
Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning

Abstract

1 Secure aggregation is a critical component in federated learning (FL), which
2 enables the server to learn the aggregate model of the users without observing
3 their local models. Conventionally, secure aggregation algorithms focus only on
4 ensuring the privacy of individual users in a *single* training round. We contend that
5 such designs can lead to significant privacy leakages over *multiple* training rounds,
6 due to partial user selection/participation at each round of FL. In fact, we show that
7 the conventional random user selection strategies in FL may lead to leaking users'
8 individual models within a number of rounds that is linear in the number of users.
9 To address this challenge, we introduce a secure aggregation framework, Multi-
10 RoundSecAgg, with multi-round privacy guarantees. In particular, we introduce a
11 new metric to quantify the privacy guarantees of FL over multiple training rounds,
12 and develop a structured user selection strategy that guarantees the long-term
13 privacy of each user (over any number of training rounds). Our framework also
14 carefully accounts for the fairness and the average number of participating users at
15 each round. Our experiments on MNIST, CIFAR-10 and CIFAR-100 datasets in
16 the IID and the non-IID settings demonstrate the performance improvement over
17 the baselines, both in terms of privacy protection and test accuracy.

18 1 Introduction

19 Federated learning (FL) enables collaborative
20 training of machine learning models over the
21 data collected and stored locally by multiple
22 data-owners. The training in FL is typically
23 coordinated by a central server who maintains a
24 global model that is updated locally by the users.
25 The local updates are then aggregated by the
26 server to update the global model. Throughout
27 the training process, the users never share their
28 data with the server, i.e., the data is always kept
29 on device, rather, they only share their local
30 updates. However, as has been shown recently,
31 the local models may still reveal substantial
32 information about the local datasets, and the
33 private training data can be reconstructed from
34 the local models through inference or inversion
35 attacks (see e.g., [11, 26, 42, 12]).

36 To prevent such information leakage, *secure aggregation* protocols are proposed (e.g., [4, 31, 15,
37 40, 2, 38, 30]) to protect the privacy of the local models, both from the server and the other users,
38 while still allowing the server to learn their aggregate. More specifically, the secure aggregation
39 protocols ensure that, at any given round, the server can only learn the aggregate model of the users,
40 and beyond that no further information is revealed about the individual model.

41 Secure aggregation protocols, however, only ensure the privacy of the individual users in a *single*
42 *training round*, and do not consider their privacy over multiple training rounds [4, 2, 31, 32]. On
43 the other hand, due to partial user selection [7, 5, 6, 28], the server may be able to reconstruct the
44 individual models of some users using the aggregated models from the previous rounds. In fact, we
45 show that after a sufficient number of rounds, all local models can be recovered with a high accuracy
46 if the server uniformly chooses a random subset of the users to participate at every round. As shown

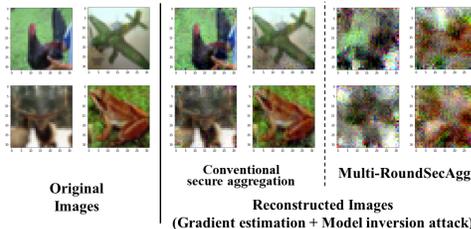


Figure 1: A qualitative comparison of the reconstructed images in two settings is shown. The first setting corresponds to the case that model privacy with random user selection (e.g., FedAvg [25]) is protected by conventional secure aggregation schemes as [4] at each round. In the second setting, our proposed method ensures the long-term privacy of individual models over any number of rounds, and hence model inversion attack cannot work well. This reconstruction process is described in detail in Appendix H.

47 in Fig.1, performing model inversion attack [12] with the recovered local models yields reconstructed
48 images with a similar quality as the original images.

49 **Contributions.** As such motivated, we study long-term user privacy in FL. Specifically, our
50 contributions are as follows.

- 51 1. [We introduce a new metric to capture long-term privacy guarantees for secure aggregation protocols](#)
52 [in FL for the first time.](#) This long-term privacy requires that the server cannot reconstruct any
53 individual model using the aggregated models from any number of training rounds. Using this
54 metric, we show that the conventional random selection schemes can result in leaking the local
55 models after a sufficient number of rounds, even if secure aggregation is employed at each round.
- 56 2. We propose Multi-RoundSecAgg, a privacy-preserving structured user selection strategy that
57 ensures the long-term privacy of the individual users over any number of training rounds. This
58 strategy also takes into account the fairness of the selection process and the average number of
59 participating users at each round.
- 60 3. We demonstrate that Multi-RoundSecAgg creates a trade-off between the long-term privacy
61 guarantee and the average number of participating users. In particular, as the average number of
62 participating users increases, the long-term privacy guarantee becomes weaker.
- 63 4. We provide the convergence analysis of Multi-RoundSecAgg, which shows that the long-term
64 privacy guarantee and the average number of participating users control the convergence rate. The
65 convergence rate is maximized when the average number of participating users is maximized.
66 (e.g., the random user selection strategy maximizes the average number of participating users at
67 the expense of not providing long-term privacy guarantees). As we require stronger long-term
68 privacy guarantees, the average number of participating users decreases and a larger number of
69 training rounds is required to achieve the same level of accuracy as the random selection strategy.
- 70 5. Finally, our experiments in both IID and non-IID settings on MNIST, CIFAR-10 and CIFAR-100
71 datasets demonstrate that Multi-RoundSecAgg achieves almost the same test accuracy compared
72 to the random selection scheme while providing better long-term privacy guarantees.

73 2 Related Work

74 The underlying principle of the secure aggregation protocol in [4] is that each pair of users exchange a
75 pairwise secret key which they can use to mask their local models before sharing them with the server.
76 The pairwise masks cancel out when the server aggregates the masked models, allowing the server to
77 aggregate the local models. These masks also ensure that the local models are kept private, i.e., no
78 further information is revealed beyond the aggregate of the local models. This protocol, however,
79 incurs a significant communication cost due to exchanging and reconstructing the pairwise keys.

80 Recently, several works have developed computation and communication-efficient protocols [31, 15,
81 2, 35, 8, 10, 38], which are complementary to and can be combined with our work. Another line of
82 work focused on designing partial user selection strategies to overcome the communication bottleneck
83 in FL while speeding up the convergence by selecting the users based on their local loss [7, 5, 6, 28].

84 Previous works, either on secure aggregation or on partial user selection, however, do not consider
85 mitigating the potential privacy leakage as a result of partial user participation and the server observing
86 the aggregated models across multiple training rounds. While [27] pointed out to the privacy leakage
87 of secure aggregation, mitigating this leakage has not been considered and our work is the first [secure](#)
88 [aggregation protocol](#) to address this challenge.

89 [Differential privacy \(DP\)](#), in which each user adds artificial noises to the local models, can be one of
90 [the potential solution to protect the privacy leakage over the multiple rounds](#) [9, 1, 37, 3, 16]. In DP,
91 [however, the privacy guarantee sacrifices the model performance, which is known as a privacy-utility](#)
92 [trade-off. It is worth noting that secure aggregation and DP are complementary, i.e., all the benefits](#)
93 [of DP can be applied to our approach by adding noise to the local models](#) [3]. In this paper, our
94 [objective is to understand the secure aggregation itself.](#)

95 3 System Model

96 In this section, we first describe the basic federated learning model in Section 3.1. Next, we introduce
97 the multi-round secure aggregation problem for federated learning and define the key metrics to
98 evaluate the performance of a multi-round secure aggregation protocol in Section 3.2.

99 3.1 Basic Federated Learning Model

100 We consider a cross-device federated learning setup consisting of a server and N users. User $i \in [N]$
 101 has a local dataset \mathcal{D}_i consisting of $m_i = |\mathcal{D}_i|$ data samples. The users are connected to each other
 102 through the server, i.e., all communications between the users goes through the server [24, 4, 17].
 103 The goal is to collaboratively learn a global model \mathbf{x} with dimension d , using the local datasets that
 104 are generated, stored, and processed locally by the users. The training task can be represented by
 105 minimizing a global loss function,

$$\min_{\mathbf{x}} L(\mathbf{x}) \text{ s.t. } L(\mathbf{x}) = \frac{1}{\sum_{i=1}^N w_i} \sum_{i=1}^N w_i L_i(\mathbf{x}), \quad (1)$$

106 where L_i is the loss function of user i and $w_i \geq 0$ is a weight parameter assigned to user i to specify
 107 the relative impact of that user. A common choice for the weight parameters is $w_i = m_i$ [17]. We
 108 define the optimal model parameters \mathbf{x}^* and \mathbf{x}_i^* as $\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^d} L(\mathbf{x})$ and $\mathbf{x}_i^* = \arg \min_{\mathbf{x} \in \mathbb{R}^d} L_i(\mathbf{x})$.
 109 **Federated Averaging with Partial User Participation.** To solve (1), the most common algorithm
 110 is the *FedAvg* (federated averaging) algorithm [24]. *FedAvg* is an iterative algorithm, where the model
 111 training is done by repeatedly iterating over individual local updates. At the beginning of training
 112 round t , the server sends the current state of the global model, denoted by $\mathbf{x}^{(t)}$, to the users. Each
 113 round consists of two phases, local training and aggregation. In the local training phase, user $i \in [N]$
 114 updates the global model by carrying out E (≥ 1) local stochastic gradient descent (SGD) steps and
 115 sends the updated local model $\mathbf{x}_i^{(t)}$ to the server. One of key features of cross-device FL is partial
 116 device participation. Due to various reasons such as unreliable wireless connectivity, or battery issues,
 117 at any given round, only a fraction of the users are available to participate in the protocol. We refer
 118 to such users as *available* users throughout the paper. In the aggregation phase, the server selects
 119 $K \leq N$ users among the available users if this is possible and aggregates their **local updates**. The
 120 server updates the global model as follows

$$\mathbf{x}^{(t+1)} = \sum_{i \in \mathcal{S}^{(t)}} w'_i \mathbf{x}_i^{(t)} = \mathbf{X}^{(t)\top} \mathbf{p}^{(t)}, \quad (2)$$

121 where $\mathcal{S}^{(t)}$ is the set of participating users at round t , $w'_i = \frac{w_i}{\sum_{i \in \mathcal{S}^{(t)}} w_i}$, and $\mathbf{p}^{(t)} \in \{0, 1\}^N$ is the
 122 corresponding characteristic vector. That is, $\mathbf{p}^{(t)}$ denotes a participation vector at round t whose i -th
 123 entry is 0 when user i is not selected and 1 otherwise. $\mathbf{X}^{(t)}$ denotes the concatenation of the weighted
 124 local models at round t , i.e., $\mathbf{X}^{(t)} = [w'_1 \mathbf{x}_1^{(t)}, \dots, w'_N \mathbf{x}_N^{(t)}]^\top \in \mathbb{R}^{N \times d}$. Finally, the server broadcasts
 125 the updated global model $\mathbf{x}^{(t+1)}$ to the users for the next round.

126 **Threat Model.** Similar to the prior works on secure aggregation as [4, 15, 31], we consider the
 127 honest-but-curious model. All participants follow the protocol honestly in this model, but try to learn
 128 as much as possible about the users. At each round, the privacy of individual model $\mathbf{x}_i^{(t)}$ in (2) is
 129 protected by secure aggregation such that the server only learns the aggregated model $\sum_{i \in \mathcal{S}^{(t)}} w'_i \mathbf{x}_i^{(t)}$.
 130

131 3.2 Multi-round Secure Aggregation

132 Conventional secure aggregation protocols only consider the privacy guarantees over a single training
 133 round. While secure aggregation protocols have provable privacy guarantees at any single round,
 134 in the sense that no information is leaked beyond the aggregate model at each round, the privacy
 135 guarantees do not extend to attacks *that span multiple training rounds*. Specifically, by using the
 136 aggregate models and participation information across multiple rounds, an individual model may be
 137 reconstructed. For instance, consider the following user participation strategy across three training
 138 rounds, $\mathbf{p}^{(1)} = [1, 1, 0]^\top$, $\mathbf{p}^{(2)} = [0, 1, 1]^\top$, and $\mathbf{p}^{(3)} = [1, 0, 1]^\top$. Assume a scenario where the local
 139 updates do not change significantly over time (e.g., models start to converge, or the server fixes the
 140 global model over consecutive rounds), i.e., $\mathbf{x}_i = \mathbf{x}_i^{(t)}$ for all $i \in [3]$ and $t \in [3]$. Then, the server can
 141 single out individual model, e.g., $\mathbf{x}_1 = (\mathbf{x}^{(1)} + \mathbf{x}^{(3)} - \mathbf{x}^{(2)})/2$. Similarly, the server can single out all
 142 individual models \mathbf{x}_i , even if a secure aggregation protocol is employed at each round.

143 In this paper, we study secure aggregation protocols with long-term privacy guarantees (which we
 144 term *multi-round secure aggregation*) for the cross-device FL setup which has not been studied before.

145 We assume that user $i \in [N]$ drops from the protocol at each round with probability p_i . $\mathcal{U}^{(t)}$ denotes
 146 the index set of available users at round t and $\mathbf{u}^{(t)} \in \{0, 1\}^N$ is a vector indicating the available users
 147 such that $\{\mathbf{u}^{(t)}\}_j = \mathbb{1}\{j \in \mathcal{U}^{(t)}\}$, where $\{\mathbf{u}\}_j$ is j -th entry of \mathbf{u} and $\mathbb{1}\{\cdot\}$ is the indicator function.
 148 The server selects K users from $\mathcal{U}^{(t)}$, if $|\mathcal{U}^{(t)}| \geq K$, based on the history of selected users in previous
 149 rounds. If $|\mathcal{U}^{(t)}| < K$, the server skips this round. The local models of the selected users are then
 150 aggregated via a secure aggregation protocol (i.e., by communicating masked models), at the end of
 151 which the server learns the aggregate of the local models of the selected users. Our goal is to design a
 152 user selection algorithm $\mathcal{A}^{(t)} : \{0, 1\}^{t \times N} \times \{0, 1\}^N \rightarrow \{0, 1\}^N$,

$$\mathcal{A}^{(t)}(\mathbf{P}^{(t)}, \mathbf{u}^{(t)}) = \mathbf{p}^{(t)} \text{ such that } \|\mathbf{p}^{(t)}\|_0 \in \{0, K\}, \quad (3)$$

153 to prevent the potential information leakage over multiple rounds, where $\mathbf{p}^{(t)} \in \{0, 1\}^N$ is the
 154 participation vector defined in (2), $\|\mathbf{x}\|_0$ denotes the L_0 -“norm” of a vector \mathbf{x} and K denotes the
 155 number of selected users. We note that $\mathcal{A}^{(t)}$ can be a random function. $\mathbf{P}^{(t)}$ is a matrix representing
 156 the user participation information up to round t , and is termed the *participation matrix*, given by

$$\mathbf{P}^{(t)} = [\mathbf{p}^{(0)}, \mathbf{p}^{(1)}, \dots, \mathbf{p}^{(t-1)}]^\top \in \{0, 1\}^{t \times N}. \quad (4)$$

157 **Key Metrics.** A multi-round secure aggregation protocol can be represented by $\mathcal{A} = \{\mathcal{A}^{(t)}\}_{t \in [J]}$,
 158 where $\mathcal{A}^{(t)}$ is the user selection algorithm at round t defined in (3) and J is the total number of rounds.
 159 The inputs of $\mathcal{A}^{(t)}$ are a random vector $\mathbf{u}^{(t)}$, which indicates the available users at round t , and the
 160 participation matrix $\mathbf{P}^{(t)}$ defined in (4) which can be a random matrix. Given the participation matrix
 161 $\mathbf{P}^{(J)}$, we evaluate the performance of the corresponding multi-round secure aggregation protocol
 162 through the following metrics.

163 1. **Multi-round Privacy Guarantee.** The secure aggregation protocols ensure that the server can
 164 only learn the sum of the local models of some users in each single round, but they do not consider
 165 what the server can learn over the long run. Our multi-round privacy definition extends the
 166 guarantees of the secure aggregation protocols from one round to all rounds by requiring that the
 167 server can only learn a sum of the local models even if the server exploits the aggregate models
 168 of all rounds. That is, our multi-round privacy guarantee is a natural extension of the privacy
 169 guarantee provided by the secure aggregation protocols considering a single training round.
 170 Specifically, a multi-round privacy guarantee T requires that any non-zero partial sum of the
 171 local models that the server can reconstruct, through any linear combination $\mathbf{X}^\top \mathbf{P}^{(J)\top} \mathbf{z}$, where
 172 $\mathbf{z} \in \mathbb{R}^J \setminus \{\mathbf{0}\}$, must be of the form¹

$$\mathbf{X}^\top \mathbf{P}^{(J)\top} \mathbf{z} = \sum_{i \in [n]} a_i \sum_{j \in \mathcal{S}_i} \mathbf{x}_j = a_1 \sum_{j \in \mathcal{S}_1} \mathbf{x}_j + a_2 \sum_{j \in \mathcal{S}_2} \mathbf{x}_j + \dots + a_n \sum_{j \in \mathcal{S}_n} \mathbf{x}_j, \quad (5)$$

173 where $|\mathcal{S}_i| \geq T$, $a_i \neq 0, \forall i \in [n]$ and $n \in \mathbb{Z}^+$. Here all the sets \mathcal{S}_i , the number of sets n , and each
 174 a_i could all depend on \mathbf{z} . In equation (5), we consider the worst-case scenario, where the local
 175 models do not change over the rounds. That is, $\mathbf{X}^{(t)} = \mathbf{X}, \forall t \in [J]$. Intuitively, this guarantee
 176 ensures that the best that the server can do is to reconstruct a partial sum of T local models which
 177 corresponds to the case where $n = 1$. When $T \geq 2$, this condition implies that the server cannot
 178 get any user model from the aggregate models of all training rounds (the best it can obtain is the
 179 sum of two local models).

180 **Remark 1.** (Weaker Privacy Notion). It is worth noting that, a weaker privacy notion would
 181 require that $\|\mathbf{P}^{(J)\top} \mathbf{z}\|_0 \geq T$ when $\mathbf{P}^{(J)\top} \mathbf{z} \neq \mathbf{0}$. When $T = 2$, this definition requires that the server
 182 cannot reconstruct any individual model (the best it can do is to obtain a linear combination of
 183 two local models). This notion, however, allows constructions in the form of $a\mathbf{x}_i + b\mathbf{x}_j$ for any
 184 $a, b \in \mathbb{R} \setminus \{0\}$. When $a \gg b$, however, this is almost the same as recovering \mathbf{x}_i perfectly, hence
 185 this privacy criterion is weaker than that of (5).

186 **Remark 2.** (Multi-round Privacy of Random Selection). In Section 6, we empirically show that a
 187 random selection strategy in which K available users are selected uniformly at random at each
 188 round does not ensure multi-round privacy even with respect to the weaker definition of Remark
 189 1. Specifically, the local models can be reconstructed within a number of rounds that is linear in
 190 N . We also show theoretically in Appendix H that when $\min(N - K, K) \geq cN$, where $c > 0$ is a
 191 constant, then the probability that the server can reconstruct all local models after N rounds is

¹We assume that $w_i = \frac{1}{N}, \forall i \in [N]$ in this paper.

at least $1 - 2e^{-c'N}$ for a constant c' that depends on c . Finally, we show that a random selection scheme in which the users are selected in an i.i.d fashion according to $\text{Bern}(\frac{K}{N(1-p)})$ reveals all local models after N rounds with probability that converges to 1 exponentially fast.

Remark 3. (Worst-Case Assumption). In (5), we considered the worst-case assumption where the models do not change over time. When the local models change over rounds, the multi-round privacy guarantee becomes even stronger as the number of unknowns increases. In Fig. 1 and Appendix H, we empirically show that the conventional secure aggregation schemes leak extensive information of training data even in the realistic settings where the models change over the rounds.

2. **Aggregation Fairness Gap.** The average aggregation fairness gap quantifies the largest gap between any two users in terms of the expected relative number of rounds each user has participated in training. Formally, the average aggregation fairness gap is defined as follows

$$F = \max_{i \in [N]} \limsup_{J \rightarrow \infty} \frac{1}{J} \mathbb{E} \left[\sum_{t=0}^{J-1} \mathbb{1} \{ \{ \mathbf{p}^{(t)} \}_i = 1 \} \right] - \min_{i \in [N]} \liminf_{J \rightarrow \infty} \frac{1}{J} \mathbb{E} \left[\sum_{t=0}^{J-1} \mathbb{1} \{ \{ \mathbf{p}^{(t)} \}_i = 1 \} \right], \quad (6)$$

where $\{ \mathbf{p}^{(t)} \}_i$ is i -th entry of the vector $\mathbf{p}^{(t)}$ and the expectation is over the randomness of the user selection algorithm \mathcal{A} and the user availability. The main intuition behind this definition is that when $F = 0$, all users participate on average on the same number of rounds. This is important to take the different users into consideration equally and our experiments show that the accuracy of the schemes with small F are much higher than the schemes with high F .

3. **Average Aggregation Cardinality.** The aggregation cardinality quantifies the expected number of models to be aggregated per round. Formally, it is defined as

$$C = \liminf_{J \rightarrow \infty} \frac{\mathbb{E} \left[\sum_{t=0}^{J-1} \| \mathbf{p}^{(t)} \|_0 \right]}{J}, \quad (7)$$

where the expectation is over the randomness in \mathcal{A} and the user availability. Intuitively, less number of rounds are needed to converge as more users participate in the training. In fact, as we show in Section 5.2, C directly controls the convergence rate.

3.3 Baseline Schemes

In this subsection, we introduce three baseline schemes for multi-round secure aggregation.

Random Selection. In this scheme, at each round, the server selects K users at random from the set of available users if this is possible.

Random Weighted Selection. This scheme is a modified version of random selection to reduce F when the dropout probabilities of the users are not equal. Specifically, K users are selected at random from the available users with the minimum frequency of participation in the previous rounds.

User Partitioning (Grouping). In this scheme, the users are partitioned into $G = N/K$ equal-sized groups denoted as $\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_G$. At each round, the server selects one of the groups if none of the users in this group has dropped out. If multiple groups are available, to reduce the aggregation fairness gap, the server selects a group including a user with the minimum frequency of participation in previous rounds. If no group is available, the server skips this round.

4 Proposed Scheme: Multi-RoundSecAgg

In this section, we present Multi-RoundSecAgg, which has two components as follows.

- The first component designs a family of sets of users that satisfy the multi-round privacy requirement. The inputs of the first component are the number of users (N), the number of selected users at each round (K), and the desired multi-round privacy guarantee (T). The output is a family of sets of K users satisfying the multi-round privacy guarantee T , termed as a *privacy-preserving family*. This family is represented by a matrix \mathbf{B} , where the rows are the characteristic vectors of these user sets.
- The second component selects a set from this designed family to satisfy the fairness guarantee. The inputs to the second component are the family \mathbf{B} , the set of available users at round t , $\mathcal{U}^{(t)}$, and the frequency of participation of each user. The output is the set of users that will participate at round t .

We now describe these two components in detail.

236 **Component 1 (Batch Partitioning (BP) of the users to guarantee multi-round privacy).** The
 237 first component designs a family of R_{BP} sets, where R_{BP} is the size of the set, satisfying the multi-
 238 round privacy requirement T . We denote the $R_{BP} \times N$ binary matrix corresponding to these sets by
 239 $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_{R_{BP}}]^\top$, where $\|\mathbf{b}_i\|_0 = K, \forall i \in [R_{BP}]$. That is, the rows of \mathbf{B} are the characteristic
 240 vectors of those sets. The main idea of our scheme is to restrict certain sets of users of size T , denoted
 241 as batches, to either participate together or not participate at all. This guarantees a multi-round privacy
 242 T as we show in Section 5.

243 To construct a family of sets with this property, the users are first partitioned into N/T
 244 batches. At any given round, either all or none of the users of a particular batch participate
 245 in training. The server can choose K/T batches to participate in training, provided that all
 246 users in any given selected batch are available. Since there are $\binom{N/T}{K/T}$ possible sets with
 247 this property, then the size of this privacy-preserving family of sets is given by $R_{BP} \stackrel{\text{def}}{=} \binom{N/T}{K/T}^2$.

248 In the extreme case of $T = 1$, this strategy specializes to random
 249 selection where the server can choose any K possible users. In
 250 the other extreme case of $T = K$, this strategy specializes to
 251 the partitioning strategy where there are N/K possible sets. We
 252 next provide an example to illustrate the construction of \mathbf{B} .

$$\mathbf{B} = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}$$

253 **Example 1** ($N = 8, K = 4, T = 2$). In this example, the users
 254 are partitioned into 4 batches as $\mathcal{G}_1 = \{1, 2\}, \mathcal{G}_2 = \{3, 4\}, \mathcal{G}_3 =$
 255 $\{5, 6\}$ and $\mathcal{G}_4 = \{7, 8\}$ as given in Fig. 2. The server can choose any two batches out of these 4
 256 batches, hence we have $R_{BP} = \binom{4}{2} = 6$ possible sets. This ensures a multi-round privacy $T = 2$.

Figure 2: Example of our construction with $N = 8, K = 4$ and $T = 2$.

257 **Component 2 (Available batch selection to guarantee fairness).** At round t , user $i \in [N]$ is
 258 available to participate in the protocol with a probability $1 - p_i \in (0, 1]$. The frequency of participation
 259 of user i before round t is denoted by $f_i^{(t)} \stackrel{\text{def}}{=} \sum_{j=0}^{t-1} \mathbb{1} \{ \{p^{(j)}\}_i = 1 \}$. Given the set of available users at
 260 round t , $\mathcal{U}^{(t)}$, and the frequencies of participation $\mathbf{f}^{(t-1)} = (f_1^{(t-1)}, \dots, f_N^{(t-1)})$, the server selects
 261 K users. To do so, the server first finds the submatrix of \mathbf{B} denoted by $\mathbf{B}^{(t)}$ corresponding to $\mathcal{U}^{(t)}$.
 262 Specifically, the i -th row of \mathbf{B} denoted by \mathbf{b}_i^\top is included in $\mathbf{B}^{(t)}$ provided that $\text{supp}(\mathbf{b}_i) \subseteq \mathcal{U}^{(t)}$. If
 263 $\mathbf{B}^{(t)}$ is an empty matrix, then the server skips this round. Otherwise, the server selects a row from $\mathbf{B}^{(t)}$
 264 uniformly at random if $p_i = p, \forall i \in [N]$. If the users have different p_i , the server selects a row from
 265 $\mathbf{B}^{(t)}$ that includes the user with the minimum frequency of participation $\ell_{\min}^{(t-1)} \stackrel{\text{def}}{=} \arg \min_{i \in \mathcal{U}^{(t)}} f_i^{(t-1)}$.
 266 If there are many such rows, then the server selects one of them uniformly at random.

267 **Remark 4.** (Necessity of the Second Component). The second component is necessary to guarantee
 268 that the aggregation fairness gap goes to zero as we show in Theorem 1 and Section 6.

269 Overall, the algorithm first designs a privacy-preserving family of sets to ensure the multi-round
 270 privacy guarantee T . Then specific sets are selected from this family to ensure fairness. We describe
 271 the two components of Multi-RoundSecAgg in detail in Algorithm 1 and Algorithm 2 in Appendix D.

272 5 Theoretical Results

273 In this section, we provide the theoretical guarantees of Multi-RoundSecAgg in Section 5.1 and the
 274 convergence analysis of Multi-RoundSecAgg in Section 5.2.

275 5.1 Theoretical Guarantees of Multi-RoundSecAgg

276 In this subsection, we establish the theoretical guarantees of Multi-RoundSecAgg in terms of the
 277 multi-round privacy guarantee, the aggregation fairness gap and the average aggregation cardinality.

278 **Theorem 1.** Multi-RoundSecAgg with parameters N, K, T ensures a multi-round privacy guarantee
 279 of T , an aggregation fairness gap $F = 0$, and an average aggregation cardinality given by

$$C = K \left(1 - \sum_{i=N/T-K/T+1}^{N/T} \binom{N/T}{i} q^i (1-q)^{N/T-i} \right),$$

²We assume for simplicity that N/T and K/T are integers.

280 where $q = 1 - (1 - p)^T$, when all users have the dropout probability p .

281 We provide the proof of Theorem 1 in Appendix A.

282 **Remark 5.** (Trade-off between ‘‘Multi-round Privacy Guarantee’’ and ‘‘Average Aggregation
283 Cardinality’’). Theorem 1 indicates a trade-off between the multi-round privacy and the average
284 aggregation cardinality since as T increases, C decreases which slows down the convergence as we
285 show in Sec. 5.2. We show this trade-off in Fig. 3.

286 **Remark 6.** (Necessity of Batch Partitioning (BP)). We show that any strategy that satisfies the privacy
287 guarantee in Equation (5) must have a batch partitioning structure, and for given $N, K, T, K \leq N/2$,
288 the largest number of distinct user sets in any strategy is at most $\binom{N/T}{K/T}$, which is achieved in our
289 design in Section 4. We provide the proof in Appendix C.

290 **Remark 7.** (Non-linear Reconstructions of Aggregated Models). The privacy criterion in Eq. (5)
291 considers linear reconstructions of the aggregated models. One may also consider more general
292 non-linear reconstructions. The long-term privacy guarantees of batch partitioning hold even under
293 such reconstructions as the users in the same batch always participate together or do not participate
294 at all. Hence, the server cannot separate individual models within the same batch even through
295 non-linear operations.

296 5.2 Convergence Analysis of Multi-RoundSecAgg

297 For convergence analysis of Multi-RoundSecAgg, we
298 first introduce a few common assumptions [23, 39].

299 **Assumption 1.** L_1, \dots, L_N in (1) are all ρ -smooth:
300 for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ and $i \in [N]$, $L_i(\mathbf{a}) \leq L_i(\mathbf{b}) + (\mathbf{a} -$
301 $\mathbf{b})^\top \nabla L_i(\mathbf{b}) + \frac{\rho}{2} \|\mathbf{a} - \mathbf{b}\|^2$.

302 **Assumption 2.** L_1, \dots, L_N in (1) are all μ -strongly
303 convex: for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ and $i \in [N]$, $L_i(\mathbf{a}) \geq$
304 $L_i(\mathbf{b}) + (\mathbf{a} - \mathbf{b})^\top \nabla L_i(\mathbf{b}) + \frac{\mu}{2} \|\mathbf{a} - \mathbf{b}\|^2$.

305 **Assumption 3.** Let $\xi_i^{(t)}$ be a sample uniformly selected from the dataset \mathcal{D}_i . The variance of the
306 stochastic gradients at each user is bounded, i.e., $\mathbb{E}\|\nabla L_i(\mathbf{x}_i^{(t)}, \xi_i^{(t)}) - \nabla L_i(\mathbf{x}_i^{(t)})\|^2 \leq \sigma_i^2$ for $i \in [N]$.

307 **Assumption 4.** The expected squared norm of the stochastic gradients is uniformly bounded, i.e.,
308 $\mathbb{E}\|\nabla L_i(\mathbf{x}_i^{(t)}, \xi_i^{(t)})\|^2 \leq G^2$ for all $i \in [N]$.

309 We now state the convergence guarantees of Multi-RoundSecAgg.

310 **Theorem 2.** Consider a FL setup with N users to train a machine learning model from (1). Assume
311 K users are selected by Multi-RoundSecAgg with average aggregation cardinality C defined in (7) to
312 update the global model from (2), and all users have the same dropout rate, hence Multi-RoundSecAgg
313 selects a random set of K users uniformly from the set of available user sets at each round. Then, the
314 following is satisfied

$$\mathbb{E}[L(\mathbf{x}^{(J)})] - L^* \leq \frac{\rho}{\gamma + \frac{C}{K}EJ - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2} \mathbb{E}\|\mathbf{x}^{(0)} - \mathbf{x}^*\|^2 \right), \quad (8)$$

315 where $\alpha = \frac{1}{N} \sum_{i=1}^N \sigma_i^2 + 6\rho\Gamma + 8(E-1)^2G^2$, $\beta = \frac{4(N-K)E^2G^2}{K(N-1)}$, $\Gamma = L^* - \sum_{i=1}^N L_i^*$, and $\gamma = \max\left\{\frac{8\rho}{\mu}, E\right\}$.

316 We provide the proof of Theorem 2 in Appendix B.

317 **Remark 8.** (The average aggregation cardinality controls the convergence rate.) Theorem 2 shows
318 how the average aggregation cardinality affects the convergence. When the average aggregation
319 cardinality is maximized, i.e., $C = K$, the convergence rate in Theorem 2 equals that of the random
320 selection algorithm provided in Theorem 3 of [23]. In (8), we have the additional term E (number of
321 local epochs) in front of J compared to Theorem 3 of [23] as we use global round index t instead of
322 using step index of local SGD. As the average aggregation cardinality decreases, a greater number of
323 training rounds is required to achieve the same level of accuracy.

324 **Remark 9.** (General Convex and Non-Convex Convergence Rates). Theorem 2 considers the
325 strongly-convex case, but we consider the general convex and the non-convex cases in Appendix I.

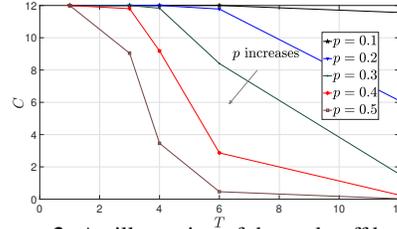


Figure 3: An illustration of the trade-off between the multi-round privacy guarantee T and the average aggregation cardinality C . In this example, $N = 120$ and $K = 12$.

326 **Remark 10.** (Different Dropout Rates). When the dropout probabilities of the users are not the same,
 327 characterizing the convergence guarantees of Multi-RoundSecAgg is challenging. This is due to the
 328 fact that batch selection based on the frequency of participation breaks the conditional unbiasedness
 329 of the user selection, which is required for the convergence guarantee. In experiments, however, we
 330 empirically show that Multi-RoundSecAgg guarantees the convergence with different dropout rates.

331 6 Experiments

332 Our experiments consist of two parts. We first numerically demonstrate the performance of Multi-
 333 RoundSecAgg compared to the baselines of Section 3.3 in terms of the key metrics of Section 3.2.
 334 Next, we implement convolutional neural networks (CNNs) for image classification with MNIST [21],
 335 CIFAR-10, and CIFAR-100 [20] to investigate how the key metrics affect the test accuracy.

336 **Setup.** We consider a FL setting with $N = 120$ users, where the server aims to choose $K = 12$ users
 337 at every round. We study two settings for partitioning the CIFAR-100 dataset across the users.

- 338 • **IID Setting.** 50000 training samples are shuffled and partitioned uniformly across $N = 120$ users.
- 339 • **Non-IID Setting.** We distribute the dataset using a Dirichlet distribution [13], which samples
 340 $\mathbf{d}_c \sim \text{Dir}(\beta = 0.5)$ which specifying the prior class distribution over 100 classes, and allocate a
 341 portion $d_{c,i}$ of the class c to user i . The parameter β controls the heterogeneity of the distributions
 342 at each user, where $\beta \rightarrow \infty$ results in IID setting.

343 We implement a VGG-11 [29], which is sufficient for our needs, as our goal is to evaluate various
 344 schemes, not to achieve the best accuracy. The hyperparameters are provided in Appendix F.

345 **Modeling dropouts.** To model heterogeneous system, users have different dropout probability p_i
 346 selected from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. At each round, user $i \in [N]$ drops with probability p_i .

347 **Implemented Schemes.** For the benchmarks, we
 348 implement the three baselines introduced in Sec. 3.3,
 349 referred to as *Random*, *Weighted Random*, and *Partition*.
 350 For Multi-RoundSecAgg, we construct three privacy-
 351 preserving families with different target multi-round
 352 privacy guarantees, $T = 6$, $T = 4$, and $T = 3$ which
 353 we refer to as Multi-RoundSecAgg ($T = 6$), Multi-
 354 RoundSecAgg ($T = 4$), and Multi-RoundSecAgg ($T =$
 355 3), respectively. One can view the Random and Partition as extreme cases of Multi-RoundSecAgg
 356 with $T = 1$ and $T = K$, respectively. Table 1 summarizes the family size R defined in Section 4.

Scheme	Family size (= R)
Random selection	$\sim 10^{16}$
Weighted random selection	$\sim 10^{16}$
User partition	10
Multi-RoundSecAgg, $T=6$	190
Multi-RoundSecAgg, $T=4$	4060
Multi-RoundSecAgg, $T=3$	91389

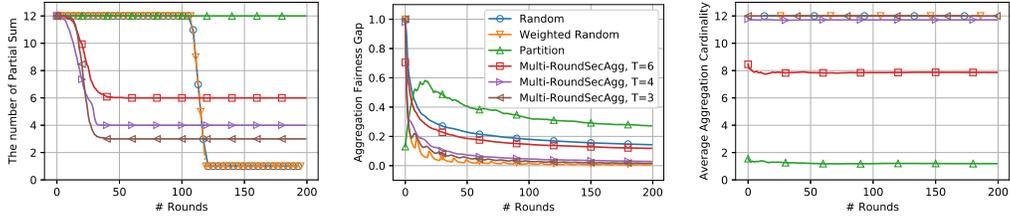
Table 1: Family size with $N = 120$, $K = 12$.

357 **Key Metrics.** To numerically demonstrate the performance of the six schemes in terms of the key
 358 metrics defined in Sec. 3.2, at each round, we measure the following metrics.

- 359 • For the multi-round privacy guarantee, we measure the number of models in the partial sum that
 360 the server can reconstruct, which is given by $T^{(t)} := \min_{\mathbf{z} \in \mathbb{R}^J} \|\mathbf{z}^\top \mathbf{P}^{(t)}\|_0$, s.t. $\mathbf{P}^{(t)\top} \mathbf{z} \neq \mathbf{0}$. This
 361 corresponds to the weaker privacy definition of Remark 1. We use this weaker privacy definition
 362 as the random selection and the random weighted selection strategies provide the worst privacy
 363 guarantee even with this weaker definition, as demonstrated later. On the other hand, Multi-
 364 RoundSecAgg provides better privacy guarantees with both the strong and the weaker definitions.
- 365 • For the aggregation fairness gap, we measure the instantaneous fairness gap, $F^{(t)} :=$
 366 $\max_{i \in [N]} F_i^{(t)} - \min_{i \in [N]} F_i^{(t)}$ where $F_i^{(t)} = \frac{1}{t+1} \sum_{l=0}^t \mathbb{1}\{\{\mathbf{p}^{(l)}\}_i = 1\}$.
- 367 • We measure the instantaneous aggregation cardinality as $C^{(t)} := \frac{1}{t+1} \sum_{l=0}^t \|\mathbf{p}^{(l)}\|_0$.

368 We demonstrate these key metrics in Figure 4. We make the following key observations.

- 369 • Multi-RoundSecAgg achieves better multi-round privacy guarantee than both the random selection
 370 and random weighted selection strategies, while user partitioning achieves the best multi-round
 371 privacy guarantee, $T = K = 12$. However, the partitioning strategy has the worst aggregation
 372 cardinality, which results in the lowest convergence rate as demonstrated later.
- 373 • Figure 5 demonstrates the trade-off between the multi-round privacy guarantee T and the average
 374 aggregation cardinality C . Interestingly, Multi-RoundSecAgg when $T = 3$ or $T = 4$ achieves better
 375 multi-round privacy guarantee than both the random selection and the weighted random selection
 376 strategies while achieving almost the same average aggregation cardinality.



(a) Multi-round privacy guarantee. (b) Aggregation fairness gap. (c) Average aggregation cardinality.

Figure 4: The key metrics with $N = 120$ (number of users), $K = 12$ (number of selected users at each round).

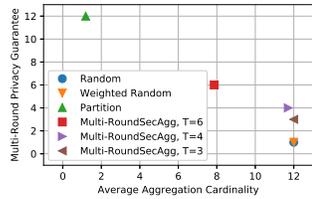
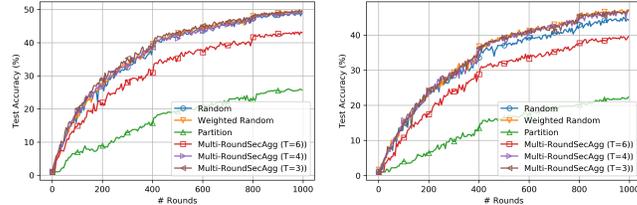


Figure 5: Trade-off between multi-round privacy and average aggregation cardinality with $N = 120$, $K = 12$.



(a) IID data distribution. (b) Non-IID data distribution.

Figure 6: Training rounds versus test accuracy of VGG11 in [29] on the CIFAR-100 with $N = 120$ and $K = 12$.

377 **Remark 11.** (Multi-round Privacy of Random and Weighted Random). The multi-round privacy
 378 guarantees of Random and Weighted Random drop sharply as shown in Fig. 4(a) as the participating
 379 matrix $\mathbf{P}^{(t)} \in \{0, 1\}^{t \times N}$ becomes full rank with high probability when $t \geq N$, and hence the server can
 380 reconstruct the individual models by utilizing a pseudo inversion of the matrix $\mathbf{P}^{(t)}$. More precisely,
 381 Theorem 3 in Appendix H shows this *thresholding phenomenon*, where the probability that the server
 382 can reconstruct individual models after certain number of rounds converges to 1 exponentially fast.

383 **Key Metrics versus Test Accuracy.** To investigate how the key metrics affect the test accuracy, we
 384 measure the test accuracy of the six schemes in the two settings, the IID and the non-IID settings.
 385 Our results are demonstrated in Figure 6. We make the following key observations.

- 386 • In the IID setting, the Multi-RoundSecAgg schemes show test accuracies that are comparable to the
 387 random selection and random weighted selection schemes while the Multi-RoundSecAgg schemes
 388 provide higher levels of privacy. Specifically, the Multi-RoundSecAgg schemes achieve $T = 3, 4, 6$
 389 based on the privacy-preserving family design while the random selection and random weighted
 390 selection schemes have $T = 1$, i.e., the server can learn an individual local model.
- 391 • In the non-IID setting, Multi-RoundSecAgg not only outperforms the random selection scheme but
 392 also achieves a smaller aggregation fairness gap as demonstrated in Fig. 4(b).
- 393 • In both IID and non-IID settings, the user partitioning scheme has the worst accuracy as its average
 394 aggregation cardinality is much smaller than the other schemes as demonstrated in Fig. 4(c).

395 We also implement additional experiments on MNIST and CIFAR-10 datasets in Appendix E and
 396 present ablation study for various settings of (N, K, T) in Appendix G

397 7 Conclusion

398 Partial user participation may breach user privacy in federated learning, even if secure aggregation is
 399 employed at every training round. To address this challenge, we introduced the notion of long-term
 400 privacy, which ensures that the privacy of individual models are protected over all training rounds. We
 401 developed Multi-RoundSecAgg, a structured user selection strategy that guarantees long-term privacy
 402 while taking into account the fairness in user selection and average number of participating users,
 403 and showed that Multi-RoundSecAgg provides a trade-off between long-term privacy and average
 404 number of participating users (hence the convergence rate). Our experiments on the CIFAR-100,
 405 CIFAR-10, and MNIST datasets on both the IID and non-IID settings show that Multi-RoundSecAgg
 406 achieves comparable accuracy to the random selection strategy (which does not ensure long-term
 407 privacy), while ensuring long-term privacy guarantees.

408 **References**

- 409 [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar,
410 and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC*
411 *conference on computer and communications security*, pages 308–318, 2016.
- 412 [2] James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova.
413 Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020*
414 *ACM SIGSAC Conference on Computer and Communications Security*, pages 1253–1269, 2020.
- 415 [3] Kallista Bonawitz, Peter Kairouz, Brendan McMahan, and Daniel Ramage. Federated learning
416 and privacy: Building privacy-preserving systems for machine learning and data science on
417 decentralized data. *Queue*, 19(5):87–114, 2021.
- 418 [4] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan,
419 Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for
420 privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on*
421 *Computer and Communications Security*, pages 1175–1191, 2017.
- 422 [5] Wenlin Chen, Samuel Horvath, and Peter Richtarik. Optimal client sampling for federated
423 learning. *arXiv preprint arXiv:2010.13723*, 2020.
- 424 [6] Yae Jee Cho, Samarth Gupta, Gauri Joshi, and Osman Yağan. Bandit-based communication-
425 efficient client selection strategies for federated learning. *arXiv preprint arXiv:2012.08009*,
426 2020.
- 427 [7] Yae Jee Cho, Jianyu Wang, and Gauri Joshi. Client selection in federated learning: Convergence
428 analysis and power-of-choice selection strategies. *arXiv preprint arXiv:2010.01243*, 2020.
- 429 [8] Beongjun Choi, Jy-yong Sohn, Dong-Jun Han, and Jaekyun Moon. Communication-
430 computation efficient secure aggregation for federated learning. *arXiv preprint*
431 *arXiv:2012.05433*, 2020.
- 432 [9] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy.
433 *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- 434 [10] Ahmed Roushdy Elkordy and A Salman Avestimehr. Secure aggregation with heterogeneous
435 quantization in federated learning. *arXiv preprint arXiv:2009.14388*, 2020.
- 436 [11] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit
437 confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC*
438 *Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- 439 [12] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients—
440 how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*,
441 2020.
- 442 [13] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical
443 data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.
- 444 [14] Vishesh Jain, Ashwin Sah, and Mehtaab Sawhney. Singularity of discrete random matrices ii.
445 *arXiv preprint arXiv:2010.06554*, 2020.
- 446 [15] Swanand Kadhe, Nived Rajaraman, O Ozan Koyluoglu, and Kannan Ramchandran. Fastsecagg:
447 Scalable secure aggregation for privacy-preserving federated learning. *arXiv preprint*
448 *arXiv:2009.11248*, 2020.
- 449 [16] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis,
450 Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings,
451 et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine*
452 *Learning*, 14(1–2):1–210, 2021.

- 453 [17] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis,
454 Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings,
455 et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*,
456 2019.
- 457 [18] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and
458 Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In
459 *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.
- 460 [19] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly,
461 and Neil Houlsby. Big transfer (bit): General visual representation learning. In *Computer Vision–
462 ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part
463 V 16*, pages 491–507. Springer, 2020.
- 464 [20] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images.
465 Technical report, Citeseer, 2009.
- 466 [21] Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database. [http://yann.
467 lecun.com/exdb/mnist](http://yann.lecun.com/exdb/mnist), 2010.
- 468 [22] Yann LeCun, Patrick Haffner, Léon Bottou, and Yoshua Bengio. Object recognition with
469 gradient-based learning. In *Shape, contour and grouping in computer vision*, pages 319–345.
470 Springer, 1999.
- 471 [23] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zihua Zhang. On the convergence
472 of fedavg on non-iid data. In *International Conference on Learning Representations*, 2019.
- 473 [24] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas.
474 Communication-efficient learning of deep networks from decentralized data. In *Int. Conf. on
475 Artificial Int. and Stat. (AISTATS)*, pages 1273–1282, 2017.
- 476 [25] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially
477 private recurrent language models. *Int. Conf. on Learning Representations (ICLR)*, 2018.
- 478 [26] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep
479 learning: Passive and active white-box inference attacks against centralized and federated
480 learning. In *2019 IEEE symposium on security and privacy (SP)*, pages 739–753. IEEE, 2019.
- 481 [27] Balázs Pejó and Gergely Biczók. Quality inference in federated learning with secure aggregation.
482 *arXiv preprint arXiv:2007.06236*, 2020.
- 483 [28] Monica Ribero and Haris Vikalo. Communication-efficient federated learning via optimal client
484 sampling. *arXiv preprint arXiv:2007.15197*, 2020.
- 485 [29] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale
486 image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- 487 [30] Jinhyun So, Ramy E Ali, Başak Güler, and A Salman Avestimehr. Secure aggregation for
488 buffered asynchronous federated learning. *arXiv preprint arXiv:2110.02177*, 2021.
- 489 [31] Jinhyun So, Başak Güler, and A Salman Avestimehr. Turbo-aggregate: Breaking the quadratic
490 aggregation barrier in secure federated learning. *IEEE Journal on Selected Areas in Information
491 Theory*, 2(1):479–489, 2021.
- 492 [32] Jinhyun So, Corey J Nolet, Chien-Sheng Yang, Songze Li, Qian Yu, Ramy E Ali, Basak Guler,
493 and Salman Avestimehr. Lightsecagg: a lightweight and versatile design for secure aggregation
494 in federated learning. *Proceedings of Machine Learning and Systems*, 4:694–720, 2022.
- 495 [33] Sebastian U Stich. Local sgd converges fast and communicates little. *arXiv preprint
496 arXiv:1805.09767*, 2018.
- 497 [34] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural
498 networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.

- 499 [35] Minxue Tang, Xuefei Ning, Yitu Wang, Yu Wang, and Yiran Chen. Fedgp: Correlation-based
500 active client selection for heterogeneous federated learning. *arXiv preprint arXiv:2103.13822*,
501 2021.
- 502 [36] Tuan Tran. The smallest singular value of random combinatorial matrices. *arXiv preprint*
503 *arXiv:2007.06318*, 2020.
- 504 [37] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS
505 Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and
506 performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–
507 3469, 2020.
- 508 [38] Chien-Sheng Yang, Jinhyun So, Chaoyang He, Songze Li, Qian Yu, and Salman Avestimehr.
509 Lightsecagg: Rethinking secure aggregation in federated learning. *arXiv preprint*
510 *arXiv:2109.14236*, 2021.
- 511 [39] Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less
512 communication: Demystifying why model averaging works for deep learning. In *Proceedings*
513 *of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5693–5700, 2019.
- 514 [40] Yizhou Zhao and Hua Sun. Information theoretic secure aggregation with user dropouts. *arXiv*
515 *preprint arXiv:2101.07750*, 2021.
- 516 [41] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated
517 learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- 518 [42] Ligeng Zhu and Song Han. Deep leakage from gradients. In *Federated Learning*, pages 17–31.
519 Springer, 2020.