

Trading Off Privacy, Utility and Efficiency in Federated Learning

XIAOJIN ZHANG, Hong Kong University of Science and Technology, Hong Kong YAN KANG, Webank, China KAI CHEN, Hong Kong University of Science and Technology, Hong Kong LIXIN FAN, Webank, China QIANG YANG^{*}, WeBank and Hong Kong University of Science and Technology, Hong Kong

Federated learning (FL) enables participating parties to collaboratively build a global model with boosted utility without disclosing private data information. Appropriate protection mechanisms have to be adopted to fulfill the opposing requirements in preserving *privacy* and maintaining high model *utility*. In addition, it is a mandate for a federated learning system to achieve high *efficiency* in order to enable large-scale model training and deployment. We propose a unified federated learning framework that reconciles horizontal and vertical federated learning. Based on this framework, we formulate and quantify the trade-offs between privacy leakage, utility loss, and efficiency reduction, which leads us to the No-Free-Lunch (NFL) theorem for the federated learning system. NFL indicates that it is unrealistic to expect an FL algorithm to simultaneously provide excellent privacy, utility, and efficiency in certain scenarios. We then analyze the lower bounds for the privacy leakage, utility loss, and efficiency reduction mechanisms, including *Randomization, Homomorphic Encryption, Secret Sharing* and *Compression*. Our analysis could serve as a guide for selecting protection parameters to meet particular requirements.

 $\label{eq:CCS Concepts: Computing methodologies} \rightarrow \mbox{Artificial Intelligence; } \bullet \mbox{Machine Learning; } \bullet \mbox{Distributed methodologies;}$

Additional Key Words and Phrases: federated learning, privacy, utility, efficiency, trade-off, divergence, optimization

1 INTRODUCTION

The rapid expansion of large-scale datasets has sparked a demand for distributed learning. With the enforcement of data privacy regulations such as the General Data Protection Regulation (GDPR), the data owned by one company is not allowed to be disclosed to others. Federated learning (FL) [24, 25, 28, 31, 32] meets this requirement by allowing multiple parties to train a machine learning model collaboratively without sharing private data. In recent years, FL has achieved significant progress in developing privacy-preserving machine learning systems. It has been extended from the conventional *horizontal federated learning* (HFL) to the *vertical federated learning* (VFL) scenarios [18, 22, 46]. HFL typically involves a large amount of parties with different samples but share the same feature space. While VFL typically involves a handful of parties that own distinct features of the same set

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. 2157-6904/2023/5-ART \$15.00 https://doi.org/10.1145/3595185

^{*}Corresponding author

Authors' addresses: Xiaojin Zhang, xiaojinzhang@ust.hk, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong; Yan Kang, yangkang@webank.com, Webank, Shenzhen, China; Kai Chen, kaichen@cse.ust.hk, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong; Lixin Fan, lixinfan@webank.com, Webank, Shenzhen, China; Qiang Yang, qyang@cse.ust.hk, WeBank and Hong Kong University of Science and Technology, Hong Kong.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

of sample instances. HFL is popular in healthcare and mobile applications [2, 21, 27, 36], while VFL is widely adopted in finance and advertisement [4, 23, 39].

Preserving privacy is of immense practical importance when federating across different parties. Although the private data of each client is not shared with other collaborators, the private information might still be reconstructed by semi-honest parties upon observing the shared model information [14, 15, 26, 47, 51]. The fundamental requirement for maintaining privacy is to keep potential *privacy leakage* below an acceptable level. This is accomplished by reducing the dependence between shared model information and private data. To protect private data of the participants, many protection mechanisms have been proposed, such as *Randomization Mechanism* [1, 17, 41], *Secret Sharing* [5, 6, 38], *Homomorphic Encryption (HE)* [16, 48], and *Compression Mechanism* [33]. However, the adoption of these protection mechanisms might result in a certain amount of *utility loss* and *efficiency reduction*, as compared with a federated model trained without any protection [11, 23]. As a result, theoretical analysis of the trade-off between privacy leakage, utility loss, and efficiency reduction is critical for guiding FL practitioners to choose better protection parameters or design smarter FL algorithms.

Motivated by this goal, the work [50] proposed a statistical framework to analyze the privacy-utility trade-off in FL on a rigorous theoretical foundation. However, its trade-off analysis did not involve efficiency, which is a crucial factor in designing FL algorithms. Besides, [50] did not consider the VFL setting, which has a broad range of applications in finance and advertisement. These gaps inspire us to investigate the following crucial open problem: *is it possible to design a protection mechanism that simultaneously achieves infinitesimal privacy leakage, utility loss, and efficiency reduction?* Our main finding is crystallized in a No-Free-Lunch (NFL) Theorem (Theorem 4.6) that provides a negative answer for this question in certain scenarios under a unified FL framework that reconciles both HFL and VFL settings. For the outline of our work, please refer to Figure 6.

Our contributions are as follows:

- We propose a unified FL framework (see Section 5) for HFL and VFL. The unified FL framework provides a conceptual view of the relationship between privacy leakage, utility loss, and efficiency reduction through the lens of the protector and the adversary. Our No-Free-Lunch theorem is formulated based on this framework, and thus it applies to both HFL and VFL.
- We provide a No-Free-Lunch theorem (**Theorem 4.6**) for federated learning, which quantifies the tradeoff between privacy leakage (Def. 3.3), utility loss (Def. 3.4) and efficiency reduction (Def. 3.5). This quantification indicates that the weighted summation of the privacy leakage, utility loss, and efficiency reduction is greater than a problem-dependent non-zero constant. It characterizes the amount of utility and efficiency that is inevitable to lose in the case that the privacy leakage budget is unduly low.
- We apply the NFL theorem to analyze trade-offs between privacy leakage, utility loss, and efficiency reduction for widely-adopted protection mechanisms, including *Randomization*, *Homomorphic Encryption*, *Secret Sharing* and *Compression*. The trade-off is characterized via the lower bounds derived for privacy leakage, utility loss, and efficiency reduction.

2 RELATED WORK

The related work for attacking mechanisms and protection mechanisms in federated learning and privacy-utilityefficiency trade-off are briefly reviewed in this section.

2.1 Attacking Mechanisms and Protection Mechanisms in Federated Learning

We focus on *semi-honest* adversaries who faithfully follow the federated learning protocol but may infer the private information of other participants based on the exposed model information.

In HFL, [15, 47, 51, 52] demonstrate that adversaries could exploit gradient information to restore the private image data to pixel-level accuracy, with distinct settings of prior distributions and conditional distributions.

A variety of protection mechanisms have been proposed in HFL to prevent private data from being deduced by adversarial participants. The most popular ones are *Homomorphic Encryption (HE)* [16, 48], *Randomization Mechanism* [1, 17, 41], *Secret Sharing* [5, 6, 38] and *Compression Mechanism* [33]. Another school of FL [19, 20] tries to protect privacy by splitting a neural network into private and public models and sharing only the public one [19, 23].

VFL has two kinds of privacy leakage: feature leakage and label leakage. It is challenging for the adversary to infer the features of other parties because, in production VFL, participating parties typically have black-box knowledge about each other. The literature has proposed mainly two kinds of label inference attacks in VFL: the gradient-based [26] and the model-based [14] attacks. [26] also demonstrated three noise-based protections that can prevent gradient-based attacks. [53] proposed a data encoding protection mechanism called CoAE that can thwart model-based attacks effectively in some scenarios. Crypto-based protections are widely adopted in federated logistic regression and XGBoost. However, they are seldom applied to VNN that involves complex neural networks for their high communication and computational cost.

2.2 Privacy-Utility-Efficiency Trade-off

In the past decade, there has been wide interest in understanding the privacy-utility trade-off [10, 30, 35, 37, 42, 44, 50]. [10, 30] quantified the privacy-utility trade-off using the solution of the optimization problem. [37] provided a privacy-utility trade-off region for the special case with i.i.d. data sources and known distribution. [35] illustrated that the optimal privacy-utility trade-off could be solved using a standard linear program and provided a closed-form solution for the special case when the data to be released is a binary variable. [44] measured distortion using the expected Hamming distance between the input and output databases and measured privacy leakage using identifiability, differential privacy, and mutual-information privacy separately. [42] provided a trade-off when utility and privacy were evaluated using χ^2 -based information measures. [43] analyzed the trade-off between the speed of error convergence and the wall-clock time for distributed SGD. [7] analyzed the trade-off between communication, privacy, and accuracy for distributed statistical tasks. The accuracy is measured using statistical mean estimation, frequency estimation, and distribution estimation separately, and the privacy is measured using differential privacy. However, none of these works focus on federated learning and measure utility using model performance.

The work [50] proposed a quantitative trade-off between utility and privacy in horizontal federated learning by exploiting some key properties of the privacy leakage and the triangle inequality of the divergence. [29] evaluated the accuracy-privacy-cost trade-off for federated learning empirically. In this work, we offer a general theoretical analysis of the trade-off between privacy, utility, and efficiency that applies to both HFL and VFL.

3 A UNIFIED FEDERATED LEARNING FRAMEWORK

In this section, we first introduce general notations used throughout this work. Then, we propose a unified federated learning framework with a conceptual view of the relationship between privacy leakage, utility loss, and efficiency reduction. Next, we provide formal definitions for key components of this framework, including protection and attacking mechanisms, privacy leakage, utility loss, and efficiency reduction. We then formulate the goal of the protector as a constrained optimization problem and put forward a critical question with which we are concerned. Our No-Free-Lunch theorem provides an answer for this problem formulated based on this framework, and thus it applies to both HFL and VFL.

Notation	Meaning
ϵ_p	Privacy leakage (Def. 3.3)
ϵ_u	Utility loss (Def. 3.4)
ϵ_{e}	Efficiency reduction (Def. 3.5)
D	Private information, including private data and statistical information
$W_{\rm fed}$	parameter for the federated model
W_{k}^{O}	Unprotected model information of client k
W_{k}^{S}	Protected model information of client k
P_{ν}^{δ}	Distribution of unprotected model information of client k
P_{l}^{S}	Distribution of protected model information of client <i>k</i>
W_{i}^{S}	Support of P_{L}^{S}
W_{L}^{k}	Support of P_{L}^{δ}
$\hat{W_k}$	Union of the supports of P_L^S and P_L^O
$F_{l_{1}}^{\mathcal{B}}$	Adversary's prior belief distribution about the private information of client k
$F_{t}^{\mathcal{A}}$	Adversary's belief distribution about client k after observing the protected private information
F_{L}^{KO}	Adversary's belief distribution about client k after observing the unprotected private information
$\operatorname{IS}(\cdot \cdot)$	Jensen-Shannon divergence between two distributions
$TV(\cdot \cdot)$	Total variation distance between two distributions

Table 1. Table of Notation

3.1 Notations

We adhere to the tradition of using uppercase letters to indicate the random variables, such as D, and lowercase letters to denote the specific values they take on. We represent [K] as $\{1, 2, \dots, K\}$. We use lowercase letters such as f and p to denote probability density functions and uppercase letters such as F and P to denote distributions.

The probability density function f at value d is represented by $f_{D_k}(d)$, and the subindex represents the random variable. The conditional density function is denoted by the notation $f_{D_k|W_k}(d|w)$. Let p and q represent the probability densities (or probability masses) of P and Q for distributions P and Q over \mathbb{R}^n . The Jensen-Shannon divergence, which is a smoothed variation of the Kullback-Leibler divergence, is defined as $JS(P||Q) = \frac{1}{2} [KL(P, M) + KL(Q, M)]$ with M = (P + Q)/2. Let TV(P||Q) denote the total variation distance between P and Q, which is defined as $TV(P||Q) = \sup_{A \subset \mathbb{R}^n} |P(A) - Q(A)|$. The detailed description of notations is illustrated in Table 1.

3.2 The Conceptual View of the Unified Federated Learning Framework

In this section, we propose a unified federated learning framework that unifies HFL and VFL through a conceptual view of the relationship between privacy leakage, utility loss, and efficiency reduction via the lens of the protector and the adversary. We first introduce HFL and VFL and unify their terminologies and notations. Then, we elaborate on this conceptual view.

HFL has two representative aggregation implementations: FedAvg and FedSGD, which are mathematically equivalent. Our framework applies to both aggregations. For illustrative purposes, we use FedAvg to explain the procedure of secure horizontal federated learning:

(1) With the global model state from the server, each client k trains its local model using its private data set D_k , and obtains the local model θ_k^O .

- (2) In order to prevent the semi-honest adversaries from inferring other clients' private information D_k according to θ_k^O , each client k adopts a protection mechanism M to convert model θ_k^O to protected model θ_k^S , and sends θ_k^S to the server.
- (3) The server aggregates θ_k^S , $k = 1, \dots, K$ to generate a new global model θ_{agg}^S .
- (4) Each client k downloads the global model θ_{agg}^{S} and uses it to update its local model.

The processes (1)-(4) iterate until the utility of the aggregated model θ_{agg}^{S} does not improve.



Fig. 1. Illustration of the HFL procedure (left) and VFL procedure (right). For the HFL setting, we consider the server as the adversary who aims to infer individual clients' private data. For the VFL setting, we consider client 1 (has no labels) as the adversary who wants to recover the labels owned by client 2.

In VFL, without loss of generality, we consider a 2-client scenario where client 1 (has no labels) is the adversary who wants to recover the labels owned by client 2. The overall secure vertical federated learning procedure is illustrated on the right panel of Fig. 1 and summarized as follows:

- (1) Client 1 and 2 compute their model outputs u_1 and u_2 based on their local bottom model information using their corresponding local data sets, respectively.
- (2) Client 1 sends u_1 to client 2.
- (3) Client 2 aggregates u_1 and u_2 and then feeds the aggregated result to its top model $\theta_{2,t}^O$ to generate the predicted labels, which further is used to compute loss against the ground truth label y.
- (4) Client 2 updates its top model $\theta_{2,t}^O$ and continues to compute the derivatives g_1 and g_2 of the loss L w.r.t. u_1 and u_2 , respectively.
- (5) In order to prevent the adversarial client 1 from inferring labels of client 2 based on g_1 , client 2 adopts protection mechanism M to convert g_1 to protected g_1^S , and sends g_1^S to client 1.
- 6 Clients 1 and 2 update their local bottom model $\theta_{1,b}^{S}$ and $\theta_{2,b}^{O}$ based on g_{1}^{S} and g_{2} , respectively.

The processes (1–6) iterate until the utility of the joint model θ_{joint}^S does not improve.

For the convenience of our analysis, we unify the terminologies and notations used for HFL and VFL (summarized in Table 2). The final trained model in HFL is typically called *aggregated model*, while the one in VFL is called *joint model*. We refer to them as *federated model*. The private information of HFL typically involves private features, labels, and statistical information; the one we study in this work for VFL is private labels. We refer to them all as *private information*. In HFL, the information exploited by the attacker to infer the private information of protector k is θ_k^S if the attacker is the server, and is θ_{agg}^S if the attacker is the client (see Figure 1 left). In VFL, the exploited information is referred to as g_k^S (see Figure 1 right). We unify them as W_k^S and W_{fed}^S . The utility of the federated model of HFL is usually defined as $\frac{1}{K} \sum_{k=1}^{K} U_k(\theta_{agg}^S)$, while the one of VFL is $U(\theta_{joint}^S)$. We unify them as $U(W_{fed}^S)$.

	Horizontal FL	Vertical FL	Unified FL
global model	aggregated model	joint model	federated model
privacy	private data or statistical information	private labels	private information
information exposed to the attacker	$ \begin{array}{c} \theta_k^{\mathcal{S}} \text{ (attacker: server)} \\ \text{or } \theta_{\text{agg}}^{\mathcal{S}} \text{ (attacker: client)} \end{array} $	$g_k^{\mathcal{S}}$	$W_k^{\mathcal{S}}$ or $W_{\text{fed}}^{\mathcal{S}}$
utility	$\frac{1}{K}\sum_{k=1}^{K}U_k(\theta_{\mathrm{agg}}^{\mathcal{S}})$	$U(\theta_{\text{joint}}^{\mathcal{S}})$	$U(W_{\rm fed}^{\mathcal{S}})$

Table 2. Unified terminologies and notations for HFL and VFL

Horizontal and vertical federated learning are the two primary forms of federated learning. While they have distinct training procedures, their privacy-preserving problem can be boiled down to the competition between the adversary and the protector. The former tries to learn as much private information on *D* as possible via privacy attacks. At the same time, the latter applies protection mechanisms to mitigate privacy leakage while maintaining utility loss and efficiency reduction below an acceptable level.

Figure 2 gives our proposed unified federated learning framework that illustrates the relationship between the privacy leakage ϵ_p , utility loss ϵ_e and efficiency reduction ϵ_u through the lens of the adversary and the protector. Specifically, to protect the privacy of its local data D, the protector converts the plain-text message msg to protected one msg^S exposed to the adversary aiming to mitigate the privacy leakage while maintaining the utility loss below an acceptable level. On the other hand, the adversary launches privacy attacks on protected model information W^S aiming to infer as much information on D as possible. Since protection mechanisms may jeopardize model utility and reduce efficiency, the protector needs to control the strength of the applied protection mechanism to strike a balance between private leakage, utility loss, and efficiency reduction.

The unified FL framework provides a conceptual view of the relationship between ϵ_p , ϵ_u , and ϵ_e through the lens of the protector and the adversary, regardless of the underlying FL architecture. Thus, it also reconciles other FL architectures, such as peer-to-peer (P2P) FL. In this work, we propose the No-Free-Lunch (NFL) theorem that quantifies the trade-off between ϵ_p , ϵ_u , and ϵ_e from the perspective of the protector of our unified federated learning framework. Thus, the NFL theorem applies to HFL, VFL, and P2P FL (e.g., Swarm Learning [45]).

3.3 Protection and Privacy Attacking Mechanisms

3.3.1 Protection Mechanisms. The data protector applies certain protection mechanisms to the exposed model information to mitigate privacy leakage. We formally define the protection mechanism as follows:

Definition 3.1 (Protection Mechanism and Protected Distribution). The **protection mechanism** $M : \mathbb{R}^m \to \mathbb{R}^m$ maps the original model information W_k^O , which follows a distribution P_k^O , to its protected (or distorted) counterpart W_k^S , which follows a distribution P_k^S . The objective of M is to protect private data so that the

Trading Off Privacy, Utility and Efficiency in Federated Learning • 7



Fig. 2. The unified federated learning framework (including HFL and VFL) illustrates the relationship between privacy leakage ϵ_p , utility loss ϵ_u , and efficiency reduction ϵ_e . D denotes the private data, which can be either features or labels. *msg* denotes the message encoded with knowledge of D. It is sent from the data owner to the adversary. In HFL, *msg* can be model parameters, model gradients, and model outputs sent from individual clients to the server. In VFL, *msg* is typically the intermediate gradient (see Figure 1 (right)) sent from the client with labels to adversarial clients. W denotes any model information derived from *msg* that the privacy attack can leverage to infer D. In HFL, W typically includes model parameters and model gradients. In VFL, W includes intermediate gradients, model parameters, and model gradients. We use superscripts S and O to distinguish the protected information and the plain-text one, respectively.

dependency between W_k^S and D_k is reduced, compared to the dependency between the unprotected information W_k^O and D_k . The distribution P_k^S is referred to as the **protected distribution** of client k.

EXAMPLE: We take the randomization mechanism as an illustrative example and introduce the *protected* distribution. Assume that $W_k^O \sim P_k^O = \mathcal{N}(\mu_0, \Sigma_0)$, and $\epsilon_k \sim \mathcal{N}(0, \Sigma_{\epsilon})$, where $\Sigma_0 = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$, $\Sigma_{\epsilon} = \text{diag}(\sigma_{\epsilon}^2, \dots, \sigma_{\epsilon}^2)$. Then the protected parameter $W_k^S = W_k^O + \epsilon_k \sim \mathcal{N}(\mu_0, \Sigma_0 + \Sigma_{\epsilon})$. That is, the protected distribution $P_k^S = \mathcal{N}(\mu_0, \Sigma_0 + \Sigma_{\epsilon})$. Please refer to Section D for more details.

In this work, we consider widely-adopted protection mechanisms, including Randomization Mechanism, Homomorphic Encryption, Secret Sharing and Compression Mechanism.

3.3.2 Privacy Attacking Mechanisms. The requirement of adopting specific protection mechanisms depends on the adversary's threat model. We consider the adversary to be *semi-honest*, he/she faithfully adheres to the federated learning protocol, yet may execute privacy attacks on exposed data to deduce the private information of other participants.

Let D_k denote private features or labels in horizontal federated learning and denote labels in vertical federated learning. Let W_k represent the model information exposed by client k. Let $f_{D_k|W_k}$ represent the probability density function of the posterior distribution $F_{D_k|W_k}$. Next, we introduce the Bayesian Inference Attack that estimates the private information by maximizing the posterior belief $f_{D_k|W_k}$.

Definition 3.2 (Bayesian Inference Attack). Given the distorted model information W_k^S , the **Bayesian Inference Attack** finds data *d* that maximizes the posterior belief:

$$\arg\max_{d} \operatorname{IH}(d|w) = \arg\max_{d} [\operatorname{I}(w|d) + \operatorname{H}(d)], \tag{1}$$

where $IH(d|w) = \log f_{D_k|W_k^S}(d|w)$ corresponds to the logarithm of the posterior belief, $I(w|d) = \log f_{W_k^S|D_k}(w|d)$ measures the logarithm of the likelihood based on the observed model information w, and $H(d) = \log f_{D_k}(d)$ represents the logarithm of the prior belief $f_{D_k}(d)$. According to Bayes' theorem, maximizing the logarithm of posterior $f_{D_k|W_k^S}(d|w)$ on D_k involves maximizing summation of $\log(f_{W_k^S|D_k}(w|d))$ and $\log(f_{D_k}(d))$.

Below, we introduce some representative privacy attacks (see Table 3) that fall under the family of the Bayesian Inference Attack.

Table 3. Privacy attacking mechanisms that fit the Bayesian inference attack. For Gradient Inversion attacks, g denotes the observed gradient, and w corresponds to θ_k^S in the left panel of Figure 1. For Label Inference attacks, w corresponds to g_1^S in the right panel of Figure 1).

Attack	Work	I(w d)	H(d)	Туре
Gradient Inversion	DLG [52]	$C - \frac{1}{2\sigma^2} \ g - \nabla_w \mathcal{L}(d, w)\ _2^2$	Constant	Horizontal
	Inverting Gradients [15]	$\frac{26}{\ \nabla_{\mathbf{w}}\mathcal{L}(d,\mathbf{w}),g\rangle}$	$\mathrm{TV}(d)$	Horizontal
	Improved DLG [51]	$C - \ g - \nabla_w \mathcal{L}(d, w)\ _2^2$	Label(d)	Horizontal
	GradInversion [47]	$C - \frac{1}{2\sigma^2} \ g - \nabla_w \mathcal{L}(d, w)\ _2^2$	$\operatorname{Group}(d)$	Horizontal
Label Inference	Norm-based Scoring [26]	$N(\sum_{i=1}^{\psi} \mathbb{1}\{r(w) \in [l_i, u_i]\} \cdot \mathbb{1}\{d \in C_i\})$	Constant	Vertical
	Direct Label Inference [14]	$N(\sum_{i=1}^{\psi} \mathbb{1}\{w \in [l_i, u_i]\} \cdot \mathbb{1}\{d \in C_i\})$	Constant	Vertical

Privacy Attacks in HFL. The following attacking mechanisms proposed for HFL fit into the Bayesian inference attack framework formulated in Eq. (1). These attacks infer private data *d* through maximizing the similarity between the observed gradient *g* and the estimated gradient $\nabla_w \mathcal{L}(d, w)$ with different prior on *d* (*w* here corresponds to θ_k^S in the left panel of Figure 1).

- Deep Gradient Leakage (DLG) [52]: H(d) is constant; I(w|d) is the negative ℓ_2 distance between the observed gradient g and the estimated gradient $\nabla_w \mathcal{L}(d, w)$.
- Inverting Gradients [15]: H(d) is the TV loss of estimated data, denoted as TV(d); I(w|d) corresponds to the cosine similarity between the observed gradient g and the estimated gradient ∇_wL(d, w).
- Improved DLG [51]): H(d) is the prior with the label information of d, denoted as Label(d); I(w|d) corresponds to the negative ℓ₂ distance between the observed gradient q and the estimated gradient ∇_w L(d, w).
- GradInversion [47]: H(d) is the group consistency of estimated data, denoted as Group(d); I(w|d) corresponds to the negative ℓ₂ distance between the observed gradient g and the estimated gradient ∇_wL(d, w).

Privacy Attacks in VFL. We focus on label inference attacks because labels may contain sensitive user information, and they are typically valuable assets in real-world VFL applications such as finance, healthcare, and advertisement. More specifically, we consider gradient scoring attack and direct label inference attack, which comply with the Bayesian inference attack framework. Both attacks are mounted by client 1, aiming to infer labels owned by client 2 based on back-propagated gradient *w* (*w* here corresponds to g_1^S in the right panel of Figure 1).

- Gradient Scoring Attack. I(w|d) = N (∑^ψ_{i=1} 1{r(w) ∈ [l_i, v_i]} · 1{d ∈ C_i}), where r : ℝ^m → [0, 1] denotes a scoring function; d represents one data point; l_i and v_i denotes the lower bound and upper bound, respectively; ψ represents the total number of classes; C_i represents class i; N(·) computes the logarithm of the normalized input. H(d) = constant. The gradient scoring attack typically applies to the binary classification task. The attacker needs to design a scoring function satisfying that r(w) ∈ [l₁, v₁] if the corresponding data of w belongs to the negative class, and r(w) ∈ [l₂, v₂] if the corresponding data of w belongs to the negative class, and r(w) ∈ [l₂, v₂] if the corresponding data of w attack are generally larger than that of the negative ones and formulates the attacking problem as the classification problem.
- Direct Label Inference Attack (DLI) [14]. DLI is tailored to the VFL scenario where the top model owned by client 2 is an activation function (e.g., softmax) [28], the adversary has access to the gradients of the final activation function, and thus it can infer labels through the signs of back-propagated gradients. In DLI, I(w|d) = N (∑_{i=1}^ψ 11{w_i ∈ ([l_i, v_i]} · 11{d ∈ C_i}), where l_i = -∞ and v_i = 0 if d ∈ C_i; l_i = 0 and v_i = +∞ otherwise; H(d) = constant. Thus, if w_i < 0 its corresponding data point d belongs to class C_i and if w_i>0 the data point d does not belong to class C_i.

Remark: Attacking mechanisms, such as the model completion attack [14], that use the cumulative information over rounds and is beyond the scope of our article.

3.4 Privacy Leakage, Utility Loss and Efficiency Reduction

In this section, we formally define privacy leakage, utility loss, and efficiency reduction (depicted in Figure 2).

The privacy leakage (Definition 3.3) measures the variation between the adversary's prior and posterior beliefs on private information. The adversary obtains the posterior belief by mounting a Bayesian inference attack (Definition 3.2) on the protected model information. Thus, we formally call the privacy leakage the *Bayesian* privacy leakage. The protected model information may lead to the federated model with lower utility (Definition 3.4) and the federated training with less efficiency (Definition 3.5) in certain scenarios.

Let $F_k^{\mathcal{A}}$, F_k^O and $F_k^{\mathcal{B}}$ represent the attacker's belief distributions about D_k upon observing the protected information, the original information and without observing any information, respectively, and the probability density functions of which are $f_{D_k}^{\mathcal{A}}$, $f_{D_k}^O$, and $f_{D_k}^{\mathcal{B}}$. Specifically, $f_{D_k}^{\mathcal{A}}(d) = \int_{W_k} f_{D_k|W_k}(d|w)dP_k^{\mathcal{S}}(w)$, $f_{D_k}^O(d) = \int_{W_k} f_{D_k|W_k}(d|w)dP_k^O(w)$, and $f_{D_k}^{\mathcal{B}}(d) = f_{D_k}(d)$.

We use JS divergence to measure the privacy leakage instead of KL divergence. The advantage of JS divergence over KL divergence is that it is symmetrical, and its square root satisfies the triangle inequality [12]. This property facilitates the quantification of the trade-offs.

Definition 3.3 (Bayesian Privacy Leakage). Let $\epsilon_{p,k}$ represent the privacy leakage of client k, which is defined as:

$$\epsilon_{p,k} = \sqrt{\mathrm{JS}(F_k^{\mathcal{A}}||F_k^{\mathcal{B}})},\tag{2}$$

where $JS(F_k^{\mathcal{A}}||F_k^{\mathcal{B}}) = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d) + \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{B}}(d) \log \frac{f_{D_k}^{\mathcal{B}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d), F_k^{\mathcal{A}} \text{ and } F_k^{\mathcal{B}} \text{ represent the attacker's belief distribution about } D_k \text{ upon observing the protected information and without observing any information, respectively, and <math>f_{D_k}^{\mathcal{M}}(d) = \frac{1}{2}(f_{D_k}^{\mathcal{A}}(d) + f_{D_k}^{\mathcal{B}}(d)).$

Furthermore, the Bayesian privacy leakage in FL resulting from releasing the protected model information is defined as:

$$\epsilon_p = \frac{1}{K} \sum_{k=1}^{K} \epsilon_{p,k}.$$

Remark:

(1) The local model information W_k represents the *model parameters, model gradients* and *model outputs*, all of which may optionally be exchanged or get exposed to semi-honest adversaries.

(2) If the private information D_k is continuous, then

$$JS(F_k^{\mathcal{A}}||F_k^{\mathcal{B}}) = \frac{1}{2} \left[KL\left(F_k^{\mathcal{A}}, F_k^{\mathcal{M}}\right) + KL\left(F_k^{\mathcal{B}}, F_k^{\mathcal{M}}\right) \right]$$
$$= \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d) + \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{B}}(d) \log \frac{f_{D_k}^{\mathcal{B}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d).$$

If the private information D_k is discrete, then

$$JS(F_k^{\mathcal{A}}||F_k^{\mathcal{B}}) = \frac{1}{2} \left[KL\left(F_k^{\mathcal{A}}, F_k^{\mathcal{M}}\right) + KL\left(F_k^{\mathcal{B}}, F_k^{\mathcal{M}}\right) \right]$$
$$= \frac{1}{2} \sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} + \frac{1}{2} \sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{B}}(d) \log \frac{f_{D_k}^{\mathcal{B}}(d)}{f_{D_k}^{\mathcal{M}}(d)}$$

When evaluating the utility loss, we consider the scenario when each protector is assigned a private key (if one exists).

Definition 3.4 (Utility Loss). The utility loss of client *k* (denoted as $\epsilon_{u,k}$) measures the variation in utility of client *k* with the federated model drawn from unprotected distribution P_{fed}^O and the utility of the federated model drawn from protected distribution P_{fed}^S :

$$\epsilon_{u,k} = \mathbb{E}_{W^{O}_{\text{fed}} \sim P^{O}_{\text{fed}}} [U_k(W^{O}_{\text{fed}})] - \mathbb{E}_{W^{S}_{\text{fed}} \sim P^{S}_{\text{fed}}} [U_k(W^{S}_{\text{fed}})],$$

where U_k represents the utility function of client k. Furthermore, the utility loss in FL system is defined as:

$$\epsilon_{u} = \mathbb{E}_{W^{O}_{\text{fed}} \sim P^{O}_{\text{fed}}}[U(W^{O}_{\text{fed}})] - \mathbb{E}_{W^{S}_{\text{fed}} \sim P^{S}_{\text{fed}}}[U(W^{S}_{\text{fed}})],$$

where U represents the utility function of the FL system.

Efficiency reduction is another criterion we consider in this work, and defined as follows.

Definition 3.5 (Efficiency Reduction). Let $\epsilon_{e,k}$ represent the efficiency reduction of client k. The efficiency reduction of client k measures the variation in efficiency with the models drawn from the unprotected and protected distributions P_k^O and P_k^S , which is defined as:

$$\epsilon_{e,k} = \mathbb{E}_{W_k^S \sim P_k^S}[C(W_k^S)] - \mathbb{E}_{W_k^O \sim P_k^O}[C(W_k^O)],$$

where *C* denotes a mapping from the model information to the efficiency measured in terms of the communication cost (e.g., the transmitted bits) or the training cost. Furthermore, the efficiency reduction in FL system is defined as:

$$\epsilon_e = \frac{1}{K} \sum_{k=1}^{K} \epsilon_{e,k}.$$

3.5 The Competition Between The Protector and The Adversary

The Goal of the Protector The goal of protector k is formulated as an optimization problem that aims at finding a protected distribution P_k^S achieving the minimum utility loss and efficiency reduction under the privacy constraint χ_k , which is formally expressed as

$$\min_{\substack{P_k^S}} \eta_u \cdot \epsilon_{u,k} + \eta_e \cdot \epsilon_{e,k},\tag{3}$$

subject to
$$\epsilon_{p,k} \leq \chi_k$$
. (4)

where η_u represents the preference of the protector towards model utility, and η_e represents the preference of the protector towards efficiency.

The Goal of the Adversary Let D_k be the client k's private information. Let W_k^S represent the client k's model information exposed to the adversary. The goal of the adversary is formally expressed as

$$\max_{F_{D_k|W_k^{\mathcal{S}}}} \epsilon_{p,k} \coloneqq \mathrm{JS}(F_k^{\mathcal{A}}||F_k^{\mathcal{B}}), \tag{5}$$

where $F_{D_k|W_k^S}$ represents the adversary's posterior distribution of the private information D_k upon observing the exposed model information W_k^S .

Competition Between the Protector and the Adversary Given the goal of the protector and the adversary formally stated in Eq. (3) and Eq. (5), the competition between the protector and the adversary is formally expressed as

$$\begin{array}{ll}
\min_{\substack{P_k^S \\ k}} & \eta_u \cdot \epsilon_{u,k} + \eta_e \cdot \epsilon_{e,k}, \\
\text{subject to} & \max_{\substack{P_{D_k \mid W_k^S} \\ k \neq k}} \epsilon_{p,k} \le \chi_k.
\end{array}$$
(6)

where χ_k is the privacy level required by client *k*.

Remark: This optimization problem can be extended to satisfy the personalized requirement of the federated learning system. For example, if the efficiency reduction of the federated learning system is required not to exceed φ_k , then the optimization problem is expressed as

$$\min_{\substack{P_k^S \\ \text{subject to}}} \epsilon_{u,k},$$

$$\max_{F_{D_k \mid W_i^S}} \epsilon_{p,k} \le \chi_k, \epsilon_{e,k} \le \varphi_k.$$
(7)

A natural question is: *is it possible to design a protection mechanism that simultaneously achieves infinitesimal privacy leakage, utility loss, and efficiency reduction?* In the following section, we provide a negative answer for this question in certain scenarios, which leads to our No-Free-Lunch Theorem (Theorem 4.6).

4 NO FREE LUNCH THEOREM IN FEDERATED LEARNING

In this section, we propose the No-Free-Lunch (NFL) theorem (Theorem 4.6), stating that it is unrealistic to expect a privacy-preserving FL algorithm to simultaneously achieve the infinitesimal privacy leakage, utility loss, and efficiency reduction in certain scenarios.

Before elaborating on the NFL theorem, we introduce Assumption 4.1 and Assumption 4.2, under which the federated learning scenarios we analyze guarantee the existence of trade-offs between privacy, utility, and efficiency.

Assumption 4.1 is proposed in [50] and illustrated in Figure 3. We present it here for completeness. Intuitively, it states that the cumulative density of the near-optimal parameters is bounded, which rules out scenarios where the utility is constant and most parameters are near-optimal parameters.

Assumption 4.1. let W_{fed}^{S} represent the support of the protected distribution P_{fed}^{S} of the federated model information. Let W_{fed}^{*} represent the set of parameters achieving the maximum utility

$$W_{\text{fed}}^* = \operatorname*{argmax}_{w \in W_{\text{fed}}} U(w)$$

where *U* represents the utility of the federated learning system. Given a non-negative constant *c*, the *near-optimal parameters* is defined as

$$\mathcal{W}_{c} = \left\{ w \in \mathcal{W}_{\text{fed}}^{\mathcal{S}} : |U(w^{*}) - U(w)| \leq c, \forall w^{*} \in \mathcal{W}_{\text{fed}}^{*} \right\}.$$

Let $c = \Delta$ be the *maximum* constant that satisfies:

$$\int_{\mathcal{W}_{\text{fed}}^{S}} \mathbb{1}\{w \in \mathcal{W}_{\Delta}\} p_{\mathcal{W}_{\text{fed}}}^{S}(w) dw \leq \frac{\text{TV}(P_{\text{fed}}^{O} || P_{\text{fed}}^{S})}{2},\tag{8}$$

where $p_{W_{\text{fed}}}^{S}$ represents the probability density function of the protected federated model information W_{fed}^{S} . We assume that Δ is positive, i.e., $\Delta > 0$.



Fig. 3. Illustration of Assumption 4.1. The blue area represents the probability that the realization of the protected model information W_{fed}^S falls inside the *near-optimal parameters* W_{Δ} . It is referred to as the cumulative probability density of the near-optimal parameters. Note that we use model parameters with one dimension in this figure for illustrative purposes.

Remark: The Eq. (8) states that the cumulative density of the near-optimal distorted parameters is upperbounded by $\frac{\text{TV}(P_{\text{fed}}^{O}||P_{\text{fed}}^{S})}{2}$, which is at most $\frac{1}{2}$. Therefore, the utility is not constant, and most distorted parameters are sub-optimal.

Remark: If the distribution of the protected model information is continuous and $TV(P_{fed}^O || P_{fed}^S) = 0$, then there does not exist a positive constant Δ satisfying Eq. (8), which implies that Assumption 4.1 does not hold. From Lemma C.3 of [50], the utility equals 0.

From Lemma C.3 of [50], the utility equals 0. Let $p_{W_k}^O$ represent the probability density of unprotected model information of client k, and $p_{W_k}^S$ represent the probability density of protected model information of client k. The following Assumption 4.2 states that the distorted parameters have higher communication cost than that of the original parameters with high probability. Therefore, the expected communication cost of the distorted parameter is higher than that of the original parameter. This assumption rules out scenarios where the communication cost is constant, and the communication cost of most distorted parameters is smaller than that of the original parameters. We provide an example for illustration in Figure 4.

Assumption 4.2. Let $w_{\max}^O = \arg \max_{w \in W_k^O} C(w)$. Let Ξ_k denote the minimum non-negative constant satisfying that $C(w_{\max}^O) - \Xi_k \leq C(w)$, $\forall w \in W_k^S$. Let \mathcal{U}_k represent the set of distorted parameters with improved probability density. Specifically, $\mathcal{U}_k = \{w \in W_k^S : p_{W_k}^S(w) \geq p_{W_k}^O(w)\}$. We denote \mathcal{W}_{k,Ξ_k}^+ as the set of distorted parameters with high probability density, and the communication cost of which is larger than that of the original parameters with a gap of at least $2 \cdot \Xi_k$. We denote \mathcal{W}_{k,Ξ_k}^+ as:

$$\mathcal{W}_{k,\Xi_k}^+ = \left\{ w \in \mathcal{U}_k : C(w_{\max}^O) + 2 \cdot \Xi_k \le C(w) \quad \text{and} \quad p_{W_k}^S(w) \ge 2 \cdot p_{W_k}^O(w) \right\}.$$

We denote $\mathcal{W}_{k\Xi_k}^- = \mathcal{U}_k \setminus \mathcal{W}_{k\Xi_k}^+$. Let Γ_k denote the maximum constant satisfying that

$$\underbrace{\int_{\mathcal{U}_{k}} \mathbb{1}\{w \in \mathcal{W}_{k,\Xi_{k}}^{+}\} p_{W_{k}}^{\mathcal{S}}(w) dw}_{\text{term 1}} - \underbrace{\int_{\mathcal{U}_{k}} \mathbb{1}\{w \in \mathcal{W}_{k,\Xi_{k}}^{-}\} p_{W_{k}}^{\mathcal{S}}(w) dw}_{\text{term 2}} \ge \Gamma_{k} \cdot \text{TV}(P_{k}^{\mathcal{O}}||P_{k}^{\mathcal{S}}), \tag{9}$$

where $p_{W_k}^S$ denotes the probability density function of the protected model information and $p_{W_k}^O$ denotes the probability density function of the original model information. We assume that Ξ_k and Γ_k are positive, i.e., $\Xi_k > 0, \Gamma_k > 0, \forall k \in [K]$.



Fig. 4. Illustration of Assumption 4.2. The left panel shows the relationship between the sets W_{k,Ξ_k}^+ , W_{k,Ξ_k}^- , U_k , and W_k^S . The right panel illustrates Eq. (9). In the right panel, the blue area (term 1 in Eq.(9)) and red area (term 2 in Eq.(9)) represent the probability that the realization of the protected model information falls inside W_{k,Ξ_k}^+ and W_{k,Ξ_k}^- , respectively, and the blue area is larger than the red area. Note that we use model parameters with one dimension for illustrative purposes.

Remark: The Eq. (9) states that the difference between the distorted parameters' cumulative density in W_{k,Ξ_k}^+ and W_{k,Ξ_k}^- is lower-bounded by a non-negative value. Therefore, it conveys that the expected communication cost of the distorted parameter is higher than that of the original parameter. Since $\Xi_k > 0$, the communication cost is not constant.

In the next section, we introduce the quantitative trade-offs between privacy and efficiency and between privacy and utility, respectively. These two trade-offs lead to the main conclusion of our No-Free-Lunch theorem formulated in Theorem 4.6.

4.1 Trade-off between Privacy and Efficiency

The following Lemma 4.1 bounds efficiency reduction using the distortion measured by total variation distance $TV(P_{fed}^{O}||P_{fed}^{S})$. The intuition is that a larger distortion leads to a higher efficiency reduction. The analysis is deferred to Lemma B.3.

Lemma 4.1. Let Assumption 4.2 hold, and ϵ_e be defined in Definition 3.5. Let P_{fed}^O and P_{fed}^S represent the distribution of the aggregated parameter before and after being protected. Let $\Xi = \min_{k \in [K]} \Xi_k$, and $\Gamma = \min_{k \in [K]} \Gamma_k$. Then, we have:

$$\epsilon_{e} \geq \Xi \cdot \Gamma \cdot \frac{1}{K} \sum_{k=1}^{K} \text{TV}(P_{k}^{O} || P_{k}^{S}).$$

The following Lemma 4.2 illustrates the quantitative relationship between ϵ_p and total variation distance. Lemma 4.2 applies to both discrete (see proof in [50]) and continuous private information (see Lemma B.1).

Lemma 4.2. Let $\epsilon_{p,k}$ be defined in Definition 3.3, P_k^O and P_k^S represent the distribution of the parameter of client k before and after being protected. Then, we have:

$$\epsilon_p \ge C_1 - \frac{1}{K} \sum_{k=1}^K C_2 \cdot \mathrm{TV}(P_k^O || P_k^S), \tag{10}$$

where $C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^O || F_k^B)}$, and $C_2 = \frac{1}{2}(e^{2\xi} - 1)$, $\xi_k = \max_{w \in W_k, d \in D_k} \left| \log \left(\frac{f_{D_k | W_k}(d | w)}{f_{D_k}(d)} \right) \right|$, and $\xi = \max_{k \in [K]} \xi_k$ represents the maximum privacy leakage over all possible information w released by the clients, and $[K] = \{1, 2, \dots, K\}$.

Intuitively, a larger distortion would increase communication costs while decreasing privacy leaks. The Lemma 4.3 demonstrates how the summation of privacy leakage and efficiency reduction is lower bounded by a constant that depends on the nature of the scenario.

Lemma 4.3 (No free lunch theorem (NFL) for privacy and efficiency). Let ϵ_p be defined in Def. 3.3, and let ϵ_e be defined in Def. 3.5, with Assumption 4.2 we have:

$$\epsilon_p + C_x \cdot \epsilon_e \ge C_1,$$

where $C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^O || F_k^{\mathcal{B}})}$ and $C_x = \frac{1}{2\Xi\Gamma} (e^{2\xi} - 1)$.

4.2 Trade-off between Privacy and Utility

The Lemma 4.4 demonstrates how the total variation distance between the protected and unprotected distributions lower bounds utility loss.

Lemma 4.4. Let Assumption 4.1 hold, and ϵ_u be defined in Definition 3.4. Let P_{fed}^O and P_{fed}^S represent the distribution of the federated model information before and after being protected. Then, we have:

$$\epsilon_u \geq \frac{\Delta}{2} \cdot \mathrm{TV}(P_{\mathrm{fed}}^O || P_{\mathrm{fed}}^S).$$

With Lemma 4.2 and Lemma 4.4, it is now natural to provide a quantitative relationship between the utility loss and the privacy leakage.

Lemma 4.5 (No free lunch theorem (NFL) for privacy and utility). Let ϵ_p be defined in Def. 3.3, and let ϵ_u be defined in Def. 3.4 at the convergence step, with Assumption 4.1 we have:

$$\epsilon_p + C_d \cdot \epsilon_u \ge C_1,\tag{11}$$

where $\xi_k = \max_{w \in \mathcal{W}_k, d \in \mathcal{D}_k} \left| \log \left(\frac{f_{D_k | W_k(d|w)}}{f_{D_k}(d)} \right) \right|$, and $\xi = \max_{k \in [K]} \xi_k$ represents the maximum privacy leakage over all possible information w released by client k, and $[K] = \{1, 2, \dots, K\}$, $C_1 = \frac{1}{K} \sum_{k=1}^K \sqrt{JS(F_k^O ||F_k^{\mathcal{B}})}$, $C_d = \frac{\gamma}{4\Delta} (e^{2\xi} - 1)$, and $\gamma = \frac{\frac{1}{K} \sum_{k=1}^K TV(P_k^O ||F_k^{\mathcal{B}})}{TV(P_{k-1}^O ||F_k^{\mathcal{B}})}$.

Remark: If the distribution of the protected model information is continuous, and $TV(P_{fed}^O||P_{fed}^S) = 0$, then Assumption 4.1 does not hold, and the utility loss is equal to 0 (from Lemma C.3 of [50]).

4.3 Trade-off between Privacy, Utility and Efficiency

The Theorem 4.6 illustrates the quantitative trade-off between privacy leakage, utility loss, and efficiency reduction, which shows that the weighted summation of these three metrics is larger than a problem-dependent constant (also exhibited in Figure 5). It implies that simultaneously achieving infinitesimal privacy leakage, utility loss, and efficiency reduction is unrealistic in certain scenarios. The complete analysis is deferred to Appendix C.

Theorem 4.6 (No free lunch theorem (NFL) for privacy, utility and efficiency). Let ϵ_e be defined in Def. 3.5, ϵ_p be defined in Def. 3.3, and let ϵ_u be defined in Def. 3.4 at the convergence step. Let Assumption 4.1 and Assumption 4.2 hold, then we have that:

$$\epsilon_p + \frac{C_d}{2} \cdot \epsilon_u + \frac{C_x}{2} \cdot \epsilon_e \ge C_1, \tag{12}$$

where $\xi_k = \max_{w \in \mathcal{W}_k, d \in \mathcal{D}_k} \left| \log \left(\frac{f_{D_k | W_k(d|w)}}{f_{D_k}(d)} \right) \right|, \xi = \max_{k \in [K]} \xi_k, C_1 = \frac{1}{K} \sum_{k=1}^K \sqrt{JS(F_k^O | | F_k^B)}, C_d = \frac{Y}{4\Delta} (e^{2\xi} - 1), C_x = \frac{1}{2\Xi\Gamma} (e^{2\xi} - 1), \gamma = \frac{\sum_{k=1}^K TV(P_k^O | | P_k^S)}{TV(P_{od}^O | | P_{sd}^S)}.$



Fig. 5. Privacy-Utility-Efficiency Trade-offs

Note that C_1 measures the discrepancy between the prior belief and the posterior belief of the attacker and is typically a positive constant in real-world applications. Theorem 4.6 states that the weighted summation of the privacy leakage, utility loss, and efficiency reduction is greater than or equal to C_1 in certain scenarios, indicating that a protection mechanism cannot be expected to achieve exceptional privacy, utility, and efficiency simultaneously in scenarios where Assumption 4.1 and Assumption 4.2 hold.

5 APPLICATIONS

In this section, we apply the No-Free-Lunch theorem to four well-adopted protection mechanisms for quantifying their trade-offs between privacy leakage, utility loss, and efficiency reduction.

We measure efficiency in terms of the communication cost and denote $\delta > 0$ as a small constant. $C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^O || F_k^B)}$ and $C_2 = (e^{2\xi} - 1)/2$ are two constants independent of the protection mechanisms adopted, and *m* is the dimension of the parameter.

Table 4. Comparison of homomorphic encryption, secret sharing, randomization, and compression in terms of privacy-utilityefficiency trade-off characterized via lower bounds for ϵ_p , ϵ_u , and ϵ_e . LB denotes lower bound.

	Paillier HE	Secret Sharing	Randomization	Compression
$\epsilon_p(LB)$	$C_1 - C_2 \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right]$	$C_1 - \frac{C_2}{K} \sum_{k=1}^{K} \left(1 - \prod_{j=1}^{m} \frac{2\delta}{b_k^j + r_k^j} \right)$	$1 - \frac{3C_2}{2} \cdot \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum_{i=1}^{m} \frac{1}{\sigma_i^4}}\right\}$	$C_1 - C_2 \cdot \left(1 - \prod_{i=1}^m \rho_i\right)$
$\epsilon_u(LB)$	0 (Assumption 4.1 does not hold)	0 (Assumption 4.1 does not hold)	$rac{\Delta}{100} \cdot \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum\limits_{i=1}^m rac{1}{\sigma_i^4}} ight\}$	$\frac{\underline{\Lambda}}{2} \cdot \left(1 - \prod_{i=1}^{m} \rho_i\right)$
$\epsilon_e(LB)$	$\Xi \cdot \Gamma \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right]$	N/A (Assumption 4.2 does not hold)	$\frac{\Xi \cdot \Gamma}{100} \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum_{i=1}^{m} \frac{1}{\sigma_i^4}}\right\}$	$\Xi \cdot \Gamma \cdot \left(1 - \prod_{i=1}^{m} \rho_i\right)$

5.1 Randomization Mechanism

Let W_k^O be the parameter sampled from distribution $P_k^O = \mathcal{N}(\mu_0, \Sigma_0)$, where $\mu_0 \in \mathbb{R}^n, \Sigma_0 = \operatorname{diag}(\sigma_1^2, \dots, \sigma_m^2)$ is a diagonal matrix. The distorted parameter $W_k^S = W_k^O + \epsilon_k$, where $\epsilon_k \sim \mathcal{N}(0, \Sigma_{\epsilon})$ and $\Sigma_{\epsilon} = \operatorname{diag}(\sigma_{\epsilon}^2, \dots, \sigma_{\epsilon}^2)$. The bounds for privacy leakage, utility loss and efficiency reduction for randomization mechanism is illustrated in the following theorem. The full proof is deferred to Appendix D.

Theorem 5.1. For randomization mechanism, the privacy leakage is bounded by

$$\epsilon_p \ge C_1 - \frac{C_2}{100} \cdot \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

The utility loss is bounded by

$$\epsilon_u \geq \frac{\Delta}{200} \cdot \min\left\{1, \sigma_\epsilon^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

The efficiency reduction is bounded by

$$\epsilon_e \geq \frac{\Xi\Gamma}{100} \min\left\{1, \sigma_\epsilon^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

5.2 Paillier Homomorphic Encryption

The **Paillier** encryption mechanism proposed by [34] is an asymmetric additive homomorphic encryption mechanism, which was widely applied in FL [3, 8, 13, 40, 49]. Paillier encryption contains three parts: key generation, encryption, and decryption. Let *h* represent the plaintext, and *c* represent the ciphertext. Let (*n*, *g*) represent the public key, and (λ , μ) represent the private key. Note that the primes *p* and *q* are rather large.

Encryption. Randomly select *r* and encode *h* as:

$$c = g^h \cdot r^n \bmod n^2,$$

where $n = p \cdot q$, p and q are two selected primes. Note that g is an integer selected randomly, and $g \in \mathbb{Z}_{n^2}^*$. Therefore, n can divide the order of g.

Decryption. Using the private key (λ, μ) to decrypt the ciphertext *c* as:

$$h = L(c^{\lambda} \bmod n^2) \cdot \mu \bmod n,$$

where $L(x) = \frac{x-1}{n}$, $\mu = (L(g^{\lambda} \mod n^2))^{-1} \mod n$, $\lambda = lcm(p-1, q-1)$, and *lcm* represents the least common multiple.

Let $W_k^O = (a_k^1, \dots, a_k^m)$ represent the plaintext parameter that follows a degenerate distribution. Let W_k^S represent the ciphertext parameter that follows a uniform distribution over $[0, n^2 - 1]^m$. To facilitate the calculation of the total variation distance between the unprotected distribution of W_k^O and the protected distribution of W_k^S , we relax the discrete unprotected distribution to a continuous distribution using a parameter $\delta > 0$. Assuming that W_k^O follows a uniform distribution over $[a_k^1 - \delta, a_k^1 + \delta] \times [a_k^2 - \delta, a_k^2 + \delta] \times \cdots \times [a_k^m - \delta, a_k^m + \delta]$, we can calculate the total variation distance between the distribution of W_k^O and the distribution of W_k^S . The following Theorem 5.2 provides bounds for privacy leakage and efficiency reduction using the magnitude of the ciphertext *n*. The bound for privacy leakage decreases with *n*, which could guide the selection of *n* to adapt to privacy and efficiency requirements. The full proof is deferred to Appendix E.

Theorem 5.2. For Paillier algorithm, the utility loss is 0, the privacy leakage is bounded by:

$$\epsilon_p \ge C_1 - C_2 \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right].$$

 $\epsilon_e \ge \Xi \cdot \Gamma \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right].$

5.3 Secret Sharing Mechanism

The efficiency reduction is bounded by

Various privacy-preserving protocols based on MPC (especially secret sharing) have been proposed to build secure machine learning models, including linear regression, logistic regression, and recommendation systems. [5, 6, 38] were proposed to distribute secrets between participants.

Note that $TV(P_{fed}^{O}||P_{fed}^{S}) = 0$ for secret sharing mechanism. From Lemma C.3 of [50], the utility loss equals 0. The communication cost for the parameter of the secret sharing mechanism is not guaranteed to satisfy Assumption 4.2, and the analysis of the lower bound of the efficiency reduction is beyond the scope of our article.

Assumption 4.2, and the analysis of the lower bound of the efficiency reduction is beyond the scope of our article. Let W_k^O represent the original model information that follows a uniform distribution over $[a_k^1 - \delta, a_k^1 + \delta] \times [a_k^2 - \delta, a_k^2 + \delta] \cdots \times [a_k^m - \delta, a_k^m + \delta]$. Let W_k^S represent the distorted model information that follows uniform distribution over $[a_k^1 - b_k^1, a_k^1 + r_k^1] \times [a_k^2 - b_k^2, a_k^2 + r_k^2] \cdots \times [a_k^m - b_k^m, a_k^m + r_k^m]$. The following Theorem 5.3 measures utility loss and provides the lower bound for privacy leakage. The full proof is deferred to Appendix F.

Theorem 5.3. For secret sharing mechanism, the utility loss $\epsilon_u = 0$, the privacy leakage is bounded by

$$\epsilon_p \ge C_1 - C_2 \cdot \frac{1}{K} \sum_{k=1}^K \left(1 - \prod_{j=1}^m \frac{2\delta}{b_k^j + r_k^j} \right).$$

5.4 Compression Mechanism

For the compression mechanism, the client does not transfer all the parameters to the server. Let $W_k^O(i)$ ($W_k^S(i)$) denote dimension *i* of W_k^O (W_k^S). Let b_i represent a random variable that follows the Bernoulli distribution. Specifically, b_i takes the value 1 with probability ρ_i , and 0 with probability $1 - \rho_i$. Each dimension *i* of the distorted parameter $W_k^S(i)$ is defined as

$$W_k^{\mathcal{S}}(i) = \begin{cases} W_k^{\mathcal{O}}(i) & \text{if } b_i = 1, \\ 0, & \text{if } b_i = 0. \end{cases}$$

The following Theorem 5.4 provides bounds for privacy leakage, utility loss and efficiency reduction using the compression probability. The full proof is deferred to Appendix G.

Theorem 5.4. For compression mechanism, the privacy leakage is lower bounded by

$$\epsilon_p \ge C_1 - C_2 \cdot \left(1 - \prod_{i=1}^m \rho_i\right)$$

The utility loss is bounded by

$$\epsilon_u \geq \frac{\Delta}{2} \cdot \left(1 - \prod_{i=1}^m \rho_i\right).$$

The efficiency reduction ϵ_e is bounded by

$$\epsilon_e \geq \Xi \cdot \Gamma \cdot \left(1 - \prod_{i=1}^m \rho_i\right).$$

6 CONCLUSION AND FUTURE WORK

In this work, we propose a unified federated learning (FL) framework that reconciles both HFL and VFL. Under this unified FL framework, we provide the No-Free-Lunch (NFL) theorem that quantifies the trade-off between privacy leakage, utility loss, and efficiency reduction for well-defined scenarios. We then leverage our proposed NFL theorem to analyze the lower bounds of the widely-adopted protection mechanisms, including the randomization mechanism, Paillier mechanism, secret sharing mechanism, and compression mechanism.

Lots of problems are worth investigating. In this work, we use the Jensen-Shannon (JS) divergence instead of the commonly-used KL divergence to measure privacy leakage. The JS divergence is expressed as the summation of two KL divergences. Thus, it inherits the advantages of KL divergence. Besides, JS divergence satisfies the triangle inequality, which facilitates the theoretical analysis. One potential question is, if we use a generalized version of JS-divergence weighted using a hyperparameter α [9], are we still able to bound the privacy leakage and quantify the trade-off analysis?

The optimal privacy-utility-efficiency trade-off is cast as a constrained optimization problem in which the *utility loss* and *efficiency reduction* are minimized subject to a predefined constraint for *privacy leakage*. The optimization problem and the derived lower bounds provide an avenue for proposing meta-algorithms that search the optimal hyperparameter characterizing the protection mechanism. Designing a meta-algorithm that seeks to determine the ideal protection hyperparameter at each communication round is an intriguing problem. Whether it is possible to design an algorithm that can learn the hyperparameter adaptively is another promising research direction.

ACKNOWLEDGMENTS

We would like to thank Di Chai, Xiaodian Cheng, Tao Fan, Shaofeng Jiang, Yilun Jin, Liu Yang, Ke Yi, Junxue Zhang for helpful discussions. This work was partially supported by the National Key Research and Development Program of China under Grant 2018AAA0101100 and Hong Kong RGC TRS T41-603/20-R.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In <u>Proceedings of the 2016 ACM SIGSAC conference on computer and communications security</u>. ACM, New York, NY, USA, 308–318.
- [2] Rodolfo Stoffel Antunes, Cristiano André da Costa, Arne Küderle, Imrana Abdullahi Yari, and Björn Eskofier. 2022. Federated Learning for Healthcare: Systematic Review and Architecture Proposal. <u>ACM Trans. Intell. Syst. Technol.</u> 13, 4, Article 54 (may 2022), 23 pages. https://doi.org/10.1145/3501813
- [3] Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shiho Moriai, et al. 2017. Privacy-preserving deep learning via additively homomorphic encryption. IEEE Transactions on Information Forensics and Security 13, 5 (2017), 1333–1345.
- [4] Kyohei Atarashi and Masakazu Ishihata. 2021. Vertical Federated Learning for Higher-Order Factorization Machines. In <u>Advances in Knowledge Discovery and Data Mining</u>, Kamal Karlapalem, Hong Cheng, Naren Ramakrishnan, R. K. Agrawal, P. Krishna Reddy, Jaideep Srivastava, and Tanmoy Chakraborty (Eds.). Springer International Publishing, Cham, 346–357.
- [5] G.R. Blakley. 1979. Safeguarding cryptographic keys. In <u>Proceedings of the 1979 AFIPS National Computer Conference</u>. AFIPS Press, Monval, NJ, USA, 313–317.
- [6] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (Dallas, Texas, USA) (CCS '17). Association for Computing Machinery, New York, NY, USA, 1175–1191. https://doi.org/10.1145/3133956.3133982
- [7] Wei-Ning Chen, Peter Kairouz, and Ayfer Ozgur. 2020. Breaking the communication-privacy-accuracy trilemma. <u>Advances in Neural</u> <u>Information Processing Systems</u> 33 (2020), 3312–3324.
- [8] Kewei Cheng, Tao Fan, Yilun Jin, Yang Liu, Tianjian Chen, Dimitrios Papadopoulos, and Qiang Yang. 2021. SecureBoost: A Lossless Federated Learning Framework. IEEE Intelligent Systems 36, 6 (2021), 87–98. https://doi.org/10.1109/MIS.2021.3082561
- [9] Jacob Deasy, Nikola Simidjievski, and Pietro Lió. 2020. Constraining Variational Inference with Geometric Jensen-Shannon Divergence. In Advances in Neural Information Processing Systems, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 10647–10658.
- [10] Flávio du Pin Calmon and Nadia Fawaz. 2012. Privacy against statistical inference. In <u>2012 50th annual Allerton conference on</u> communication, control, and computing (Allerton). IEEE, 1401–1408.
- [11] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. <u>Foundations and Trends in Theoretical</u> Computer Science 9, 3-4 (2014), 211–407.
- [12] Dominik Maria Endres and Johannes E Schindelin. 2003. A new metric for probability distributions. <u>IEEE Transactions on Information</u> theory 49, 7 (2003), 1858–1860.
- [13] Haokun Fang and Quan Qian. 2021. Privacy preserving machine learning with homomorphic encryption and federated learning. <u>Future</u> Internet 13, 4 (2021), 94.
- [14] Chong Fu, Xuhong Zhang, Shouling Ji, Jinyin Chen, Jingzheng Wu, Shanqing Guo, Jun Zhou, Alex X. Liu, and Ting Wang. 2022. Label Inference Attacks Against Vertical Federated Learning. In <u>31st USENIX Security Symposium (USENIX Security 22)</u>. USENIX Association, Boston, MA, 1397–1414. https://www.usenix.org/conference/usenixsecurity22/presentation/fu-chong
- [15] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting Gradients How Easy is It to Break Privacy in Federated Learning? (NIPS'20). Curran Associates Inc., Red Hook, NY, USA, Article 1421, 11 pages.
- [16] Craig Gentry. 2009. A fully homomorphic encryption scheme. Stanford university.
- [17] Robin C Geyer, Tassilo Klein, and Moin Nabi. 2017. Differentially private federated learning: A client level perspective. <u>arXiv preprint</u> <u>arXiv:1712.07557</u> (2017).
- [18] Bin Gu, Zhiyuan Dang, Xiang Li, and Heng Huang. 2020. Federated Doubly Stochastic Kernel Learning for Vertically Partitioned Data (KDD '20). Association for Computing Machinery, New York, NY, USA, 2483–2493.
- [19] Hanlin Gu, Lixin Fan, Bowen Li, Yan Kang, Yuan Yao, and Qiang Yang. 2021. Federated Deep Learning with Bayesian Privacy. <u>arXiv</u> preprint arXiv:2109.13012 (2021).
- [20] Otkrist Gupta and Ramesh Raskar. 2018. Distributed learning of deep neural network over multiple agents. Journal of Network and Computer Applications 116 (2018), 1–8.
- [21] Andrew Hard, Chloé M Kiddon, Daniel Ramage, Francoise Beaufays, Hubert Eichner, Kanishka Rao, Rajiv Mathews, and Sean Augenstein. 2018. Federated Learning for Mobile Keyboard Prediction. https://arxiv.org/abs/1811.03604

- [22] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista A. Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D'Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaïd Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Hang Qi, Daniel Ramage, Ramesh Raskar, Mariana Raykova, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. 2021. Advances and Open Problems in Federated Learning. Found. Trends Mach. Learn. 14, 1-2 (2021), 1–210. https://doi.org/10.1561/220000083
- [23] Yan Kang, Yuanqin He, Jiahuan Luo, Tao Fan, Yang Liu, and Qiang Yang. 2022. Privacy-preserving Federated Adversarial Domain Adaptation over Feature Groups for Interpretability. <u>IEEE Transactions on Big Data</u> (2022), 1–12. https://doi.org/10.1109/TBDATA.2022. 3188292
- [24] Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. 2016. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. <u>CoRR</u> abs/1610.02527 (2016). arXiv:1610.02527 http://arxiv.org/abs/1610.02527
- [25] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492 (2016).
- [26] Oscar Li, Jiankai Sun, Xin Yang, Weihao Gao, Hongyi Zhang, Junyuan Xie, Virginia Smith, and Chong Wang. 2022. Label Leakage and Protection in Two-party Split Learning. In <u>International Conference on Learning Representations</u>. https://openreview.net/forum?id= cOtBRgsf2fO
- [27] Wei Yang Bryan Lim, Nguyen Cong Luong, Dinh Thai Hoang, Yutao Jiao, Ying-Chang Liang, Qiang Yang, Dusit Niyato, and Chunyan Miao. 2020. Federated learning in mobile edge networks: A comprehensive survey. <u>IEEE Communications Surveys & Tutorials</u> 22, 3 (2020), 2031–2063.
- [28] Yang Liu, Yan Kang, Tianyuan Zou, Yanhong Pu, Yuanqin He, Xiaozhou Ye, Ye Ouyang, Ya-Qin Zhang, and Qiang Yang. 2022. Vertical Federated Learning. <u>arXiv preprint arXiv:2211.12814</u> (2022).
- [29] Hanlin Lu, Changchang Liu, Ting He, Shiqiang Wang, and Kevin S Chan. 2020. Sharing models or coresets: A study based on membership inference attack. arXiv preprint arXiv:2007.02977 (2020).
- [30] Ali Makhdoumi and Nadia Fawaz. 2013. Privacy-utility tradeoff under statistical uncertainty. In 2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 1627–1634.
- [31] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54), Aarti Singh and Jerry Zhu (Eds.). PMLR, 1273–1282. https://proceedings.mlr. press/v54/mcmahan17a.html
- [32] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In Artificial Intelligence and Statistics. PMLR, 1273–1282.
- [33] Milad Khademi Nori, Sangseok Yun, and Il-Min Kim. 2021. Fast federated learning by balancing communication trade-offs. <u>IEEE</u> Transactions on Communications 69, 8 (2021), 5168–5182.
- [34] Pascal Paillier. 1999. Public-key cryptosystems based on composite degree residuosity classes. In <u>International conference on the theory</u> and applications of cryptographic techniques. Springer, 223–238.
- [35] Borzoo Rassouli and Deniz Gündüz. 2019. Optimal utility-privacy trade-off with total variation distance as a privacy measure. <u>IEEE</u> Transactions on Information Forensics and Security 15 (2019), 594–603.
- [36] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. 2020. The future of digital health with federated learning. NPJ digital medicine 3, 1 (2020), 1–7.
- [37] Lalitha Sankar, S Raj Rajagopalan, and H Vincent Poor. 2013. Utility-privacy tradeoffs in databases: An information-theoretic approach. IEEE Transactions on Information Forensics and Security 8, 6 (2013), 838–852.
- [38] Adi Shamir. 1979. How to Share a Secret. Commun. ACM 22, 11 (nov 1979), 612–613. https://doi.org/10.1145/359168.359176
- [39] Ben Tan, Bo Liu, Vincent Zheng, and Qiang Yang. 2020. <u>A Federated Recommender System for Online Services</u>. Association for Computing Machinery, New York, NY, USA, 579–581. https://doi.org/10.1145/3383313.3411528
- [40] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. 2019. A hybrid approach to privacy-preserving federated learning. In Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security. 1–11.
- [41] Stacey Truex, Ling Liu, Ka-Ho Chow, Mehmet Emre Gursoy, and Wenqi Wei. 2020. LDP-Fed: Federated learning with local differential privacy. In Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and Networking. 61–66.
- [42] Hao Wang and Flavio P Calmon. 2017. An estimation-theoretic view of privacy. In <u>2017 55th Annual Allerton Conference on</u> Communication, Control, and Computing (Allerton). IEEE, 886–893.
- [43] Jianyu Wang and Gauri Joshi. 2019. Adaptive communication strategies to achieve the best error-runtime trade-off in local-update SGD. Proceedings of Machine Learning and Systems 1 (2019), 212–229.

- [44] Weina Wang, Lei Ying, and Junshan Zhang. 2016. On the relation between identifiability, differential privacy, and mutual-information privacy. IEEE Transactions on Information Theory 62, 9 (2016), 5018–5029.
- [45] Stefanie Warnat-Herresthal, Hartmut Schultze, Krishnaprasad Lingadahalli Shastry, Sathyanarayanan Manamohan, Saikat Mukherjee, Vishesh Garg, Ravi Sarveswara, Kristian Händler, Peter Pickkers, N Ahmad Aziz, et al. 2021. Swarm learning for decentralized and confidential clinical machine learning. Nature 594, 7862 (2021), 265–270.
- [46] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated machine learning: Concept and applications. <u>ACM Transactions</u> on Intelligent Systems and Technology (TIST) 10, 2 (2019), 1–19.
- [47] Hongxu Yin, Arun Mallya, Arash Vahdat, Jose M Alvarez, Jan Kautz, and Pavlo Molchanov. 2021. See through Gradients: Image Batch Recovery via GradInversion. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 16337–16346.
- [48] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. 2020. BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning. In 2020 USENIX Annual Technical Conference (USENIX ATC 20). USENIX Association, 493–506. https://www.usenix.org/conference/atc20/presentation/zhang-chengliang
- [49] Jiale Zhang, Bing Chen, Shui Yu, and Hai Deng. 2019. PEFL: A privacy-enhanced federated learning scheme for big data analytics. In 2019 IEEE Global Communications Conference (GLOBECOM). IEEE, 1–6.
- [50] Xiaojin Zhang, Hanlin Gu, Lixin Fan, Kai Chen, and Qiang Yang. 2022. No Free Lunch Theorem for Security and Utility in Federated Learning. ACM Trans. Intell. Syst. Technol. 14, 1, Article 1 (nov 2022), 35 pages. https://doi.org/10.1145/3563219
- [51] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. idlg: Improved deep leakage from gradients. arXiv preprint arXiv:2001.02610 (2020).
- [52] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep Leakage from Gradients. In <u>Advances in Neural Information Processing Systems</u>, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. https: //proceedings.neurips.cc/paper/2019/file/60a6c4002cc7b29142def8871531281a-Paper.pdf
- [53] T. Zou, Y. Liu, Y. Kang, W. Liu, Y. He, Z. Yi, Q. Yang, and Y. Zhang. 2022. Defending Batch-Level Label Inference and Replacement Attacks in Vertical Federated Learning. <u>IEEE Transactions on Big Data</u> (jul 2022), 1–12. https://doi.org/10.1109/TBDATA.2022.3192121

A OUTLINE OF OUR WORK



B LOWER BOUNDS FOR PRIVACY LEAKAGE, UTILITY LOSS AND EFFICIENCY REDUCTION

When the private information D is discrete, the privacy leakage could be lower bounded by the total variation distance between P_k^O and P_k^S , as is shown in the following lemma.

Lemma B.1. Let $\epsilon_{p,k}$ be defined in Definition 3.3. Let P_k^O and P_k^S represent the distribution of the parameter of client *k* before and after being protected. Then for any client *k*, we have that

$$\epsilon_{p,k} \ge \sqrt{\mathrm{JS}(F_k^O||F_k^{\mathcal{B}})} - \frac{1}{2}(e^{2\xi} - 1)\mathrm{TV}(P_k^O||P_k^S).$$
(13)

Furthermore, we have that

$$\epsilon_p \geq \frac{1}{K} \sum_{k=1}^K \sqrt{\mathrm{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}).$$

To provide a lower bound for privacy leakage when the private information *D* is discrete, we first show the following lemma.

Lemma B.2. Let P_k^O and P_k^S represent the distribution of the parameter of client *k* before and after being protected. Let $F_k^{\mathcal{A}}$ and F_k^O represent the belief of client *k* about *S* after observing the protected and original parameter. Then we have

$$JS(F_k^{\mathcal{A}}||F_k^{O}) \le \frac{1}{4}(e^{2\xi} - 1)^2 TV(P_k^{O}||P_k^{S})^2.$$

Trading Off Privacy, Utility and Efficiency in Federated Learning • 23

PROOF. Recall that $JS(F_k^{\mathcal{A}}||F_k^{\mathcal{B}}) = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d) + \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}^{\mathcal{B}}(d) \log \frac{f_{D_k}^{\mathcal{B}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d)$ $= \frac{1}{2} \int_{\mathcal{D}_k} \int_{W_k^{\mathcal{S}}} f_{D_k|W_k}(d|w) dP_k^{\mathcal{S}}(w) \log \frac{\int_{W_k^{\mathcal{S}}} f_{D_k|W_k}(d|w) dP_k^{\mathcal{S}}(w)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d) + \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d). \text{ Denote } F_k^{\mathcal{M}} = \frac{1}{2} \int_{\mathcal{D}_k} f_{D_k}(d) \log \frac{f_{D_k}(d)}{f_{D_k}^{\mathcal{M}}(d)} \mathbf{d}\mu(d).$ $\frac{1}{2}(F_k^{\mathcal{A}} + F_k^O)$, we have

$$\begin{aligned} \mathrm{JS}(F_k^{\mathcal{A}}||F_k^{\mathcal{O}}) &= \frac{1}{2} \left[KL\left(F_k^{\mathcal{A}}, F_k^{\mathcal{M}}\right) + KL\left(F_k^{\mathcal{O}}, F_k^{\mathcal{M}}\right) \right] \\ &= \frac{1}{2} \left[\sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} + \sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{O}}(d) \log \frac{f_{D_k}^{\mathcal{O}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \right] \\ &= \frac{1}{2} \left[\sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{A}}(d) \log \frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} - \sum_{d \in \mathcal{D}_k} f_{D_k}^{\mathcal{O}}(d) \log \frac{f_{D_k}^{\mathcal{M}}(d)}{f_{D_k}^{\mathcal{O}}(d)} \right] \\ &\leq \frac{1}{2} \sum_{d \in \mathcal{D}_k} \left| f_{D_k}^{\mathcal{A}}(d) - f_{D_k}^{\mathcal{O}}(d) \right| \left| \log \frac{f_{D_k}^{\mathcal{M}}(d)}{f_{D_k}^{\mathcal{O}}(d)} \right|, \end{aligned}$$

where the inequality is due to $\frac{f_{D_k}^{\mathcal{A}}(d)}{f_{D_k}^{\mathcal{M}}(d)} \leq \frac{f_{D_k}^{\mathcal{M}}(d)}{f_{D_k}^{\mathcal{O}}(d)}$. **Bounding** $\left| f_{D_k}^{\mathcal{A}}(d) - f_{D_k}^{\mathcal{O}}(d) \right|$. Let $\mathcal{U}_k = \{ w \in \mathcal{W}_k : dP_k^{\mathcal{S}}(w) - dP_k^{\mathcal{O}}(w) \geq 0 \}$, and $\mathcal{V}_k = \{ w \in \mathcal{W}_k : dP_k^{\mathcal{S}}(w) - dP_k^{\mathcal{O}}(w) \geq 0 \}$. Then we have

$$\begin{aligned} \left| f_{D_{k}}^{\mathcal{A}}(d) - f_{D_{k}}^{\mathcal{O}}(d) \right| &= \left| \int_{\mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) [dP_{k}^{\mathcal{S}}(w) - dP_{k}^{\mathcal{O}}(w)] \right| \\ &= \left| \int_{\mathcal{U}_{k}} f_{D_{k}|W_{k}}(d|w) [dP_{k}^{\mathcal{S}}(w) - dP_{k}^{\mathcal{O}}(w)] + \int_{\mathcal{V}_{k}} f_{D_{k}|W_{k}}(d|w) [dP_{k}^{\mathcal{S}}(w) - dP_{k}^{\mathcal{O}}(w)] \right| \\ &\leq \left(\sup_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) - \inf_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) \right) \int_{\mathcal{U}_{k}} [dP_{k}^{\mathcal{S}}(w) - dP_{k}^{\mathcal{O}}(w)]. \end{aligned}$$
(14)

Notice that

$$\sup_{w \in \mathcal{W}_k} f_{D_k | W_k}(d | w) - \inf_{w \in \mathcal{W}_k} f_{D_k | W_k}(d | w) = \inf_{w \in \mathcal{W}_k} f_{D_k | W_k}(d | w) \left| \frac{\sup_{w \in \mathcal{W}_k} f_{D_k | W_k}(d | w)}{\inf_{w \in \mathcal{W}_k} f_{D_k | W_k}(d | w)} - 1 \right|.$$

From the definition of ξ , we know that for any $w \in \mathcal{W}_k$,

$$e^{-\xi} \leq \frac{f_{D_k|W_k}(d|w)}{f_{D_k}(d)} \leq e^{\xi}$$

Therefore, for any pair of parameters $w, w' \in \mathcal{W}_k$, we have

$$\frac{f_{D_k|W_k}(d|w)}{f_{D_k|W_k}(d|w')} = \frac{f_{D_k|W_k}(d|w)}{f_{D_k}(d)} / \frac{f_{D_k|W_k}(d|w')}{f_{D_k}(d)} \le e^{2\xi}$$

Therefore, the first term of Eq. (14) is bounded by

$$\sup_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) - \inf_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) \le \inf_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w)(e^{2\xi} - 1).$$
(15)

From the definition of total variation distance, we have

$$\int_{\mathcal{U}_k} \left[dP_k^{\mathcal{S}}(w) - dP_k^{\mathcal{O}}(w) \right] = \mathrm{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}).$$
(16)

Combining Eq. (15) and Eq. (16), we have

$$|f_{D_{k}}^{\mathcal{A}}(d) - f_{D_{k}}^{O}(d)| = \left(\sup_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w) - \inf_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w)\right) \int_{\mathcal{U}_{k}} [dP_{k}^{S}(w) - dP_{k}^{O}(w)]$$

$$\leq \inf_{w \in \mathcal{W}_{k}} f_{D_{k}|W_{k}}(d|w)(e^{2\xi} - 1)\mathrm{TV}(P_{k}^{O}||P_{k}^{S}).$$
(17)

Bounding $\left| \log \left(\frac{f_{D_k}^{\mathcal{M}}(d)}{f_{D_k}^{\mathcal{O}}(d)} \right) \right|$. Since $\left| \log \left(\frac{a}{b} \right) \right| \le \frac{|a-b|}{\min\{a,b\}}$ for any two positive numbers *a* and *b*, we have that

$$\left|\log \frac{f_{D_{k}}^{\mathcal{M}}(d)}{f_{D_{k}}^{\mathcal{O}}(d)}\right| \leq \frac{|f_{D_{k}}^{\mathcal{M}}(d) - f_{D_{k}}^{\mathcal{O}}(d)|}{\min\{f_{D_{k}}^{\mathcal{M}}(d), f_{D_{k}}^{\mathcal{O}}(d)\}}$$

$$= \frac{|f_{D_{k}}^{\mathcal{M}}(d) - f_{D_{k}}^{\mathcal{O}}(d)|}{2\min\{f_{D_{k}}^{\mathcal{M}}(d), f_{D_{k}}^{\mathcal{O}}(d)\}}$$

$$\leq \frac{\inf_{w \in W_{k}} f_{D_{k}|W_{k}}(d|w)(e^{2\xi} - 1)\operatorname{TV}(P_{k}^{\mathcal{O}}||P_{k}^{S})}{2\min\{f_{D_{k}}^{\mathcal{M}}(d), f_{D_{k}}^{\mathcal{O}}(d)\}}$$

$$\leq \frac{1}{2}(e^{2\xi} - 1)\operatorname{TV}(P_{k}^{\mathcal{O}}||P_{k}^{S}), \qquad (18)$$

where the third inequality is due to $\min\{f_{D_k}^{\mathcal{M}}(d), f_{D_k}^{\mathcal{O}}(d)\} \ge \min\{f_{D_k}^{\mathcal{A}}(d), f_{D_k}^{\mathcal{O}}(d)\} \ge \inf_{w \in \mathcal{W}_k} f_{D_k|W_k}(d|w)$. Combining Eq. (17) and Eq. (18), we have

$$\begin{split} \mathsf{JS}(F_{k}^{\mathcal{A}}||F_{k}^{O}) &\leq \frac{1}{2} \left[\sum_{d \in \mathcal{D}_{k}} \left| (f_{D_{k}}^{\mathcal{A}}(d) - f_{D_{k}}^{O}(d)) \right| \left| \log \frac{f_{D_{k}}^{\mathcal{M}}(d)}{f_{D_{k}}^{O}(d)} \right| \right] \\ &\leq \frac{1}{4} (e^{2\xi} - 1)^{2} \mathrm{TV}(P_{k}^{O}||P_{k}^{S})^{2} \sum_{d \in \mathcal{D}_{k}} \inf_{w \in W_{k}^{O}} f_{D_{k}|W_{k}}(d|w) \\ &\leq \frac{1}{4} (e^{2\xi} - 1)^{2} \mathrm{TV}(P_{k}^{O}||P_{k}^{S})^{2}. \end{split}$$

With the above lemma, now we are ready to show Lemma B.1 when the private information D is discrete, the analysis is similar to the continuous scenario.

PROOF. The square root of JS divergence satisfies triangle inequality, which implies that

$$\sqrt{\mathrm{JS}(F_k^O||F_k^{\mathcal{B}})} - \sqrt{\mathrm{JS}(F_k^{\mathcal{A}}||F_k^{\mathcal{B}})} \le \sqrt{\mathrm{JS}(F_k^{\mathcal{A}}||F_k^O)},$$

where $\sqrt{\mathrm{JS}(F_k^O||F_k^{\mathcal{B}})}$ is a problem-dependent constant, $\sqrt{\mathrm{JS}(F_k^{\mathcal{A}}||F_k^{\mathcal{B}})}$ represents the privacy leakage, and $\sqrt{\mathrm{JS}(F_k^{\mathcal{A}}||F_k^O)} \leq \frac{1}{2}(e^{2\xi}-1)\mathrm{TV}(P_k^O||P_k^S)$ from Lemma B.2. Therefore, $\forall k \in [K]$, we have that

$$\sqrt{\mathrm{JS}(F_k^O||F_k^{\mathcal{B}})} \le \epsilon_{p,k} + \frac{1}{2}(e^{2\xi} - 1)\mathrm{TV}(P_k^O||P_k^{\mathcal{S}}).$$

Lemma B.3. Let Assumption 4.2 hold, and ϵ_e be defined in Definition 3.4. Let P_k^O and P_k^S represent the distribution of the federated model information before and after being protected. Let $\Xi = \min_{k \in [K]} \Xi_k$, and $\Gamma = \min_{k \in [K]} \Gamma_k$. Then we have,

$$\epsilon_e \ge \Xi \cdot \Gamma \cdot \frac{1}{K} \sum_{k=1}^{K} \text{TV}(P_k^O || P_k^S).$$
(19)

PROOF. Let $\mathcal{U}_k = \{w \in \mathcal{W}_k : dP_k^{\mathcal{S}}(w) - dP_k^{\mathcal{O}}(w) \ge 0\}$, and $\mathcal{V}_k = \{w \in \mathcal{W}_k : dP_k^{\mathcal{S}}(w) - dP_k^{\mathcal{O}}(w) < 0\}$, where \mathcal{W}_k represents the union of the supports of $P_k^{\mathcal{S}}$ and $P_k^{\mathcal{O}}$. For any $w \in \mathcal{V}_k$, the definition of \mathcal{V}_k implies that $dP_k^{\mathcal{O}}(w) > dP_k^{\mathcal{S}}(w) \ge 0$. Therefore, w belongs to the support of $P_k^{\mathcal{O}}$, which is denoted as $\mathcal{W}_k^{\mathcal{O}}$. Therefore we have that

$$\mathcal{V}_k \subset \mathcal{W}_k^O. \tag{20}$$

Similarly, we have that

$$\mathcal{U}_k \subset \mathcal{W}_k^{\mathcal{S}}.\tag{21}$$

Recall that P_k^S represents the distribution of the aggregated parameter after being protected, and $p_{W_k}^S(w)$ represents the corresponding probability density function.

 W_{k,Ξ_k}^+ is defined as:

$$\mathcal{W}_{k,\Xi_k}^+ = \left\{ w \in \mathcal{U}_k : C(w_{\max}^O) + 2 \cdot \Xi_k \le C(w) \text{ and } p_{W_k}^S(w) \ge 2 \cdot p_{W_k}^O(w) \right\},\$$

where $w_{\max}^{O} = \arg \max_{w \in W_{k}^{O}} C(w)$. We denote $W_{k,\Xi_{k}}^{-} = \mathcal{U}_{k} \setminus W_{k,\Xi_{k}}^{+}$. From Assumption 4.2, Γ_{k} represents the maximum constant satisfying that

$$\int_{\mathcal{U}_k} (\mathbb{1}\{w \in \mathcal{W}_{k,\Xi_k}^+\} - \mathbb{1}\{w \in \mathcal{W}_{k,\Xi_k}^-\}) p_{W_k}^S(w) dw \ge \Gamma_k \cdot \mathrm{TV}(P_k^O || P_k^S),$$
(22)

where $p_{W_k}^{S}$ denotes the probability density function of the protected model information and $p_{W_k}^{O}$ denotes the probability density function of the original model information. Then we have

$$\begin{split} &\varepsilon_{e,k} = \left[\mathbb{E}_{w \sim P_{k}^{S}} [C(w)] - \mathbb{E}_{w \sim P_{k}^{O}} [C(w)] \right] \\ &= \left[\int_{\mathcal{W}_{k}} C(w) dP_{k}^{S}(w) - \int_{\mathcal{W}_{k}} C(w) dP_{k}^{O}(w) \right] \\ &= \left[\int_{\mathcal{U}_{k}} C(w) [dP_{k}^{S}(w) - dP_{k}^{O}(w)] - \int_{\mathcal{V}_{k}} C(w) [dP_{k}^{O}(w) - dP_{k}^{S}(w)] \right] \\ &= \int_{\mathcal{U}_{k}} C(w) \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{+} \} [dP_{k}^{S}(w) - dP_{k}^{O}(w)] + \int_{\mathcal{U}_{k}} C(w) \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{-} \} [dP_{k}^{S}(w) - dP_{k}^{O}(w)] \\ &- \int_{\mathcal{V}_{k}} C(w) \mathbb{1} \{ w \in \mathcal{W}_{k}^{+} \} [dP_{k}^{O}(w) - dP_{k}^{S}(w)] \\ &\geq 2\Xi_{k} \cdot \int_{\mathcal{U}_{k}} \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{+} \} (p_{W_{k}}^{S}(w) - p_{W_{k}}^{O}(w)) dw - \Xi_{k} \cdot \int_{\mathcal{U}_{k}} \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{-} \} (p_{W_{k}}^{S}(w) - p_{W_{k}}^{O}(w)) dw \\ &\geq 2\Xi_{k} \cdot \int_{\mathcal{U}_{k}} \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{+} \} \frac{p_{W_{k}}^{S}(w)}{2} dw - \Xi_{k} \cdot \int_{\mathcal{U}_{k}} \mathbb{1} \{ w \in \mathcal{W}_{k,\Xi_{k}}^{-} \} p_{W_{k}}^{S}(w) dw \\ &\geq \Xi_{k} \cdot \Gamma_{k} \cdot \mathrm{TV}(P_{k}^{O} || P_{k}^{S}), \end{split}$$

where the first inequality is due to $C(w_{\max}^{O}) + 2\Xi_k \leq C(w), \forall w \in \mathcal{W}_{k,\Xi_k}^+$, and $C(w_{\max}^{O}) - \Xi_k \leq C(w), \forall w \in \mathcal{U}_k$, the second inequality is due to the $p_{W_k}^{\mathcal{S}}(w) - p_{W_k}^{O}(w) \ge \frac{p_{W_k}^{\mathcal{S}}(w)}{2}, \forall w \in \mathcal{W}_{k,\Xi_k}^+$, and the third inequality is due to $\int_{\mathcal{U}_k} (\mathbbm{1}\{w \in \mathcal{W}_{k,\Xi_k}^+\} - \mathbbm{1}\{w \in \mathcal{W}_{k,\Xi_k}^-\}) p_{W_k}^{\mathcal{S}}(w) dw \ge \Gamma_k \cdot \mathrm{TV}(P_k^{O}||P_k^{\mathcal{S}})$ from Assumption 4.2. Therefore, we have that

$$\epsilon_{e} = \frac{1}{K} \sum_{k=1}^{K} \epsilon_{e,k} \ge \frac{1}{K} \sum_{k=1}^{K} \Xi_{k} \cdot \Gamma_{k} \cdot \operatorname{TV}(P_{k}^{O} || P_{k}^{S})$$
$$\ge \Xi \cdot \Gamma \cdot \frac{1}{K} \sum_{k=1}^{K} \operatorname{TV}(P_{k}^{O} || P_{k}^{S}).$$

TRADE-OFF BETWEEN PRIVACY, UTILITY AND EFFICIENCY С

First, we quantify the trade-off between privacy and efficiency. The following lemma illustrates that the summation of the privacy leakage and the efficiency reduction is lower bounded by a problem-dependent constant.

Lemma C.1 (No free lunch theorem (NFL) for privacy and efficiency). Let ϵ_p be defined in Def. 3.3, and let ϵ_e be defined in Def. 3.5, with Assumption 4.2 we have:

$$\epsilon_p + C_x \cdot \epsilon_e \ge C_1,$$

where $C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^{\mathcal{O}} || F_k^{\mathcal{B}})}$ and $C_x = \frac{1}{2\Xi\Gamma} (e^{2\xi} - 1)$.

PROOF. First, from Lemma 4.2 we have that

$$\epsilon_p \geq \frac{1}{K} \sum_{k=1}^{K} \sqrt{\mathrm{JS}(F_k^O||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^{K} \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_k^O||P_k^{\mathcal{S}}),$$

From Lemma 4.1, we have

$$\epsilon_{e} \geq \Xi \cdot \Gamma \cdot \frac{1}{K} \sum_{k=1}^{K} \text{TV}(P_{k}^{O} || P_{k}^{S})$$

Combining the above two equations, we have that

$$\frac{1}{K}\sum_{k=1}^{K}\sqrt{\mathrm{JS}(F_{k}^{O}||F_{k}^{\mathcal{B}})} \leq \epsilon_{p} + \frac{1}{2\Xi\Gamma}(e^{2\xi}-1)\epsilon_{e},$$

where $\epsilon_p = \frac{1}{K} \sum_{k=1}^{K} \epsilon_{p,k}$. The above equation could be further simplified as

$$C_1 \leq \epsilon_p + C_x \epsilon_e$$
,

where $C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^O || F_k^{\mathcal{B}})}$ and $C_x = \frac{1}{2 \Xi \Gamma} (e^{2\xi} - 1)$.

The trade-off between privacy and utility was the main result shown in [50], the analysis of which is applicable to a more general definition for utility loss (Definition 3.4). The following lemma illustrates that utility loss is lower bounded by the total variation distance between the unprotected and protected distributions.

Lemma C.2. [50] Let Assumption 4.1 hold, and ϵ_u be defined in Definition 3.4. Let P_{fed}^O and P_{fed}^S represent the distribution of the aggregated parameter before and after being protected. Then we have,

$$\epsilon_u \geq \frac{\Delta}{2} \cdot \mathrm{TV}(P_{\mathrm{fed}}^O || P_{\mathrm{fed}}^S).$$

With Lemma B.1 and Lemma C.2, it is now natural to provide a quantitative relationship between the utility loss and the privacy leakage (Theorem C.3).

Theorem C.3. [50] No free lunch theorem (NFL) for privacy and utility] Let ϵ_p be defined in Def. 3.3, and let ϵ_u be defined in Def. 3.4 at the convergence step, with Assumption 4.1 we have:

$$\epsilon_p + C_d \cdot \epsilon_u \ge C_1,\tag{23}$$

where $\xi_k = \max_{w \in W_k, d \in D_k} \left| \log \left(\frac{f_{D_k | W_k(d | w)}}{f_{D_k(d)}} \right) \right|$, and $\xi = \max_{k \in [K]} \xi_k$ represents the maximum privacy leakage over all possible information *w* released by client *k*, and $[K] = \{1, 2, \dots, K\}, C_1 = \frac{1}{K} \sum_{k=1}^{K} \sqrt{JS(F_k^O || F_k^{\mathcal{B}})}, C_d = \frac{\gamma}{4\Lambda} (e^{2\xi} - 1), C_k = \frac{1}{4\Lambda} (e^{2\xi} - 1)$ $\frac{1}{K}\sum_{k=1}^{K} \mathrm{TV}(P_{k}^{O}||P_{k}^{S})$

and
$$\gamma = \frac{K - 2_{K} = 1}{\text{TV}(P_{\text{fed}}^{O} || P_{\text{fed}}^{S})}$$
.

With Theorem C.1 and Theorem C.3, our main result Theorem 4.6 is proven.

ANALYSIS FOR RANDOMIZATION MECHANISM D

Randomization mechanism adds random noise such as Gaussian noise to model gradients [1, 17, 41].

- Let W_k^O be the model information sampled from distribution $P_k^O = \mathcal{N}(\mu_0, \Sigma_0)$, where $\mu_0 \in \mathbb{R}^n, \Sigma_0 =$
- Let W_k² be the model information charges and the set of the protected model information W_k^S = W_k^O + ε_k, where ε_k ~ N(0, Σ_ε) and Σ_ε = diag(σ_ε², · · · , σ_ε²). Therefore, W_k^S follows the distribution P_k^S = N(μ₀, Σ₀ + Σ_ε).
 The protected model information W_{fed}^S = ¹/_K Σ_{k=1}^K (W_k^O + ε_k) follows distribution P_{fed}^S = N(μ₀, Σ₀/K + Σ_ε/K).

The following lemmas establish bounds for the privacy leakage and utility loss, and efficiency reduction using the variance of the noise σ_{ϵ}^2 .

ACM Trans. Intell. Syst. Technol.

Lemma D.1. For randomization mechanism, the privacy leakage is bounded by

$$\epsilon_p \ge C_1 - \frac{3C_2}{2} \cdot \min\left\{1, \sigma_\epsilon^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\},\$$

and the efficiency reduction is

$$\epsilon_e \geq \frac{\Xi\Gamma}{100} \min\left\{1, \sigma_e^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

PROOF. From Lemma 4.2 we have that

$$\epsilon_p \geq \frac{1}{K} \sum_{k=1}^K \sqrt{\mathrm{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}).$$
we have that

From Lemma C.2 of [50], we have that

$$\frac{1}{100}\min\left\{1,\sigma_{\epsilon}^{2}\sqrt{\sum_{i=1}^{m}\frac{1}{\sigma_{i}^{4}}}\right\} \leq \operatorname{TV}(P_{k}^{O}||P_{k}^{S}) \leq \frac{3}{2}\min\left\{1,\sigma_{\epsilon}^{2}\sqrt{\sum_{i=1}^{m}\frac{1}{\sigma_{i}^{4}}}\right\}.$$
(24)

Then we have that

$$\epsilon_p \ge C_1 - \frac{3C_2}{2} \cdot \min\left\{1, \sigma_\epsilon^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}$$

From Lemma 4.1 and Eq. (24), we have that

$$\epsilon_{e} \geq \frac{1}{K} \sum_{k=1}^{K} \Xi \cdot \Gamma \cdot \text{TV}(P_{k}^{O} || P_{k}^{S})$$
$$\geq \frac{\Xi\Gamma}{100} \min\left\{1, \sigma_{e}^{2} \sqrt{\sum_{i=1}^{m} \frac{1}{\sigma_{i}^{4}}}\right\}$$

Lemma D.2. For randomization mechanism, the utility loss is bounded by

$$\epsilon_u \geq \frac{\Delta}{200} \cdot \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

PROOF. From Lemma C.2 of [50], we have that

$$\mathrm{TV}(P_{\mathrm{fed}}^{O}||P_{\mathrm{fed}}^{S}) \ge \frac{1}{100} \min\left\{1, \sigma_{\epsilon}^{2} \sqrt{\sum_{i=1}^{m} \frac{1}{\sigma_{i}^{4}}}\right\}.$$
(25)

From Lemma 4.4, we have that

$$\epsilon_u \ge \frac{\Delta}{2} \cdot \text{TV}(P_{\text{fed}}^O || P_{\text{fed}}^S).$$
 (26)

Combining the Eq. (25) and Eq. (26), we have

$$\epsilon_u \geq \frac{\Delta}{200} \cdot \min\left\{1, \sigma_{\epsilon}^2 \sqrt{\sum_{i=1}^m \frac{1}{\sigma_i^4}}\right\}.$$

E ANALYSIS FOR PAILLIER HOMOMORPHIC ENCRYPTION

The **Paillier** encryption mechanism was proposed by [34] is an asymmetric additive homomorphic encryption mechanism, which was widely applied in FL [3, 8, 40, 49]. We first introduce the basic definition of Paillier algorithm in federated learning [13]. Paillier encyption contains three parts including key generation, encryption and decryption. Let *h* represent the plaintext, and *c* represent the ciphertext.

Key Generation. Let (n, q) represent the public key, and (λ, μ) represent the private key. Select two primes p and *q* that are rather large, satisfying that gcd(pq, (p-q)(q-1)) = 1. Select *g* randomly satisfying that $g \in \mathbb{Z}_{n^2}^*$. Let $n = p \cdot q, \lambda = \text{lcm}(p - 1, q - 1), \text{ and } \mu = (L(q^{\lambda} \mod n^2))^{-1} \mod n.$

Encryption. Randomly select *r* and encode *h* as:

$$c = q^h \cdot r^n \mod n^2$$
,

where $n = p \cdot q$, p and q are two selected primes. Note that g is an integer selected randomly, and $g \in \mathbb{Z}_{n^2}^*$. Therefore, *n* can divide the order of *q*.

Decryption. Using the private key (λ, μ) to decrypt the ciphertext *c* as:

 $h = L(c^{\lambda} \bmod n^2) \cdot \mu \bmod n,$

where $L(x) = \frac{x-1}{n}$, $\mu = (L(g^{\lambda} \mod n^2))^{-1} \mod n$, and $\lambda = \text{lcm}(p-1, q-1)$.

Let *m* represent the number of dimension of the parameter.

- Let W^O_k represent the plaintext that follows a uniform distribution over [a¹_k − δ, a¹_k + δ] × [a²_k − δ, a²_k + δ] × ···× [a^m_k − δ, a^m_k + δ].
 Assume that the ciphertext W^S_k follows a uniform distribution over [0, n² − 1]^m.
 Let W^O_{fed} represent the federated plaintext that follows a uniform distribution over [a¹ − δ, a¹ + δ] × [a²_k − δ, a²_k + δ].
- $\delta, \bar{a}^2 + \delta] \cdots \times [\bar{a}^m \delta, \bar{a}^m + \delta], \text{ where } \bar{a}_i = \sum_{k=1}^K a_k^i.$ The federated ciphertext W_{fed}^S follows a uniform distribution over $[0, n^2 1]^m$.

Intuitively, longer ciphertext should theoretically increase efficiency and decrease privacy leakage. The following lemma provides lower bounds for privacy leakage and efficiency reduction for Paillier mechanism.

Lemma E.1. For Paillier mechanism, the privacy leakage is bounded by

$$\epsilon_p \ge C_1 - C_2 \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right],$$
(27)

and the efficiency reduction is bounded by

$$\epsilon_e \geq \Xi \cdot \Gamma \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m\right].$$

PROOF. Let W_k^O represent the plaintext *m*, and W_k^S represent the ciphertext *c*. Recall for encryption, we have that

$$c = q^m \cdot r^n \mod n^2$$

The ciphertext W_k^S follows a uniform distribution over $[0, n^2 - 1]^m$, and the plaintext W_k^O follows a uniform distribution over $[a_k^1 - \delta, a_k^1 + \delta] \times [a_k^2 - \delta, a_k^2 + \delta] \cdots \times [a_k^m - \delta, a_k^m + \delta]$, and $a_k^i \in [0, n^2 - 1], \forall i = 1, 2, \cdots, m$. Then we have that

$$\begin{aligned} \operatorname{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}) &= \int_{[a_k^1 - \delta, a_k^1 + \delta]} \int_{[a_k^2 - \delta, a_k^2 + \delta]} \cdots \int_{[a_k^m - \delta, a_k^m + \delta]} \left(\left(\frac{1}{2\delta}\right)^m - \left(\frac{1}{n^2}\right)^m \right) dw_1 dw_2 \cdots dw_m \\ &= \left[\left(\frac{1}{2\delta}\right)^m - \left(\frac{1}{n^2}\right)^m \right] \cdot (2\delta)^m. \end{aligned}$$

emma 4.2 we have that
$$\epsilon_p \geq \frac{1}{K} \sum_{k=1}^K \sqrt{\operatorname{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \operatorname{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}). \end{aligned}$$

hing the above two equations, we have that

From Lemma 4.2 we have that

$$\epsilon_{p} \geq \frac{1}{K} \sum_{k=1}^{K} \sqrt{\mathrm{JS}(F_{k}^{O} || F_{k}^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^{K} \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_{k}^{O} || P_{k}^{S}).$$

Combining the above two equations, we have that

$$\begin{split} \epsilon_p &\geq \frac{1}{K} \sum_{k=1}^K \sqrt{\mathrm{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}) \\ &= \sqrt{\mathrm{JS}(F_k^{\mathcal{B}}||F_k^{\mathcal{O}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \left[\left(\frac{1}{2\delta}\right)^m - \left(\frac{1}{n^2}\right)^m \right] \cdot (2\delta)^m \\ &= C_1 - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \left[\left(\frac{1}{2\delta}\right)^m - \left(\frac{1}{n^2}\right)^m \right] \cdot (2\delta)^m \\ &= C_1 - C_2 \cdot \left[1 - \left(\frac{2\delta}{n^2}\right)^m \right]. \end{split}$$

From Lemma 4.1, we have that

$$\epsilon_{e} \geq \frac{1}{K} \sum_{k=1}^{K} \Xi \cdot \Gamma \cdot \mathrm{TV}(P_{k}^{O} || P_{k}^{S})$$
$$\geq \Xi \cdot \Gamma \cdot \left[1 - \left(\frac{2\delta}{n^{2}}\right)^{m} \right].$$

For Paillier mechanism, the distorted parameter given secret key becomes the original parameter. The following lemma shows that the utility loss for Paillier mechanism is 0.

Lemma E.2. For Paillier mechanism, the utility loss $\epsilon_u = 0$.

PROOF. Let P_{fed}^{S} represent the distribution of the distorted parameter which is decrypted by the client. Note that $\text{TV}(P_{\text{fed}}^{O}||P_{\text{fed}}^{S}) = 0$. From Lemma C.3 of [50], the utility loss is equal to 0.

ANALYSIS FOR SECRET SHARING MECHANISM F

Many MPC-based protocols (particularly secret sharing) are used to build secure machine learning models, such as linear regression, logistic regression, recommend systems, and so on. [5, 6, 38] were developed to distribute a secret among a group of participants.

Let *m* represent the number of dimension of the model information.

- Let W^O_k represent the original model information that follows a uniform distribution over [a¹_k − δ, a¹_k + δ] × [a²_k − δ, a²_k + δ] ··· × [a^m_k − δ, a^m_k + δ].
 Let W^S_k represent the distorted model information that follows a uniform distribution over [a¹_k − b¹_k, a¹_k + r¹_k] × [a²_k − b²_k, a²_k + r²_k] ··· × [a^m_k − b^m_k, a^m_k + r^m_k].

Lemma F.1. For secret sharing mechanism, the lower bound for privacy leakage is

$$\epsilon_p \ge C_1 - C_2 \cdot \frac{1}{K} \sum_{k=1}^K \left(1 - \prod_{j=1}^m \frac{2\delta}{b_k^j + r_k^j} \right)$$

PROOF. Notice that W_k^S follows a uniform distribution over $[a_k^1 - b_k^1, a_k^1 + r_k^1] \times [a_k^2 - b_k^2, a_k^2 + r_k^2] \cdots \times [a_k^m - b_k^m, a_k^m + r_k^m]$, and W_k^O follows a uniform distribution over $[a_k^1 - \delta, a_k^1 + \delta] \times [a_k^2 - \delta, a_k^2 + \delta] \cdots \times [a_k^m - \delta, a_k^m + \delta]$, and $\delta < b_k^m, r_k^m, \forall i = 1, 2, \cdots, m$. Then we have that

$$TV(P_k^O||P_k^S) = \int_{[a_k^1 - \delta, a_k^1 + \delta]} \int_{[a_k^2 - \delta, a_k^2 + \delta]} \cdots \int_{[a_k^m - \delta, a_k^m + \delta]} \left(\left(\frac{1}{2\delta}\right)^m - \prod_{j=1}^m \left(\frac{1}{b_k^j + r_k^j}\right) \right) dw_1 dw_2 \cdots dw_m$$
$$= \left(\left(\frac{1}{2\delta}\right)^m - \prod_{j=1}^m \left(\frac{1}{b_k^j + r_k^j}\right) \right) \cdot (2\delta)^m.$$

Therefore, we have that

$$\begin{aligned} \epsilon_p &\geq \frac{1}{K} \sum_{k=1}^K \sqrt{\mathsf{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \mathsf{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}) \\ &= \frac{1}{K} \sum_{k=1}^K \sqrt{\mathsf{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \left(\left(\frac{1}{2\delta}\right)^m - \prod_{j=1}^m \left(\frac{1}{b_k^j + r_k^j}\right) \right) \cdot (2\delta)^m. \end{aligned}$$

where the first inequality is due to Eq. (10) in Lemma 4.2. Then we have that

$$\epsilon_p \ge C_1 - C_2 \cdot \frac{1}{K} \sum_{k=1}^K \left(1 - \prod_{j=1}^m \frac{2\delta}{b_k^j + r_k^j} \right)$$

Lemma F.2. For secret sharing mechanism, the utility loss $\epsilon_u = 0$.

PROOF. For secret sharing mechanism, the federated model information does not change after being protected, which implies that $P_{\text{fed}}^{O} = P_{\text{fed}}^{S}$. Therefore, we have that

$$\begin{split} \epsilon_{u} &= \frac{1}{K} \sum_{k=1}^{K} \epsilon_{u,k} \\ &= \frac{1}{K} \sum_{k=1}^{K} [U_{k}(P_{\text{fed}}^{O}) - U_{k}(P_{\text{fed}}^{S})] \\ &= \frac{1}{K} \sum_{k=1}^{K} \left[\mathbb{E}_{w \sim P_{\text{fed}}^{O}} [U_{k}(w)] - \mathbb{E}_{w \sim P_{\text{fed}}^{S}} [U_{k}(w)] \right] \\ &= 0. \end{split}$$

The communication cost for the model information of secret sharing mechanism is not guaranteed to satisfy Assumption 4.2, and the analysis of the lower bound of the efficiency reduction is beyond the scope of our article.

G ANALYSIS FOR COMPRESSION MECHANISM

To facilitate the analysis, we simplify the compression mechanism as follows. Let b_i be a random variable sampled from Bernoulli distribution. The probability that b_i is equal to 1 is ρ_i .

$$W_k^{\mathcal{S}}(i) = \begin{cases} W_k^{\mathcal{O}}(i) & \text{if } b_i = 1, \\ 0, & \text{if } b_i = 0. \end{cases}$$
(28)

Let *m* represent the number of dimension of the model information.

- Let W_k^O represent the original model information that follows a uniform distribution over $[a_k^1 \delta, a_k^1 + \delta] \times [a_k^2 \delta, a_k^2 + \delta] \cdots \times [a_k^m \delta, a_k^m + \delta].$
- Each dimension *i* of the distorted model information $W_k^{\mathcal{S}}(i)$ takes the value identical with that of $W_k^{\mathcal{O}}$ with
- Probability ρ_i , and 0 with probability $1 \rho_i$. Let W_{fed}^O represent the federated plaintext that follows a uniform distribution over $[\tilde{a}^1 \delta, \tilde{a}^1 + \delta] \times [\tilde{a}^2 \delta, \tilde{a}^2 + \delta] \cdots \times [\tilde{a}^m \delta, \tilde{a}^m + \delta]$, where $\tilde{a}^i = \frac{1}{K} \sum_{k=1}^{K} a_k^i$. Each dimension *i* of the distorted model information $W_{\text{fed}}^S(i)$ takes the value identical with that of W_{fed}^O
- with probability ρ_i , and 0 with probability $1 \rho_i$.

Lemma G.1. For compression mechanism, the privacy leakage is lower bounded by

$$\epsilon_p \ge C_1 - C_2 \cdot \left(1 - \prod_{i=1}^m \rho_i\right).$$

The efficiency reduction is lower bounded by

$$\epsilon_e \geq \Xi \cdot \Gamma \cdot \left(1 - \prod_{i=1}^m \rho_i\right).$$

PROOF. The original model information W_k^O follows a uniform distribution over $[a_k^1 - \delta, a_k^1 + \delta] \times [a_k^2 - \delta, a_k^2 + \delta] \cdots \times [a_k^m - \delta, a_k^m + \delta]$, where *m* represents the dimension of W_k^O . Besides,

Trading Off Privacy, Utility and Efficiency in Federated Learning • 33

$$W_k^{\mathcal{S}}(i) = \begin{cases} W_k^{\mathcal{O}}(i) & \text{with probability } \rho_i, \\ 0, & \text{with probability } 1 - \rho_i. \end{cases}$$
(29)

Then we have that

$$\operatorname{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}) = \int_{[a_k^1 - \delta, a_k^1 + \delta]} \int_{[a_k^2 - \delta, a_k^2 + \delta]} \cdots \int_{[a_k^m - \delta, a_k^m + \delta]} \left(\left(\frac{1}{2\delta}\right)^m - \prod_{i=1}^m \left(\frac{\rho_i}{2\delta}\right) \right) dw_1 dw_2 \cdots dw_m \\ = \left(\left(\frac{1}{2\delta}\right)^m - \prod_{i=1}^m \left(\frac{\rho_i}{2\delta}\right) \right) \cdot (2\delta)^m.$$

From Lemma 4.2 we have that

that

$$\epsilon_{p} \geq \frac{1}{K} \sum_{k=1}^{K} \sqrt{\mathrm{JS}(F_{k}^{O}||F_{k}^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^{K} \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_{k}^{O}||P_{k}^{\mathcal{S}}).$$
be equations, we have that

Combining the above two equations, we have that

$$\begin{split} \epsilon_p &\geq \frac{1}{K} \sum_{k=1}^K \sqrt{\mathrm{JS}(F_k^{\mathcal{O}}||F_k^{\mathcal{B}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \mathrm{TV}(P_k^{\mathcal{O}}||P_k^{\mathcal{S}}) \\ &= \sqrt{\mathrm{JS}(F_k^{\mathcal{B}}||F_k^{\mathcal{O}})} - \frac{1}{K} \sum_{k=1}^K \frac{1}{2} (e^{2\xi} - 1) \cdot \left(\left(\frac{1}{2\delta} \right)^m - \prod_{i=1}^m \left(\frac{\rho_i}{2\delta} \right) \right) \cdot (2\delta)^m \\ &= C_1 - C_2 \cdot \left(1 - \prod_{i=1}^m \rho_i \right). \end{split}$$

From Lemma 4.1, we have that

$$\begin{aligned} \epsilon_{e} &\geq \frac{1}{K} \sum_{k=1}^{K} \Xi \cdot \Gamma \cdot \mathrm{TV}(P_{k}^{O} || P_{k}^{S}) \\ &\geq \Xi \cdot \Gamma \cdot \left(\left(\frac{1}{2\delta} \right)^{m} - \prod_{i=1}^{m} \left(\frac{\rho_{i}}{2\delta} \right) \right) \cdot (2\delta)^{m} \\ &= \Xi \cdot \Gamma \cdot \left(1 - \prod_{i=1}^{m} \rho_{i} \right). \end{aligned}$$

Lemma G.2. For compression mechanism, the utility loss is bounded by

$$\epsilon_u \geq \frac{\Delta}{2} \cdot \left(1 - \prod_{i=1}^m \rho_i\right).$$

PROOF. Recall that W_{fed}^O follows a uniform distribution over $[\tilde{a}^1 - \delta, \tilde{a}^1 + \delta] \times [\tilde{a}^2 - \delta, \tilde{a}^2 + \delta] \cdots \times [\tilde{a}^m - \delta, \tilde{a}^m + \delta]$, where $\tilde{a}^i = \frac{1}{K} \sum_{k=1}^K a_k^i$. Besides,

$$W_{\text{fed}}^{\mathcal{S}}(i) = \begin{cases} W_{\text{fed}}^{\mathcal{O}}(i) & \text{with probability } \rho_i, \\ 0, & \text{with probability } 1 - \rho_i. \end{cases}$$
(30)

Then we have that

$$\operatorname{TV}(P_{\text{fed}}^{O}||P_{\text{fed}}^{S}) = \int_{[\widetilde{a}^{1}-\delta,\widetilde{a}^{1}+\delta]} \int_{[\widetilde{a}^{2}-\delta,\widetilde{a}^{2}+\delta]} \cdots \int_{[\widetilde{a}^{m}-\delta,\widetilde{a}^{m}+\delta]} \left(\left(\frac{1}{2\delta}\right)^{m} - \prod_{i=1}^{m} \left(\frac{\rho_{i}}{2\delta}\right) \right) dw_{1}dw_{2} \cdots dw_{m}$$
$$= \left(\left(\frac{1}{2\delta}\right)^{m} - \prod_{i=1}^{m} \left(\frac{\rho_{i}}{2\delta}\right) \right) \cdot (2\delta)^{m}.$$

From Lemma 4.4, we have

$$\epsilon_{u} \geq \frac{\Delta}{2} \cdot \operatorname{TV}(P_{\text{fed}}^{O} || P_{\text{fed}}^{S}).$$

$$\frac{\Delta}{2} \cdot \left(\left(\frac{1}{2\delta} \right)^{m} - \prod_{i=1}^{m} \left(\frac{\rho_{i}}{2\delta} \right) \right) \cdot (2\delta)^{m}$$

$$\frac{\Delta}{2} \cdot \left(1 - \prod_{i=1}^{m} \rho_{i} \right).$$

Therefore,

$$\epsilon_{u} \geq \frac{\Delta}{2} \cdot \left(\left(\frac{1}{2\delta} \right)^{m} - \prod_{i=1}^{m} \left(\frac{\rho_{i}}{2\delta} \right) \right) \cdot (2\delta)^{m}$$
$$= \frac{\Delta}{2} \cdot \left(1 - \prod_{i=1}^{m} \rho_{i} \right).$$

(31)