one can quantize each element of the data points onto  $\mathbb{F}_q$ , for some sufficiently large prime  $q$ , such that overflow does not occur to any computation during the execution of the Lloyd’s algorithm. See Appendix A for one such quantization method.

### 4.1 Protocol Description

**Secure data sharing.** In the first phase of SecFC, each client mixes its local data with random noises to generate secret shares of the data, and communicates the shares with the other clients.

For each  $j \in [n]$ , and each  $\mathbf{x}_i \in \mathcal{D}_j$ , client  $j$  first evenly partitions  $\mathbf{x}_i \in \mathbb{F}_q^d$  into  $\ell$  segments  $\mathbf{x}_{i,1}, \dots, \mathbf{x}_{i,\ell} \in \mathbb{F}_q^{\frac{d}{\ell}}$ , for some design parameter  $\ell$ . Next, client  $j$  randomly samples  $t$  noise vectors  $\mathbf{z}_{i,\ell+1}, \dots, \mathbf{z}_{i,\ell+t}$  from  $\mathbb{F}_q^{\frac{d}{\ell}}$ . Then, for a set of distinct parameters  $\{\beta_1, \dots, \beta_{\ell+t}\}$  from  $\mathbb{F}_q$  that are agreed upon among all clients and the server, following the data encoding scheme of Lagrange Coded Computing (LCC) [19], client  $j$  performs Lagrange interpolation on the points  $(\beta_1, \mathbf{x}_{i,1}), \dots, (\beta_\ell, \mathbf{x}_{i,\ell}), (\beta_{\ell+1}, \mathbf{z}_{i,\ell+1}), \dots, (\beta_{\ell+t}, \mathbf{z}_{i,\ell+t})$  to obtain the

following polynomial

$$\mathbf{f}_i(\alpha) = \sum_{u=1}^{\ell} \mathbf{x}_{i,u} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha - \beta_v}{\beta_u - \beta_v} + \sum_{u=\ell+1}^{\ell+t} \mathbf{z}_{i,u} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha - \beta_v}{\beta_u - \beta_v}. \quad (2)$$

Note that  $\mathbf{f}_i(\beta_u) = \mathbf{x}_{i,u}$  for  $u = 1, \dots, \ell$ , and  $\mathbf{f}_i(\beta_u) = \mathbf{z}_{i,u}$  for  $u = \ell + 1, \dots, \ell + t$ .

Finally, for another set of public parameters  $\{\alpha_1, \dots, \alpha_n\}$  that are pair-wise distinct, and  $\{\beta_1, \dots, \beta_{\ell+t}\} \cap \{\alpha_1, \dots, \alpha_n\} = \emptyset$ , client  $j$  evaluates  $\mathbf{f}_i(\alpha)$  at  $\alpha_{j'}$  to compute a secret share  $\tilde{\mathbf{x}}_{i,j'}$  of  $\mathbf{x}_i$  for client  $j'$ , for all  $j' \in [n]$  (i.e.,  $\tilde{\mathbf{x}}_{i,j'} = \mathbf{f}_i(\alpha_{j'})$ ), and sends  $\tilde{\mathbf{x}}_{i,j'}$  to client  $j'$ .

By the end of secure data sharing, each client  $j$  possesses a secret share of the entire dataset  $\mathbf{X}$ , i.e.,  $\{\tilde{\mathbf{x}}_{i,j}\}_{i=1}^m$ .

**Coded center update.** The proposed SecFC runs in an iterative manner until the centers converge. The server starts by randomly sampling a  $k$ -clustering assignment  $\mathcal{S}$ . In the subsequent iterations, the computations of the distances from the data points to the centers, and the center updates are all performed on the secret shares of the original data.

In each iteration, given the current clustering  $(\mathcal{S}_1, \dots, \mathcal{S}_k)$ , each client  $j$  updates the secret shares of the centers as  $\tilde{\boldsymbol{\mu}}_{h,j} = \sum_{i \in \mathcal{S}_h} \tilde{\mathbf{x}}_{i,j}$ , for all  $h \in [k]$ . Then, for each  $h \in [k]$  and each  $i \in [m]$ , client  $j$  computes the *coded* distance from data point  $i$  to center  $h$  as  $\tilde{d}_{i,h,j} = \|\tilde{\boldsymbol{\mu}}_{h,j} - |\mathcal{S}_h| \cdot \tilde{\mathbf{x}}_{i,j}\|_2^2$ , and sends it to the server. For each  $(i, h) \in [m] \times [k]$ , the coded distance  $\tilde{d}_{i,h,j}$  can be viewed as the evaluation of a polynomial  $\phi_{i,h}(\alpha) = \|\sum_{i' \in \mathcal{S}_h} \mathbf{f}_{i'}(\alpha) - |\mathcal{S}_h| \cdot \mathbf{f}_i(\alpha)\|_2^2$  at point  $\alpha_j$ . As  $\phi_{i,h}(\alpha)$  has degree  $2(\ell + t - 1)$ , the server can interpolate  $\phi_{i,h}(\alpha)$  from the computation results of  $2\ell + 2t - 1$  clients. Having recovered  $\phi_{i,h}(\alpha)$ , the server computes  $d'_{i,h} = \sum_{u=1}^{\ell} \phi_{i,h}(\beta_u) = \sum_{u=1}^{\ell} \|\sum_{i' \in \mathcal{S}_h} \mathbf{x}_{i',u} - |\mathcal{S}_h| \cdot \mathbf{x}_{i,u}\|_2^2$ , and normalizes it by the square of the size of cluster  $h$  to obtain the distance from data point  $i$  to cluster center  $h$ , i.e.,  $d_{i,h} = \|\boldsymbol{\mu}_h - \mathbf{x}_i\|_2^2 = \frac{d'_{i,h}}{|\mathcal{S}_h|^2}$ .<sup>1</sup> Having obtained the distances between all centers and all data points, the server updates the clustering assignment by first identifying the nearest center for each data point, and assigning the data point to the corresponding cluster. We summarize the proposed SecFC protocol in Algorithm 1.

## 4.2 Theoretical Analysis

We theoretically analyze the performance, privacy, and complexity of the proposed SecFC protocol.

**Theorem 1.** *For the federated clustering problem with  $n$  clients, the proposed SecFC scheme with parameter  $\ell$  achieves universal performance and data  $t$ -privacy, when  $2t + 2\ell - 1 \leq n$ .*

*Proof.* The data  $t$ -privacy of SecFC against colluding clients follows from the  $t$ -privacy of LCC encoding. As shown in Theorem 1 of [19], each data point  $\mathbf{x}_i$  is

<sup>1</sup>Here the division is performed in the real field.

---

**Algorithm 1: SecFC**

---

**Data:** Datasets  $\mathcal{D}_1, \dots, \mathcal{D}_n$  on  $n$  clients  
**Input:** Number of clusters  $k$   
**Output:**  $k$ -clustering  $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_k)$

- 1 *Phase 1: Secure data sharing;*
- 2 **for**  $j \leftarrow 1$  **to**  $n$  **do**
- 3     **foreach**  $\mathbf{x}_i \in D_j$  **do** Client  $j$  generates a secret share  $\tilde{\mathbf{x}}_{i,j'}$ , and sends it to client  $j'$ ,  $\forall j' \in [n]$ ;
- 4 **end**
- 5 *Phase 2: Coded center update;*
- 6 Server randomly samples an initial clustering  $\mathcal{S}$  and broadcasts to all clients;
- 7 **while**  $\mathcal{S} = (\mathcal{S}_1, \dots, \mathcal{S}_k)$  *is different from last iteration* **do**
- 8     **for**  $j \leftarrow 1$  **to**  $n$  **do**
- 9         **foreach**  $h \in [k]$  **do** Client  $j$  updates local coded center  
            $\tilde{\boldsymbol{\mu}}_{h,j} = \sum_{i \in \mathcal{S}_h} \tilde{\mathbf{x}}_{i,j}$ ;
- 10         **for**  $(i, h) \in [m] \times [k]$  **do**
- 11             Client  $j$  computes a coded distance  
            $\tilde{d}_{i,h,j} = \|\tilde{\boldsymbol{\mu}}_{h,j} - |\mathcal{S}_h| \cdot \tilde{\mathbf{x}}_{i,j}\|_2^2$ , and sends it to the server
- 12         **end**
- 13     **end**
- 14     **for**  $(i, h) \in [m] \times [k]$  **do**
- 15         Server recovers  $d'_{i,h} = \|\boldsymbol{\mu}_h - |\mathcal{S}_h| \cdot \mathbf{x}_i\|_2^2$  from  $(\tilde{d}_{i,h,1}, \dots, \tilde{d}_{i,h,n})$  via interpolating and evaluating a polynomial of degree  $2t + 2\ell - 2$ ;
- 16         Server computes the actual distance  $d_{i,h} = \frac{d'_{i,h}}{|\mathcal{S}_h|^2}$  from  $\mathbf{x}_i$  to  $\boldsymbol{\mu}_h$ ;
- 17     **end**
- 18     Server associates each data point to its nearest center and updates the clustering  $\mathcal{S}$ .
- 19 **end**

---

information-theoretical security against  $t$  colluding clients. In other words, for any subset  $\mathcal{C} \subset [n]$  of size not larger than  $t$ , the mutual information  $I(\mathbf{x}_i; (\tilde{\mathbf{x}}_{i,j})_{j \in \mathcal{C}}) = 0$ .

As explained in the above protocol description, for a given clustering  $(\mathcal{S}_1, \dots, \mathcal{S}_k)$ , as long as  $n \geq 2\ell + 2t - 1$ , the server is able to exactly reconstruct the polynomial  $\phi_{i,h}(\alpha)$  from the computation results of the  $n$  clients, and subsequently compute the exact distance  $d_{i,h}$  from data point  $\mathbf{x}_i$  to the cluster center  $\boldsymbol{\mu}_h = \frac{\sum_{\mathbf{x}_{i'} \in \mathcal{S}_h} \mathbf{x}_{i'}}{|\mathcal{S}_h|}$ . Therefore, the clustering update of SecFC exactly follows that of a centralized Lloyd's algorithm, and achieves identical clustering performance. It is easy to see that this performance is achieved universally regardless of the clients' knowledge about the numbers of local clusters, and the distribution of data points.

Finally, with the computation results received from the clients, the server is



only able to decode the distances between the data points and the current cluster centers, but not their individual values. This achieves data privacy against a curious server.  $\square$

**Remark 1.** Taking advantage of a key property of the Lloyd’s algorithm where only the distances between points and centers are needed to carry out the center update step, SecFC is able to protect clients’ data privacy against the server, without sacrificing any clustering performance.

**Remark 2.** SecFC achieves the same clustering performance as the original Lloyd’s algorithm, which is nevertheless not guaranteed to produce the optimal clustering, and its performance relies heavily on the clustering initialization [44]. It was proposed in [24] to improve center initialization by first projecting the dataset onto the subspace spanned by its top  $k$  singular vectors, and applying an approximation algorithm (see, e.g., [23, 45]) on the projected dataset to obtain the initial centers. How one can incorporate these initialization tricks into SecFC without compromising data privacy remains a compelling research question.

Theorem 1 reveals a trade-off between privacy and efficiency for SecFC. As a larger  $t$  allows for data privacy against more curious colluding clients, a larger  $\ell$  reduces both the communication loads among the clients, and the computation complexities at the clients and the server.

**Complexity analysis.** We analyze the communication and computation complexities of SecFC. Let  $m_j = |\mathcal{D}_j|$  be the number local data points at client  $j$ , and  $s$  be the number of iterations.

*Communication complexity of client  $j$ .* The communication load of client  $j$  consists of two parts: 1) The communication cost of client  $j$  for data sharing is  $\mathcal{O}(\frac{m_j d n}{\ell})$ ; 2) In each iteration, client  $j$  communicates  $km$  computation results to the server, incurring a cumulative communication cost of  $\mathcal{O}(kms)$ . The total communication cost of client  $j$  during the execution of SecFC is  $\mathcal{O}(\frac{m_j d n}{\ell} + kms)$ .

*Computation complexity of client  $j$ .* The computation performed by client  $j$  consists of two parts: 1) Utilizing fast polynomial interpolation and evaluation [46], client  $j$  can generate the secret shares of its local data with complexity  $\mathcal{O}(\frac{m_j d}{\ell} n \log^2 n)$ ; 2) Within each iteration, client  $j$  needs to first compute a secret share for each of the  $k$  centers, and then compute the distance from each data point to each center in the coded domain. This takes a total of  $\mathcal{O}(\frac{k m d s}{\ell})$  operations. Hence, the total computation complexity of client  $j$  is  $\mathcal{O}(\frac{m_j d}{\ell} n \log^2 n + \frac{k m d s}{\ell})$ .

*Computation complexity of the server.* The computation occurs on server in each iteration consists of two parts. 1) The decoding and recovery of actual pair-wise distances take  $\mathcal{O}(km(\ell+t) \log^2(\ell+t))$  operations; 2) The cost of cluster assignment on each data point is  $\mathcal{O}(k)$ . Thus, the total computation complexity of the server is  $\mathcal{O}(kms(\ell+t) \log^2(\ell+t))$ .

## 5 Experiments

In this section, we empirically evaluate the performance and complexity of the proposed SecFC protocol, for federated clustering of synthetic and real datasets. We compare SecFC with centralized implementation of Lloyd’s heuristic and the  $k$ -FED protocol as baselines. All experiments are carried out on a machine using Intel(R) Xeon(R) Gold 5118 CPU @ 2.30GHz, with 12 cores of 48 threads.

The *center separation* technique [24] is applied to all three algorithms, for clustering initialization to improve the performance. It is a one-time method to generate  $k$  groups of data  $\mathcal{S}'_1, \dots, \mathcal{S}'_k$  from a given clustering assignment. Given the current cluster centers  $\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k$ , each  $\mathcal{S}'_h$  is selected such that  $\mathcal{S}'_h \leftarrow \{i : \|\mathbf{x}_i - \boldsymbol{\mu}_h\|_2 \leq \frac{1}{3} \|\mathbf{x}_i - \boldsymbol{\mu}_s\|_2, \forall s \in [k]\}$ . In this way, the data in  $\mathcal{S}'_h$  is close to center  $\boldsymbol{\mu}_h$  and far from all other centers. Then, the means of  $\mathcal{S}'_1, \dots, \mathcal{S}'_k$  are used as the initial centers for the subsequent iterative execution of the algorithms.

We cluster data points with ground truth labels, and use *classification accuracy* as the performance measure. For each  $i \in [m]$ , let  $y_i$  and  $y'_i$  respectively denote the ground truth label of data point  $\mathbf{x}_i$ , and the index of the cluster it belongs to after the clustering process, such that  $y_i, y'_i \in [k]$ . As the clustering operation does not directly predict  $\mathbf{x}_i$ ’s label, the classification of  $\mathbf{x}_i$  is subject to a potential permutation  $\pi$  of the label set  $[k]$ . We use the Hungarian method [47] to find the optimal permutation  $\pi_0$  that maximizes  $\sum_{i \in [m]} \mathbb{1}(y_i = \pi(y'_i))$ , where  $\mathbb{1}(\cdot)$  is the indicator function. Then, the accuracy of a  $k$ -means clustering algorithm is given by  $\frac{1}{m} \sum_{i \in [m]} \mathbb{1}(y_i = \pi_0(y'_i))$ .

### 5.1 Separating Mixture of Gaussians

We first consider a synthetic dataset, where the data is generated by a mixture of  $k$  Gaussians. That is, we generate isotropic Gaussian clusters comprised of a total of  $m$  data points each with  $d$  features. Each data point  $\mathbf{x}_i$  is sampled from an average of normal distributions with density  $p(\mathbf{x}) = \sum_{h=1}^k \frac{1}{k} \mathcal{N}(\boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h)$ . The centers  $\{\boldsymbol{\mu}_h\}_{h=1}^k$  are randomly chosen. The covariance matrix  $\boldsymbol{\Sigma}_h$  is diagonal with covariance  $\sigma_h^2$ . The standard deviations are the same for all  $h$ , i.e.,  $\sigma_h = \sigma$ ,  $h \in [k]$ . The ground truth label of each data point is given by the index of the closest center to that particular point.

To compare SecFC with  $k$ -FED, we generate clients’ datasets with different  $k'$ , where the parameter  $k'$  is the maximum number of clusters a client can have data points in and indicates the level of data heterogeneity across clients. The  $k$ -FED scheme assumes that this number is known to the clients [13], which is not required for our SecFC. We generate local datasets for the cases of  $k' \leq \sqrt{k}$ , which is required by  $k$ -FED to have performance guarantees, and also for the case of  $k' = k$ , which corresponds to i.i.d. data distribution across clients.

We consider two Gaussian mixtures of dimension  $d = 100$ , with standard deviation  $\sigma = 1$  and  $\sigma = 20$  respectively. For each value of  $\sigma$ , we consider a setting of  $k = 4$  clusters,  $n = 10$  clients, and  $m = 10000$  data points are generated; and another setting of  $k = 16$  clusters,  $n = 16$  clients, and  $m = 16384$

data points are generated. For each of the datasets, the centralized Lloyd,  $k$ -FED, and SecFC are executed for 10 runs, and their performance are shown in Table 1 and Table 2.

Table 1: Clustering accuracy (%) for a Gaussian mixture with  $\sigma = 1$ .

	$k = 4$			$k = 16$		
	$k' = 1$	$k' = 2$	$k' = 4$	$k' = 2$	$k' = 4$	$k' = 16$
Centralized Lloyd	100.0 $\pm$ 0.0	100.0 $\pm$ 0.0	100.0 $\pm$ 0.0	88.2 $\pm$ 5.5	88.2 $\pm$ 5.5	88.2 $\pm$ 5.5
$k$ -FED	75.0 $\pm$ 0.1	99.7 $\pm$ 2.1	25.8 $\pm$ 0.6	80.3 $\pm$ 5.3	78.7 $\pm$ 3.7	7.6 $\pm$ 0.1
SecFC	100.0 $\pm$ 0.0	100.0 $\pm$ 0.0	100.0 $\pm$ 0.0	86.0 $\pm$ 1.5	86.0 $\pm$ 1.5	86.0 $\pm$ 1.5

For each value of  $\sigma$ , and each combination of  $k$  and  $k'$ , SecFC consistently achieves almost identical performance with that of the centralized Lloyd, which is better than the accuracy of  $k$ -FED. For a fixed  $k$ , the accuracy of SecFC almost stays the same for different values of  $k'$ , demonstrating its universal superiority. In sharp contrast, the performance of  $k$ -FED varies significantly with  $k'$ , and drops drastically for the case of  $k' = k$ . With a larger  $\sigma$ , the Gaussian clusters are more scattered, hence increasing the difficulty for clustering algorithms to correctly classify data points. This is empirically verified by the performance drop when increasing  $\sigma$  from 1 to 20, for all three algorithms.

**Complexity evaluation.** We also evaluate the execution overheads of SecFC, under different combinations of system parameters. By default, the clustering is done for  $k = 4$  clusters on  $n = 10$  clients, for a dataset generated by parameters  $m = 1000$ ,  $d = 20$ , and  $\sigma = 1$ . For SecFC, as the computations at both the clients and the server are independent across centers and data points, we parallelize the computations onto different processes to speed up the execution of SecFC. Based on the capacity of the system’s computing power, the distance computations of the  $m$  data points at each client are parallelized onto multiple processes, and the distance decodings of the  $mk$  (data, center) pairs at the server are also parallelized onto multiple processes. We present the run times of  $k$ -FED and SecFC, as functions of the parameters  $n$ ,  $d$ , and  $m$  respectively in Figure 1. Here we measure the average run time of SecFC spent in each client-server iteration, and the overall run time of  $k$ -FED as it performs one-shot communication from clients to the server. We omit the communication times in all scenarios as they are negligible compared with computation times.

We observe in Figure 1 that, as expected, SecFC is slower than  $k$ -FED (in fact, one iteration of SecFC is already slower than the entire execution of  $k$ -FED). However, the execution time of SecFC in each iteration is rather short ( $< 16s$  in the worst case), which demonstrates its potential for practical secure clustering tasks. For different client number  $n$ , the security parameter  $t$  is kept to be  $t = \lceil \frac{1}{3}n \rceil$  to provide privacy against up to  $\frac{n}{3}$  colluding clients. The client run time is independent of  $n$  because all clients compute on the coded data of the entire dataset in SecFC. Server complexity increases slightly with  $n$  as the increased  $t$  contributes to a higher decoding complexity. For different data dimension  $d$ , the server run time remains almost unchanged as it only decodes the distance between data points and centers. The run times of the clients

Table 2: Clustering accuracy (%) for a Gaussian mixture with  $\sigma = 20$ .

	$k = 4$			$k = 16$		
	$k' = 1$	$k' = 2$	$k' = 4$	$k' = 2$	$k' = 4$	$k' = 16$
Centralized Lloyd	$96.3 \pm 0.1$	$96.3 \pm 0.1$	$96.3 \pm 0.1$	$84.0 \pm 0.2$	$84.0 \pm 0.2$	$84.0 \pm 0.2$
$k$ -FED	$46.7 \pm 0.0$	$86.6 \pm 5.5$	$25.6 \pm 0.1$	$42.0 \pm 2.2$	$41.7 \pm 1.2$	$8.4 \pm 0.1$
SecFC	$96.3 \pm 0.0$	$96.3 \pm 0.0$	$96.3 \pm 0.0$	$84.0 \pm 0.1$	$84.0 \pm 0.1$	$84.0 \pm 0.1$

in SecFC and the overall run time of  $k$ -FED both increase with  $d$  as they are operating on data points with a higher dimension. Finally, when the number of data points  $m$  is increased, the run times of server and clients in SecFC, and the overall run time of  $k$ -FED all grow almost linearly.




Figure 1: Run time of SecFC and  $k$ -FED under different system parameters.

## 5.2 Experiments on Real Data

We also evaluate the clustering performance of SecFC on the MNIST handwritten digit dataset [48]. As in [12], we use rotations to construct a clustered dataset such that different rotation angles correspond to data generated by different distributions. For the set of images in MNIST labelled by a certain digit, we rotate the images by a degree of 0, 90, 180, 270 respectively to construct a dataset with  $k = 4$  clusters. We execute SecFC and the two baseline algorithms on the datasets generated for digits 2 and 3 over  $n = 10$  clients, and present their performance in Table 3.

Table 3: Clustering accuracy (%) for Rotated MNIST datasets.

$k = 4$	Digit 2			Digit 3		
	$k' = 1$	$k' = 2$	$k' = 4$	$k' = 1$	$k' = 2$	$k' = 4$
Centralized Lloyd	$99.5 \pm 0.1$	$99.5 \pm 0.1$	$99.5 \pm 0.1$	$98.1 \pm 0.0$	$98.1 \pm 0.0$	$98.1 \pm 0.0$
$k$ -FED	$50.5 \pm 0.0$	$90.3 \pm 10.4$	$25.8 \pm 0.1$	$97.7 \pm 0.0$	$98.1 \pm 0.0$	$68.2 \pm 9.4$
SecFC	$99.5 \pm 0.1$	$99.5 \pm 0.1$	$99.5 \pm 0.1$	$98.1 \pm 0.0$	$98.1 \pm 0.0$	$98.1 \pm 0.0$

Similar to the synthetic data, SecFC achieves almost identical performance with centralized Lloyd in all tasks of clustering digit 2 and digit 3 images, which is insensitive to the value of  $k'$ .  $k$ -FED achieves a lower accuracy in general, which also varies substantially with  $k'$ .

## 6 Strengthening SecFC with Membership Privacy

One shortcoming of SecFC is that distribution of the global dataset, i.e.,  $(\mathcal{I}_1, \dots, \mathcal{I}_n)$ , where  $\mathcal{I}_j$  denotes the set of indices of the local data points at client  $j$ , will become publicly known during data sharing. It would be desirable for client  $j$  to also keep  $\mathcal{I}_j$  private, so that from the final clustering result  $(\mathcal{S}_1, \dots, \mathcal{S}_k)$ , no client or the server can infer which data point belongs to which client.

**Private ID alignment.** This can be achieved by adding an additional step of private ID alignment before data sharing in SecFC. We assume that each data point has a unique ID (e.g., its hash value), and we denote the set of IDs of the data points in  $\mathcal{D}_j$  at client  $j$  as  $\mathcal{E}_j$ . Before the clustering task starts, both  $\mathcal{D}_j$  and  $\mathcal{E}_j$  are kept private at client  $j$ , and are not known by the other clients and the server. As the first step of the clustering algorithm, all clients collectively run a private set union (PSU) protocol (see, e.g., [20, 21, 22]), over the local ID sets  $\mathcal{E}_1, \dots, \mathcal{E}_n$ , such that by the end of PSU, each client knows the set of all data IDs  $\mathcal{E} = \cup_{j \in [n]} \mathcal{E}_j$  in the network, without knowing the individual ID sets of other clients. Let  $|\mathcal{E}| = m$ . All the clients and the server agree on a bijective map between  $\mathcal{E}$  and  $[m]$ . This also induces a bijective map between  $\mathcal{E}_j$  and  $\mathcal{I}_j$  at each client  $j$ .

After private ID alignment, each client learns the total number of data points in the system. During the secure data sharing phase, to protect the privacy of  $\mathcal{I}_j \subset [m]$ , each client  $j$  secret shares a data  $\bar{\mathbf{x}}_i^{(j)}$  for each  $i = 1, \dots, m$ , as described in Section 4.1. Here  $\bar{\mathbf{x}}_i^{(j)} = \mathbf{x}_i$  for  $i \in \mathcal{I}_j$  and  $\bar{\mathbf{x}}_i^{(j)} = \mathbf{0}$  otherwise. We denote the secret share of  $\bar{\mathbf{x}}_i^{(j)}$  created by client  $j'$  for client  $j$  as  $\tilde{\mathbf{x}}_{i,j}^{(j')}$ . Unlike the original data sharing where each client  $j$  only receives one secret share  $\tilde{\mathbf{x}}_{i,j}$  from the owner of  $\mathbf{x}_i$ , for each  $i = 1, \dots, m$ , now with membership privacy, client  $j$  receives  $n$  shares  $\tilde{\mathbf{x}}_{i,j}^{(1)}, \dots, \tilde{\mathbf{x}}_{i,j}^{(n)}$ , one from each client. Each client  $j$  computes  $\tilde{\mathbf{x}}_{i,j} = \sum_{j'=1}^n \tilde{\mathbf{x}}_{i,j}^{(j')}$ . This completes the data sharing phase. Next, the clients perform the same coded center update process as in SecFC to obtain the final  $k$ -clustering.

With the extension of incorporating the private ID alignment step, SecFC can further protect the membership privacy of the data points, without sacrificing data privacy and clustering performance. We give more detailed description and analysis of this extended version of SecFC in Appendix B.

## 7 Conclusion and discussions

We propose a secure federated clustering algorithm named SecFC, for clustering data points arbitrarily distributed across a federated learning network. SecFC's design builds upon the well-known Lloyd's algorithm, and SecFC universally achieves identical performance to a centralized Lloyd's execution. Moreover, SecFC protects the data privacy of each client via secure data sharing, and computation over coded data throughout the clustering process. It guarantees information-theoretic privacy against colluding clients, and no explicit knowledge about the data points and cluster centers at the server. Experiment results from

clustering synthetic and real data demonstrate the universal superiority of SecFC and its computational practicality under a wide range of system parameters. Finally, we propose an extended version of SecFC to further provide membership privacy for the clients’ datasets, such that, while knowing the final clustering, no party knows the owner of any data point.

While the high level of security provided by SecFC is desirable for many applications, it cannot be directly applied to scenarios where the centers of clusters (e.g., in mean estimation [49, 50]), or the membership of specific data points (e.g., in clustering-based outlier detection [51, 52]) need to be known explicitly. These problems can be resolved respectively by having  $\ell + t$  clients to collectively reveal the cluster centers, and having each client publish the IDs of their local data.

## References

- [1] 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*, pages 1273–1282. PMLR, 2017.
- [2] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving Google keyboard query suggestions. *arXiv preprint arXiv:1812.02903*, 2018.
- [3] Amir Jalalirad, Marco Scavuzzo, Catalin Capota, and Michael Sprague. A simple and efficient federated recommender system. In *IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, pages 53–58, 2019.
- [4] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R. Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus Maier-Hein, et al. The future of digital health with federated learning. *NPJ digital medicine*, 3(1):1–7, 2020.
- [5] Canh T. Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with Moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020.
- [6] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948*, 2020.
- [7] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pages 6357–6368. PMLR, 2021.

- [8] 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.
- [9] 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.
- [10] Xiaomin Ouyang, Zhiyuan Xie, Jiayu Zhou, Jianwei Huang, and Guoliang Xing. ClusterFL: a similarity-aware federated learning system for human activity recognition. In *International Conference on Mobile Systems, Applications, and Services*, pages 54–66, 2021.
- [11] Yeongwoo Kim, Ezeddin Al Hakim, Johan Haraldson, Henrik Eriksson, José Mairton B. da Silva, and Carlo Fischione. Dynamic clustering in federated learning. In *ICC 2021-IEEE International Conference on Communications*, pages 1–6. IEEE, 2021.
- [12] 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.
- [13] Don Kurian Dennis, Tian Li, and Virginia Smith. Heterogeneity for the win: One-shot federated clustering. In *International Conference on Machine Learning*, pages 2611–2620. PMLR, July 2021.
- [14] Stuart Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [15] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in Neural Information Processing Systems*, 32, 2019.
- [16] 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.
- [17] Zhibo Wang, Mengkai Song, Zhifei Zhang, Yang Song, Qian Wang, and Hairong Qi. Beyond inferring class representatives: User-level privacy leakage from federated learning. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 2512–2520. IEEE, 2019.
- [18] Liam Fowl, Jonas Geiping, Wojtek Czaja, Micah Goldblum, and Tom Goldstein. Robbing the Fed: Directly obtaining private data in federated learning with modified models. *arXiv preprint arXiv:2110.13057*, 2021.
- [19] Qian Yu, Songze Li, Netanel Raviv, Seyed Mohammadreza Mousavi Kalan, Mahdi Soltanolkotabi, and Salman A. Avestimehr. Lagrange coded computing: Optimal design for resiliency, security, and privacy. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1215–1225. PMLR, 2019.

- [20] 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*. Springer, 2012.
- [21] Sivakanth Gopi, Pankaj Gulhane, Janardhan Kulkarni, Judy Hanwen Shen, Milad Shokouhi, and Sergey Yekhanin. Differentially private set union. In *International Conference on Machine Learning*. PMLR, 2020.
- [22] Keith Frikken. Privacy-preserving set union. In Jonathan Katz and Moti Yung, editors, *Applied Cryptography and Network Security*. Springer, 2007.
- [23] Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for  $k$ -means clustering. In *Annual Symposium on Computational Geometry*, June 2002.
- [24] Pranjali Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 37–49. Springer, 2012.
- [25] Inderjit S. Dhillon and Dharmendra S. Modha. A data-clustering algorithm on distributed memory multiprocessors. In *Workshop on Large-Scale Parallel KDD Systems, SIGKDD*. August 1999.
- [26] Manasi N. Joshi. Parallel  $K$ -Means algorithm on distributed memory multiprocessors. *Technical report, University of Minnesota*, 2003.
- [27] Xiaowei Xu, Jochen Jäger, and Hans-Peter Kriegel. A fast parallel clustering algorithm for large spatial databases. In *High performance data mining*, pages 263–290. Springer, 1999.
- [28] Hillol Kargupta, Weiyun Huang, Krishnamoorthy Sivakumar, and Erik Johnson. Distributed clustering using collective principal component analysis. *Knowledge and Information Systems*, 3(4):422–448, 2001.
- [29] Eshref Januzaj, Hans-Peter Kriegel, and Martin Pfeifle. Towards effective and efficient distributed clustering. In *Workshop on Clustering Large Data Sets*, 2003.
- [30] Maria-Florina F. Balcan, Steven Ehrlich, and Yingyu Liang. Distributed  $k$ -means and  $k$ -median clustering on general topologies. *Advances in Neural Information Processing Systems*, 26, 2013.
- [31] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. *Advances in neural information processing systems*, 30, 2017.
- [32] 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.



- [33] Ravikumar Balakrishnan, Tian Li, Tianyi Zhou, Nageen Himayat, Virginia Smith, and Jeff Bilmes. Diverse client selection for federated learning via submodular maximization. In *International Conference on Learning Representations*, 2022.
- [34] Sudipto Mukherjee, Himanshu Asnani, Eugene Lin, and Sreeram Kannan. Clustergan: Latent space clustering in generative adversarial networks. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 4610–4617, 2019.
- [35] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via self-conditioned gans. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14286–14295, 2020.
- [36] Jaideep Vaidya and Chris Clifton. Privacy-preserving  $k$ -means clustering over vertically partitioned data. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 206–215, 2003.
- [37] Geetha Jagannathan and Rebecca N Wright. Privacy-preserving distributed  $k$ -means clustering over arbitrarily partitioned data. In *ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, pages 593–599, 2005.
- [38] Paul Bunn and Rafail Ostrovsky. Secure two-party  $k$ -means clustering. In *ACM Conference on Computer and Communications Security*, pages 486–497, 2007.
- [39] Sankita Patel, Sweta Garasia, and Devesh Jinwala. An efficient approach for privacy preserving distributed  $k$ -means clustering based on shamir’s secret sharing scheme. In *IFIP International Conference on Trust Management*, pages 129–141. Springer, 2012.
- [40] Jiawei Yuan and Yifan Tian. Practical privacy-preserving mapreduce based  $k$ -means clustering over large-scale dataset. *IEEE Transactions on Cloud Computing*, 7(2):568–579, 2017.
- [41] Payman Mohassel, Mike Rosulek, and Ni Trieu. Practical privacy-preserving  $k$ -means clustering. *Cryptology ePrint Archive*, 2019.
- [42] Rafail Ostrovsky, Yuval Rabani, Leonard J. Schulman, and Chaitanya Swamy. The effectiveness of lloyd-type methods for the  $k$ -means problem. *Journal of the ACM*, 59(6):1–22, 2013.
- [43] David Arthur and Sergei Vassilvitskii. How slow is the  $k$ -means method? In *Symposium on Computational Geometry*, pages 144–153, 2006.
- [44] M. Emre Celebi, Hassan A. Kingravi, and Patricio A. Vela. A comparative study of efficient initialization methods for the  $k$ -means clustering algorithm. *Expert Systems with Applications*, 40(1):200–210, 2013.

- [45] Sariel Har-Peled and Soham Mazumdar. On coresets for  $k$ -means and  $k$ -median clustering. In *ACM Symposium on Theory of Computing*, pages 291–300, 2004.
- [46] Kiran S. Kedlaya and Christopher Umans. Fast polynomial factorization and modular composition. *SIAM Journal on Computing*, 40(6):1767–1802, 2011.
- [47] Harold W Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- [48] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [49] Kevin A Lai, Anup B Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 665–674. IEEE, 2016.
- [50] JA Cuesta-Albertos, C Matrán, and A Mayo-Iscar. Robust estimation in the normal mixture model based on robust clustering. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(4):779–802, 2008.
- [51] Antonio Loureiro, Luis Torgo, and Carlos Soares. Outlier detection using clustering methods: a data cleaning application. In *Proceedings of KDNNet Symposium on Knowledge-based systems for the Public Sector*. Springer Bonn, 2004.
- [52] Sheng-yi Jiang and Qing-bo An. Clustering-based outlier detection method. In *2008 Fifth international conference on fuzzy systems and knowledge discovery*, volume 2, pages 429–433. IEEE, 2008.

## Appendix

### A Data Quantization

In SecFC, the operations of data sharing and distance computation are carried out in a finite field  $\mathbb{F}_q$ , for some large prime  $q$ . Hence, for a data point  $\mathbf{x}_i$  from the real field, one needs to first quantize it onto  $\mathbb{F}_q$ . Specifically, for some scaling factor  $\lambda$  that determines the quantization precision, we first scale  $\mathbf{x}_i$  by  $\lambda$ , and embed the scaled value onto  $\mathbb{F}_q$  as follows.

$$\bar{\mathbf{x}}_i = \mathcal{Q}(\mathbf{x}_i, \lambda) = \begin{cases} \lfloor \lambda \mathbf{x}_i \rfloor, & \text{if } \mathbf{x}_i \geq 0 \\ \lfloor q + \lambda \mathbf{x}_i \rfloor, & \text{if } \mathbf{x}_i < 0 \end{cases}, \quad (3)$$

where  $\lfloor x \rfloor$  denotes the largest integer less than or equal to  $x$ , and the quantization function  $\mathcal{Q}$  is applied element-wise. Assuming each element of  $\lambda \mathbf{x}_i$  is within  $[-\eta, \eta]$  for some  $\eta > 0$ , then on the range of  $\mathcal{Q}$ ,  $[0, \eta)$  is the image of the positive part, and  $[q - \eta, q)$  is the image of the negative part. While  $\lambda$  controls the precision loss of quantization, it also needs to be chosen so that overflow does not occur during computations of SecFC. A larger  $\eta$  requires a larger finite field size  $q$  to avoid computation overflow, and a smaller  $\eta$  leads to a higher precision loss between  $\mathbf{x}_i$  and  $\bar{\mathbf{x}}_i$ . We can choose a proper scaling factor  $\lambda$  to preserve enough precision while keeping the field size practical.

To avoid computation overflow, we should choose  $q$  such that all the intermediate computation results on the scaled data  $\lambda \mathbf{x}_i$  are within the range  $(-\frac{q}{2}, \frac{q}{2})$ . In the worst case, the largest distance across data points and cluster centers  $D \triangleq \max_{i \in [m], h \in [k]} \|\boldsymbol{\mu}_h - |\mathcal{S}_h| \cdot \mathbf{x}_i\|_2^2$  results in the largest output value. Therefore, we should choose  $q$  that is at least  $2\lambda^2 D$ .

### B Extension of SecFC with Membership Privacy

After the step of private ID alignment, each client learns the total number of data points in the system. Each client  $j$  knows the set of IDs of its private data points as  $\mathcal{I}_j$ , but does not know the sets of data IDs of the other clients.

During the secure data sharing phase, to protect the privacy of  $\mathcal{I}_j \subset [m]$ , each client  $j$  first creates a data point  $\bar{\mathbf{x}}_i^{(j)}$  for *each*  $i = 1, \dots, m$ , where

$$\bar{\mathbf{x}}_i^{(j)} = \begin{cases} \mathbf{x}_i, & i \in \mathcal{I}_j \\ \mathbf{0}, & \text{otherwise} \end{cases}. \quad (4)$$

Next, for each  $i \in [m]$ , and some design parameter  $\ell$ , as described in Section 4.1, using Lagrange interpolation and evaluation, each client  $j'$  encodes  $\bar{\mathbf{x}}_i^{(j')}$ , and creates a secret share  $\tilde{\mathbf{x}}_{i,j}^{(j')}$  for each  $j \in [n]$  as

$$\tilde{\mathbf{x}}_{i,j}^{(j')} = \sum_{u=1}^{\ell} \bar{\mathbf{x}}_{i,u}^{(j')} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v} + \sum_{u=\ell+1}^{\ell+t} \mathbf{z}_{i,u}^{(j')} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v}, \quad (5)$$

where  $\bar{\mathbf{x}}_{i,1}^{(j')}, \dots, \bar{\mathbf{x}}_{i,\ell}^{(j')} \in \mathbb{F}_q^{\frac{d}{\ell}}$  are segments of  $\bar{\mathbf{x}}_i^{(j')}$ , and  $\mathbf{z}_{i,\ell+1}^{(j')}, \dots, \mathbf{z}_{i,\ell+t}^{(j')} \in \mathbb{F}_q^{\frac{d}{\ell}}$  are noises randomly generated by client  $j'$  for each  $i \in [m]$ .

Unlike the original data sharing of SecFC where each client  $j$  only receives one secret share  $\tilde{\mathbf{x}}_{i,j}$  from the owner of  $\mathbf{x}_i$ , for each  $i \in [m]$ , client  $j$  receives  $n$  shares  $\tilde{\mathbf{x}}_{i,j}^{(1)}, \dots, \tilde{\mathbf{x}}_{i,j}^{(n)}$ , one from each client. Each client  $j$  computes  $\tilde{\mathbf{x}}_{i,j} = \sum_{j'=1}^n \tilde{\mathbf{x}}_{i,j}^{(j')}$ . We can see from (4) and (5) that

$$\begin{aligned} \tilde{\mathbf{x}}_{i,j} &= \sum_{u=1}^{\ell} \left( \sum_{j'=1}^n \bar{\mathbf{x}}_{i,u}^{(j')} \right) \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v} + \sum_{u=\ell+1}^{\ell+t} \left( \sum_{j'=1}^n \mathbf{z}_{i,u}^{(j')} \right) \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v} \\ &= \sum_{u=1}^{\ell} \mathbf{x}_{i,u} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v} + \sum_{u=\ell+1}^{\ell+t} \bar{\mathbf{z}}_{i,u} \cdot \prod_{v \in [\ell+t] \setminus \{u\}} \frac{\alpha_j - \beta_v}{\beta_u - \beta_v}, \end{aligned} \quad (6)$$

where  $\bar{\mathbf{z}}_{i,u} = \sum_{j'=1}^n \mathbf{z}_{i,u}^{(j')}$  is still uniformly distributed in  $\mathbb{F}_q^{\frac{d}{\ell}}$ .

This completes the phase of secure data sharing.

As the data share in (6) resembles the algebraic structure of  $\mathbf{f}_i(\alpha_j)$  in (2), for all  $i \in [m]$  and  $j \in [n]$ , we can carry out the same coded center update process as in the original SecFC to reach the final  $k$ -clustering.

The data  $t$ -privacy of this extended version of SecFC is guaranteed by the Lagrange secure data sharing in (5). Also, as mentioned above, it achieves the same clustering performance as the original SecFC, which demonstrates its universal performance according to Theorem 1.