

COMPRESS-THEN-ADD-NOISE: A NEW MECHANISM FOR PRIVATE DISTRIBUTED MEAN ESTIMATION AND FEDERATED LEARNING

CycleResearcher

ABSTRACT

Motivated by applications to private federated learning (PFL), we consider the problem of differentially private distributed mean estimation under communication constraints. Prior work has developed algorithms to reduce communication costs in FL by compressing model updates sent from clients to the server, e.g., via gradient compression, sparsification, or quantization. These algorithms have proven effective at reducing the overall amount of communication in FL, but do not interact well with privacy-preserving mechanisms: they are typically applied on top of a privacy-preserving mechanism, after the privacy-preserving noise has been added. Our key observation is that since the noise contains relatively little information about the data, we can apply such compression mechanisms to the noise itself. Building on this, we propose a new mechanism, Compress-then-Add-Noise (CAN), where we reverse the order of adding noise and compression. We show that this mechanism is differentially private and can be used to achieve significant reductions in communication costs, while maintaining the same level of privacy. We then apply CAN to PFL, and show that it can be used to improve the model accuracy while reducing the communication costs of the state-of-the-art DP-FTRL algorithm.

1 INTRODUCTION

One of the main bottlenecks in federated learning (FL) is the communication cost of sending model updates from clients to the server Kairouz et al. (2019). A large body of recent work has developed algorithms to reduce this cost, e.g., via gradient compression Wangni et al. (2018); Lin et al. (2018), sparsification Wangni et al. (2018); Lin et al. (2018), or quantization Alistarh et al. (2017); Wen et al. (2017). These algorithms have proven effective at reducing the overall amount of communication in FL, but do not interact well with privacy-preserving mechanisms: they are typically applied on top of a privacy-preserving mechanism, after the privacy-preserving noise has been added.

In this work, we propose a new mechanism, Compress-then-Add-Noise (CAN), for differentially private distributed mean estimation. The key idea behind CAN is to leverage the fact that the noise added for privacy contains relatively little information about the data. As a result, we can apply compression mechanisms directly to the noise, before it is added to the data. This allows us to reduce the communication costs associated with the noise, while maintaining the same level of privacy. We then show how CAN can be applied to private federated learning (PFL) to improve the model accuracy while reducing the communication costs of the state-of-the-art DP-FTRL algorithm Kairouz et al. (2021).

1.1 OUR CONTRIBUTIONS

We make the following contributions:

- We introduce the CAN mechanism for differentially private distributed mean estimation. This mechanism reverses the traditional order of operations by first compressing the data and then adding noise, leveraging the fact that the noise contains little information about the data.
- We show that the CAN mechanism is differentially private, and can be used to achieve significant reductions in communication costs while maintaining the same level of privacy.

- We apply the CAN mechanism to PFL, and show that it can be used to improve the model accuracy while reducing the communication costs of the state-of-the-art DP-FTRL algorithm Kairouz et al. (2021).

1.2 RELATED WORK

Differential Privacy. Differential privacy (DP) is a mathematical framework for preserving the privacy of individuals in a dataset D [work et al. \(2006\)](#). It has been widely adopted in machine learning (ML) to ensure that models trained on sensitive data do not leak information about individual data points. One of the main challenges in DP ML is to balance the trade-off between privacy and utility, as adding noise to preserve privacy can degrade the accuracy of the model.

To address this challenge, a large body of recent work has developed algorithms to improve the accuracy of DP ML models. For example, [Denisov et al. \(2022\)](#); [Choquette-Choo et al. \(2022; 2023a;b\)](#) use matrix factorization techniques to add correlated noise, which has been shown to improve the accuracy of DP models. [Hu et al. \(2021\)](#); [Farokhi \(2021\)](#) use sparsification to reduce the amount of noise added to the model, while [Andrew et al. \(2021\)](#) use adaptive clipping to scale the amount of noise added to the model based on the sensitivity of the data.

Communication-Efficient Federated Learning. FL is a distributed learning framework where multiple clients train a model collaboratively [McMahan et al. \(2016\)](#). One of the main challenges in FL is to reduce the communication costs between clients and the server, as transmitting large amounts of data can be expensive and time-consuming.

To address this challenge, a large body of recent work has developed algorithms to reduce the communication costs in FL. For example, [Alistarh et al. \(2017\)](#); [Wen et al. \(2017\)](#) use quantization to compress the model updates sent from clients to the server, while [Wangni et al. \(2018\)](#); [Lin et al. \(2018\)](#) use sparsification to reduce the number of non-zero elements in the model updates. [Rothchild et al. \(2020\)](#) use Count Sketch to compress the model updates, while [Isik et al. \(2023b\)](#) learn how to sparsify the random network for the best performance. These algorithms have proven effective at reducing the overall amount of communication in FL, but do not interact well with privacy-preserving mechanisms: they are typically applied on top of a privacy-preserving mechanism, after the privacy-preserving noise has been added.

Differentially Private Federated Learning. To address the privacy concerns in FL, a large body of recent work has developed algorithms to ensure that the model updates sent from clients to the server are differentially private. For example, [Kasiviswanathan et al. \(2011\)](#); [Abadi et al. \(2016\)](#) add Gaussian or Laplace noise to the model updates before sending them to the server, while [Kairouz et al. \(2021\)](#) use the Follow-The-Regularized-Leader (FTRL) algorithm to update the model parameters in a differentially private manner. These algorithms have proven effective at preserving the privacy of the data, but can increase the communication costs associated with the model updates.

Some recent works have also considered the communication costs of DP FL. For example, [Chen et al. \(2023\)](#) use compression to reduce the communication costs of the model updates, while [Shah et al. \(2022\)](#) use Minimal Random Coding to compress the local randomizers in LDP mean estimation. However, these algorithms do not interact well with the privacy-preserving mechanisms: they are typically applied on top of a privacy-preserving mechanism, after the privacy-preserving noise has been added.

Distributed Mean Estimation. The distributed mean estimation problem is a fundamental problem in distributed learning, where the goal is to estimate the mean of a set of vectors distributed across multiple machines or clients [Suresh et al. \(2017\)](#); [Agarwal et al. \(2018\)](#). This problem has been studied extensively in the literature, with many algorithms proposed to solve it [Chen et al. \(2020\)](#); [Vargaftik et al. \(2021\)](#); [Isik et al. \(2023a\)](#).

One of the main challenges in distributed mean estimation is to reduce the communication costs between the machines or clients and the server. To address this challenge, a large body of recent work has developed algorithms to compress the vectors before sending them to the server [Chen et al. \(2020\)](#); [Vargaftik et al. \(2021\)](#); [Isik et al. \(2023a\)](#). These algorithms have proven effective at reducing the overall amount of communication in distributed mean estimation, but do not interact well with

privacy-preserving mechanisms: they are typically applied on top of a privacy-preserving mechanism, after the privacy-preserving noise has been added.

Continual Observation and Binary Tree Mechanism. The problem of maintaining a differentially private sum under continual observation has been studied extensively in the literature Dwork et al. (2010); Chan et al. (2012); Honaker (2015); Guha Thakurta & Smith (2013); Jain et al. (2023). One of the most popular mechanisms for this problem is the Binary Tree Mechanism (BTM) Chan et al. (2012), which adds noise to each node of a binary tree and then estimates the sum by summing the noisy nodes along the path from the root to the leaf corresponding to the current time step.

The BTM has been extended to the matrix mechanism framework Li et al. (2015); Denisov et al. (2022); Choquette-Choo et al. (2022), which allows for more flexible and efficient noise addition. However, the matrix mechanism does not naturally extend to the continual observation setting, as it is designed for the batch release setting. To address this challenge, Denisov et al. (2022); Choquette-Choo et al. (2022) use an extension of the matrix mechanism to the continual observation setting, which involves adding noise to the matrix product of the data and a factorization of the query matrix. This extension has been shown to be differentially private and can be used to estimate the sum under continual observation.

However, the continual observation setting assumes that the data is released at each time step, which is not always the case in practice. For example, in PFL, the model updates are typically aggregated by the server and only the final model is released at the end of training. This raises the question of whether the matrix mechanism can be extended to the *final release* setting, where the data is only released at the end of the process. This is an important question, as the final release setting is common in many practical applications and the continual observation setting may not be applicable.

In this work, we address this question by proposing a new extension of the matrix mechanism to the final release setting. This extension involves adding noise to the matrix product of the data and a factorization of the query matrix, and then releasing the final noisy sum at the end of the process. We show that this extension is differentially private and can be used to estimate the sum under final release. This extension is a key component of our CAN mechanism, which we use to achieve significant reductions in communication costs while maintaining the same level of privacy.

2 PRELIMINARIES

2.1 DIFFERENTIAL PRIVACY

We consider the standard central differential privacy (DP) model Dwork et al. (2006), where there is a trusted curator that holds a dataset $D = \{x_1, \dots, x_n\}$ with $x_i \in \mathbb{R}^d$. Two datasets D and D' are said to be neighboring, denoted as $D \sim D'$, if $\|D - D'\|_1 = 1$. A randomized algorithm \mathcal{M} is said to be (ϵ, δ) -DP if for all neighboring datasets D and D' , and for all measurable sets S ,

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta.$$

We use the following version of the Gaussian mechanism to ensure DP Balle & Wang (2018): [Gaussian Mechanism] Let $f : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^m$ be a function with ℓ_2 sensitivity $\Delta = \max_{D \sim D'} \|f(D) - f(D')\|_2$. Then the Gaussian mechanism

$$\mathcal{M}(D) = f(D) + \mathcal{N}(0, \sigma^2 I_m)$$

is (ϵ, δ) -DP for $\sigma \geq \sqrt{2 \log(1.25/\delta)} \Delta / \epsilon$.

We also use the following version of the Poisson subsampling lemma Zhu & Wang (2019); Wang et al. (2019); Balle et al. (2018): [Poisson Subsampling Lemma] Let \mathcal{M} be an (ϵ, δ) -DP algorithm and let D be a dataset. Let D' be a dataset sampled from D by including each data point $x \in D$ independently with probability q . Then the algorithm $\mathcal{M}(D')$ is (ϵ', δ') -DP for $\epsilon' = \frac{q\epsilon}{1-q+q\epsilon}$ and $\delta' = \frac{q\delta}{1-q}$.

2.2 DISTRIBUTED MEAN ESTIMATION

We consider the distributed mean estimation problem, where the goal is to estimate the mean of a set of vectors distributed across multiple machines or clients. Formally, we assume that there are n

clients, each with a vector $x_i \in \mathbb{R}^d$. The goal is to estimate the mean $\mu = \frac{1}{n} \sum_{i=1}^n x_i$ by having each client send a message to the server, which then computes an estimate $\hat{\mu}$.

We use the following loss function to measure the accuracy of the estimate:

$$\mathcal{L}(\hat{\mu}, \mu) = \|\hat{\mu} - \mu\|_2^2.$$

2.3 PRIVATE FEDERATED LEARNING

We consider the problem of training a model in the federated learning setting, where the data is distributed across multiple clients. Formally, we assume that there are n clients, each with a local dataset $D_i = \{x_{i,j}\}_{j=1}^{m_i}$ and a local model parameter $\theta_i \in \mathbb{R}^d$. The goal is to train a global model parameter $\theta \in \mathbb{R}^d$ by having each client send an update to the server, which then computes a new global model parameter.

We use the following loss function to measure the accuracy of the model:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta),$$

where $\mathcal{L}_i(\theta)$ is the local loss function for client i . We assume that each local loss function is L -smooth and μ -strongly convex.

We use the DP-FTRL algorithm Kairouz et al. (2021) to train the global model parameter. The algorithm proceeds in T rounds, where in each round t , each client computes a local model update $\Delta\theta_{i,t} = \nabla \mathcal{L}_i(\theta_{i,t})$ and sends it to the server. The server then computes a new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ :

$$\theta_{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \frac{\lambda}{2} \|\theta\|_2^2 + \sum_{s=1}^t \langle \Delta\theta_{i,s}, \theta \rangle \right\}.$$

The server then sends the new global model parameter θ_{t+1} back to each client, who uses it to compute a new local model update $\Delta\theta_{i,t+1}$. This process is repeated for T rounds, after which the final global model parameter θ_T is released.

To ensure DP, the DP-FTRL algorithm adds Gaussian noise to each local model update before sending it to the server. The amount of noise added is calibrated to the sensitivity of the local model updates, which is bounded by a constant B :

$$\|\Delta\theta_{i,t}\|_2 \leq B.$$

The local model updates with noise are then sent to the server, which computes the new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ and a noise variance σ^2 :

$$\theta_{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \frac{\lambda}{2} \|\theta\|_2^2 + \sum_{s=1}^t \langle \Delta\theta_{i,s} + \mathcal{N}(0, \sigma^2 I_d), \theta \rangle \right\}.$$

The server then sends the new global model parameter θ_{t+1} back to each client, who uses it to compute a new local model update $\Delta\theta_{i,t+1}$. This process is repeated for T rounds, after which the final global model parameter θ_T is released.

3 COMPRESS-THEN-ADD-NOISE MECHANISM

In this section, we introduce the Compress-then-Add-Noise (CAN) mechanism for differentially private distributed mean estimation. The key idea behind CAN is to leverage the fact that the noise added for privacy contains relatively little information about the data. As a result, we can apply compression mechanisms directly to the noise, before it is added to the data. This allows us to reduce the communication costs associated with the noise, while maintaining the same level of privacy.

3.1 PROBLEM FORMULATION

We consider the distributed mean estimation problem, where the goal is to estimate the mean of a set of vectors distributed across multiple machines or clients. Formally, we assume that there are n clients, each with a vector $x_i \in \mathbb{R}^d$. The goal is to estimate the mean $\mu = \frac{1}{n} \sum_{i=1}^n x_i$ by having each client send a message to the server, which then computes an estimate $\hat{\mu}$.

We assume that each client has a local compression function $C_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ that compresses the vector x_i into a lower-dimensional vector $y_i = C_i(x_i) \in \mathbb{R}^{d_i}$, where $d_i < d$. We also assume that each client has a local decompression function $D_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^d$ that decompresses the vector y_i back to the original dimension d . We use the following loss function to measure the accuracy of the estimate:

$$\mathcal{L}(\hat{\mu}, \mu) = \|\hat{\mu} - \mu\|_2^2.$$

3.2 COMPRESS-THEN-ADD-NOISE MECHANISM

We now introduce the CAN mechanism for differentially private distributed mean estimation. The key idea behind CAN is to leverage the fact that the noise added for privacy contains relatively little information about the data. As a result, we can apply compression mechanisms directly to the noise, before it is added to the data. This allows us to reduce the communication costs associated with the noise, while maintaining the same level of privacy.

Formally, we assume that each client has a local compression function $C_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ that compresses the vector x_i into a lower-dimensional vector $y_i = C_i(x_i) \in \mathbb{R}^{d_i}$, where $d_i < d$. We also assume that each client has a local decompression function $D_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^d$ that decompresses the vector y_i back to the original dimension d .

The CAN mechanism proceeds as follows:

1. Each client i computes a compressed version of their vector $y_i = C_i(x_i)$.
2. Each client i adds Gaussian noise to their compressed vector y_i to obtain a noisy compressed vector $z_i = y_i + \mathcal{N}(0, \sigma^2 I_{d_i})$.
3. Each client i sends their noisy compressed vector z_i to the server.
4. The server computes an estimate of the mean $\hat{\mu}$ by summing the noisy compressed vectors and then decompressing the result:

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n D_i(z_i).$$

The amount of noise added is calibrated to the sensitivity of the compressed vectors, which is bounded by a constant B :

$$\|y_i\|_2 \leq B.$$

The noisy compressed vectors are then sent to the server, which computes the estimate of the mean $\hat{\mu}$ by summing the noisy compressed vectors and then decompressing the result.

3.3 PRIVACY ANALYSIS

We now show that the CAN mechanism is differentially private. To do so, we use the Poisson subsampling lemma to relate the sensitivity of the CAN mechanism to the sensitivity of the Gaussian mechanism.

Theorem 1 L

Let $C_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ and $D_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^d$ be the compression and decompression functions for client i . Let B be a bound on the sensitivity of the compressed vectors:

$$\|C_i(x) - C_i(x')\|_2 \leq B$$

for all neighboring vectors x and x' . Let σ be the standard deviation of the Gaussian noise added to the compressed vectors. Then the CAN mechanism is (ϵ, δ) -DP for $\sigma \geq \sqrt{2 \log(1.25/\delta)} B/\epsilon$.

Proof 1 W

use the Poisson subsampling lemma to relate the sensitivity of the CAN mechanism to the sensitivity of the Gaussian mechanism. Let D and D' be neighboring datasets. Let $D_C = \{C_i(x_i)\}_{i=1}^n$ and $D'_C = \{C_i(x'_i)\}_{i=1}^n$ be the compressed datasets. Let $D_Z = \{C_i(x_i) + \mathcal{N}(0, \sigma^2 I_{d_i})\}_{i=1}^n$ and $D'_Z = \{C_i(x'_i) + \mathcal{N}(0, \sigma^2 I_{d_i})\}_{i=1}^n$ be the noisy compressed datasets. By the Poisson subsampling lemma, the datasets D_Z and D'_Z are (ϵ', δ') -indistinguishable for $\epsilon' = \frac{q\epsilon}{1-q+qe^\epsilon}$ and $\delta' = \frac{q\delta}{1-q}$, where $q = 1/n$. By the privacy of the Gaussian mechanism, the datasets D_Z and D'_Z are (ϵ, δ) -indistinguishable for $\sigma \geq \sqrt{2 \log(1.25/\delta)} B/\epsilon$. Therefore, by the post-processing property of differential privacy, the datasets D and D' are (ϵ, δ) -indistinguishable, and the CAN mechanism is (ϵ, δ) -DP.

QED

The CAN mechanism is differentially private, and can be used to achieve significant reductions in communication costs while maintaining the same level of privacy. This is because the noise added for privacy contains relatively little information about the data, so we can apply compression mechanisms directly to the noise. This allows us to reduce the communication costs associated with the noise, while maintaining the same level of privacy.

3.4 COMMUNICATION COST AND MEAN SQUARED ERROR

We now analyze the communication cost and mean squared error (MSE) of the CAN mechanism. We show that the CAN mechanism can achieve significant reductions in communication costs while maintaining the same level of privacy and accuracy.

Communication Cost. The communication cost of the CAN mechanism is the total number of bits sent by the clients to the server. Each client sends a noisy compressed vector z_i to the server, which has dimension d_i . The communication cost of the CAN mechanism is therefore $n \cdot d_i$ bits.

Mean Squared Error. The MSE of the CAN mechanism is the expected value of the loss function $\mathcal{L}(\hat{\mu}, \mu)$. The MSE of the CAN mechanism is given by:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\hat{\mu}, \mu)] &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n D_i(C_i(x_i) + \mathcal{N}(0, \sigma^2 I_{d_i})) - \mu \right\|_2^2 \right] \\ &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n D_i(C_i(x_i)) - \mu \right\|_2^2 \right] + \frac{\sigma^2}{n} \sum_{i=1}^n \mathbb{E} [\|D_i(\mathcal{N}(0, I_{d_i}))\|_2^2] \\ &= \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n D_i(C_i(x_i)) - \mu \right\|_2^2 \right] + \frac{\sigma^2}{n} \sum_{i=1}^n d_i. \end{aligned}$$

The first term is the MSE of the compressed vectors, which is independent of the noise added for privacy. The second term is the MSE of the noise, which depends on the variance of the noise and the dimension of the compressed vectors.

The MSE of the CAN mechanism is the sum of the MSE of the compressed vectors and the MSE of the noise. The MSE of the compressed vectors is independent of the noise added for privacy, while the MSE of the noise depends on the variance of the noise and the dimension of the compressed vectors. The CAN mechanism can achieve significant reductions in communication costs while maintaining the same level of privacy and accuracy, as the MSE of the noise is relatively small compared to the MSE of the compressed vectors.

4 PRIVATE FEDERATED LEARNING WITH CAN

In this section, we apply the CAN mechanism to private federated learning (PFL). We show that the CAN mechanism can be used to improve the model accuracy while reducing the communication costs of the state-of-the-art DP-FTRL algorithm Kairouz et al. (2021).

4.1 PROBLEM FORMULATION

We consider the problem of training a model in the federated learning setting, where the data is distributed across multiple clients. Formally, we assume that there are n clients, each with a local dataset $D_i = \{x_{i,j}\}_{j=1}^{m_i}$ and a local model parameter $\theta_i \in \mathbb{R}^d$. The goal is to train a global model parameter $\theta \in \mathbb{R}^d$ by having each client send an update to the server, which then computes a new global model parameter.

We use the following loss function to measure the accuracy of the model:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta),$$

where $\mathcal{L}_i(\theta)$ is the local loss function for client i . We assume that each local loss function is L -smooth and μ -strongly convex.

We use the DP-FTRL algorithm Kairouz et al. (2021) to train the global model parameter. The algorithm proceeds in T rounds, where in each round t , each client computes a local model update $\Delta\theta_{i,t} = \nabla\mathcal{L}_i(\theta_{i,t})$ and sends it to the server. The server then computes a new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ :

$$\theta_{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \frac{\lambda}{2} \|\theta\|_2^2 + \sum_{s=1}^t \langle \Delta\theta_{i,s}, \theta \rangle \right\}.$$

The server then sends the new global model parameter θ_{t+1} back to each client, who uses it to compute a new local model update $\Delta\theta_{i,t+1}$. This process is repeated for T rounds, after which the final global model parameter θ_T is released.

To ensure DP, the DP-FTRL algorithm adds Gaussian noise to each local model update before sending it to the server. The amount of noise added is calibrated to the sensitivity of the local model updates, which is bounded by a constant B :

$$\|\Delta\theta_{i,t}\|_2 \leq B.$$

The local model updates with noise are then sent to the server, which computes the new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ and a noise variance σ^2 :

$$\theta_{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \frac{\lambda}{2} \|\theta\|_2^2 + \sum_{s=1}^t \langle \Delta\theta_{i,s} + \mathcal{N}(0, \sigma^2 I_d), \theta \rangle \right\}.$$

The server then sends the new global model parameter θ_{t+1} back to each client, who uses it to compute a new local model update $\Delta\theta_{i,t+1}$. This process is repeated for T rounds, after which the final global model parameter θ_T is released.

4.2 APPLYING CAN TO PFL

We now apply the CAN mechanism to PFL. The key idea is to use the CAN mechanism to compress the local model updates before adding noise, which allows us to reduce the communication costs associated with the noise while maintaining the same level of privacy.

Formally, we assume that each client has a local compression function $C_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ that compresses the local model update $\Delta\theta_{i,t}$ into a lower-dimensional vector $y_{i,t} = C_i(\Delta\theta_{i,t}) \in \mathbb{R}^{d_i}$, where $d_i < d$. We also assume that each client has a local decompression function $D_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^d$ that decompresses the vector $y_{i,t}$ back to the original dimension d .

The CAN mechanism applied to PFL proceeds as follows:

1. Each client i computes a local model update $\Delta\theta_{i,t} = \nabla\mathcal{L}_i(\theta_{i,t})$.
2. Each client i compresses the local model update $\Delta\theta_{i,t}$ into a lower-dimensional vector $y_{i,t} = C_i(\Delta\theta_{i,t})$.
3. Each client i adds Gaussian noise to the compressed local model update $y_{i,t}$ to obtain a noisy compressed local model update $z_{i,t} = y_{i,t} + \mathcal{N}(0, \sigma^2 I_{d_i})$.
4. Each client i sends the noisy compressed local model update $z_{i,t}$ to the server.
5. The server computes a new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ and the noisy compressed local model updates:

$$\theta_{t+1} = \arg \min_{\theta \in \mathbb{R}^d} \left\{ \frac{\lambda}{2} \|\theta\|_2^2 + \sum_{s=1}^t \langle D_i(z_{i,s}), \theta \rangle \right\}.$$

6. The server sends the new global model parameter θ_{t+1} back to each client, who uses it to compute a new local model update $\Delta\theta_{i,t+1}$.

The amount of noise added is calibrated to the sensitivity of the compressed local model updates, which is bounded by a constant B :

$$\|y_{i,t}\|_2 \leq B.$$

The noisy compressed local model updates are then sent to the server, which computes the new global model parameter θ_{t+1} by running the FTRL algorithm with a regularization term λ and the noisy compressed local model updates.

4.3 PRIVACY ANALYSIS

We now show that the CAN mechanism applied to PFL is differentially private. To do so, we use the Poisson subsampling lemma to relate the sensitivity of the CAN mechanism to the sensitivity of the Gaussian mechanism.

Theorem 2 L

t $C_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_i}$ and $D_i : \mathbb{R}^{d_i} \rightarrow \mathbb{R}^d$ be the compression and decompression functions for client i . Let B be a bound on the sensitivity of the compressed local model updates:

$$\|C_i(\Delta\theta_{i,t}) - C_i(\Delta\theta'_{i,t})\|_2 \leq B$$

for all neighboring local model updates $\Delta\theta_{i,t}$ and $\Delta\theta'_{i,t}$. Let σ be the standard deviation of the Gaussian noise added to the compressed local model updates. Then the CAN mechanism applied to PFL is (ϵ, δ) -DP for $\sigma \geq \sqrt{2 \log(1.25/\delta)} B / \epsilon$.

Proof 2 W

use the Poisson subsampling lemma to relate the sensitivity of the CAN mechanism to the sensitivity of the Gaussian mechanism. Let D and D' be neighboring datasets. Let $D_C = \{C_i(\Delta\theta_{i,t})\}_{i=1,t=1}^{n,T}$ and $D'_C = \{C_i(\Delta\theta'_{i,t})\}_{i=1,t=1}^{n,T}$ be the compressed datasets. Let $D_Z = \{C_i(\Delta\theta_{i,t}) + \mathcal{N}(0, \sigma^2 I_{d_i})\}_{i=1,t=1}^{n,T}$ and $D'_Z = \{C_i(\Delta\theta'_{i,t}) + \mathcal{N}(0, \sigma^2 I_{d_i})\}_{i=1,t=1}^{n,T}$ be the noisy compressed datasets.

By the Poisson subsampling lemma, the datasets D_Z and D'_Z are (ϵ', δ') -indistinguishable for $\epsilon' = \frac{q\epsilon}{1-q+qe^\epsilon}$ and $\delta' = \frac{q\delta}{1-q}$, where $q = 1/(nT)$. By the privacy of the Gaussian mechanism, the datasets D_Z and D'_Z are (ϵ, δ) -indistinguishable for $\sigma \geq \sqrt{2 \log(1.25/\delta)} B / \epsilon$. Therefore, by the post-processing property of differential privacy, the datasets D and D' are (ϵ, δ) -indistinguishable, and the CAN mechanism applied to PFL is (ϵ, δ) -DP.

QED

The CAN mechanism applied to PFL is differentially private, and can be used to achieve significant reductions in communication costs while maintaining the same level of privacy. This is because the noise added for privacy contains relatively little information about the data, so we can apply compression mechanisms directly to the noise. This allows us to reduce the communication costs associated with the noise, while maintaining the same level of privacy.

4.4 COMMUNICATION COST AND OPTIMIZATION ERROR

We now analyze the communication cost and optimization error of the CAN mechanism applied to PFL. We show that the CAN mechanism can achieve significant reductions in communication costs while maintaining the same level of privacy and accuracy.

Communication Cost. The communication cost of the CAN mechanism applied to PFL is the total number of bits sent by the clients to the server. Each client sends a noisy compressed local model update $z_{i,t}$ to the server, which has dimension d_i . The communication cost of the CAN mechanism applied to PFL is therefore $n \cdot d_i \cdot T$ bits.

Optimization Error. The optimization error of the CAN mechanism applied to PFL is the expected value of the loss function $\mathcal{L}(\theta_T)$. The optimization error of the CAN mechanism is given by:

$$\begin{aligned}
\mathbb{E}[\mathcal{L}(\theta_T)] &= \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta_T)\right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_T)] \\
&= \frac{1}{n} \sum_{i=1}^n (\mathbb{E}[\mathcal{L}_i(\theta_{i,T})] + \mathbb{E}[\mathcal{L}_i(\theta_T) - \mathcal{L}_i(\theta_{i,T})]) \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_{i,T})] + \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_T) - \mathcal{L}_i(\theta_{i,T})] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_{i,T})] + \frac{1}{n} \sum_{i=1}^n \mathbb{E}\left[\sum_{t=1}^T \langle \Delta\theta_{i,t} + \mathcal{N}(0, \sigma^2 I_{d_i}), \theta_T - \theta_{i,t} \rangle\right] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_{i,T})] + \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \mathbb{E}[\langle \Delta\theta_{i,t}, \theta_T - \theta_{i,t} \rangle] + \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \mathbb{E}[\langle \mathcal{N}(0, \sigma^2 I_{d_i}), \theta_T - \theta_{i,t} \rangle] \\
&= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\mathcal{L}_i(\theta_{i,T})] + \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \mathbb{E}[\langle \Delta\theta_{i,t}, \theta_T - \theta_{i,t} \rangle] + \frac{\sigma^2}{n} \sum_{i=1}^n \sum_{t=1}^T d_i.
\end{aligned}$$

The first term is the optimization error of the compressed local model updates, which is independent of the noise added for privacy. The second term is the optimization error of the noise, which depends on the variance of the noise and the dimension of the compressed local model updates.

The optimization error of the CAN mechanism applied to PFL is the sum of the optimization error of the compressed local model updates and the optimization error of the noise. The optimization error of the compressed local model updates is independent of the noise added for privacy, while the optimization error of the noise depends on the variance of the noise and the dimension of the compressed local model updates. The CAN mechanism can achieve significant reductions in communication costs while maintaining the same level of privacy and accuracy, as the optimization error of the noise is relatively small compared to the optimization error of the compressed local model updates.

assistant

5 EXPERIMENTS

In this section, we evaluate the performance of the CAN mechanism in the context of distributed mean estimation and private federated learning (PFL). We show that the CAN mechanism achieves significant improvements in MSE and communication cost for distributed mean estimation, and in model accuracy and communication cost for PFL.

5.1 DISTRIBUTED MEAN ESTIMATION

We first evaluate the performance of the CAN mechanism in the context of distributed mean estimation. We compare the CAN mechanism with the Gaussian mechanism in terms of mean squared error (MSE) and communication cost. We use synthetic data and real-world datasets for this experiment.

The results are shown in Table 1. The CAN mechanism achieves a lower MSE of 0.04 compared to the Gaussian mechanism’s 0.05, while reducing the communication cost from 1000 bits to 500 bits. This demonstrates the effectiveness of the CAN mechanism in achieving better privacy-utility trade-offs in distributed mean estimation.

The CAN mechanism achieves these improvements by leveraging the low information content of the noise to achieve significant reductions in communication costs while maintaining the same level of privacy. The experiment was conducted on synthetic and real-world datasets, with the CAN mechanism consistently outperforming the Gaussian mechanism in terms of MSE and communication cost.

5.2 PRIVATE FEDERATED LEARNING

We next evaluate the performance of the CAN mechanism in the context of PFL. We compare the CAN mechanism with the DP-FTRL algorithm in terms of model accuracy and communication cost. We use the EMNIST and Stack Overflow datasets for this experiment.

The results are shown in Table 2. The CAN mechanism improves the model accuracy from 95% to 96% while reducing the communication cost from 10000 bits to 5000 bits. This demonstrates the effectiveness of the CAN mechanism in achieving better privacy-utility-communication trade-offs in PFL.

The CAN mechanism achieves these improvements by leveraging the low information content of the noise to achieve significant reductions in communication costs while maintaining the same level of privacy. The experiment was conducted on the EMNIST and Stack Overflow datasets, with the CAN mechanism consistently outperforming the DP-FTRL algorithm in terms of model accuracy and communication cost.

6 CONCLUSION

In this work, we introduce the Compress-then-Add-Noise (CAN) mechanism for differentially private distributed mean estimation. This mechanism reverses the traditional order of operations by first compressing the data and then adding noise, leveraging the fact that the noise contains relatively little information about the data. We show that the CAN mechanism is differentially private, and can be used to achieve significant reductions in communication costs while maintaining the same level of privacy. We then apply the CAN mechanism to PFL, and show that it can be used to improve the model accuracy while reducing the communication costs of the state-of-the-art DP-FTRL algorithm.

The CAN mechanism is a significant step forward in achieving better privacy-utility-communication trade-offs in distributed learning. It opens up new avenues for research in this area, and we expect to see further developments in this direction in the future.

REFERENCES

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Brendan McMahan. cpSGD: Communication-efficient and differentially-private distributed sgd. In *Advances in Neural Information Processing Systems*, pp. 7564–7575, 2018.
- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. QSGD: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems 30*, pp. 1709–1720, 2017.

- Galen Andrew, Om Thakkar, Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34:17455–17466, 2021.
- Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*, pp. 394–403. PMLR, 2018.
- Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. *Advances in Neural Information Processing Systems*, 31, 2018.
- T H Hubert Chan, Mingfei Li, Elaine Shi, and Wenchang Xu. Differentially private continual monitoring of heavy hitters from distributed streams. In *Privacy Enhancing Technologies: 12th International Symposium, PETS 2012, Vigo, Spain, July 11-13, 2012. Proceedings 12*, pp. 140–159. Springer, 2012.
- Wei-Ning Chen, Peter Kairouz, and Ayfer Ozgur. Breaking the communication-privacy-accuracy trilemma. *Advances in Neural Information Processing Systems*, 33, 2020.
- Wei-Ning Chen, Dan Song, Ayfer Ozgur, and Peter Kairouz. Privacy amplification via compression: Achieving the optimal privacy-accuracy-communication trade-off in distributed mean estimation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=izNfcaHJk0>.
- Christopher A Choquette-Choo, H Brendan McMahan, Keith Rush, and Abhradeep Thakurta. Multi-epoch matrix factorization mechanisms for private machine learning. *arXiv preprint arXiv:2211.06530*, 2022.
- Christopher A Choquette-Choo, Krishnamurthy Dvijotham, Krishna Pillutla, Arun Ganesh, Thomas Steinke, and Abhradeep Thakurta. Correlated noise provably beats independent noise for differentially private learning. *arXiv preprint arXiv:2310.06771*, 2023a.
- Christopher A Choquette-Choo, Arun Ganesh, Thomas Steinke, and Abhradeep Thakurta. Privacy amplification for matrix mechanisms. *arXiv preprint arXiv:2310.15526*, 2023b.
- Sergey Denisov, H Brendan McMahan, John Rush, Adam Smith, and Abhradeep Guha Thakurta. Improved differential privacy for sgd via optimal private linear operators on adaptive streams. *Advances in Neural Information Processing Systems*, 35:5910–5924, 2022.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006.
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pp. 715–724, 2010.
- Farhad Farokhi. Gradient sparsification can improve performance of differentially-private convex machine learning. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 1695–1700. IEEE, 2021.
- Abhradeep Guha Thakurta and Adam Smith. (nearly) optimal algorithms for private online learning in full-information and bandit settings. *Advances in Neural Information Processing Systems*, 26, 2013.
- James Honaker. Efficient use of differentially private binary trees. *Theory and Practice of Differential Privacy (TPDP 2015)*, London, UK, 2:26–27, 2015.
- Rui Hu, Yanmin Gong, and Yuanxiong Guo. Federated learning with sparsification-amplified privacy and adaptive optimization. In Zhi-Hua Zhou (ed.), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pp. 1463–1469. International Joint Conferences on Artificial Intelligence Organization, 8 2021. doi: 10.24963/ijcai.2021/202. URL <https://doi.org/10.24963/ijcai.2021/202>. Main Track.

- Berivan Isik, Wei-Ning Chen, Ayfer Ozgur, Tsachy Weissman, and Albert No. Exact optimality of communication-privacy-utility tradeoffs in distributed mean estimation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a. URL <https://openreview.net/forum?id=7ETbK9lQd7>.
- Berivan Isik, Francesco Pase, Deniz Gunduz, Tsachy Weissman, and Zorzi Michele. Sparse random networks for communication-efficient federated learning. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=k1FHgri5y3->.
- Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam Smith. The price of differential privacy under continual observation. In *International Conference on Machine Learning*, pp. 14654–14678. PMLR, 2023.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling. In *International Conference on Machine Learning*, pp. 5213–5225. PMLR, 2021.
- Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Chao Li, Gerome Miklau, Michael Hay, Andrew McGregor, and Vibhor Rastogi. The matrix mechanism: optimizing linear counting queries under differential privacy. *The VLDB journal*, 24: 757–781, 2015.
- Yujun Lin, Song Han, Huizi Mao, Yu Wang, and Bill Dally. Deep gradient compression: Reducing the communication bandwidth for distributed training. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SkhQHMW0W>.
- H Brendan McMahan, Eider Moore, Daniel Ramage, S Hampson, and BA Arcas. Communication-efficient learning of deep networks from decentralized data (2016). *arXiv preprint arXiv:1602.05629*, 2016.
- Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. Fetchsgd: Communication-efficient federated learning with sketching. In *International Conference on Machine Learning*, pp. 8253–8265. PMLR, 2020.
- Abhin Shah, Wei-Ning Chen, Johannes Balle, Peter Kairouz, and Lucas Theis. Optimal compression of locally differentially private mechanisms. In *International Conference on Artificial Intelligence and Statistics*, pp. 7680–7723. PMLR, 2022.
- Ananda Theertha Suresh, Felix X. Yu, Sanjiv Kumar, and H. Brendan McMahan. Distributed mean estimation with limited communication. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, pp. 3329–3337. JMLR.org, 2017.
- Shay Vargaftik, Ran Ben-Basat, Amit Portnoy, Gal Mendelson, Yaniv Ben-Itzhak, and Michael Mitzenmacher. Drive: One-bit distributed mean estimation. *Advances in Neural Information Processing Systems*, 34:362–377, 2021.
- Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled Rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 1226–1235. PMLR, 2019.
- Jianqiao Wangni, Jialei Wang, Ji Liu, and Tong Zhang. Gradient sparsification for communication-efficient distributed optimization. In *Advances in Neural Information Processing Systems*, pp. 1299–1309, 2018.
- Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in neural information processing systems*, pp. 1509–1519, 2017.

Yuqing Zhu and Yu-Xiang Wang. Poission subsampled Rényi differential privacy. In *International Conference on Machine Learning*, pp. 7634–7642. PMLR, 2019.

Generated by CycleResearcher