How Much Privacy Does Federated Learning with Secure Aggregation Guarantee?

Ahmed Roushdy Elkordy[1], Jiang Zhang[1], Yahya H. Ezzeldin[1], Konstantinos Psounis[1], Salman Avestimehr[1]
[1]University of Southern California

Motivation

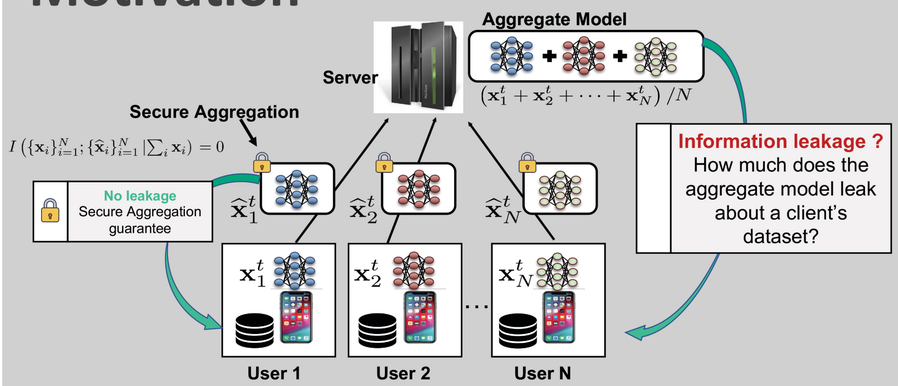

Theoretical Results

Theorem: FedSGD with Secure aggregation guarantees

(Single Round Leakage) $$I\left(\mathbf{x}_i^{(t)}; \frac{1}{N}\sum_{i\in\mathcal{N}}\mathbf{x}_i^{(t)} \middle| \sum_{i\in\mathcal{N}}\mathbf{x}_i^{(t-1)}\right) \leq \frac{\overset{\star}{C}d}{(N-1)B} + \frac{d}{2}\log\left(\frac{N}{N-1}\right)$$

(Multi-round Leakage) $$I\left(\mathcal{D}_1; \left\{\sum_{i=1}^N \mathbf{x}_i^t\right\}_{t=1}^T\right) \leq T\left[\frac{\overset{\star}{C}d}{(N-1)B} + \frac{d}{2}\log\left(\frac{N}{N-1}\right)\right]$$

Notations $\mathbf{x}_i^{(t)}$: mini-batch gradient N: number of nodes B: batch size d: model size
$d^\star$: $\text{rank}(\mathbf{K}_g^{(t)})$ $d^\star \leq d$ T : training rounds D_i : local data of node i

Main theoretical takeaways:

- Leakage decreases with increasing number of users at a rate O(N)
- Leakage decreases with increasing Batch size
- Increasing the model size does not have a linear impact on the leakage
- More training rounds can potentially cause a linear increase in leakage

Empirical Evaluation Results

We use the Mutual Information Neural Estimator (MINE) to estimate our mutual information expressions in the theory (multiple training sessions treated as samples).

[Impact of number of users and model size]

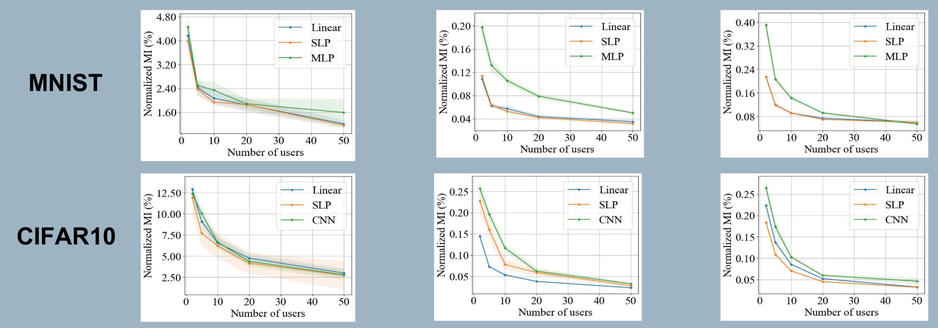

Increasing number of users decreases the leakage

Increasing the model size increases the leakage

[Impact of batch size]

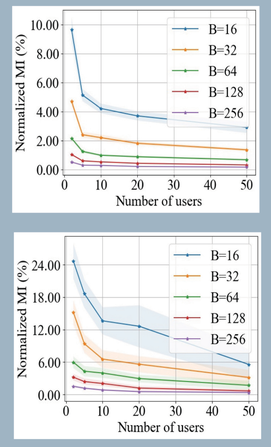

Increasing the batch size decreases the leakage

[Impact of data heterogeneity]

Dataset: FEMNIST (written by 3500 writers and include 62 classes)

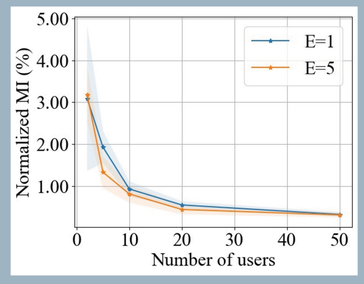

Leakage also decreases with N in the presence of heterogeneous dataset

[Privacy as a cost for performance]

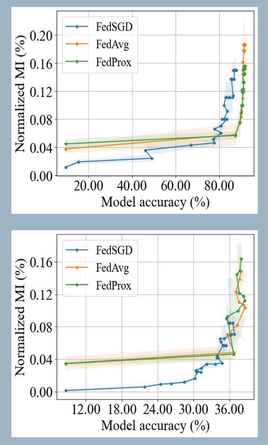

After a number of training rounds, the increase in accuracy is small per round, but accumulated leakage is still reasonable.

ProperData Frontier 2nd Annual Symposium
Dec 1-2, 2022 | Irvine, California

SaTC Frontiers: Collaborative: *1956393* [1], *1955227* [2], *2103439* [3], *1956435* [4]
Institutions: UC Irvine [1], Northeastern Univ. [2], UC Davis [3], USC [4]
Lead PIs: A. Markopoulou [1], D. Choffnes [2], Z. Shafiq [3], K. Psounis [4]