



Secure Aggregation Is Not Private Against Membership Inference Attacks

Khac-Hoang Ngo^{1(✉)}, Johan Östman², Giuseppe Durisi¹,
and Alexandre Graell i Amat¹

¹ Department of Electrical Engineering, Chalmers University of Technology,
Gothenburg, Sweden

{ngok,durisi,alexandre.graell}@chalmers.se

² AI Sweden, Gothenburg, Sweden

johan.ostman@ai.se

Abstract. Secure aggregation (SecAgg) is a commonly-used privacy-enhancing mechanism in federated learning, affording the server access only to the aggregate of model updates while safeguarding the confidentiality of individual updates. Despite widespread claims regarding SecAgg’s privacy-preserving capabilities, a formal analysis of its privacy is lacking, making such presumptions unjustified. In this paper, we delve into the privacy implications of SecAgg by treating it as a local differential privacy (LDP) mechanism for each local update. We design a simple attack wherein an adversarial server seeks to discern which update vector a client submitted, out of two possible ones, in a single training round of federated learning under SecAgg. By conducting privacy auditing, we assess the success probability of this attack and quantify the LDP guarantees provided by SecAgg. Our numerical results unveil that, contrary to prevailing claims, SecAgg offers weak privacy against membership inference attacks even in a single training round. Indeed, it is difficult to hide a local update by adding other independent local updates when the updates are of high dimension. Our findings underscore the imperative for additional privacy-enhancing mechanisms, such as noise injection, in federated learning.

Keywords: Federated learning · Secure aggregation · Differential privacy · Membership inference

1 Introduction

Federated learning (FL) [27] allows multiple clients to collaboratively train a machine learning model. In each training round, the clients share their local model updates with a central server, which then aggregates them to improve the global model. Although raw data is not shared in the clear, vanilla FL is prone to model-inversion attacks [18] and membership-inference attacks [29]. To mitigate such attacks, secure aggregation (SecAgg) [9] has been proposed, where

the clients jointly mask their local model updates so that only the aggregate is revealed to the server.

Many papers explicitly or implicitly assume that SecAgg provides strong privacy against honest-but-curious servers in a single round [16, 19, 30, 31]. However, a formal analysis of the privacy offered by SecAgg is lacking, making this presumption unjustified. SecAgg has been combined with differential privacy (DP) [14] to ensure that the server only sees the aggregate of the noisy local updates [2, 22]. However, in these works, the privacy analysis does not account for SecAgg. It remains unclear how much privacy SecAgg by itself provides for individual updates.

Main Contributions. We address the question: *how much privacy does SecAgg by itself guarantee for the local updates?* Specifically, we formally analyze the privacy of SecAgg against membership inference attacks wherein the server aims to distinguish, from two potential update vectors, the one a client submitted in a single training round of FL with SecAgg. Our approach consists in treating SecAgg as a local differential privacy (LDP) mechanism for each update, where the sum of the other clients' updates plays the role of a source of uncontrolled noise. We then characterize the privacy parameters (ϵ, δ) for SecAgg to satisfy (ϵ, δ) -LDP via the following steps.

- We show that, under some practical assumptions, as the client population grows, the sum of the clients' updates converges to a Gaussian vector (Theorem 1). We analyze the optimal privacy guarantee of the Gaussian mechanism with correlated noise (Theorem 2).
- We evaluate the optimal LDP parameters of SecAgg in some special cases (Theorem 3 and Corollary 1) and verify that these parameters are close to that of a Gaussian mechanism, even for a small number of clients (Fig. 1).
- Exploiting the similarity of SecAgg and a Gaussian mechanism, we audit the privacy of SecAgg. Specifically, we design a simple membership inference attack wherein the server regards SecAgg as a Gaussian mechanism with correlated noise. We then evaluate the achievable false negative rate (FNR) and false positive rate (FPR) of this attack and use these values to compute a lower bound on the smallest ϵ for SecAgg to satisfy (ϵ, δ) -LDP.

We apply our privacy auditing procedure to federated averaging for a classification problem on the ADULT dataset [5] and the EMNIST Digits dataset [11]. We show that both the FNR and FPR can be small simultaneously, and the audited (ϵ, δ) are high. Our results reveal that SecAgg provides weak privacy even for a single training round. Indeed, it is difficult to hide a local update by adding other independent local updates when the updates are of high dimension. Therefore, SecAgg cannot be used as a sole privacy-enhancing mechanism in FL.

2 Related Work

Secure Aggregation. Based on cryptographic multi-party computation, SecAgg ensures that the central server sees only the aggregate of the clients’ local updates, while individual updates are kept confidential. This is achieved by letting the clients jointly add randomly sampled masks to their updates via secret sharing, such that when the masked updates are aggregated, the masks cancel out [6, 9]. With SecAgg, a client’s update is obfuscated by many other clients’ updates. However, the level of privacy provided by SecAgg lacks a formal analysis. In [16], this level was measured by the mutual information between a local update and the aggregated update. However, mutual information only measures the average privacy leakage and does not capture the threat to the most vulnerable data points. Furthermore, the bound provided in [16] is not explicit, i.e., not computable.

Differential Privacy. DP is a rigorous privacy measure that quantifies the ability of an adversary to guess which dataset, out of two neighboring ones, a model was trained on [14, 15]. DP is typically achieved by adding noise to the model/gradients obtained from the dataset [1]. A variant of DP is LDP [13, 24], where the noise is added to individual data points. When applied to achieve client-level privacy in FL, LDP lets the clients add noise to their updates before sending the updates to the server.

Privacy Attacks in FL with SecAgg. Model inversion attacks [18] and membership-inference attacks [29] have been shown to seriously jeopardize the integrity of FL. When SecAgg is employed, the server can perform disaggregation attacks to learn the individual data. A malicious server performs active attacks by suppressing the updates of non-target clients [8, 17]. For an honest-but-curious server, existing passive attacks require that the server leverages the aggregated model across many rounds [25, 26]. Differently from these works, we consider a *passive* attack based only on the observation in a *single round*.

3 Preliminaries

We denote random quantities with lowercase nonitalic letters, such as a scalar x and a vector \mathbf{x} . The only exception is the privacy-loss random variable (PLRV) L , which is in uppercase. Deterministic quantities are denoted with italic letters, such as a scalar x and a vector \mathbf{x} . We denote the multidimensional normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and its probability density function (PDF) evaluated at \mathbf{x} by $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$. We denote by $\Phi(x)$ the cumulative distribution function (CDF) of the standard normal distribution $\mathcal{N}(0, 1)$, i.e., $\Phi(x) \triangleq \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$. We denote by $[m : n]$ the set of integers from m to n ; $[n] \triangleq [1 : n]$; $(\cdot)^+ \triangleq \max\{0, \cdot\}$. Furthermore, $\mathbb{1}\{\cdot\}$ denotes the indicator function, and $f(n) = o(g(n))$ means that $f(n)/g(n) \rightarrow 0$ as $n \rightarrow \infty$.

Let ϕ be a decision rule of a hypothesis test between $\{H: \text{the underlying distribution is } P\}$ and $\{H': \text{the underlying distribution is } Q\}$. Specifically, ϕ returns 0 and 1 in favor of H and H' , respectively. A false positive (resp. false negative) occurs when H (resp. H') is true but rejected. The FPR and FNR of the test are given by $\alpha_\phi \triangleq \mathbb{E}_P[\phi]$ and $\beta_\phi \triangleq 1 - \mathbb{E}_Q[\phi]$, respectively.

Definition 1 (Trade-off function). *The trade-off function $T_{P,Q}(\cdot): [0, 1] \rightarrow [0, 1]$ is the map from the FPR to the corresponding minimum FNR of the test between P and Q , i.e., $T_{P,Q}(\alpha) \triangleq \inf_{\phi: \alpha_\phi \leq \alpha} \beta_\phi$, $\alpha \in [0, 1]$.*

We also write the trade-off function for the distributions of x and y as $T_{x,y}(\cdot)$. We next state the definition of LDP.

Definition 2 (LDP [13,24]). *A mechanism M satisfies (ϵ, δ) -LDP if and only if, for every pair of data points $(\mathbf{x}, \mathbf{x}')$ and for every measurable set \mathcal{E} , we have $\mathbb{P}[M(\mathbf{x}) \in \mathcal{E}] \leq e^\epsilon \mathbb{P}[M(\mathbf{x}') \in \mathcal{E}] + \delta$.*

For a mechanism M , we define the optimal LDP curve $\delta_M(\epsilon)$ as the function that returns the smallest δ for which M satisfies (ϵ, δ) -LDP. We next define a variant of LDP that is built upon the trade-off function in a similar manner as f -DP [12, Def. 3].

Definition 3 (f -LDP). *For a function f , a mechanism M satisfies f -LDP if for every pair of data points $(\mathbf{x}, \mathbf{x}')$, we have that $T_{M(\mathbf{x}),M(\mathbf{x}')}(\alpha) \geq f(\alpha)$, $\forall \alpha \in [0, 1]$.*

For a mechanism M , we define the optimal f -LDP curve $f_M(\cdot)$ as the upper envelope of all functions f such that M satisfies f -LDP. In the paper, we regard SecAgg as a LDP mechanism and provide bounds on both its optimal LDP curve and optimal f -LDP curve.

4 Privacy Analysis of Secure Aggregation

We consider a FL scenario with $n + 1$ clients and a central server. The model update of client i can be represented as a vector $\mathbf{x}_i \in \mathbb{R}^d$. Under SecAgg, the server only learns the aggregate model update $\bar{\mathbf{x}} = \sum_{i=0}^n \mathbf{x}_i$, while the individual updates $\{\mathbf{x}_i\}_{i=0}^n$ remain confidential.

4.1 Threat Model

Server. The server is honest and follows the FL and SecAgg protocols. We assume that it observes the exact sum $\bar{\mathbf{x}}$. In practical SecAgg, the clients discretize their updates (using, e.g., randomized rounding) and the server obtains a modulo sum. These operations introduce perturbations that can improve privacy [34]. However, they do not capture the essence of SecAgg which is to use the updates of other clients to obfuscate an individual update. The rounding and modulo operations can be applied even in a setting without SecAgg. We ignore the perturbation caused by these operations to focus on the privacy obtainable by using the updates of other clients to obfuscate an update.

Clients. The clients are also honest. Client i computes the local model update \mathbf{x}_i from the global model in the previous round and its local dataset. We also assume that each vector \mathbf{x}_i , $i \in [0 : n]$, has correlated entries since these entries together describe a model, and that the vectors are mutually independent. The latter assumption holds in the first training round if the clients have independent local data. Furthermore, the independence assumption results in the best-case scenario for privacy as, if the vectors are dependent, the sum reveals more information about each vector. Therefore, the privacy level for the case of independent $\{\mathbf{x}_i\}_{i=0}^n$ acts as an upper bound on the privacy for the case of dependent vectors. So if privacy does not hold for independent data, it will also not hold when there is dependence.

Privacy Threat. The server is curious. It seeks to infer the membership of a targeted client, say client 0, from the aggregate model updates $\bar{\mathbf{x}}$. We consider a membership inference game [33] where: i) a challenger selects a pair of possible local updates $(\mathbf{x}_0, \mathbf{x}'_0)$ of client 0, one of which is used in the aggregation, and sends this pair to the server, ii) the server observes $\bar{\mathbf{x}}$ and guesses if \mathbf{x}_0 or \mathbf{x}'_0 was submitted by client 0. Note that this attack can be an entry point for the server to further infer the data of client 0. Our goal is to quantify the capability of SecAgg in mitigating this attack.

4.2 SecAgg as a Noiseless LDP Mechanism

Hereafter, we focus on client 0; the analysis for other clients follows similarly. Our key idea is to view SecAgg through the lens of noiseless DP [7], where the contribution of other clients can be seen as noise and no external (controlled) noise is added. More precisely, for client 0, SecAgg plays the role of the mechanism

$$M(\mathbf{x}_0) = \mathbf{x}_0 + \mathbf{y}, \quad (1)$$

where $\mathbf{y} = \sum_{i=1}^n \mathbf{x}_i$ is a source of uncontrolled noise.

The aforementioned membership inference game can be cast as follows: given $M(\mathbf{x}_0)$, the server guesses whether it came from $P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}_0}$ or $P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}'_0}$ for the worst-case pair $(\mathbf{x}_0, \mathbf{x}'_0)$. This game is closely related to the LDP framework. First, the tradeoff between the FPR and FNR of the server's guesses is captured by the f -LDP guarantee of M . Second, as M achieves a stronger (ϵ, δ) -LDP guarantee, the distributions $P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}_0}$ and $P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}'_0}$ become more similar, and the hypothesis test between them becomes harder. Therefore, we shall address the following question: how much LDP or f -LDP does SecAgg guarantee for client 0? Specifically, we shall establish bounds on the optimal LDP curve $\delta_M(\epsilon)$ and optimal f -LDP curve $f_M(\cdot)$ of the mechanism M .

4.3 Asymptotic Privacy Guarantee

Let us first focus on the large- n regime. The following asymptotic analysis will be used as inspiration for our privacy auditing procedure to establish a lower bound on the LDP curve in Sect. 4.5.

We assume that the ℓ_2 norm of the vectors $\{\mathbf{x}_i\}_{i=1}^n$ scales as $o(\sqrt{n})$, which holds if, e.g., d is fixed. In this case, \mathbf{y} converges to a Gaussian vector when $n \rightarrow \infty$, as stated in the next theorem.

Theorem 1 (Asymptotic noise distribution). *Assume that $\{\mathbf{x}_i\}_{i=1}^n$ are independent, $\|\mathbf{x}_i\|_2 = o(\sqrt{n})$ for $i \in [n]$, and $\frac{1}{n} \sum_{i=1}^n \text{Cov}[\mathbf{x}_i] \rightarrow \Sigma$ as $n \rightarrow \infty$. Then $\frac{1}{\sqrt{n}} (\sum_{i=1}^n \mathbf{x}_i - \mathbb{E}[\sum_{i=1}^n \mathbf{x}_i])$ converges in distribution to $\mathcal{N}(\mathbf{0}, \Sigma)$ as $n \rightarrow \infty$.*

Proof. Theorem 1 follows by applying the multivariate Lindeberg-Feller central limit Theorem [32, Prop. 2.27] to the triangular array $\{\frac{\mathbf{x}_i}{\sqrt{n}}\}_{n,i}$, upon verifying the Lindeberg condition $\lim_{n \rightarrow \infty} \sum_{i=1}^n \mathbb{E} \left[\frac{\|\mathbf{x}_i\|_2^2}{n} \mathbb{1} \left\{ \frac{\|\mathbf{x}_i\|_2}{\sqrt{n}} > \varepsilon \right\} \right] = 0, \forall \varepsilon > 0$. Since $\|\mathbf{x}_i\|_2 = o(\sqrt{n})$, i.e., $\|\mathbf{x}_i\|_2/\sqrt{n} \rightarrow 0$ as $n \rightarrow \infty$, this condition indeed holds. \square

Theorem 1 implies that, when n is large, under the presented assumptions, the mechanism $\widetilde{M}(\mathbf{x}_0) = M(\mathbf{x}_0) - \mathbb{E}[\mathbf{y}]$ behaves like a Gaussian mechanism with noise distribution $\mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{y}})$, where $\Sigma_{\mathbf{y}}$ is the covariance matrix of \mathbf{y} . Furthermore, since the map from M to \widetilde{M} is simply a shift by a fixed vector $\mathbb{E}[\mathbf{y}]$, i.e., it is a bijection, we have from the post-processing property¹ that the optimal LDP curve of M is the same as that of \widetilde{M} .

We now provide privacy guarantees for a Gaussian mechanism with correlated noise, to capture the correlation between the entries of the vectors \mathbf{x}_i . The next theorem, proved in Appendix A.1, is an extension of the optimal privacy curve of the uncorrelated Gaussian mechanism [4, Theorem 8].

Theorem 2 (Correlated Gaussian mechanism). *Consider the mechanism $G(\mathbf{x}) = \mathbf{x} + \mathbf{y}$ where \mathbf{x} belongs to a set $\mathcal{S}_d \subset \mathbb{R}^d$, and $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{y}})$. The optimal LDP curve of G is*

$$\delta_G(\epsilon) = \Phi\left(\frac{\Delta}{2} - \frac{\epsilon}{\Delta}\right) - e^\epsilon \Phi\left(-\frac{\Delta}{2} - \frac{\epsilon}{\Delta}\right) \tag{2}$$

where $\Delta = \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_d} \Delta_{\mathbf{x}, \mathbf{x}'}$ with

$$\Delta_{\mathbf{x}, \mathbf{x}'} \triangleq \sqrt{(\mathbf{x} - \mathbf{x}')^\top \Sigma_{\mathbf{y}}^{-1} (\mathbf{x} - \mathbf{x}')}. \tag{3}$$

In Sect. 4.4, we shall verify the similarity between the privacy of SecAgg and that of the Gaussian mechanism G via numerical examples.

Parameter Δ is the maximum Euclidean distance between a pair of input vectors transformed by matrix $\Sigma^{-1/2}$ (similar to the whitening transformation). It plays the same role as the ratio of the sensitivity and the noise standard deviation in the case of uncorrelated noise [4]. We remark that the privacy guarantee of G is weakened as Δ increases: for a given ϵ , $\delta_G(\epsilon)$ increases with Δ . To achieve small ϵ and δ , we need Δ to be small. The impact of Δ can also be

¹ If a mechanism M satisfies (ϵ, δ) -LDP, then so does $h \circ M$ for a mapping h that is independent of M . The proof of this result is similar to the proof of the post-processing property of DP [15, Prop. 2.1].

seen via the hypothesis test associated to the considered membership inference game. Consider an adversary that observes an output \mathbf{z} of G and tests between $\{H : \mathbf{z} \text{ came from } P_{G(\mathbf{x})|\mathbf{x}=\mathbf{x}}\}$ and $\{H' : \mathbf{z} \text{ came from } P_{G(\mathbf{x})|\mathbf{x}=\mathbf{x}'}\}$. This is effectively a test between $\mathcal{N}(\mathbf{x}, \Sigma_{\mathbf{y}})$ and $\mathcal{N}(\mathbf{x}', \Sigma_{\mathbf{y}})$. The trade-off function for this test is stated in the following proposition, which is proved in Appendix A.3.

Proposition 1. $T_{\mathcal{N}(\mathbf{x}, \Sigma_{\mathbf{y}}), \mathcal{N}(\mathbf{x}', \Sigma_{\mathbf{y}})}(\alpha) = \Phi(\Phi^{-1}(1 - \alpha) - \Delta_{\mathbf{x}, \mathbf{x}'})$, $\alpha \in [0, 1]$.

The trade-off function decreases with $\Delta_{\mathbf{x}, \mathbf{x}'}$. A large $\Delta_{\mathbf{x}, \mathbf{x}'}$ facilitates the distinguishability of the pair $(\mathbf{x}, \mathbf{x}')$, and thus weakens the privacy guarantee. Furthermore, the worst-case pair $(\mathbf{x}, \mathbf{x}')$ that minimizes the trade-off function is given by the maximizer of $\Delta_{\mathbf{x}, \mathbf{x}'}$. It follows that the optimal f -LDP curve of the Gaussian mechanism G is $f_G(\alpha) = \Phi(\Phi^{-1}(1 - \alpha) - \Delta)$, $\alpha \in [0, 1]$.

4.4 Upper Bounding $\delta_M(\epsilon)$ via Dominating Pairs of Distributions

We now upper-bound the optimal LDP curve of M in (1) for finite n . We shall then consider the case in which the upper bound is tight and verify the convergence of SecAgg to a Gaussian mechanism.

We define the hockey-stick divergence with parameter α between two probability measures P and Q as $E_{\alpha}(P\|Q) = \sup_{\mathcal{E}}(P(\mathcal{E}) - \alpha Q(\mathcal{E}))$. We also write the hockey-stick divergence between the distributions of \mathbf{x} and \mathbf{y} as $E_{\alpha}(\mathbf{x}\|\mathbf{y})$. The condition for LDP in Definition 2 is equivalent to $\sup_{\mathbf{x} \neq \mathbf{x}'} E_{e^{\epsilon}}(M(\mathbf{x})\|M(\mathbf{x}')) \leq \delta$. Therefore, the optimal LDP curve of mechanism M is given by

$$\delta_M(\epsilon) = \sup_{\mathbf{x} \neq \mathbf{x}'} E_{e^{\epsilon}}(M(\mathbf{x})\|M(\mathbf{x}')). \tag{4}$$

A pair of measures (P, Q) is called a dominating pair of distributions for M if

$$\sup_{\mathbf{x} \neq \mathbf{x}'} E_{e^{\epsilon}}(M(\mathbf{x})\|M(\mathbf{x}')) \leq E_{e^{\epsilon}}(P\|Q), \quad \forall \epsilon \geq 0. \tag{5}$$

If equality is achieved in (5) for every $\epsilon \geq 0$, then (P, Q) is said to be a tightly dominating pair of distributions for M . For each dominating pair (P, Q) , we associate a privacy-loss random variable $L \triangleq \ln \frac{dP}{dQ}(\mathbf{y})$ with $\mathbf{y} \sim P$, where $\frac{dP}{dQ}$ is the Radon-Nikodym derivative. We have that

$$E_{e^{\epsilon}}(P\|Q) = \mathbb{E} [(1 - e^{\epsilon-L})^+] \triangleq \delta_L(\epsilon). \tag{6}$$

It follows readily that $\delta_L(\epsilon)$ is an upper bound on the optimal LDP curve $\delta_M(\epsilon)$.

Without a known distribution of \mathbf{y} , it is challenging to characterize a dominating pair of distributions for the mechanism M in (1). In the next theorem, proved in Appendix B, we make some assumptions on $P_{\mathbf{y}}$ to enable such characterization.

Theorem 3 (Dominating pair of distributions). *Let $\mathbf{x}_0 = (x_{01}, x_{02}, \dots, x_{0d})$ and assume that $\underline{r}_j \leq x_{0j} \leq \bar{r}_j$, $j \in [d]$. Assume further that \mathbf{y} has independent entries, i.e., $P_{\mathbf{y}} = P_{y_1} \times \dots \times P_{y_d}$, and that the marginal probabilities $\{P_{y_j}\}$*

are log-concave and symmetric.² Then, a dominating pair of distributions for the mechanism $M(\mathbf{x}_0)$ in (1) is given by $(P_{r_1+y_1} \times \cdots \times P_{r_d+y_d}, P_{\bar{r}_1+y_1} \times \cdots \times P_{\bar{r}_d+y_d})$.

The family of log-concave distributions includes the typical noise distributions in DP, namely, the Gaussian and Laplace distributions, as well as many other common distributions, e.g., the exponential and uniform distributions [3]. If each vector $\mathbf{x}_i, i \in [n]$, has independent entries following a log-concave distribution, then so does the sum $\mathbf{y} = \sum_{i=1}^n \mathbf{x}_i$, because log-concavity is closed under convolutions [28]. Under the presented assumptions, Theorem 3 allows us to characterize an upper bound on the LDP curve of M as $\delta_L(\epsilon)$ with $L = \sum_{j=1}^d \ln \frac{P_{r_j+y_j}(z_j)}{P_{\bar{r}_j+y_j}(z_j)}$ where $z_j \sim P_{r_j+y_j}, j \in [d]$.


Corollary 1. *If the support of \mathbf{x}_0 contains (r_1, \dots, r_d) and $(\bar{r}_1, \dots, \bar{r}_d)$, the dominating pair of distributions in Theorem 3 becomes tight, and the resulting upper bound $\delta_L(\epsilon)$ is the optimal LDP curve.*

We now use Corollary 1 to evaluate the optimal LDP curve of mechanism M in (1) when each \mathbf{x}_i has independent entries. We aim to verify the convergence of SecAgg to a Gaussian mechanism implied by Theorem 1 and understand how the LDP curve depends on the model size d . We consider two cases. In the first case, the entries follow the exponential distribution with parameter 1, and thus \mathbf{y} has independent entries following the Gamma distribution with shape n and scale 1. For convenience, we further assume that \mathbf{x}_0 is truncated such that $0 \leq x_{0j} \leq 4, j \in [d]$. In the second case, the entries are uniformly distributed in $[-1/2, 1/2]$, and thus \mathbf{y} has independent entries following the shifted Irwin-Hall distribution with PDF $p_{y_i}(y) = \frac{1}{(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} [(y+n/2-k)^+]^{n-1}$. Both cases satisfy the conditions of Corollary 1. We can therefore obtain the optimal LDP curves and depict them in Fig. 1. We also show the optimal LDP curve of the Gaussian mechanism G with the same noise covariance matrix $\Sigma_{\mathbf{y}}$. We see that the optimal LDP curve of M is indeed close to that of G , even for a small value of n in the second case. Furthermore, although Theorem 1 assumes a fixed d and $n \rightarrow \infty$, Fig. 1 suggests that M behaves similarly to a Gaussian mechanism even for large d . Remarkably, for a given ϵ , the parameter δ increases rapidly with d , indicating that the privacy of SecAgg is weak for high-dimensional models.

4.5 Lower Bounding $\delta_M(\epsilon)$ and Upper Bounding $f_M(\cdot)$ via Privacy Auditing

In practical FL, the updates typically have a distribution that does not satisfy the conditions of Theorem 3 and is not known in closed form. Therefore, we now establish a numerical routine to compute a lower bound on the optimal LDP curve and an upper bound on the optimal f -LDP curve of M . The proposed

² P_{y_j} is symmetric if there exists a y^* such that $P_{y_j}(\mathcal{A}+y^*) = P_{y_j}(-\mathcal{A}+y^*)$ for every subset \mathcal{A} of the support of y_j . Here, $-\mathcal{A} \triangleq \{-y: y \in \mathcal{A}\}$.



(a) $\mathbf{x}_i, i \in [n]$, has independent entries exponentially distributed with parameter 1; $0 \leq x_{0j} \leq 4, j \in [d]$. For the Gaussian mechanism, $\mathcal{S}_d = [0, 4]^d$ and $\Sigma_{\mathbf{y}} = n\mathbf{I}_d$. (b) $\mathbf{x}_i, i \in [n]$, has independent entries uniformly distributed in $[-1/2, 1/2]$; $-1/2 \leq x_{0j} \leq 1/2, j \in [d]$. For the Gaussian mechanism, $\mathcal{S}_d = [-1/2, 1/2]^d$ and $\Sigma_{\mathbf{y}} = \frac{n}{12}\mathbf{I}_d$.

Fig. 1. The optimal LDP curve of M in (1) where each \mathbf{x}_i has independent entries, compared with the Gaussian mechanism G with the same noise covariance matrix.

numerical routine exploits the similarity of SecAgg and a Gaussian mechanism as discussed in Sects. 4.3 and 4.4. The bounds are based on the following result.

Proposition 2 (LDP via the trade-off function). *A mechanism M satisfies (ϵ, δ) -LDP if and only if for every $\alpha \in [0, 1]$,*

$$\epsilon \geq \ln \max \left\{ \frac{1 - \delta - \alpha}{\inf_{\mathbf{x} \neq \mathbf{x}'} T_{M(\mathbf{x}), M(\mathbf{x}')}(\alpha)}, \frac{1 - \delta - \inf_{\mathbf{x} \neq \mathbf{x}'} T_{M(\mathbf{x}), M(\mathbf{x}')}(\alpha)}{\alpha} \right\}. \quad (7)$$

The proof of Proposition 2 follows from similar arguments for DP in [23, Thm. 2.1]. This proposition implies that, if a pair (FPR, FNR) is achievable for some decision rule ϕ between the distributions of $M(\mathbf{x})$ and $M(\mathbf{x}')$ for some $(\mathbf{x}, \mathbf{x}')$, then the mechanism does not satisfy (ϵ, δ) -LDP for $\delta \in [0, 1]$ and

$$\epsilon < \ln \max \left\{ \frac{1 - \delta - \text{FPR}}{\text{FNR}}, \frac{1 - \delta - \text{FNR}}{\text{FPR}} \right\}. \quad (8)$$

This gives a lower bound on the optimal LDP curve of M . Furthermore, it follows readily from Definition 3 that, for an achievable pair (FPR, FNR) of the mentioned test, the mechanism does not satisfy f -LDP for any trade-off function f such that $f(\text{FPR}) < \text{FNR}$. Therefore, a collection of achievable pairs (FPR, FNR) constitutes an upper bound on the optimal f -LDP curve of M .

We shall use this result to perform *privacy auditing* [21, 33]. Specifically, following the defined membership inference game (see privacy threat in Sect. 4.1), we conduct a hypothesis test between $\{H: \mathbf{z} \text{ came from } P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}_0}\}$ and $\{H': \mathbf{z} \text{ came from } P_{M(\mathbf{x}_0)|\mathbf{x}_0=\mathbf{x}'_0}\}$ for a given output \mathbf{z} of M . That is, we select a pair $(\mathbf{x}_0, \mathbf{x}'_0)$ for the challenger and a decision rule ϕ for the server. We then

evaluate the achievable pair (FPR, FNR), and obtain therefrom a lower bound on the optimal LDP curve and an upper bound on the optimal f -LDP curve of M . To design the attack, we draw inspiration from the asymptotic analysis in Sect. 4.3 as follows.

We consider the likelihood-ratio test, i.e., the test rejects H if

$$\ln \frac{P_{M(\mathbf{x}_0) | \mathbf{x}_0}(\mathbf{z} | \mathbf{x}_0)}{P_{M(\mathbf{x}_0) | \mathbf{x}_0}(\mathbf{z} | \mathbf{x}'_0)} = \ln \frac{P_{\mathbf{y}}(\mathbf{z} - \mathbf{x}_0)}{P_{\mathbf{y}}(\mathbf{z} - \mathbf{x}'_0)} \leq \theta \quad (9)$$

for a given threshold θ . We choose the input pair $(\mathbf{x}_0, \mathbf{x}'_0)$ as the worst-case pair, i.e., $(\mathbf{x}_0, \mathbf{x}'_0) = \arg \min_{x, x'} T_{M(x), M(x')}(\alpha)$, $\forall \alpha \in [0, 1]$. However, the trade-off function is not known in closed-form in general, and thus finding the worst-case pair is challenging. Motivated by Theorem 1, we treat \mathbf{y} as a Gaussian vector with the same mean $\boldsymbol{\mu}_{\mathbf{y}}$ and covariance $\boldsymbol{\Sigma}_{\mathbf{y}}$. We thus approximate the trade-off function $T_{M(x), M(x')}(\cdot)$ by $T_{\mathcal{N}(x, \boldsymbol{\Sigma}_{\mathbf{y}}), \mathcal{N}(x', \boldsymbol{\Sigma}_{\mathbf{y}})}(\cdot)$, and choose $(\mathbf{x}_0, \mathbf{x}'_0)$ as the minimizer of the latter. Using Proposition 1, we have that

$$(\mathbf{x}_0, \mathbf{x}'_0) = \arg \max_{x, x' \in \mathcal{X}_0} \Delta_{x, x'} \quad (10)$$

where \mathcal{X}_0 is the support of \mathbf{x}_0 .

If the server does not know $P_{\mathbf{y}}(\mathbf{y})$ in closed form, we let it approximate $P_{\mathbf{y}}(\mathbf{y})$ as $\mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_{\mathbf{y}}, \boldsymbol{\Sigma}_{\mathbf{y}})$. That is, the test rejects \mathbf{x}_0 if

$$\ln \frac{\mathcal{N}(\mathbf{z} - \mathbf{x}_0; \boldsymbol{\mu}_{\mathbf{y}}, \boldsymbol{\Sigma}_{\mathbf{y}})}{\mathcal{N}(\mathbf{z} - \mathbf{x}'_0; \boldsymbol{\mu}_{\mathbf{y}}, \boldsymbol{\Sigma}_{\mathbf{y}})} \leq \theta. \quad (11)$$

Moreover, if the server does not know $\boldsymbol{\mu}_{\mathbf{y}}$ and $\boldsymbol{\Sigma}_{\mathbf{y}}$ but can generate samples from $P_{\mathbf{y}}$, we let it estimate $\boldsymbol{\mu}_{\mathbf{y}}$ and $\boldsymbol{\Sigma}_{\mathbf{y}}$ as the sample mean and sample covariance matrix, and use these estimates instead of the true values in (10) and (11).

We evaluate the FNR and FPR of the test via Monte-Carlo simulation. Specifically, we repeat the test N_s times and count the number of false negatives N_{FN} and the number false positives N_{FP} . We obtain a high-confidence upper bound on FNR using the Clopper-Pearson method [10] as $\overline{\text{FNR}} = B(1 - \gamma/2; N_{\text{FN}} + 1, N_s - N_{\text{FN}})$, where $B(x; a, b)$ is the quantile of the Beta distribution with shapes (a, b) , and $1 - \gamma$ is the confidence level. A high-confidence upper bound $\overline{\text{FPR}}$ on FPR is obtained similarly. By varying the threshold θ , we obtain an empirical trade-off curve $\overline{\text{FNR}}$ vs. $\overline{\text{FPR}}$. This curve is an upper bound on the optimal f -LDP curve of SecAgg. For a given $\delta \in [0, 1]$, we also compute a lower confidence bound on ϵ for SecAgg to satisfies (ϵ, δ) -LDP. Specifically, we use $\overline{\text{FNR}}$ and $\overline{\text{FPR}}$ in place of FNR and FPR in (8). Note that $\overline{\text{FNR}}$ and $\overline{\text{FPR}}$ are lower bounded by $B(1 - \gamma/2; 1, N_s)$ even if $N_{\text{FN}} = N_{\text{FP}} = 0$. Therefore, the estimated ϵ is upper bounded by $\bar{\epsilon}_{\delta} \triangleq \ln \frac{1 - \delta - B(1 - \gamma/2; 1, N_s)}{B(1 - \gamma/2; 1, N_s)}$. That is, it is impossible to audit an arbitrarily large ϵ with a finite number of trials.

5 Experiments and Discussion

Experimental Setting. We consider federated averaging with $n_{\text{tot}} = 100$ clients out of which $n+1$ clients are randomly selected in each round. The experiments³ are conducted for a classification problem on the ADULT dataset [5] and the EMNIST Digits dataset [11]. The ADULT dataset contains 30 162 entries with 104 features; the entries belong to two classes with 7508 positive labels and 22 654 negative labels. The EMNIST Digits dataset contains 280 000 images of size 28×28 of handwritten digits belonging to 10 balanced classes. We allocate the training samples between n_{tot} clients according to a latent Dirichlet allocation model with concentration parameter ω . Here, with $\omega \rightarrow \infty$, the training samples are distributed evenly and uniformly between the clients; with $\omega \rightarrow 0$, each client holds samples from only one class. We consider a single-layer neural network and use the cross-entropy loss and stochastic gradient descent with a learning rate of 0.01 and batch size of 64. The model size is $d = 210$ for ADULT and $d = 7850$ for EMNIST Digits.

We focus on the first round, containing one local epoch, of federated averaging and perform privacy auditing for a fixed initial model, which is known to the server. Note that performing an attack in the first round is the most challenging because, in later rounds, the server accumulates more observations. Let $\{\mathbf{x}_i\}_{i=0}^n$ be the local updates of the selected clients in the first round. The server does not know the distribution of $\{\mathbf{x}_i\}_{i=0}^n$ in closed form, but can sample from this distribution by simulating the learning scenario. Note that it is a common assumption in membership inference attacks that the adversary can sample from the population [33]. We let the server compute the sample mean $\hat{\boldsymbol{\mu}}_{\mathbf{x}}$ and sample covariance matrix $\hat{\boldsymbol{\Sigma}}_{\mathbf{x}}$ from 25 000 samples of \mathbf{x}_i , then estimate the mean and covariance matrix of $\mathbf{y} = \sum_{i=1}^n \mathbf{x}_i$ as $\hat{\boldsymbol{\mu}}_{\mathbf{y}} = n\hat{\boldsymbol{\mu}}_{\mathbf{x}}$ and $\hat{\boldsymbol{\Sigma}}_{\mathbf{y}} = n\hat{\boldsymbol{\Sigma}}_{\mathbf{x}}$. The server then uses $\hat{\boldsymbol{\mu}}_{\mathbf{y}}$ and $\hat{\boldsymbol{\Sigma}}_{\mathbf{y}}$ for privacy auditing, as described in Sect. 4.5. Following (10), we find the worst-case input pair $(\mathbf{x}_0, \mathbf{x}'_0)$ by searching for the maximizer of $\Delta_{\mathbf{x}_0, \mathbf{x}'_0}$ among 5000 and 1000 samples of \mathbf{x}_0 for the ADULT and EMNIST Digits datasets, respectively. The Clopper-Pearson confidence level is $1 - \gamma = 95\%$. For a given initial model, we consider $N_s = 5000$ trials with random data partition and batch selection. In the simulation results, we report the average of $\overline{\text{FPR}}$, $\overline{\text{FNR}}$, and the audited (ϵ, δ) over 10 and 5 initial models for the ADULT and EMNIST Digits datasets, respectively.

Homogeneous Data Partitioning. We first consider $\omega = \infty$. In Fig. 2(a), we show the trade-off between the estimated FNR and FPR for the ADULT dataset, achieved by varying the threshold θ in (11) for $n+1 \in \{60, 70, 90\}$ clients. Both the FNR and FPR can be as small as 0.005 simultaneously. Hence, the server can reliably distinguish the selected input pair, and the membership inference attack is successful. We note that a reference for the trade-off curve that represents different privacy levels is given in [12, Fig. 3]. There, the case with

³ The code is available at https://github.com/khachoang1412/SecAgg_not_private.

both FNR and FPR equal to 0.07 is already considered nonprivate. Comparing Fig. 2(a) with this reference, we conclude that SecAgg provides essentially no privacy for the ADULT dataset. Next, in Fig. 2(b), we show the average audited LDP curves for the ADULT dataset. We observe that the audited LDP curves are close to the largest auditable $(\bar{\epsilon}_\delta, \delta)$ with the considered N_s and γ . As ϵ increases, δ remains high until it drops due to the limit of the Clopper-Pearson method—even for $\epsilon = 7$, $\delta > 10^{-1}$. Furthermore, increasing n provides only a marginal privacy improvement. This shows that the privacy of SecAgg, viewed through the lens of LDP, is weak.

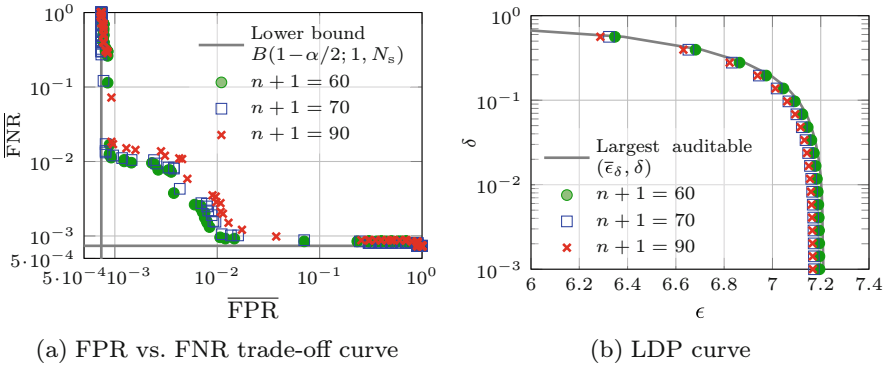


Fig. 2. The audited FPR vs. FNR trade-off and LDP curve, averaged over 10 initial models, for SecAgg in federated learning on the ADULT dataset with homogeneous data partitioning. Here, $d = 210$.

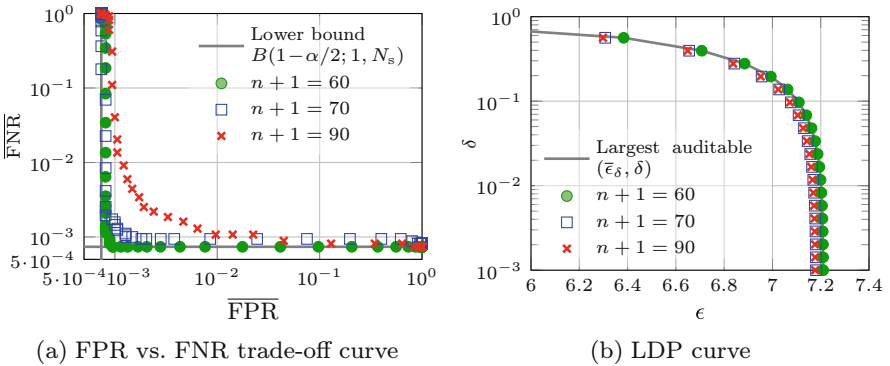


Fig. 3. Audited FPR vs. FNR trade-off and LDP curves, averaged over 5 initial models, for federated learning with SecAgg on the EMNIST Digits dataset with homogeneous data partitioning. Here, $d = 7850$.

In Fig. 3, we show the FPR vs. FNR trade-off and the audited LDP curve for the EMNIST Digits dataset. Similar conclusions hold: SecAgg provides weak

privacy. In this case, with a larger model size than the ADULT dataset, the adversary achieves even smaller FPR and FNR simultaneously.

Heterogeneous Data Partitioning. We next consider $\omega = 1$ and show the FPR vs. FNR trade-off and the audited LDP curve for the EMNIST Digits dataset in Fig. 4. In this case, the FPR and FNR are simultaneously reduced with respect to the homogeneous case, and the audited (ϵ, δ) coincide with the largest auditable values. This is because the worst-case pair $(\mathbf{x}_0, \mathbf{x}'_0)$ is better separated than in the homogeneous case and thus easier to distinguish.

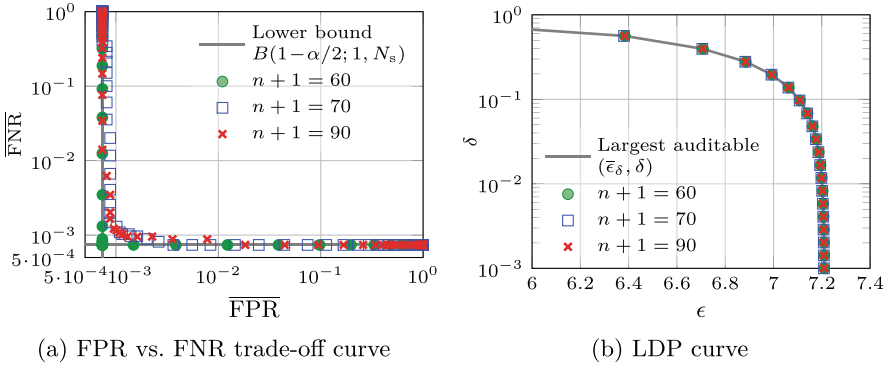


Fig. 4. Same as Fig. 3 but with heterogeneous data partitioning.

Discussion. We have seen that SecAgg is expected to perform like a correlated Gaussian mechanism (see Sect. 4.3). Why does SecAgg fail to prevent membership inference attacks, given that the Gaussian mechanism (with appropriate noise calibration) is known to be effective? We explain it as follows. We assume that the individual updates have entries with a bounded magnitude such that $\|\mathbf{x}_i\|_2 \leq r\sqrt{d}$, $i \in [0 : n]$. For large n , we expect the mechanism M in (1) to have similar privacy guarantee as G in Theorem 2 with $\mathcal{S}_d = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq r\sqrt{d}\}$ and $\Sigma_{\mathbf{y}}$ being the covariance matrix of $\sum_{i=1}^n \mathbf{x}_i$. In this case, $\Delta \geq 2\sqrt{d/n}$ (see Appendix A.2). A strong privacy guarantee requires Δ to be small, which implies that d/n must be small. This suggests that the privacy guarantee of SecAgg is weak if the vector dimension d is large compared to the number of clients n . This is, however, the case in most practical FL scenarios, as the model size is typically much larger than the number of clients. Note that reducing the ratio d/n by dividing the model into smaller chunks that are federated via SecAgg does not solve this issue, as the privacy loss composes over these chunks.

While with our results we have shown that SecAgg provides weak privacy for small models where d is in the order of 10^2 – 10^3 , we remark that the privacy guarantee is expected to further deteriorate for larger models. This is supported by the rapid degradation of the privacy guarantee with d of the Gaussian mechanism (see Fig. 1) and the similarity of SecAgg and a Gaussian mechanism.

6 Conclusions

We analyzed the privacy of SecAgg through the lens of LDP. Via privacy auditing, we showed that membership inference attacks on the output of SecAgg succeed with high probability: adding independent local updates is not sufficient to hide a local update when the model is of high dimension. While this result may not be surprising, our work fills an important gap by providing a formal analysis of the privacy of SecAgg and challenges the prevailing claims of the privacy robustness of SecAgg. Hence, it underscores that additional privacy mechanisms, such as noise addition, are needed in federated learning.

Acknowledgments. This work has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101022113. This work was also partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) and by the Swedish Research Council (VR) under grant 2020-03687. The authors would like to thank Onur Günlü and Balázs Pejó for fruitful discussions.

A Correlated Gaussian Mechanism

We present an analysis of the LDP guarantee of the Gaussian mechanism $G(\mathbf{x}) = \mathbf{x} + \mathbf{y}$, where \mathbf{x} belongs to a subset \mathcal{S}_d of \mathbb{R}^d , and $\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{y}})$.

A.1 Optimal LDP Curve: Proof of Theorem 2

We extend [4, Thm. 8] to the case of correlated noise. First, for a mechanism M and a pair \mathbf{x}, \mathbf{x}' , we define the privacy loss function as $L_{M, \mathbf{x}, \mathbf{x}'}(\mathbf{z}) \triangleq \ln \frac{P_{M(\mathbf{x})}(\mathbf{z})}{P_{M(\mathbf{x}')}(\mathbf{z})}$. The PLRV $L_{M, \mathbf{x}, \mathbf{x}'}$ is defined as the output of $L_{M, \mathbf{x}, \mathbf{x}'}$ when the input follows $P_{M(\mathbf{x})}$. The PLRV can be used to express the optimal LDP curve as

$$\delta_M(\epsilon) = \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_d} (\mathbb{P}[L_{M, \mathbf{x}, \mathbf{x}'} \geq \epsilon] - e^\epsilon \mathbb{P}[L_{M, \mathbf{x}', \mathbf{x}} \leq -\epsilon]). \quad (12)$$

Equation (12) is obtained from a similar result for DP given in [4, Thm. 5], upon modifying the notion of neighboring datasets.

For the mechanism G , we have that $P_{G(\mathbf{x})}(\mathbf{z}) = \frac{\exp(-\frac{1}{2}(\mathbf{z} - \mathbf{x})^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{z} - \mathbf{x}))}{\sqrt{(2\pi)^d |\Sigma_{\mathbf{y}}|}}$.

Therefore, the PLRV can be expressed as

$$L_{G, \mathbf{x}, \mathbf{x}'} = \frac{1}{2}(\mathbf{z} - \mathbf{x}')^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{z} - \mathbf{x}') - \frac{1}{2}(\mathbf{z} - \mathbf{x})^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{z} - \mathbf{x}) \quad (13)$$

$$= \frac{1}{2}(\mathbf{x} - \mathbf{x}')^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{x} - \mathbf{x}') + (\mathbf{x} - \mathbf{x}')^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{z} - \mathbf{x}). \quad (14)$$

With \mathbf{z} identically distributed to $G(\mathbf{x})$, i.e., $\mathbf{z} \sim \mathcal{N}(\mathbf{x}, \Sigma_{\mathbf{y}})$, we have that $L_{G, \mathbf{x}, \mathbf{x}'} \sim \mathcal{N}(\eta, 2\eta)$ with $\eta = \frac{1}{2}(\mathbf{x} - \mathbf{x}')^\top \Sigma_{\mathbf{y}}^{-1}(\mathbf{x} - \mathbf{x}')$. It follows from [4, Lemma 7] that for $\mathbf{x} \sim \mathcal{N}(\eta, 2\eta)$, $\mathbb{P}[\mathbf{x} \geq \epsilon] - e^\epsilon [\mathbf{x} \leq -\epsilon]$ is monotonically increasing in

η . By applying this result with $\mathbf{x} = \mathbf{L}_{G,x,x'}$, we obtain that the maximum in the right-hand side of (12) is achieved when η is maximized, i.e., when $\eta = \max_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_d} \frac{1}{2} (\mathbf{x} - \mathbf{x}')^\top \boldsymbol{\Sigma}_y^{-1} (\mathbf{x} - \mathbf{x}') = \Delta^2/2$. We then obtain the optimal LDP curve (2) after some simple computations.

A.2 The Case $\mathcal{S}_d = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq r\sqrt{d}\}$

In this case, for every $\mathbf{x}, \mathbf{x}' \in \mathcal{S}_d$, we have that

$$(\mathbf{x} - \mathbf{x}')^\top \boldsymbol{\Sigma}_y^{-1} (\mathbf{x} - \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2^2 \cdot \frac{(\mathbf{x} - \mathbf{x}')^\top \boldsymbol{\Sigma}_y^{-1} (\mathbf{x} - \mathbf{x}')}{\|\mathbf{x} - \mathbf{x}'\|_2 \|\mathbf{x} - \mathbf{x}'\|_2} \quad (15)$$

$$\leq (\|\mathbf{x}\|_2 + \|\mathbf{x}'\|_2)^2 \cdot \lambda_{\max}(\boldsymbol{\Sigma}_y^{-1}) \quad (16)$$

$$\leq (2r\sqrt{d})^2 \lambda_{\min}^{-1}(\boldsymbol{\Sigma}_y), \quad (17)$$

where $\lambda_{\max}(\boldsymbol{\Sigma}_y^{-1})$ is the largest eigenvalue of $\boldsymbol{\Sigma}_y^{-1}$ and $\lambda_{\min}(\boldsymbol{\Sigma}_y)$ is the smallest eigenvalue of $\boldsymbol{\Sigma}_y$. Here, (16) follows from the triangle inequality and the Rayleigh-Ritz theorem, and (17) holds because both $\|\mathbf{x}\|_2$ and $\|\mathbf{x}'\|_2$ are bounded by $r\sqrt{d}$. Equalities occur in (16) and (17) if $\mathbf{x} = -\mathbf{x}' = r\sqrt{d}\mathbf{v}_{\min}$ where \mathbf{v}_{\min} is the eigenvector of $\boldsymbol{\Sigma}_y$ corresponding to $\lambda_{\min}^{-1}(\boldsymbol{\Sigma}_y)$. Therefore, $\Delta = 2r\sqrt{d}/\lambda_{\min}(\boldsymbol{\Sigma}_y)$.

If we let $\boldsymbol{\Sigma}_y$ equal the covariance matrix of the sum of n independent vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ in \mathcal{S}_d , it holds that $\lambda_{\min}(\boldsymbol{\Sigma}_y) \leq \frac{1}{d} \text{Tr}(\boldsymbol{\Sigma}_y) = \frac{1}{d} \mathbb{E}[\text{Tr}(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top)] = \frac{1}{d} \sum_{i=1}^n \mathbb{E}[\text{Tr}(\mathbf{x}_i \mathbf{x}_i^\top)] = \frac{1}{d} \sum_{i=1}^n \mathbb{E}[\|\mathbf{x}_i\|_2^2] \leq nr^2$. As a consequence, $\Delta \geq 2\sqrt{d}/n$.

A.3 Trade-Off Function: Proof of Proposition 1

The log-likelihood ratio (LLR) for the test between $\{H: \mathbf{z}$ is generated from $\mathcal{N}(\mathbf{x}, \boldsymbol{\Sigma}_y)\}$ and $\{H': \mathbf{z}$ is generated from $\mathcal{N}(\mathbf{x}', \boldsymbol{\Sigma}_y)\}$ is given by the privacy loss function $L_{G,x,x'}(\mathbf{z})$. For a threshold θ , the FNR and FPR are given by

$$\text{FNR}(\theta) = \mathbb{P}_{\mathbf{z} \sim \mathcal{N}(\mathbf{x}', \boldsymbol{\Sigma}_y)} [L_{G,x,x'}(\mathbf{z}) \geq \theta] = \mathbb{P}[L_{G,x',x} \leq -\theta], \quad (18)$$

$$\text{FPR}(\theta) = \mathbb{P}_{\mathbf{z} \sim \mathcal{N}(\mathbf{x}, \boldsymbol{\Sigma}_y)} [L_{G,x,x'}(\mathbf{z}) < \theta] = \mathbb{P}[L_{G,x,x'} < \theta]. \quad (19)$$

In the proof of Theorem 2, we have shown that $L_{G,x,x'}(\mathbf{z}) \sim \mathcal{N}(\eta, 2\eta)$ with $\eta = \frac{1}{2} (\mathbf{x} - \mathbf{x}')^\top \boldsymbol{\Sigma}_y^{-1} (\mathbf{x} - \mathbf{x}')$. Therefore, $\text{FNR}(\theta) = \Phi(\frac{-\theta - \eta}{\sqrt{2\eta}})$ and $\text{FPR}(\theta) = \Phi(\frac{\theta - \eta}{\sqrt{2\eta}})$. To achieve $\text{FPR}(\theta) \leq \alpha$, the threshold must satisfy $\theta \leq \eta - \sqrt{2\eta}\Phi^{-1}(1 - \alpha)$. Under this constraint, the minimum FNR is given by $\Phi(\Phi^{-1}(1 - \alpha) - \sqrt{2\eta})$. This is by definition the optimal trade-off $T_{\mathcal{N}(x, \boldsymbol{\Sigma}), \mathcal{N}(x', \boldsymbol{\Sigma})}(\alpha)$.

B LDP Analysis of the Mechanism (1) in a Special Case: Proof of Theorem 3

To prove Theorem 3, we shall use the following preliminary results.

Lemma 1. *For two pairs of distributions (P_1, Q_1) and (P_2, Q_2) , if $T_{P_1, Q_1}(\alpha) \geq T_{P_2, Q_2}(\alpha)$ for every $\alpha \in [0, 1]$, then $\mathbb{E}_{e^\epsilon}(P_1 \| Q_1) \leq \mathbb{E}_{e^\epsilon}(P_2 \| Q_2)$ for every $\epsilon > 0$.*

Lemma 1 follows directly by expressing the hockey-stick divergence in terms of the trade-off function as follows.

Lemma 2. *Consider two distributions P and Q defined over \mathcal{X} . Define a random variable $L_P = \ln \frac{Q(x)}{P(x)}$, $x \sim P$, and denote its CDF by $F_P(x) = \mathbb{P}[L_P \leq x]$. It holds that $\mathbb{E}_{e^\epsilon}(P \| Q) = F_P(-\epsilon) - e^\epsilon T_{P, Q}(1 - F_P(-\epsilon))$.*

Proof of Lemma 2. Define the random variable $L_Q = \ln \frac{Q(x)}{P(x)}$, $x \sim Q$, and denote its CDF by $F_Q(x)$. Observe that $1 - F_P(\theta)$ and $F_Q(\theta)$ are the FPR and FNR of the likelihood test between P and Q with threshold θ . It follows from Definition 1 and the Neyman-Pearson lemma [20, Thm. 8.6.1] that

$$T_{P, Q}(1 - F_P(\theta)) = F_Q(\theta). \tag{20}$$

We further have that

$$\begin{aligned} \mathbb{E}_{e^\epsilon}(P \| Q) &= \int_{\mathcal{X}} (P(x) - e^\epsilon Q(x))^+ dx \\ &= \int_{\mathcal{X}} (P(x) - e^\epsilon Q(x)) \mathbb{1}\{P(x) \geq e^\epsilon Q(x)\} dx \\ &= \int_{\mathcal{X}} P(x) \mathbb{1}\{P(x) \geq e^\epsilon Q(x)\} dx - e^\epsilon \int_{\mathcal{X}} Q(x) \mathbb{1}\{P(x) \geq e^\epsilon Q(x)\} dx \\ &= F_P(-\epsilon) - e^\epsilon F_Q(-\epsilon). \end{aligned} \tag{21}$$

By substituting (20) with $\theta = -\epsilon$ into (21), we complete the proof. \square

We are now ready to prove Theorem 3. For brevity, we omit the subscript 0 in \mathbf{x}_0 and \mathbf{x}_{0j} , $j \in [d]$.

The Univariate Case. We first consider $d = 1$. Then M in (1) is written as

$$M(x) = x + y \tag{22}$$

where $\underline{r} \leq x \leq \bar{r}$ and y follows a symmetric and log-concave distribution P_y . We next show that a dominating pair of distribution for this mechanism is

$(P_{\underline{r}+y}, P_{\bar{r}+y})$. By symmetry of P_y around y_0 , we have that, for every $a, b \in \mathbb{R}$,

$$\mathbf{E}_{e^\epsilon}(a + y \| b + y) = \sup_{\mathcal{A}} P_{a+y}(\mathcal{A}) - e^\epsilon P_{b+y}(\mathcal{A}) \quad (23)$$

$$= \sup_{\mathcal{A}} P_y(\mathcal{A} - a) - e^\epsilon P_y(\mathcal{A} - b) \quad (24)$$

$$= \sup_{-\mathcal{A}+z_0} P_y(-\mathcal{A} + z_0 - a) - e^\epsilon P_y(-\mathcal{A} + z_0 - b) \quad (25)$$

$$= \sup_{\mathcal{A}} P_y(\mathcal{A} + a + z_0) - e^\epsilon P_y(\mathcal{A} + b + z_0) \quad (26)$$

$$= \sup_{\mathcal{A}'=\mathcal{A}+a+b+z_0} P_y(\mathcal{A}' - b) - e^\epsilon P_y(\mathcal{A}' - a) \quad (27)$$

$$= \sup_{\mathcal{A}} P_y(\mathcal{A} - b) - e^\epsilon P_y(\mathcal{A} - a) \quad (28)$$

$$= \mathbf{E}_{e^\epsilon}(b + y \| a + y). \quad (29)$$

Therefore, it suffices to prove that

$$\mathbf{E}_{e^\epsilon}(a + y \| b + y) \leq \mathbf{E}_{e^\epsilon}(\underline{r} + y \| \bar{r} + y), \quad \forall \underline{r} \leq a \leq b \leq \bar{r}. \quad (30)$$

To this end, we shall show that $\mathbf{E}_{e^\epsilon}(a + y \| b + y)$ increases with $b - a$, and thus maximized when $(a, b) = \arg \max_{a', b' \in [\underline{r}, \bar{r}], a' \leq b'} (b' - a') = (\underline{r}, \bar{r})$. In light of Lemma 1, it

suffices to show that $T_{a+y, b+y}(x)$ decreases with $b - a$ for all $x \in \mathbb{R}$. Indeed, this is true because $T_{a+y, b+y}(\alpha) = F_y(F_y^{-1}(1 - \alpha) - (b - a))$ for y following a log-concave distribution (this follows from [12, Prop. A.3]), and because the CDF $F_y(\cdot)$ is an increasing function.

The Multivariate Case. We now address the general case with $d \geq 1$. Using the independence assumption, we can write the mechanism (1) as $M(\mathbf{x}) = (M_1(\mathbf{x}), M_2(\mathbf{x}), \dots, M_d(\mathbf{x}))$ where $M_j(\mathbf{x}) = \mathbf{e}_j^\top \mathbf{x} + y_j$, with \mathbf{e}_j being the j th d -dimensional canonical basis vector. Therefore, $M(\mathbf{x})$ is a (nonadaptive) composition of d mechanisms $\{M_j\}_{j \in [d]}$. Observe that each mechanism M_j has the form (22). Therefore, $(P_{\underline{r}_j+y_j}, P_{\bar{r}_j+y_j})$ is a dominating pair of distributions for M_j . The proof is completed by applying [35, Thm. 10], which states that if (P_j, Q_j) is a dominating pair of distributions for mechanism M_j , $j \in [d]$, then $(P_1 \times \dots \times P_d, Q_1 \times \dots \times Q_d)$ is a dominating pair of distributions for mechanism $M(\mathbf{x}) = (M_1(\mathbf{x}), \dots, M_d(\mathbf{x}))$.

References

1. Abadi, M., et al.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 308–318. ACM, New York (2016)
2. Agarwal, N., Suresh, A.T., Yu, F., Kumar, S., McMahan, H.B.: CpSGD: communication-efficient and differentially-private distributed SGD. In: Proceedings of the International Conference in Neural Information Processing Systems (NIPS), NIPS 2018, pp. 7575–7586 (2018)

3. Bagnoli, M., Bergstrom, T.: Log-concave probability and its applications. In: Aliprantis, C.D., Matzkin, R.L., McFadden, D.L., Moore, J.C., Yannelis, N.C. (eds.) *Rationality and Equilibrium. Studies in Economic Theory*, vol. 26, pp. 217–241. Springer, Heidelberg (2006). https://doi.org/10.1007/3-540-29578-X_11
4. Balle, B., Wang, Y.X.: Improving the Gaussian mechanism for differential privacy: analytical calibration and optimal denoising. In: *Proceedings of the International Conference Machine Learning (ICML)*, pp. 394–403. PMLR (2018)
5. Becker, B., Kohavi, R.: Adult. UCI Machine Learning Repository (1996)
6. Bell, J.H., Bonawitz, K.A., Gascón, A., Lepoint, T., Raykova, M.: Secure single-server aggregation with (poly)logarithmic overhead. In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 1253–1269. ACM, New York (2020)
7. Bhaskar, R., Bhowmick, A., Goyal, V., Laxman, S., Thakurta, A.: Noiseless database privacy. In: Lee, D.H., Wang, X. (eds.) *ASIACRYPT 2011. LNCS*, vol. 7073, pp. 215–232. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_12
8. Boenisch, F., Dziedzic, A., et al.: Reconstructing individual data points in federated learning hardened with differential privacy and secure aggregation. In: *Proceedings of the European Symposium on Security and Privacy (EuroS&P)*, pp. 241–257 (2023)
9. Bonawitz, K., et al.: Practical secure aggregation for privacy-preserving machine learning. In: *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191 (2017)
10. Clopper, C.J., Pearson, E.S.: The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika* **26**(4), 404–413 (1934)
11. Cohen, G., Afshar, S., Tapson, J., van Schaik, A.: EMNIST: extending MNIST to handwritten letters. In: *Proceedings of International Joint Conference on Neural Networks (IJCNN)*, pp. 2921–2926 (2017)
12. Dong, J., Roth, A., Su, W.: Gaussian differential privacy. *J. Roy. Stat. Soc.* **84**(1), 3–37 (2021)
13. Duchi, J.C., Jordan, M.I., Wainwright, M.J.: Local privacy and statistical minimax rates. In: *Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 429–438 (2013)
14. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Halevi, S., Rabin, T. (eds.) *TCC 2006. LNCS*, vol. 3876, pp. 265–284. Springer, Heidelberg (2006). https://doi.org/10.1007/11681878_14
15. Dwork, C., Roth, A., et al.: The algorithmic foundations of differential privacy. *Found. Trends® Theoret. Comput. Sci.* **9**(3–4), 211–407 (2014)
16. Elkordy, A.R., Zhang, J., Ezzeldin, Y.H., Psounis, K., Avestimehr, S.: How much privacy does federated learning with secure aggregation guarantee? In: *Proceedings of the Privacy Enhancing Technologies Symposium (PETS)*, pp. 510–526 (2023)
17. Fowl, L.H., Geiping, J., Czaja, W., Goldblum, M., Goldstein, T.: Robbing the fed: directly obtaining private data in federated learning with modified models. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2022)
18. Geiping, J., Bauermeister, H., Dröge, H., Moeller, M.: Inverting gradients - how easy is it to break privacy in federated learning? In: *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, NeurIPS 2020 (2020)
19. Hatamizadeh, A., et al.: Do gradient inversion attacks make federated learning unsafe? *IEEE Trans. Med. Imaging* **42**(7), 2044–2056 (2023)

20. Hogg, R.V., Tanis, E.A., Zimmerman, D.: Probability and Statistical Inference, 9th edn. Pearson, Upper Saddle River (2015)
21. Jagielski, M., Ullman, J., Oprea, A.: Auditing differentially private machine learning: how private is private SGD? In: Advances in Neural Information Processing Systems (NeurIPS), vol. 33, pp. 22205–22216 (2020)
22. Kairouz, P., Liu, Z., Steinke, T.: The distributed discrete Gaussian mechanism for federated learning with secure aggregation. In: Proceedings of the International Conference on Machine Learning (ICML), pp. 5201–5212. PMLR (2021)
23. Kairouz, P., Oh, S., Viswanath, P.: The composition theorem for differential privacy. *IEEE Trans. Inf. Theory* **63**(6), 4037–4049 (2017)
24. Kasiviswanathan, S.P., Lee, H.K., Nissim, K., Raskhodnikova, S., Smith, A.: What can we learn privately? *SIAM J. Comput.* **40**(3), 793–826 (2011)
25. Kerkouche, R., Ács, G., Fritz, M.: Client-specific property inference against secure aggregation in federated learning. In: Proceedings of the Workshop Privacy in the Electronic Society, WPES 2023, pp. 45–60. ACM, New York (2023)
26. Lam, M., Wei, G.Y., Brooks, D., Reddi, V.J., Mitzenmacher, M.: Gradient disaggregation: breaking privacy in federated learning by reconstructing the user participant matrix. In: Proceedings of the International Conference on Machine Learning (ICML), pp. 5959–5968. PMLR (2021)
27. McMahan, B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.y.: Communication-efficient learning of deep networks from decentralized data. In: Singh, A., Zhu, J. (eds.) Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS). PMLR, vol. 54, pp. 1273–1282. PMLR, 20–22 April 2017
28. Merkle, M.: Convolutions of logarithmically concave functions. *Publikacije Elektrotehničkog fakulteta. Serija Matematika*, pp. 113–117 (1998)
29. Nasr, M., Shokri, R., Houmansadr, A.: Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In: Proceedings of the IEEE Symposium on Security and Privacy (SP), pp. 739–753 (2019)
30. So, J., Ali, R.E., Güler, B., Jiao, J., Avestimehr, A.S.: Securing secure aggregation: mitigating multi-round privacy leakage in federated learning. In: Proceedings of the AAAI Conference on Artificial Intelligence & Conf. Innov. App. Art. Intel. & Symp. Edu. Adv. Art. Intel. AAAI 2023/IAAI 2023/EAAI 2023. AAAI Press (2023)
31. Ullah, E., Choquette-Choo, C.A., Kairouz, P., Oh, S.: Private federated learning with autotuned compression. In: Proceedings of the International Conference on Machine Learning (ICML), ICML 2023. JMLR.org (2023)
32. Van der Vaart, A.W.: Asymptotic Statistics, vol. 3. Cambridge University Press, Cambridge (2000)
33. Ye, J., Maddi, A., Murakonda, S.K., Bindschaedler, V., Shokri, R.: Enhanced membership inference attacks against machine learning models. In: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, pp. 3093–3106 (2022)
34. Youn, Y., Hu, Z., Ziani, J., Abernethy, J.: Randomized quantization is all you need for differential privacy in federated learning. In: ICML Workshop (2023)
35. Zhu, Y., Dong, J., Wang, Y.X.: Optimal accounting of differential privacy via characteristic function. In: Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS). PMLR, vol. 151, pp. 4782–4817. PMLR, 28–30 March 2022