



PriPrune: Quantifying and Preserving Privacy in Pruned Federated Learning

TIANYUE CHU*, IMDEA Networks Institute, Madrid (Leganes), Spain and Universidad Carlos III de Madrid, Madrid (Leganes), Spain

MENGWEI YANG*, University of California Irvine, Irvine, United States

NIKOLAOS LAOUTARIS, IMDEA Networks Institute, Madrid (Leganes), Spain

ATHINA MARKOPOULOU, University of California Irvine, Irvine, United States

Model pruning has been proposed as a technique for reducing the size and complexity of Federated learning (FL) models. By making local models coarser, pruning is intuitively expected to improve protection against privacy attacks. However, the level of this expected privacy protection has not been previously characterized, or optimized jointly with utility.

In this paper, we first characterize the privacy offered by pruning. We establish information-theoretic upper bounds on the information leakage from pruned FL and we experimentally validate them under state-of-the-art privacy attacks across different FL pruning schemes. Second, we introduce *PriPrune*—a privacy-aware algorithm for pruning in FL. *PriPrune* uses defense pruning masks, which can be applied locally after *any* pruning algorithm, and adapts the defense pruning rate to jointly optimize privacy and accuracy. Another key idea in the design of *PriPrune* is *Pseudo-Pruning*: it undergoes defense pruning within the local model and only sends the pruned model to the server; while the weights pruned out by defense mask are withheld locally for future local training rather than being removed. We show that *PriPrune* significantly improves the privacy-accuracy tradeoff compared to state-of-the-art pruned FL schemes. For example, on the FEMNIST dataset, *PriPrune* improves the privacy of *PruneFL* by 45.5% without reducing accuracy.

CCS Concepts: • Security and privacy → Distributed systems security; • Computing methodologies → Machine learning; • Theory of computation → Theory and algorithms for application domains.

Additional Key Words and Phrases: Federated Learning, Privacy, Model Pruning

1 Introduction

Federated Learning (FL) has emerged as the predominant paradigm for distributed machine learning across a multitude of user devices [18, 26]. It is known to have several benefits in terms of reducing the communication, computation and storage costs for the users, training better global models, and raising the bar for privacy by not sharing the local data. FL allows users to train models locally on their (user) devices without revealing their local data but instead collaborate by sharing only model updates that can be combined to build a global model through a central server. A typical FL scenario involves numerous users and resource-constrained devices [23] making it challenging to perform resource-intensive tasks like training Deep Neural Networks (DNNs) on them.

*Both authors contributed equally to the paper.

Authors' Contact Information: Tianyue Chu, IMDEA Networks Institute, Madrid (Leganes), Spain and Universidad Carlos III de Madrid, Madrid (Leganes), Spain; e-mail: tianyue.chu@imdea.org; Mengwei Yang, University of California Irvine, Irvine, California, United States; e-mail: mengwey@uci.edu; Nikolaos Laoutaris, IMDEA Networks Institute, Madrid (Leganes), Spain; e-mail: nikolaos.laoutaris@imdea.org; Athina Markopoulou, University of California Irvine, Irvine, California, United States; e-mail: athina@uci.edu.

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.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2376-3647/2024/11-ART

<https://doi.org/10.1145/3702241>

In parallel, significant efforts have been put toward optimizing sparse DNNs to create lightweight models suitable for training on edge devices [6, 14, 24]. Model pruning, which involves the removal of a certain percentage of parameters, has gained significant attention [13, 20, 28]. Its primary objective is to derive a sparse DNNs model that decreases the computational requirements and enhances efficiency without compromising accuracy [10]. Several model pruning techniques have been developed for FL to optimize model sparsity under given constraints on the processing and communication capabilities of devices, especially under heterogeneous clients, [3, 17].

It is intuitively expected that by reducing the number of weights of FL models, the pruning method should, as a side-effect, also improve their resistance to reconstruction attacks. These attacks aim at using these weights to reverse-engineer the model and obtain information about the input data; a.k.a. gradient inversion attacks launched by the server. Prior work on model pruning focuses solely on the scalability/accuracy trade-off and thus neither considers nor quantifies the privacy impact of pruning on FL. To the best of our knowledge, our paper is the first to consider privacy in the context of FL model pruning, and to address the following key questions: **Q1.** *Can we quantify privacy in FL model pruning?* **Q2.** *How can we further optimize pruning for privacy, and jointly with accuracy?* To address these questions, we make the following two contributions.

First, we **theoretically and empirically quantify privacy leakage in model pruning**. We derive information-theoretic upper bounds on the amount of information revealed about any single user’s dataset via model updates, in *any* pruned FL scheme. This is the first theoretical characterization of privacy leakage in pruned FL models. We also conduct a comprehensive empirical evaluation that quantifies the actual amount of privacy leakage of six different pruning schemes, considering several state-of-the-art privacy attacks. Within the family of gradient inversion attacks (including Deep Leakage from Gradients [45] and Gradient Inversion (GI) [11]), we also design a novel privacy attack (referred to as Sparse Gradient Inversion, or SGI), specifically tailored to exploit vulnerabilities inherent in FL model pruning. This evaluation combined with our theoretical analysis provides valuable insights into the choices and parameters that affect the privacy provided by pruning.

Building upon these insights, we make our second contribution: we design ***PriPrune*– a privacy-preserving pruning mechanism**. *PriPrune* defends against gradient inversion attacks in pruned FL, by performing additional local (defense) pruning, after any (base) FL pruning scheme. *PriPrune* applies a personalized defense mask and adapts the defense pruning rate, so as to jointly optimize model accuracy *and* privacy, using back-propagation augmented with Gumbel Softmax Sampling. Another key idea in the design of *PriPrune* is what we refer to as *Pseudo-Pruning*: it entails pruning with the defense mask within the local model and transmitting the pruned model to the server, thereby enhancing privacy. However, the weights pruned out by the defense mask in the local model are not discarded; instead, they are retained locally for subsequent rounds of local training, thereby preserving model accuracy. We evaluate *PriPrune* via comprehensive experiments across various FL pruning schemes, a state-of-the-art attack and several benchmark datasets. We show that it achieves a significantly better privacy-accuracy tradeoff than privacy-unaware pruned FL baselines. For example, on the FEMNIST dataset, *PriPrune* improves the privacy *PruneFL* [17] by 45.5% without imposing any accuracy loss.

2 Related Work

Model Pruning in FL. State-of-the-art research for model pruning in FL utilizes two main approaches: one involves the pruning mask being decided on the server side and the client only using this mask without changing it, and the other involves the collaborative selection of the pruning mask by both the user and the server side. *PruneFL* [17] follows the first approach, aiming to reduce the size and complexity of FL models without compromising accuracy. It achieves this by removing less important weights and connections in each layer of the neural network based on their magnitude and importance scores. However, the pruning mask is only decided by the server without considering user-side local representations. *LotteryFL* [22] belongs to the second approach,




Fig. 1. FL with model pruning, under a privacy attack by an honest-but-curious server. Step 1: the server broadcasts the global model with weight w^t to users. Step 2: After local training, user i employs base pruning with pruning mask \bar{m}_i^t , resulting in a pruned model with weights \bar{w}_i^t . The user then applies a defense pruning technique, yielding a pruned model with weights \hat{w}_i^t . Step 3, the server aggregates the model updates and launches a gradient inversion attack upon receiving \hat{w}_i^t . Step 4: the server applies global pruning with pruning mask \bar{m}^t , and updates the global model accordingly. Notice *Base* pruning can be any FL pruning scheme. *Defense* pruning is introduced, in this paper, specifically to protect against privacy attacks.

introducing a method that allows users to maintain local representations by selecting a subset of the global network using personalized masks. However, the system of sparse models it produces performs well only on local datasets. *FedDST* [3] involves dynamic sparse training on both the client and server sides, extracting and training sparse sub-networks. Other recent state-of-the-art studies have introduced various gradient-based one-shot pruning algorithms at initialization. *Snip* [21] leverages connection sensitivity to retain critical connections and remove less significant ones, aiming to limit the loss due to pruning. On the other hand, *GraSP* [38] focuses on preserving the gradient flow through the network by scoring weights based on the Hessian-gradient product. *SynFlow* [34], being a data-agnostic pruning algorithm, emphasizes preserving the total flow of synaptic strengths through the neural network to achieve Maximal Critical Compression during initialization. More importantly, none of them considered privacy implications in their design, as *PriPrune* does.

Reconstruction Attacks in FL. In terms of privacy leakage, communicating gradients throughout the training process in federated learning can reveal sensitive information about the participants. *Deep leakage from gradients* (DLG) [45] demonstrated that sharing the gradients can leak private training data, including images and text. *Improved DLG* (iDLG) [44], presented an analytical procedure to extract the ground-truth labels from the shared gradients and showed the advantages of iDLG over DLG. *Inverting Gradients* [11] introduced cosine similarity as a cost function in their reconstruction attacks and employed total variation loss \mathcal{L}_{TV} as an image prior. *Inverting Gradients* [11] and *GradInversion* [41] demonstrated the capability to reconstruct high-resolution images with increased batch sizes. The reconstruction quality of the attack was improved in *Reconstruction From Obfuscated*

Gradient (ROG) [42]. ROG [42] proposed conducting the reconstruction optimization in low dimensional space and trained a neural network as a postprocessing module.

While these privacy attacks have demonstrated the privacy vulnerabilities of standard federated learning, none of them has considered pruned federated learning. Therefore, we do not currently know how effective existing attacks are against the sparser models from pruned FL.

Privacy-Preserving Techniques in FL. FL is particularly vulnerable to the aforementioned inverting-the-gradient type of reconstruction attacks, since it relies on gradient updates to operate. There exists a large body of literature on privacy-enhancing methods that protect against this type of attacks in FL. The adversary is typically an “honest-but-curious” server, *i.e.*, a server that participates correctly in the FL protocol, but tries to infer training data of individual clients based on their gradient updates. One family of defense approaches is based on differential privacy (DP) [7, 8, 12, 36, 40, 43], which can be applied locally (at the clients), centrally (at the server), or in a distributed way. In the local DP model for FL, each client adds noise to the local update before sharing with the server, which makes it difficult for the server to invert the true gradient and infer the client’s training data. However, local DP is known to significantly degrade utility, in this case model accuracy. Another defense approach applied to FL is Secure Aggregation (SecAgg) [4, 32, 33, 37, 39]. SecAgg ensures that individual updates are encrypted (thus preventing the server from inverting the gradients of individual users), while their aggregate can still be decrypted (thus enabling the server to train a good model). SecAgg is powerful but has computational overhead, although significant advances have been done recently.

Both DP and SecAgg are important but orthogonal to the underlying FL and can be combined with a number of FL protocols, including all the pruning protocols we consider as base in this study. Our proposed *PriPrune* can also be combined with DP and/or SecAgg. For instance, after applying *PriPrune* to base pruning, clients can add local DP noise to the pruned updates before sending them to the server, and/or encrypt them to enable SecAgg following standard protocols. We focus on pruning itself and the privacy-utility tradeoff that can be achieved, *before* adding orthogonal defense techniques. Future work could explore adding DP and/or SecAgg, which are out of the scope of this paper.

Recent works like *FedMap* [15] and *Fed-LTP* [30] discuss privacy specifically in the context of pruned FL. *FedMap*, which appeared after this submission, enhances privacy by training models from scratch and by avoiding the sharing of detailed pruning masks. However, it does not quantify the privacy leakage that may still occur from the pruned models. *FedMap* can be considered as a base pruning technique, and *PriPrune* can still be applied to protect it. *Fed-LTP*, on the other hand, combines Lottery Ticket Pruning with Zero-concentrated Differential Privacy (zCDP) to enhance privacy. Integrating zCDP, or any type of DP, with our proposed *PriPrune* is orthogonal to our core contribution, as discussed above, and can be explored as part of future research.

3 Problem Setup

Federated Learning (FL) with Pruning. We consider a general FL system with multiple users and one server, employing *any* model pruning scheme on the users and/or the server. We refer to pruning at this stage as *base* pruning with base pruning rate \bar{p} , and it can be any state-of-the-art pruning FL scheme [3, 17]. This is depicted in Fig. 1 and formalized in Algorithm 1.

The goal is to train a global model $F(\mathcal{D}, \mathbf{w})$ through FL. At the beginning of round t , user i with the local data \mathcal{D}_i and labels y_i has local model weights \mathbf{w}_i^t ; the server broadcasts the most recent global model $F(\mathcal{D}, \mathbf{w}^t)$ to all participating users. Then the round proceeds as follows:

Algorithm 1 FL with Pruning under Privacy Attack

```

1: Given:  $T$  number of global training rounds;  $N$  number of users indexed by  $i$ ;  $E$  number of local epochs; the server aims to
   reconstruct the local data of the target user.
2: Server executes:
3: for  $t \leftarrow 1$  to  $T$  do
4:   for each user  $i \in N$  in parallel do
5:      $\hat{\mathbf{w}}_i^t \leftarrow UserUpdate(\mathbf{w}^{t-1}, \mathbf{m}^{t-1}, i)$ 
6:     if user  $i$  is targeted then
7:        $\nabla \hat{\mathbf{w}}_i^t \leftarrow \hat{\mathbf{w}}_i^t - \hat{\mathbf{w}}_i^{t-1}$ 
8:       Privacy Attack on  $(F(\mathcal{D}_i; \hat{\mathbf{w}}_i^t), \hat{\mathbf{w}}_i^t, \nabla \hat{\mathbf{w}}_i^t)$ 
9:     end if
10:    end for
11:     $\mathbf{w}^t \leftarrow \left( \sum_{i=1}^I \frac{\|\mathcal{D}_i\|}{\|\mathcal{D}\|} \hat{\mathbf{w}}_i^t \right) \odot \bar{\mathbf{m}}^t;$ 
12:  end for
13:  function  $UserUpdate(\mathbf{w}^{t-1}, \mathbf{m}^{t-1}, i)$ :
14:    for local epoch  $e \leftarrow 1$  to  $E$  do
15:      Local Training and Base Pruning:  $\bar{\mathbf{w}}_i^t \leftarrow$  Eq. (1)
16:      Defense Pruning:  $\hat{\mathbf{w}}_i^t \leftarrow \bar{\mathbf{w}}_i^t \odot \hat{\mathbf{m}}_i^t$ 
17:    end for
18:  Return  $\hat{\mathbf{w}}_i^t$  to server

```

Local Training and Base Pruning: Each user i utilizes base pruning strategy $\mathbb{P}_i(\cdot)$ with base pruning masks $\bar{\mathbf{m}}_i^t$ on its trained local model $F(\mathcal{D}_i, \mathbf{w}_i^t)$:

$$\bar{\mathbf{w}}_i^t := \left(\mathbf{w}^t - r \frac{\partial \ell_i (F(\mathcal{D}_i, \mathbf{w}^t \odot \bar{\mathbf{m}}^t), y_i)}{\partial \mathbf{w}^t} \right) \odot \bar{\mathbf{m}}_i^t \quad (1)$$

Where $\bar{\mathbf{m}}_i^t = \mathbb{P}_i(\mathbf{w}_i^t)$, r is the learning rate and \odot is the Hadamard product. The *base* pruning mask $\bar{\mathbf{m}}_i^t$ exhibits variability across different base pruning schemes, reflecting the diverse pruning strategy $\mathbb{P}_i(\cdot)$ employed.

Server: The server receives model updates $\hat{\mathbf{w}}_i^t$ from all participating users and aggregates to update the global model, as it is typical in FL. We do not consider secure aggregation [33], differential privacy [27], or other defenses known in the literature.¹

Threat Model. We consider an *honest-but-curious* server, which correctly follows the FL protocol but uses model updates to infer private information, by launching privacy attacks against target user(s), as shown in Algorithm 1. In our case, the server has full visibility of all masked local models and, upon receiving an update from a target user, it attempts to reconstruct the user's training data. We consider reconstruction attacks that invert gradients, including Deep-Leakage-from-Gradients (DLG) [45], Gradient Inversion (GI) [11], and a custom attack SGI described in Section 4.2. We focus on single-round attacks, *i.e.*, attacks that invert gradients observed at a single round².

¹We consider those defenses as out-of-the-scope for this paper, since they are orthogonal to pruning. They can be implemented on top/combined with our pruning-based defense methods.

²To the best of our knowledge, there are currently no known *practical* multi-round inverting-the-gradient types of attacks that combine information from multiple rounds. There is theoretical analysis of upper *bounds* of multi-round information leakage, including [9, 32] and we derive our own multi-round leakage in Section 4.1.2. Designing a practical multi-round attack is orthogonal to our contribution, which focuses on quantifying the leakage and designing an add-on defense, for *existing* attacks.

Defense via Pruning: In this paper, we propose to defend against privacy attacks, via pruning itself. One could specifically design an entire pruning scheme not only for accuracy and efficiency but also for privacy, and the defense would then be optimized jointly with the particular base pruning scheme e.g., PruneFL [17] and FedDST [3].

In this paper, we take a modular and universal approach: we consider *any* given state-of-the-art base pruning scheme (with rate \bar{p}_i and mask $\bar{\mathbf{m}}_i^t$) without any modification. Then, we introduce additional local *defense* pruning, with mask $\hat{\mathbf{m}}_i^t$ and pruning rate \hat{p}_i , specifically designed for protecting each user i from privacy attacks. User i eventually sends back to the server weights after applying both base ($\bar{\mathbf{m}}_i^t$) and defense ($\hat{\mathbf{m}}_i^t$) pruning: $\mathbf{w}_i^t = \bar{\mathbf{m}}_i^t \odot \hat{\mathbf{m}}_i^t \odot \mathbf{w}_i^t$. This is depicted in Fig. 1 and Algorithm 1 (see Line 15 and 16), and elaborated upon on Fig. 7. Given a base pruning ($\bar{\mathbf{m}}_i^t$), there are many possible defense pruning strategies ($\hat{\mathbf{m}}_i^t$). In Section 5.2, we propose *PriPrune* for $\hat{\mathbf{m}}_i^t$, to optimize the accuracy-vs-privacy tradeoff.

4 Quantifying Privacy in FL Model Pruning

In this section, to address question Q1, we provide a theoretical and empirical quantification of privacy leakage in FL model pruning. We aim to assess the server’s ability to infer information about an individual user’s local dataset \mathcal{D}_i . We employ Mutual Information (MI) as the metric for quantifying privacy leakage:

$$\begin{aligned} I(X; Y) &= \int_{X \times Y} \log \frac{d\mathbb{P}_{X,Y}(x, y)}{d\mathbb{P}_X(x) \otimes \mathbb{P}_Y(y)} d\mathbb{P}_{X,Y}(x, y) \\ &= H(X) + H(Y) - H(X, Y) \end{aligned} \quad (2)$$

Where $(X; Y)$ is a pair of random variables with values over the space $X \times Y$. Their joint distribution is $\mathbb{P}_{X,Y}(x, y)$ and the marginal distributions are $\mathbb{P}_X(x)$ and $\mathbb{P}_Y(y)$.

This concept, rooted in information theory and Shannon entropy, captures the interdependence between two random variables. It is very powerful for privacy analysis, given its applicability across various domains and tasks [1, 2]. This applicability extends across various datasets, pruning schemes, and attack scenarios, as recently demonstrated in the quantification of privacy within the context of FL aggregation [9].

4.1 Theoretical Quantification

We seek to quantify how much information the server can infer about the private data \mathcal{D}_i of user i over T global training rounds, based on the pruned updates $\{\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t\}_{t \in [T]}$ submitted by user i , via:

$$I_i = I\left(\mathcal{D}_i; \{\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t\}_{t \in [T]}\right) \quad (3)$$

where I represents mutual information between \mathcal{D}_i and $\{\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t\}_{t \in [T]}$, $\mathbf{w}_i^t \in \mathbb{R}^d$ is the local model weights of user i at round t .

This quantification is based on two mild assumptions:

ASSUMPTION 1. *The random vector $\hat{g}_i(\mathbf{w}_i^t, b)$ has non-singular covariance matrix Σ_i and mean 0.*

Here, $\hat{g}_i(\mathbf{w}_i^t, b) \in \mathbb{R}^{d^*}$ is the largest sub-vector of the $g_i(\mathbf{w}_i^t, b)$, where d^* is the rank of $g_i(\mathbf{w}_i^t, b)$, $d^* \leq d$, d is the model size, b denotes the size of the random samples. $g_i(\mathbf{w}_i^t, b)$ denotes the stochastic estimate of the gradient of the local loss function for user i , computed based on a random sample b , and \mathbf{B}_i^t is a data batch of size B sampled uniformly at random from its local dataset \mathcal{D}_i .

ASSUMPTION 2. *\bar{p}_i the pruning rate of user i remains unchanged throughout the training process.*

Building upon Equation 2, Equation 3 can be written as:

$$I_i \leq \sum_{t=1}^T \left(I \left(x_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) \right) \quad (4)$$

Where $x_i^t = \frac{\partial \ell_i(\bar{\mathbf{w}}_i^{t-1}; \mathcal{D}_i)}{\partial \mathbf{w}} = \frac{1}{B} \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b)$.

To that end, we first characterize the privacy leakage for a *single round* t :

$$I \left(x_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) \quad (5)$$

Under these two mild assumptions, we obtain our main theoretical result; the proof is deferred to Appendix A.1.

4.1.1 Upper Bound for Privacy Leakage in a Single Round. We now provide an upper bound for Equation 5, i.e., for the privacy leakage in a single round.

THEOREM 4.1 (SINGLE ROUND LEAKAGE). *Under Assumption 1 and 2, we have an upper bound of I_i^t for a single round:*

$$I_i^t \leq 1 - \frac{\bar{p}_i - 1}{2 \ln 2} + 2 \log \frac{1}{B} + 2\Delta + d^* \log(2\pi e) \quad (6)$$

where $\Delta = \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right|$.

Proof. Based on Equation 2, we have

$$I \left(x_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) \leq \underbrace{H \left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right)}_{A} + \underbrace{H \left(x_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right)}_{B} \quad (7)$$

For part A, based on the chain rule of entropy, we have

$$A = H \left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) + H \left(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t]} \right) \geq 0 \quad (8)$$

Due to $H \left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \mathbf{w}_i^t, \bar{\mathbf{m}}_i^t, \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) = 0$, and $H \left(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t]} \right) \geq 0$, we have

$$H \left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) \leq \underbrace{H \left(x_i^t \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right)}_{B} + 1 - \frac{\bar{p}_i - 1}{2 \ln 2} \quad (9)$$

The proof is shown in the Appendix A.1.

$$B = H \left(\frac{1}{B} \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b) \middle| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right) \quad (10)$$

Let $X = \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b)$, $Z = \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}$, $Y = \frac{1}{B} X \cdot f_{X,Z}(x, z)$ and $f_{Y,Z}(y, z)$ are the corresponding joint probability density functions.

Then Equation 10 can be written as follows:

$$\mathbf{H}(Y|Z) = \mathbf{H} \left(\underbrace{\sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b) \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right.}_{\mathbf{C}} \right) + \log \frac{1}{B} \quad (11)$$

The proof is shown in the Appendix A.1.

In recent theoretical results for analyzing the behaviour of SGD, the SGD vector is approximated by a distribution with independent components or by a multivariate Gaussian vector [46]. Based on the ZCA whitening transformation and Assumption 1, we have $\hat{g}_i(\mathbf{w}_i^t, b) = \Sigma_i^{-\frac{1}{2}} \mathbf{z}$, where \mathbf{z} has zero mean and \mathbb{I}_{d^*} covariance matrix.

For the part C, we set $\mathbf{v} = \sum_{b \in \mathbf{B}_i^t} \mathbf{z}$ and $\mathbf{u} = \Sigma_i^{-\frac{1}{2}} \mathbf{v}$, then we have:

$$\begin{aligned} \mathbf{C} &= \mathbf{H} \left(\sum_{b \in \mathbf{B}_i^t} \hat{g}_i(\mathbf{w}_i^t, b) \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) = \mathbf{H} \left(\Sigma_i^{-\frac{1}{2}} \sum_{b \in \mathbf{B}_i^t} \mathbf{z} \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) = - \int \mathbb{P}_{\mathcal{U}}(\mathbf{u}) \log \mathbb{P}_{\mathcal{U}}(\mathbf{u}) d\mathbf{u} \\ &\stackrel{(a)}{=} - \int \frac{\mathbb{P}_{\mathcal{V}}(\mathbf{v})}{\left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right|} \log \left(\frac{\mathbb{P}_{\mathcal{V}}(\mathbf{v})}{\left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right|} \right) \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| d\mathbf{v} = \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + \mathbf{H} \left(\underbrace{\sum_{b \in \mathbf{B}_i^t} \mathbf{z} \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right.}_{\mathbf{D}} \right) \end{aligned} \quad (12)$$

where (a) is based on the transformation of random vectors. Based on the Maximum Entropy Upper Bound [35], which the continuous distribution \mathbf{X} with prescribed variance $\mathcal{V}(\mathbf{X})$ maximizing the entropy is the Gaussian distribution of same variance. Since the entropy of a multivariate Gaussian distribution with mean μ and covariance $K_{i,j}$ is

$$h(\mathcal{N}_n(\mu, K)) = \frac{1}{2} \log (2\pi e)^n |K|$$

where $|K|$ denotes the determinant of K . Hence, the maximum entropy upper bound for part D is

$$\mathbf{D} \leq \frac{1}{2} \log (2\pi e)^{d^*} |\mathbb{I}_{d^*}| = \frac{d^*}{2} \log (2\pi e) \quad (13)$$

Therefore, by combining Equation 12 and Equation 13, the upper bound for part B is

$$\mathbf{B} \leq \log \frac{1}{B} + \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + \frac{d^*}{2} \log (2\pi e) \quad (14)$$

By adding parts A and B, we establish the upper bound for privacy leakage for a single round as in Theorem 4.1.

4.1.2 Upper Bound for Privacy Leakage across Multiple Rounds. Using Equation 4, we can also obtain an upper bound for Equation 3, i.e., how much information the aggregated model over T global training rounds could leak about the private data:

COROLLARY 4.2 (MULTIPLE ROUNDS LEAKAGE). *Continuing with Theorem 4.1, we can upper bound \mathbf{I}_i after T global training rounds as follows:*

$$\mathbf{I}_i \leq T \left(1 - \frac{\bar{p}_i - 1}{2 \ln 2} + 2 \log \frac{1}{B} + 2\Delta + d^* \log (2\pi e) \right) \quad (15)$$

Algorithm 2 Sparse Gradient Inversion (SGI) Attack

1: **Input:** $F(\mathcal{D}_i; \hat{\mathbf{w}}_i^t)$: model at round t from targeted user i ; learning rate η for inverting gradient optimizer; S : max iterations for attack; τ : regularization term for cosine loss in inverting gradient attack; \mathbf{m}_{rec} : the recovery mask generated by server; d : the model size

2: **Output:** reconstructed training data (\mathcal{D}_i, y_i) at round t

$$\mathbf{m}_{rec}[j]_{j \in d} = \begin{cases} 1 & \text{if } \hat{\mathbf{w}}_i^t[j] \neq 0 \\ 0 & \text{if } \hat{\mathbf{w}}_i^t[j] = 0 \end{cases}$$

3: $\nabla \hat{\mathbf{w}}_i^t \leftarrow (\hat{\mathbf{w}}_i^t - \hat{\mathbf{w}}_i^{t-1}) \odot \mathbf{m}_{rec}$

4: **Initialize** $\mathcal{D}'_0 \leftarrow N(0,1)$, $y'_0 \leftarrow Randint(0, max(y))$

5: **for** $s \leftarrow 0$ **to** $S - 1$ **do**

6: $\nabla \mathbf{w}'_s \leftarrow \partial \ell(F(\mathcal{D}'_s, \hat{\mathbf{w}}_i^t), y'_s) / \partial \hat{\mathbf{w}}_i^t$

7: $\mathcal{L}'_s \leftarrow 1 - \frac{\nabla \hat{\mathbf{w}}_i^t \cdot \nabla \mathbf{w}'_s}{\|\nabla \hat{\mathbf{w}}_i^t\| \|\nabla \mathbf{w}'_s\|} + \tau$

8: $\mathcal{D}'_{s+1} \leftarrow \mathcal{D}'_s - \eta \nabla_{\mathcal{D}'_s} \mathcal{L}'_s$, $y'_{s+1} \leftarrow y'_s - \eta \nabla_{y'_s} \mathcal{L}'_s$

9: **end for**

10: **Return** \mathcal{D}'_S, y'_S

This shows that increasing the number of global training rounds (T) leads to a proportional rise in the upper bound of information leakage from the user's local training model.

Importance of the Upper Bounds: The derived upper bounds are important for several reasons.

- They provide a theoretical upper limit to the amount of information that can potentially be leaked from model updates in pruned models. To the best of our knowledge, this has not been previously characterized for pruned models.
- Because they are information-theoretical in nature, these bounds are universally applicable. They apply to *any* learning task, *any* data distribution, *any* pruning scheme, and *any* privacy attack. The attacker is the server that sees the pruned updates $\{\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t\}_{t \in [T]}$ from client i , and infers information about the local training Data \mathcal{D}_i of that client. Regardless of the exact reconstruction strategy, the inferred information is at most the mutual information I_i between the update and the local data, quantified in Equation 3. The derivation of the bounds assumes the existence of a pruning mask but does not depend on how this mask was generated, *e.g.*, whether through magnitude-based pruning, gradient-based pruning, or any other pruning method.
- Theorem 1 expresses the bound in terms of important parameters (*e.g.*, the pruning rate p , model size d , etc), which provides insights into the choices and parameters that affect the privacy provided by pruning in practical algorithms. However, the bounds may or may not be tight compared to the actual information leakage, depending on the real-world scenario. In fact, these bounds are necessarily conservative since they are universal. This is why, in the next section, we complement the theoretical analysis with experimental evaluations in specific settings (*i.e.*, considering specific base pruning, attack and defense schemes, specific data sets, and learning tasks).

4.2 Empirical Quantification

We perform an experimental evaluation to quantify the exact amount of privacy loss in concrete settings. This allows us to compare the theoretical bounds to experimental results and show that they are qualitatively aligned. It also allows us to obtain insights (in Section 5.1) that inform the design of defense pruning (in Section 5.2). More evaluation results are provided in Section 6.




Fig. 2. Impact of varying pruning rate on privacy leakage using key metrics: Normalized Mutual Information (NMI) displayed on (a) and Peak Signal-to-Noise Ratio (PSNR) shown on (b) under a Sparse Gradient Inversion attack (SGI) in FEMNIST for five pruning methods under varying the base pruning rate.

4.2.1 Privacy Metrics. To quantify the extent of privacy leakage, we use the Normalized Mutual Information (NMI) between the training data and the reconstructed data. For two vectors $U, V \in \mathcal{R}^N$, NMI is computed as follows. The entropy of U is calculated by $H(U) = -\sum_{i=1}^{|U|} P(i) \log(P(i))$, where $P(i) = |U_i|/N$. The Mutual Information between U and V is given by $I(U; V) = \sum_{i=1}^{|U|} \sum_{j=1}^{|V|} \frac{|U_i \cap V_j|}{N} \log \left(\frac{N|U_i \cap V_j|}{|U_i||V_j|} \right)$. NMI is then measured as $\frac{I(U; V)}{\text{mean}(H(U), H(V))}$.

For image data, prior research [9] has shown that NMI aligns well with the Peak Signal-to-Noise Ratio (PSNR). This is also the case here: we see an agreement between NMI and PSNR with increasing pruning rates in Fig. 2 and Fig. A1, both showing a similar decreasing trend. In the rest of the paper, we use both NMI and PSNR as metrics to quantify the success of a reconstruction attack. Intuitively, the higher the NMI, the higher the privacy leakage, the more successful the reconstruction attack. The higher the PSNR, the closer the original and reconstructed images and the more successful the attack. Hence, Fig. 2 and Fig. A1 compare the NMI and PSNR metrics and show that that increasing pruning rate reduces the success of the privacy attack.

4.2.2 Reconstruction Attack Algorithms. We first implement the classic Gradient Inversion (GI) attack [11] from the DLG attack family. However, due to the unique sparsity patterns introduced by model pruning in FL, we identify the potential for further optimizing this attack to specifically exploit the sparsity. Consequently, we introduce an advanced attack tailored for the context of model pruning in FL, the Sparse Gradient Inversion (SGI) attack. This attack is designed to recover the pruning mask within the pruned model. Moreover, its adaptability allows seamless integration with existing privacy attacks in FL, thereby amplifying their effectiveness.

The SGI attack is provided in Algorithm 2. In each iteration t , the SGI algorithm proceeds as follows: (i) Based on the user i 's local model weights \mathbf{w}_i^t , the server attempts to recover the pruning mask of user i (Line 2). The recovered mask is then integrated into the calculation of gradient, allowing the server to obtain a gradient $\nabla \hat{\mathbf{w}}_i^t$ that closely approximates the true gradients of user i ; (ii) The SGI attack randomly initializes a set of dummy data, comprising dummy inputs \mathcal{D}'_0 and dummy labels y'_0 . (iii) After the dummy gradient \mathbf{w}' is acquired, the server then updates the dummy data in the direction that minimizes the cosine distance between the dummy gradient and the gradient $\nabla \hat{\mathbf{w}}_i^t$.

Method	Pruning Criteria	Pruning Cycles	Pruning at Server	Pruning at Client
SNIP [21]	gradient-based	one shot	✓	✗
SynFlow [34]	gradient-based	one shot	✓	✗
Random Pruning	magnitude-based	iterative	✓	✓
FedDST [3]	magnitude-based	iterative	✓	✓
PruneFL [17]	magnitude-based	iterative	✓	✓

Table 1. Comparison for Evaluated Methods




Fig. 3. Comparison between the SGI (Sparse Gradient Inversion) attack and the GI (Gradient Inversion) attack on the FEMNIST dataset for five pruning methods under varying the base pruning rate. In addition to the NMI metric, we show the corresponding recovered images, which show the degradation with an increasing pruning rate.

4.2.3 Base Pruning Schemes under Attack. We subject a set of well-established FL pruning methods to the privacy attack to assess their vulnerabilities and effectiveness in protecting user privacy. The pruning methods under evaluation include Random pruning, Snip [21], SynFlow [34], FedDST [3], and PruneFL [17], all of which are discussed in Related Work (Section 2). We present an overview of the pruning criteria, cycles, and schedules employed by the six evaluated methods. These specifics are elaborated in Table 1. By scrutinizing the table, we discern notable trends: three of the methods execute pruning exclusively at the server side, while the remaining methods engage in pruning activities at both the server and client sides. The latter approaches adopt an iterative pruning strategy throughout the course of the FL process.

We compare the effectiveness of the Sparse Gradient Inversion (SGI) attack with the Gradient Inversion attack under varying the base pruning rate. Our evaluation is based on the normalized mutual information (NMI) metric. Specifically, we analyze the NMI achieved by each attack for different pruning methods. Figure 3 demonstrates that the SGI attack consistently outperforms the Gradient Inversion attack by inferring a larger amount of information (as captured by NMI) at each pruning rate level, across all tested pruning methods. This is because the SGI attack exploits the sparsity inherent to model pruning, ultimately allowing better reconstruction attacks. In the rest of the paper, we use the SGI attack as a baseline, with a learning rate of 0.01 and 10,000 attack iterations.




Fig. 4. Impact of varying base pruning rate \bar{p} and batch size B on privacy leakage. We consider the Normalized Mutual Information (NMI) as privacy metric. We launch a Sparse Gradient Inversion (SGI) attack against 5 base pruning methods (Random, SNIP, SynFlow, FedDST, and PruneFL), over two datasets (FEMNIST in the First Column and CIFAR-10 in the Second Column).

4.2.4 Attack Performance. We launch the SGI attack on the aforementioned FL pruning schemes, on the FEMNIST and CIFAR10 datasets, to assess how FL parameters influence privacy leakage. The experimental results in Fig. 4 and Fig. 5 demonstrate that the impact of the parameters \bar{p} , B , d , and T qualitatively align with our theoretical findings.

Impact of Pruning Rate (\bar{p}). Fig. 4a and Fig. 4b illustrate that a higher pruning rate generally leads to a decrease in NMI. This is intuitively expected as pruning removes information and makes it more difficult to attack the original model. The effect is clearer for large pruning rates (e.g., above 0.5). Image reconstruction is clearer in FEMNIST than CIFAR, consistently with prior work.

Impact of Batch Size (B). Fig. 4c and Fig. 4d show increasing the batch size B contributes to reducing information leakage from the local training model. Larger batch sizes enable more data to be added during model training, which can help obscure single data and increase privacy protection.

Impact of Model Size (d). We consider 3 architecture of the models for FEMNIST dataset: Conv-1, Conv-2, Conv-4, featuring 1, 2, 4 convolutional layers, along with 2 linear layers. These models encompass 626720, 6601504 and 13077392 parameters, respectively. We consider 3 different model architectures for CIFAR 10 dataset: VGG-4, VGG-6, and VGG-11. These models consist of 1 convolutional layer and 2 linear layers, 4 convolutional layers




Fig. 5. Impact of varying model size d and training rounds T on privacy leakage using Normalized Mutual Information (NMI) under a Sparse Gradient Inversion (SGI) attack for 5 base pruning methods (Random, SNIP, SynFlow, FedDST, and PruneFL), over two datasets (FEMNIST in the First Column and CIFAR-10 in the Second Column).

and 2 linear layers, and 8 convolutional layers with 3 linear layers, respectively. These models encompass 268464, 624048 and 9747136 parameters, respectively. Figure 5a and Figure 5b depicts the relationship between the model size (d) and the corresponding information leakage. The x-axis in the figure represents the number of layers in each model. Fig. 5a and Fig. 5b show the upper bounds on information leakage increase as d^* increases. Interestingly, the impact of model size d on information leakage is not linear. In the case of over-parameterised models, some parameters may not contribute to the model's performance or information retention. Therefore, increasing the size of such models may not proportionally affect information leakage.

Impact of Global Training Rounds (T). Fig. 5c and Fig. 5d show that increasing the global training rounds (T) leads to a proportional rise in the upper bound of information leakage. As the training progresses over multiple rounds, the user's model updates are repeatedly exposed to the central server. This repeated exposure increases the risk of potential memorization of private training information by the server.




Fig. 6. Tradeoff of privacy vs. accuracy with three defense strategies: *Largest* (red), *Random* (green), and *Mix* (yellow), and their enhanced versions incorporating *Pseudo-Pruning* (*Pseudo*). *Pseudo* could improve the overall tradeoff by elevating model accuracy and *Largest* offers the highest level of privacy protection among all defense methods.

5 Design of Defense Pruning

5.1 Insights from Privacy Quantification

As shown above, pruning designed only with model accuracy and model size in mind, does not sufficiently protect against privacy attacks. Next, we describe three insights gained in the process of privacy quantification that inform our design of privacy-aware FL model pruning (*PriPrune* described in Section 5.2).

Insight 1: Largest Weights Matter. Pruning weights with large gradients improves privacy the most, but also hurts model accuracy the most. The intuition is that (base) pruning criteria, as used e.g., in *PruneFL* and *FedDST*, avoid pruning model weights with the largest gradients, in order to preserve the most valuable information for the model. These gradients can then be exploited by gradient inversion attacks. Fig. 4a and Fig. 4b confirm empirically that, as more weights with large gradients are pruned, the attack is less successful.

This implies that the defense should strategically prune certain weights with large gradients, in addition to those pruned by the base pruning scheme. We explored three defense strategies for weights to prune: (1) weights with the top- k largest gradients (denoted as *Largest*), (2) random weight pruning (denoted as *Random*), and (3) a hybrid approach combining the *Largest* and *Random* (denoted as *Mix*). We compared the three strategies for the same defense pruning rate \hat{p}_i (set 0.3) on top of the base method (*PruneFL*) and details are provided in Appendix C.1. The results in Fig. 6, within the Pruning oval, show the tradeoff achieved between privacy and accuracy: the “*Largest*” method achieves the best privacy and the worst model accuracy. This is because pruning weights with large gradients, removes valuable information, thus improving privacy but also significantly reducing model accuracy. Therefore, to improve this tradeoff and preserve privacy without harming accuracy, we need to introduce an additional mechanism.

Insight 2: *Pseudo-Pruning*. Users can prune weights with large magnitude gradients when they communicate with the server (which helps privacy) *and* still keep these weights for their local training (which maintains accuracy).

We refer to this idea as *Pseudo-Pruning*, and illustrate it in Fig. 7. Before sending the weights \bar{w}_i^t to the server user i conducts *Pseudo-Pruning* with defense mask \hat{m}_i^t . The weights pruned out by defense pruning with large gradients \tilde{w}_i^t ($= \bar{w}_i^t \odot \hat{m}_i^t$) are withheld locally, thus preserving accuracy. Meanwhile, the weights remaining after the defense pruning without large gradients, \hat{w}_i^t ($= \bar{w}_i^t \odot \hat{m}_i^t$), are transmitted to the server, thus preserving privacy.




Fig. 7. Illustration of local pruning at user i with *Pseudo-Pruning* defense: First user applies base pruning with base mask $\bar{\mathbf{m}}_i$ to derive a pruned model $\bar{\mathbf{w}}_i^t$. Then, defense pruning is conducted with mask $\hat{\mathbf{m}}_i$. Defense pruning is implemented as *Pseudo-Pruning*, where the remaining weights $\tilde{\mathbf{w}}_i^t$ ($= \bar{\mathbf{w}}_i^t \odot \hat{\mathbf{m}}_i^t$) are sent to the server; while the pruned weights $\bar{\mathbf{w}}_i^t$ ($= \bar{\mathbf{w}}_i^t \odot \tilde{\mathbf{m}}_i^t$) are locally saved for future local training. Notably, these weights pruned by defense are not actually pruned, hence the term "Pseudo-Pruning."

Here, $\tilde{\mathbf{m}}$ is the bit-wise complement matrix of $\hat{\mathbf{m}}$. In the next round, after user i receives the global model \mathbf{w}^{t+1} , the locally saved weights $\tilde{\mathbf{w}}_i^t$ are loaded into the global model, serving as the local initial model.

We combine *Pseudo-Pruning* with each of the three strategies (*Largest*, *Random*, and *Mix*) in Insight 1, and we show the evaluation results within the *Pseudo-Pruning* dashed oval in Fig. 6. Fig. 6 shows that the model accuracy is notably better with *Pseudo-Pruning* than real pruning, with the same mask. Furthermore, combining the two insights into *Largest (Pseudo)* offers the most effective privacy guarantee among all defenses discussed above. Hence, we adopt *Largest (Pseudo)* as our defense strategy. Additional evaluation details are available in Appendix C. Still, a remaining shortcoming of the defense considered so far, is that their pruning rate is manually selected and remains fixed.

Insight 3: Adapt.

The defense pruning rate (\hat{p}_i) should adapt to the model weights, to jointly optimize privacy and model accuracy. A fixed defense rate is unable to capture the evolving dynamics of the training process. The defense strategy should adapt to the changing dynamics of the training process by employing an adaptive defense pruning rate, enabling a more effective optimization of the tradeoff between privacy and accuracy.

5.2 The *PriPrune* Mechanism

Combining all aforementioned insights, we introduce the defense strategy, *PriPrune*, which dynamically adapts the *Pseudo-Pruning* defense mask $\hat{\mathbf{m}}$ to jointly optimize accuracy and privacy throughout the FL training process.

Loss Function. The objective of user i is to minimize a loss function that combines privacy and accuracy:

$$\mathcal{L}_{local} = \lambda_{acc} \mathcal{L}_{acc} + \lambda_{pri} \mathcal{L}_{pri} + \lambda_{sha} \sum_{l \in L, j \in d_l} \alpha_{i(l,j)}^t \quad (16)$$

Here, \mathcal{L}_{acc} represents the local model training loss, \mathcal{L}_{pri} represents the local privacy loss, and $\sum_{l \in L, j \in d_l} \alpha_{i(l,j)}^t$ serves as the privacy regularization term, where $\alpha_{i(l,j)}^t$ represents the probability of user i not sharing the j -th parameter

in l -th layer with the server at round t . L represents the total number of layers in the model, and d_l refers to the number of parameters in the l -th layer. Additionally, λ_{acc} , λ_{pri} , and λ_{sha} act as weights for accuracy loss, privacy loss, and the privacy regularization term, respectively.

Next, we explain the terms of the loss function of the user in Equation 16 and their rationale. First, the accuracy loss is given by: $\mathcal{L}_{acc} = \mathcal{L}(\mathcal{D}_i; \hat{\mathbf{w}}_i^t | \mathbf{w}_{i,init}^t)$. The initial local weight is computed as: $\mathbf{w}_{i,init}^t = \mathbf{w}^t \odot \hat{\mathbf{m}}_i^{t-1} + \bar{\mathbf{w}}_i^{t-1} \odot \tilde{\mathbf{m}}_i^{t-1}$, where $\tilde{\mathbf{m}}$ denotes the bit-wise complement matrix of $\hat{\mathbf{m}}$. It is the composition of locally saved weights with global weights through the utilization of the defense mask $\hat{\mathbf{m}}$.

Then, the privacy loss \mathcal{L}_{pri} is given by:

$$\mathcal{L}_{pri} = \sum_{l \in L, j \in d_l} -\frac{N_l}{\sum_{l \in L} N_l} \frac{|\mathbf{g}_{l,j}|}{\sum_{j \in d_l} |\mathbf{g}_{l,j}|} \log \alpha_{i(l,j)}^t \quad (17)$$

N_l indicates the number of weights in l -th layer, $|\mathbf{g}_{l,j}|$ represents the magnitude of the gradient for the j -th parameter in l -th layer. The privacy loss is designed to discourage sharing weights with large gradients inspired by Insight 1, aiming to mitigate information leakage. Therefore, we assign higher weights to parameters with larger gradients in the loss function. Through this process, during the optimization of the privacy loss, the algorithm updates a higher α_i^t associated with the larger weight term. This higher α_i^t signifies a higher probability of not sharing the parameters with larger gradients with the server, thus promoting the sharing of less valuable information with the server. Moreover, we also distribute constraints to l -th layer based on its parameter count N_l , in proportion to the total model weights, which ensures uniform layer-wise sparsity.

The term $\sum_{l \in L, j \in d_l} \alpha_{i(l,j)}^t$ introduces a privacy penalty based on the magnitude of α_i^t . The optimization process aims to minimize the sum of α_i^t , resulting in a decrease in the overall probability of not sharing parameters. This, in turn, leads to an increase in the total probability of sharing parameters. Thus, while the privacy loss discourages the sharing of parameters with larger gradients with the server, this term encourages the sharing of parameters with smaller gradients, thus preserving accuracy while maintaining privacy.

Optimization. In *PriPrune*, we optimize the defense mask $\hat{\mathbf{m}}$ and model weights $\hat{\mathbf{w}}_i$ jointly through gradient descent based on our designed loss function. To overcome the discrete and non-differentiable nature of the defense mask $\hat{\mathbf{m}}$, we employ Gumbel-Softmax sampling [16, 25] to substitute the original non-differentiable sample with a differentiable sample. Specifically, $\hat{\mathbf{m}}$ is parameterized by a distribution vector $\pi_{l,j} = [\alpha_{i(l,j)}^t, 1 - \alpha_{i(l,j)}^t]$. $v_{l,j}$ represents the soft *Pseudo-Pruning* decision for the j -th parameter in layer l . The reparameterization trick is used for differentiable training:

$$v_{l,j}(k) = \frac{\exp((\log \pi_{l,j}(k) + G_{l,j}(k))/\tau)}{\sum_{k \in \{0,1\}} \exp((\log \pi_{l,j}(k) + G_{l,j}(k))/\tau)} \quad (18)$$

where $G_{l,j} = -\log(-\log U_{l,j})$ is Gumbel distribution with $U_{l,j}$ sampled from a uniform i.i.d. distribution $\text{Unif}(0, 1)$, τ is the temperature of the softmax and $k \in \{0, 1\}$. We apply the hard sample trick introduced by PyTorch document [29] to acquire the defense mask $\hat{\mathbf{m}}$ and the main trick is to do $y_{hard} - f_{stopGradient}(y_{soft}) + y_{soft}$.

PriPrune. The *PriPrune* algorithm is described in Algorithm 3, which updates the *UserUpdate* function of Algorithm 1. As depicted in Algorithm 3, in each FL round, after user i receives the latest global model \mathbf{w}_i^t from the server, they sample their defense mask $\hat{\mathbf{m}}_i$ using Equation 18. The defense mask $\hat{\mathbf{m}}_i$ is then employed to combine the global model \mathbf{w}^t and the user's locally saved $\tilde{\mathbf{w}}_i^{t-1}$ to generate the local initial model $\mathbf{w}_{i,init}^t$. Following this composition step, the user initiates local training based on Equation 16, leading to updates in $\hat{\mathbf{w}}_i^t$, $\tilde{\mathbf{w}}_i^t$, and α_i^t .

Algorithm 3 *PriPrune*

```

1: function UserUpdate ( $\mathbf{w}^t, \mathbf{m}^t, i$ ):
2: Load  $\tilde{\mathbf{w}}_i^t$  from local storage of user  $i$ 
3: Initialize  $\boldsymbol{\alpha}_i^t = \boldsymbol{\alpha}_i^{t-1}$ 
4: for local epoch  $e \in E$  do
5:   Sample defense mask  $\hat{\mathbf{m}}_i^t$  by Equation 18 with  $\boldsymbol{\alpha}_i^t$ 
6:   Combine local and global model by:
7:   
$$\mathbf{w}_i^t[j]_{j \in d} = \begin{cases} \tilde{\mathbf{w}}_i^t[j] & \text{if } \hat{\mathbf{m}}_i^t[j] = 0 \\ \mathbf{w}_i^t[j] & \text{if } \hat{\mathbf{m}}_i^t[j] = 1 \end{cases}$$

8:   Base Pruning:  $\bar{\mathbf{w}}_i^t \leftarrow \mathbf{w}_i^t \odot \tilde{\mathbf{m}}_i^t$ 
9:    $\mathcal{L}_{local} \leftarrow \text{Equation16.}$ 
10:  Model parameter update:
11:  
$$\bar{\mathbf{w}}_i^t \leftarrow \bar{\mathbf{w}}_i^t - \eta \nabla \mathcal{L}_{local}(\bar{\mathbf{w}}_i^t, \boldsymbol{\alpha}_i^t)$$

12:  Defense mask parameter update:
13:  
$$\boldsymbol{\alpha}_i^t \leftarrow \boldsymbol{\alpha}_i^t - \eta \nabla \mathcal{L}_{local}(\bar{\mathbf{w}}_i^t, \boldsymbol{\alpha}_i^t)$$

14: end for
15:  $\hat{\mathbf{m}}_i^t \leftarrow \text{argmax}(\boldsymbol{\alpha}_i^t, 1 - \boldsymbol{\alpha}_i^t)$ 
16: Defense Pruning:  $\hat{\mathbf{w}}_i^t \leftarrow \bar{\mathbf{w}}_i^t \odot \hat{\mathbf{m}}_i^t; \tilde{\mathbf{w}}_i^t \leftarrow \bar{\mathbf{w}}_i^t \odot \tilde{\mathbf{m}}_i^t$ 
17: Return  $\hat{\mathbf{w}}_i^t$  to server, locally save  $\tilde{\mathbf{w}}_i^t$ 

```

Then user i conducts defense pruning, retaining the non-shared weights $\tilde{\mathbf{w}}_i^t$ locally, and transmits the shared weights $\hat{\mathbf{w}}_i^t$ to the server.

6 Experimental Evaluation

6.1 Experimental Setup

Datasets and Models. We evaluate models Conv-2 [5] and VGG-11 [31] along with their corresponding datasets FEMNIST [5] and CIFAR-10 [19]), which are commonly employed in FL studies.

Implementation details. In each round of training, we randomly select 10 clients from a pool of 193 users for the FEMNIST and 100 clients for the CIFAR-10. The training process involves 20,000 iterations, with a batch size of 20 for the FEMNIST and CIFAR-10. We employ 1 local training and initialize the learning rate at 0.25 with the base pruning rate set to 0.3. The details regarding the data split, resource and configurations are provided in Appendix B.1.

Performance Metrics. To assess the effectiveness of the evaluated pruning methods, we employ the privacy metrics and utility metrics. For privacy assessment, we use the Normalized Mutual Information (NMI) between the training data and the reconstructed data as the privacy metric. Higher values of NMI indicate a higher risk of privacy leakage. Additionally, the Peak Signal-to-Noise Ratio (PSNR) is employed as a well-established metric for image quality evaluation. The comparison of NMI and PSNR can be found in Fig. 2. See Appendix B.2 for additional evaluation results. The utility metric relies on model accuracy (ACC) to gauge model performance. Ideally, we want improved privacy, without significantly compromising accuracy.

Hyper-parameter Selection. We conducted a search for our main parameter, the defense rate \hat{p}_i , from 0.1 to 0.6, as depicted in Fig. 8c. With regards to our hyperparameters in loss functions, we search λ_{acc} in the list of [1, 5, 10], λ_{pri} in the list of [1, 10, 15] and λ_{sha} in the list of [2e-06, 2e-05, 2e-04]. Detailed results are deferred to Appendix C.3.




Fig. 8. Performance evaluation of *PriPrune* with adaptive defense rate and *Largest (Pseudo)* with 6 fixed defense rates, as defense strategies, integrated with *PruneFL* (as the base pruning method), conducted over 20,000 training rounds across the FEMNIST dataset. (a) Comparison of ACC over training rounds under these different defense strategies. (b) Comparison of Privacy (NMI) over training rounds under these different defense strategies. (c) The Privacy-Utility tradeoff, measured in terms of NMI and ACC, respectively. (d) Change in the adaptive defense rate of *PriPrune* over training rounds.

6.2 Performance of *PriPrune*

Recall that in Section 5.1, Insight 1, we initially assessed three defense strategies atop the base pruning method, namely: *Random*, *Largest*, *Random and Mix*. The results in Figure 6 indicate that while these defense strategies enhance privacy, they also degrade utility. Consequently, in Insight 2, we introduce *Pseudo-Pruning* to optimize this tradeoff. We evaluate the aforementioned three defense strategies both with and without *Pseudo-Pruning*, revealing that their integration with *Pseudo-Pruning* leads to improved accuracy and enhanced privacy. Notably, the *Largest + Pseudo-Pruning* strategy achieves the optimal balance between privacy and utility, thus emerging as our preferred defense mechanism. In Insight 3, we tackle the challenge of a fixed defense rate by proposing *PriPrune*, which incorporates *Largest + Pseudo-Pruning* (*Largest (Pseudo)*) with an adaptive defense rate. The evaluation of the privacy improvement achieved by each individual insight, in Section 5.1, essentially provides an ablation study. In this section, we conduct comprehensive experiments to meticulously evaluate the performance of the entire *PriPrune* across various pruning mechanisms and datasets.

First, we present a performance comparison between the *Largest (Pseudo)* with 6 fixed defense rates and *PriPrune* with the adaptive defense rate, both are integrated with *PruneFL* (as the base pruning method).

Comparable Utility. Fig. 8a illustrates how the utility metrics, ACC, change over FL rounds. It shows the convergence curve of *PriPrune*, with its model accuracy closely aligned with that of *PruneFL*. This alignment




Fig. 9. Comparison of tradeoffs between privacy-vs-accuracy among four base pruning mechanisms, both with and without the integration of *PriPrune*, across the FEMNIST and CIFAR-10 datasets.

indicates that *PriPrune* maintains utility comparable to PruneFL. However, for the *Largest (Pseudo)* defense, increasing its fixed defense rate from 10% to 60% results in decreased utility, rendering it unable to match the ACC achieved by PruneFL.

Enhanced Privacy. Figure 8b illustrates the change in the privacy metric NMI across FL rounds. It shows that *PriPrune* consistently maintains a low NMI throughout all FL training rounds, significantly outperforming PruneFL. Moreover, its NMI is comparable to that of *Largest (Pseudo)* with a fixed defense rate of 60%.

Optimal Tradeoff. Figure 8c illustrates the comparison of the privacy and utility tradeoffs among the mentioned defenses. It shows that increasing the defense rate of *Largest (Pseudo)* enhances privacy protection but at the expense of model accuracy. Conversely, *PriPrune* achieves the optimal tradeoff between privacy and accuracy. This is attributed to the adaptive nature of *PriPrune*'s defense rate \hat{p} , which adjusts to the model weights and is optimized jointly for privacy and model accuracy. The dynamic change of the defense rate of *PriPrune* is depicted in Fig. 8d.

Communication Efficiency and Training Speed. We conduct experiments to compare communication and training in PruneFL with and without *PriPrune*; details are deferred to Appendix C.2. When *PriPrune* is applied, the number of model parameters is reduced by $32.20\% \pm 0.91\%$ for FEMNIST and by $15.51\% \pm 0.17\%$ for CIFAR-10, thus significantly reducing the communication cost between clients and the server. This benefit is achieved because of *Pseudo-Pruning*, which retains pseudo-pruned weights locally, but transmits fewer model updates. When *PriPrune* is applied, the training time is increased by $0.18s \pm 0.01s$ for FEMNIST and by $0.50s \pm 0.01s$ for CIFAR-10, this is a negligible increase in the training speed (due to updating the defense pruning rate).

We have demonstrated that *PriPrune* stands out as the best defense solution among other strategies. Subsequently, we present the privacy performance of *PriPrune* when integrated with various state-of-the-art base pruning schemes in FL, such as SNIP, SynFlow, FedDST, and PruneFL, across both the FEMNIST and CIFAR-10 datasets.

PriPrune with State-of-the-Art Pruned FL. Figure 9 shows the results of using *PriPrune* for defense after a number of state-of-the-art base pruning schemes. Across all base pruning methods, this integration consistently enhances privacy while preserving accuracy. This highlights *PriPrune*'s effectiveness and versatility in achieving an optimal tradeoff between privacy and accuracy in diverse scenarios. This outcome is primarily due to the nature of *PriPrune*, which not only actually prunes but also selectively retains certain critical weights for local training (through *Pseudo-Pruning*). This mechanism allows the model to maintain or even improve accuracy while

reducing the information leakage that could be exploited in privacy attacks, as also shown in Fig.6. Specifically, *PriPrune* focuses on *Pseudo-Pruning* weights with large gradients that are most likely to leak sensitive information and keep them for local training. This selective pruning can lead to a model that is more robust against privacy attacks without a significant loss in accuracy. Additionally, the adaptive pruning rate in *PriPrune* is optimized for the tradeoff between privacy and accuracy, allowing the model to balance these two objectives effectively.

Methods	FEMNIST		CIFAR-10	
	Base	Base + <i>PriPrune</i>	Base	Base + <i>PriPrune</i>
PruneFL	0.22 ± 0.008	0.12 ± 0.007	0.29 ± 0.009	0.24 ± 0.009
FedDST	0.21 ± 0.025	0.13 ± 0.007	0.25 ± 0.008	0.18 ± 0.006
SynFlow	0.22 ± 0.009	0.15 ± 0.005	0.25 ± 0.012	0.21 ± 0.018
SNIP	0.23 ± 0.012	0.13 ± 0.014	0.24 ± 0.004	0.22 ± 0.008

Table 2. Comparison of privacy levels (NMI) between Base and Base + *PriPrune* for FL pruning methods. A lower NMI indicates a reduced risk of privacy leakage.

While Fig. 9 compares the tradeoffs between privacy-vs-accuracy, Table 2 further elaborates on the privacy improvement achieved by using *PriPrune* as defense after any base scheme, reporting results over three experiments. Compared to the base approach alone, adding *PriPrune* as defense consistently reduces NMI, thus improving privacy. For instance, on the FEMNIST dataset, *PriPrune* improves the privacy of *PruneFL* by 45.5% without compromising accuracy; on the CIFAR-10 dataset, *PriPrune* improves privacy of *FedDST* by 28% while maintaining accuracy.

7 Conclusion

In this paper, we revisit federated learning with model pruning, from the point of view of privacy. First, we quantify the information leakage and privacy gain offered for model pruning in FL, both theoretically and experimentally. Inspired by insights obtained from our extensive privacy quantification, we design *PriPrune*—an adaptive privacy-preserving local pruning mechanism in FL, that is jointly optimized for privacy and model performance. Our proposed mechanism achieves the best tradeoff between privacy and accuracy across diverse pruning methods and datasets under privacy attacks. One direction for our future work is to combine our pruning-based defense with classic, orthogonal defenses in FL such as differential privacy and secure aggregation, to enhance privacy further.

Acknowledgment

Tianyue Chu and Nikolaos Laoutaris were supported by the MLEDGE project (REGAGE22e00052829516), funded by the Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU/PRTR. Mengwei Yang and Athina Markopoulou were supported by NSF awards 1956393 and 1900654, and a UC Noyce Foundation gift. This work started during a long-term visit of Tianyue to the ProperData project, also supported by NSF Award 1956393

References

[1] Philip Bachman, R Devon Hjelm, and William Buchwalter. 2019. Learning Representations by Maximizing Mutual Information Across Views. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc., Vancouver Convention Centre, Canada. https://proceedings.neurips.cc/paper_files/paper/2019/file/ddf354219aac374f1d40b7e760ee5bb7-Paper.pdf

- [2] Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. 2018. Mutual Information Neural Estimation. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholmsmässan, Stockholm Sweden, 531–540. <https://proceedings.mlr.press/v80/belghazi18a.html>
- [3] Sameer Bibikar, Haris Vikalo, Zhangyang Wang, and Xiaohan Chen. 2022. Federated Dynamic Sparse Training: Computing Less, Communicating Less, Yet Learning Better. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 6 (6 2022), 6080–6088. <https://doi.org/10.1609/aaai.v36i6.20555>
- [4] 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 *CCS '17: 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>
- [5] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. Leaf: A Benchmark for Federated Settings.
- [6] Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020. The lottery ticket hypothesis for pre-trained bert networks. *Advances in neural information processing systems (NeurIPS) 33* (2020), 15834–15846.
- [7] Min Du, Ruoxi Jia, and Dawn Song. 2020. Robust anomaly detection and backdoor attack detection via differential privacy. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, Addis Ababa, Ethiopia. <https://openreview.net/forum?id=SJx0q1rtvS>
- [8] Cynthia Dwork. 2006. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II* 33. Springer, Berlin, Heidelberg, 1–12.
- [9] Ahmed Roushdy Elkordy, Jiang Zhang, Yahya H. Ezzeldin, Konstantinos Psounis, and Salman Avestimehr. 2023. How Much Privacy Does Federated Learning with Secure Aggregation Guarantee? *Proceedings on Privacy Enhancing Technologies (PoPETs) 2023*, 1 (2023), 510–526. <https://doi.org/10.56553/POPETS-2023-0030>
- [10] Jonathan Frankle and Michael Carbin. 2019. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, New Orleans, LA, USA. <https://openreview.net/forum?id=rJl-b3RcF7>
- [11] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems (NeurIPS) 33* (2020), 16937–16947.
- [12] Antonious Giris, Deepesh Data, Suhas Diggavi, Peter Kairouz, and Ananda Theertha Suresh. 2021. Shuffled Model of Differential Privacy in Federated Learning. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*, Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, Virtual Event, 2521–2529. <https://proceedings.mlr.press/v130/giris21a.html>
- [13] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Network with Pruning, Trained Quantization and Huffman Coding. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). OpenReview.net, San Juan, Puerto Rico. <http://arxiv.org/abs/1510.00149>
- [14] Song Han, Jeff Pool, John Tran, and William J. Dally. 2015. Learning both weights and connections for efficient neural networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1* (Montreal, Canada) (NIPS'15). MIT Press, Cambridge, MA, USA, 1135–1143.
- [15] Alexander Herzog, Robbie Southam, Ioannis Mavromatis, and Aftab Khan. 2024. FedMap: Iterative Magnitude-Based Pruning for Communication-Efficient Federated Learning.
- [16] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations*. OpenReview.net, Toulon, France. <https://openreview.net/forum?id=rkE3y85e>
- [17] Yuang Jiang, Shiqiang Wang, Victor Valls, Bong Jun Ko, Wei-Han Lee, Kin K. Leung, and Leandros Tassiulas. 2023. Model Pruning Enables Efficient Federated Learning on Edge Devices. *IEEE Trans. Neural Networks Learn. Syst.* 34, 12 (2023), 10374–10386. <https://doi.org/10.1109/TNNLS.2022.3166101>
- [18] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. *Learning Multiple Layers of Features from Tiny Images*. Technical Report. University of Toronto, Toronto, ON, Canada. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [20] Yann LeCun, John S. Denker, and Sara A. Solla. 1989. Optimal Brain Damage. In *Advances in Neural Information Processing Systems 2, [NIPS Conference, Denver, Colorado, USA, November 27-30, 1989]*, David S. Touretzky (Ed.). Morgan Kaufmann, Denver, Colorado, USA, 598–605. <http://papers.nips.cc/paper/250-optimal-brain-damage>

[21] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. 2019. Snip: single-Shot Network Pruning based on Connection sensitivity. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, Louisiana, USA, May 6-9, 2019*. OpenReview.net, New Orleans, Louisiana. <https://openreview.net/forum?id=B1VZqjAcYX>

[22] Ang Li, Jingwei Sun, Binghui Wang, Lin Duan, Sicheng Li, Yiran Chen, and Hai Li. 2020. Lotteryfl: Personalized and communication-efficient federated learning with lottery ticket hypothesis on non-iid datasets. arXiv:2008.03371

[23] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* 37, 3 (2020), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>

[24] Xiaolong Ma, Geng Yuan, Xuan Shen, Tianlong Chen, Xuxi Chen, Xiaohan Chen, Ning Liu, Minghai Qin, Sijia Liu, Zhangyang Wang, et al. 2021. Sanity checks for lottery tickets: Does your winning ticket really win the jackpot? *Advances in Neural Information Processing Systems (NeurIPS) 34* (2021), 12749–12760.

[25] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, Toulon, France. <https://openreview.net/forum?id=S1jE5L5gl>

[26] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Xiaojin (Jerry) Zhu (Eds.). PMLR, Fort Lauderdale, FL, USA, 1273–1282. <http://proceedings.mlr.press/v54/mcmahan17a.html>

[27] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, Vancouver, BC, Canada. <https://openreview.net/forum?id=Bj0hF1Z0b>

[28] Jun-Hyung Park, Yeachan Kim, Junho Kim, Joon-Young Choi, and SangKeun Lee. 2023. Dynamic Structure Pruning for Compressing CNNs. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, Washington, DC, USA, 9408–9416. <https://doi.org/10.1609/AAAI.V37I8.26127>

[29] Pytorch. n.d. GUMBEL SOFTMAX In Pytorch documentation. https://pytorch.org/docs/stable/generated/torch.nn.functional.gumbel_softmax.html. [Online].

[30] Yifan Shi, Kang Wei, Li Shen, Jun Li, Xueqian Wang, Bo Yuan, and Song Guo. 2024. Efficient Federated Learning With Enhanced Privacy via Lottery Ticket Pruning in Edge Computing. *IEEE Transactions on Mobile Computing* 23, 10 (Feb. 2024), 9946–9958. <https://doi.org/10.1109/TMC.2024.3370967>

[31] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). OpenReview.net, San Diego, CA, USA. <http://arxiv.org/abs/1409.1556>

[32] Jinhyun So, Ramy E. Ali, Basak Güler, Jiantao Jiao, and Amir Salman Avestimehr. 2023. Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI Washington, DC, USA, February 7-14, 2023*, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, Washington, DC, USA, 9864–9873. <https://doi.org/10.1609/AAAI.V37I8.26177>

[33] Jinhyun So, Chaoyang He, Chien-Sheng Yang, Songze Li, Qian Yu, Ramy E Ali, Basak Guler, and Salman Avestimehr. 2022. Lightsecagg: a lightweight and versatile design for secure aggregation in federated learning. *Proceedings of Machine Learning and Systems* 4 (2022), 694–720.

[34] Hidenori Tanaka, Daniel Kunin, Daniel L. K. Yamins, and Surya Ganguli. 2020. Pruning neural networks without any data by iteratively conserving synaptic flow. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 535, 13 pages.

[35] MTCAJ Thomas and A Thomas Joy. 2006. Elements of information theory.

[36] Aleksei Triastcyn and Boi Faltings. 2019. Federated Learning with Bayesian Differential Privacy. In *2019 IEEE International Conference on Big Data (IEEE BigData)*, Chaitanya K. Baru, Jun Huan, Latifur Khan, Xiaohua Hu, Ronay Ak, Yuanyuan Tian, Roger S. Barga, Carlo Zaniolo, Kisung Lee, and Yanfang (Fanny) Ye (Eds.). IEEE, Los Angeles, CA, USA, 2587–2596. <https://doi.org/10.1109/BIGDATA47090.2019.9005465>

[37] Georgia Tsaloli, Bei Liang, Carlo Brunetta, Gustavo Banegas, and Aikaterini Mitrokotsa. 2021. sf DEVA: Decentralized, Verifiable Secure Aggregation for Privacy-Preserving Learning. In *Information Security - 24th International Conference, ISC 2021, Virtual Event, November 10-12, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 13118)*, Joseph K. Liu, Sokratis K. Katsikas, Weizhi Meng, Willy Susilo, and Rolly Intan (Eds.). Springer, Virtual Event, 296–319. https://doi.org/10.1007/978-3-030-91356-4_16

[38] Chaoqi Wang, Guodong Zhang, and Roger B. Grosse. 2020. Picking Winning Tickets Before Training by Preserving Gradient Flow. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, Addis Ababa, Ethiopia. <https://openreview.net/forum?id=SkgsACVKPH>

- [39] Danye Wu, Miao Pan, Zhiwei Xu, Yujun Zhang, and Zhu Han. 2020. Towards Efficient Secure Aggregation for Model Update in Federated Learning. In *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*. IEEE, Taipei, Taiwan, 1–6. <https://doi.org/10.1109/GLOBECOM42002.2020.9347960>
- [40] Xiyuan Yang, Wenke Huang, and Mang Ye. 2023. Dynamic personalized federated learning with adaptive differential privacy. *Advances in Neural Information Processing Systems* 36 (2023), 72181–72192.
- [41] Hongxu Yin, Arun Mallya, Arash Vahdat, Jose M. Alvarez, Jan Kautz, and Pavlo Molchanov. 2021. See through Gradients: Image Batch Recovery via GradInversion. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 16332–16341. <https://doi.org/10.1109/CVPR46437.2021.01607>
- [42] Kai Yue, Richeng Jin, Chau-Wai Wong, Dror Baron, and Huaiyu Dai. 2023. Gradient Obfuscation Gives a False Sense of Security in Federated Learning. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 6381–6398. <https://www.usenix.org/conference/usenixsecurity23/presentation/yue>
- [43] Xinwei Zhang, Xiangyi Chen, Mingyi Hong, Zhiwei Steven Wu, and Jinfeng Yi. 2022. Understanding clipping for federated learning: Convergence and client-level differential privacy. In *International Conference on Machine Learning, ICML 2022*, Vol. 162. PMLR, Baltimore, Maryland, USA, 26048–26067.
- [44] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. idlg: Improved deep leakage from gradients. [arXiv:2001.02610 \[cs.LG\]](https://arxiv.org/abs/2001.02610)
- [45] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. *Deep leakage from gradients*. Curran Associates Inc., Red Hook, NY, USA, 11.
- [46] Zhanxing Zhu, Jingfeng Wu, Bing Yu, Lei Wu, and Jinwen Ma. 2019. The Anisotropic Noise in Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regularization Effects. *Proceedings of the 36th International Conference on Machine Learning* 97 (09–15 Jun 2019), 7654–7663. <https://proceedings.mlr.press/v97/zhu19e.html>

Appendix

This appendix extends the main paper, providing supplemental materials, including additional details and results, which could not be included in the main paper, due to lack of space.

A Proof of Theorem

This section provides the proof of Theorem 4.1 in Section 4.1.

A.1 Proof

Equation 2 is the definition of mutual information, which is:

$$\begin{aligned}
 I(X; Y) &= \int_{X \times Y} \log \frac{d\mathbb{P}_{X,Y}(x, y)}{d\mathbb{P}_X(x) \otimes \mathbb{P}_Y(y)} d\mathbb{P}_{X,Y}(x, y) \\
 &= H(X) - H(X|Y) \\
 &= H(X) + H(Y) - H(X, Y)
 \end{aligned}$$

Hence, Equation 3 can be written as:

$$\begin{aligned}
\mathbf{I}\left(\mathcal{D}_i; \{\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t\}_{t \in [T]}\right) &\stackrel{(a)}{\leq} \sum_{t=1}^T \mathbf{I}\left(\mathcal{D}_i; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \\
&\stackrel{(b)}{=} \sum_{t=1}^T \left(\mathbf{I}\left(\mathbf{w}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) + \sum_{t=1}^T \left(\mathbf{I}\left(\mathcal{D}_i; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \mathbf{w}_i^t, \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) \\
&\quad - \sum_{t=1}^T \left(\mathbf{I}\left(\mathcal{D}_i; \mathbf{w}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t]}\right) \right) \\
&\stackrel{(c)}{\leq} \sum_{t=1}^T \left(\mathbf{I}\left(\mathbf{w}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) \\
&= \sum_{t=1}^T \left(\mathbf{H}\left(\mathbf{w}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) + \sum_{t=1}^T \left(\mathbf{H}\left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) \\
&\quad - \sum_{t=1}^T \left(\mathbf{H}\left(\mathbf{w}_i^t, \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right) \\
&\stackrel{(d)}{=} \sum_{t=1}^T \left(\mathbf{I}\left(\mathbf{x}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \right)
\end{aligned} \tag{19}$$

Where (a) and (b) are from the chain rule; (c) is from data processing inequality $\mathcal{D}_i \rightarrow \mathbf{w}_i^t \rightarrow \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t$ that $\mathbf{I}\left(\mathcal{D}_i; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \mathbf{w}_i^t, \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) = 0$; (d) is from zero conditional entropy and invariant of mutual information, where

$$\mathbf{x}_i^t = \frac{\partial \ell_i(\bar{\mathbf{w}}_i^{t-1}; \mathcal{D}_i)}{\partial \mathbf{w}} = \frac{1}{B} \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b) \tag{20}$$

b denotes the size of the random samples.

Therefore, the privacy leakage for a single round t is Equation 5:

$$\mathbf{I}\left(\mathbf{x}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right)$$

We first prove the upper bound for Equation 5.

Based on Equation 2, we have

$$\begin{aligned}
\mathbf{I}\left(\mathbf{x}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) &= \mathbf{H}\left(\mathbf{x}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) + \mathbf{H}\left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \\
&\quad - \mathbf{H}\left(\mathbf{x}_i^t, \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) \\
&\leq \underbrace{\mathbf{H}\left(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right)}_{\mathbf{A}} + \underbrace{\mathbf{H}\left(\mathbf{x}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right)}_{\mathbf{B}}
\end{aligned}$$

For part A, based on the chain rule of entropy, we have

$$\begin{aligned} 0 \leq H(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t, \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) &= H(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) + H(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \mathbf{w}_i^t, \bar{\mathbf{m}}_i^t, \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) \\ &= H(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) + H(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t]}) \end{aligned} \quad (21)$$

Due to $H(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \mathbf{w}_i^t, \bar{\mathbf{m}}_i^t, \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) = 0$, and $H(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t]}) \geq 0$, we have

$$\begin{aligned} H(\mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) &\leq H(\mathbf{w}_i^t, \bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) \\ &\stackrel{(a)}{\leq} H(\mathbf{w}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) + H(\bar{\mathbf{m}}_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) \\ &\stackrel{(b)}{=} H(x_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) + 1 - \frac{1}{2 \ln 2} \sum_{n=1}^{\infty} \frac{(2\bar{p}_i - 1)^{2n}}{n(2n-1)} \\ &\leq \underbrace{H(x_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]})}_{B} + 1 - \frac{\bar{p}_i - 1}{2 \ln 2} \end{aligned}$$

(a) is from conditioning reduces entropy; (b) is from zero conditional entropy and the Taylor series of the binary entropy function in a neighbourhood of 0.5 with the base pruning rate \bar{p}_i .

For part B,

$$H(x_i^t \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}) = H\left(\frac{1}{B} \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b) \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right)$$

Let $X = \sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b)$, $Z = \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}$, $Y = \frac{1}{B}X$. $f_{X,Z}(x, z)$ and $f_{Y,Z}(y, z)$ are the corresponding joint probability density functions.

Then part B can be written as

$$\begin{aligned} H(Y \mid Z) &= - \int_{Y, Z} f_{Y,Z}(y, z) \log \frac{f_{Y,Z}(y, z)}{f_Z(z)} dy dz = - \int_{X, Z} |B| f_{X,Z}(By, z) \log |B| \frac{f_{X,Z}(By, z)}{f_Z(z)} dy dz \\ &= - \int_{X, Z} f_{X,Z}(x, z) \log \frac{f_{X,Z}(x, z)}{f_Z(z)} dx dz + \log \frac{1}{|B|} = H\left(\underbrace{\sum_{b \in \mathbf{B}_i^t} g_i(\mathbf{w}_i^t, b)}_{C} \mid \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]}\right) + \log \frac{1}{B} \end{aligned}$$

In recent theoretical results for analyzing the behaviour of SGD, they approximate the SGD vector by a distribution with independent components or by a multivariate Gaussian vector.

We define $\hat{g}_i(\mathbf{w}_i^t, b) \in \mathbb{R}^{d^*}$ is the largest sub-vector of the $g_i(\mathbf{w}_i^t, b)$ with non-singular covariance matrices, where $d^* \leq d$. Based on the ZCA whitening transformation and Assumption 1, we have $\hat{g}_i(\mathbf{w}_i^t, b) = \Sigma_i^{-\frac{1}{2}} \mathbf{z}$, where \mathbf{z} has zero mean and \mathbb{I}_{d^*} covariance matrix.

Then for the part C, we set $\mathbf{v} = \sum_{b \in \mathbf{B}_i^t} \mathbf{z}$ and $\mathbf{u} = \Sigma_i^{-\frac{1}{2}} \mathbf{v}$, then we have:

$$\begin{aligned} \mathbf{C} &= \mathbf{H} \left(\sum_{b \in \mathbf{B}_i^t} \hat{g}_i(\mathbf{w}_i^t, b) \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) = \mathbf{H} \left(\Sigma_i^{-\frac{1}{2}} \sum_{b \in \mathbf{B}_i^t} \mathbf{z} \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) = - \int \mathbb{P}_{\mathcal{U}}(\mathbf{u}) \log \mathbb{P}_{\mathcal{U}}(\mathbf{u}) d\mathbf{u} \\ &\stackrel{(a)}{=} - \int \frac{\mathbb{P}_{\mathcal{V}}(\mathbf{v})}{\left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right|} \log \left(\frac{\mathbb{P}_{\mathcal{V}}(\mathbf{v})}{\left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right|} \right) \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| d\mathbf{v} = \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + \underbrace{\mathbf{H} \left(\sum_{b \in \mathbf{B}_i^t} \mathbf{z} \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right)}_{\mathbf{D}} \end{aligned} \quad (22)$$

(a) is based on the transformation of random vectors. Based on the Maximum Entropy Upper Bound [35], which the continuous distribution \mathbf{X} with prescribed variance $\mathcal{V}(\mathbf{X})$ maximizing the entropy is the Gaussian distribution of same variance. Since the entropy of a multivariate Gaussian distribution with mean μ and covariance $K_{i,j}$ is

$$\mathbf{h}(\mathcal{N}_n(\mu, K)) = \frac{1}{2} \log(2\pi e)^n |K|$$

where $|K|$ denotes the determinant of K . Hence, the maximum entropy upper bound for part \mathbf{D} is

$$\mathbf{H} \left(\sum_{b \in \mathbf{B}_i^t} \mathbf{z} \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) \leq \frac{1}{2} \log(2\pi e)^{d^*} |\mathbb{I}_{d^*}| = \frac{d^*}{2} \log(2\pi e) \quad (23)$$

Therefore, by combining Equation 22 and Equation 23, the upper bound for part \mathbf{B} is

$$\mathbf{H} \left(\mathbf{x}_i^t \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) \leq \log \frac{1}{B} + \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + \frac{d^*}{2} \log(2\pi e) \quad (24)$$

After summing part \mathbf{A} and part \mathbf{B} , we derive the upper bound for the privacy leakage for the single round, which is Equation 5:

$$\mathbf{I} \left(\mathbf{x}_i^t; \mathbf{w}_i^t \odot \bar{\mathbf{m}}_i^t \left| \{\mathbf{w}_i^k \odot \bar{\mathbf{m}}_i^k\}_{k \in [t-1]} \right. \right) \leq 1 - \frac{\bar{p}_i - 1}{2 \ln 2} + 2 \log \frac{1}{B} + 2 \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + d^* \log(2\pi e)$$

Based on Equation 4, we have the upper bound for objective function Equation 3, standing how much information the aggregated model over T global training rounds could leak about the private data,

$$\mathbf{I}_i \leq T - \frac{T(\bar{p}_i - 1)}{2 \ln 2} + 2T \log \frac{1}{B} + 2T \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + d^* T \log(2\pi e)$$

For PruneFL, the upper bound for the single-round leakage is

$$\begin{aligned} \mathbf{I}_i^t &\leq 1 + \frac{\bar{\mathcal{P}} \cup \mathcal{A}}{2d \ln 2} + 2 \log \frac{1}{B} + 2 \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + d^* \log(2\pi e) \\ &\leq 1 + \frac{\bar{\mathcal{P}}}{2d \ln 2} + 2 \log \frac{1}{B} + 2 \log \left| \det \left(\Sigma_i^{-\frac{1}{2}} \right) \right| + d^* \log(2\pi e) \end{aligned} \quad (25)$$




Fig. A1. Impact of varying pruning rate in FEMNIST on privacy leakage using key metrics: Normalized Mutual Information (NMI) displayed on the left y-axis and Peak Signal-to-Noise Ratio (PSNR) shown on the right y-axis under a Gradient Inversion attack for PruneFL.

B Privacy Attacks

B.1 Extra Implementation Details

Our algorithms are implemented by Pytorch and we implement experiments on two NVIDIA RTX A5000 and two Xeon Silver 4316. Across all the examined methods, we employ 1 local training and initialize the learning rate to 0.25 and The baseline pruning rate is set at 0.3 for all datasets.

Data Partitioning. Our data partitioning methodology is aligned with the FEMNIST dataset setting, utilizing its inherent 193-user partition. As for the CIFAR-10 dataset, we divided it into 100 equal-sized, non-overlapping users, following the methodology used in PruneFL [17].

B.2 Evaluation on both NMI and PSNR Metrics

To quantify the extent of privacy leakage, we utilized the Normalized Mutual Information (NMI) metric, which has been demonstrated in prior research to align well with the Peak Signal-to-Noise Ratio (PSNR) [9] as an effective measure of privacy leakage.

We examine the relationship between NMI and PSNR. The observed trend, depicted in Fig. A1, illustrates a consistent decrease in both NMI and PSNR metrics with higher pruning rates. This result is consistent with previous research, confirming the reliability of NMI as a measure of privacy leakage, which is comparable to the widely accepted PSNR metric.

B.3 Image Reconstructions under Privacy Attacks

Figure A2 shows the recovered images under the SGI (Sparse Gradient Inversion) attack in the FEMNIST dataset, for different base pruning methods and base pruning rates.

C Defense Details

C.1 Results of Insights

Within the scope of the three defense methods in the Insights from Evaluation Section, specific configurations have been established. These settings are detailed below. For all three defense methods, the total pruning rate (\hat{p}) is uniformly set at 0.3, mirroring the original PruneFL's pruning rate.

For each defense method:




Fig. A2. Impact of varying pruning rate on image reconstruction in FEMNIST using different base pruning methods.




Fig. A3. Comparison of Privacy-Accuracy Trade-Off with Varying Defense Rate (0% to 50%) in Largest Method based on PruneFL using NMI and Attack Loss Privacy Metrics

- PruneFL + Largest: The defense strategy involves pruning based on the weights with top- k largest gradients, with the defense pruning rate set to $\hat{p}_{largest} = 0.3$.
- PruneFL + Random: The defense strategy involves random pruning, with the defense pruning rate (\hat{p}) set to $\hat{p}_{random} = 0.3$.
- PruneFL + Mix: This is a hybrid approach that combines both the largest gradient-based pruning and random pruning. The defense pruning rate is determined as the sum of $\hat{p}_{largest} = 0.15$ and $\hat{p}_{random} = 0.15$.

Figure A3 shows Privacy-Accuracy Trade-Off with Varying Defense Rate (0% to 50%) in Largest Method based on PruneFL using NMI and Attack Loss Privacy Metrics.

The outcomes are indicative of a notable trend: as more weights associated with larger gradients are pruned, privacy is enhanced at the expense of accuracy. This substantiates our **Insight 1**: Largest Weights Matter. Pruning weights with large gradients improves privacy the most but also hurts model accuracy the most.

C.2 Communication Efficiency and Training Speed

Figure A4 demonstrates the communication cost in terms of the number of model parameters across 1000 iterations of FL. In *PriPrune*, we employ *Pseudo-Pruning*, where the pseudo-pruned weights are retained locally for subsequent training rounds. Thus, *PriPrune* directly reduces communication costs by transmitting fewer




Fig. A4. Comparison of the communication costs of PruneFL with and without *PriPrune* during the training process with (a) representing FL training using the Conv-2 model on the FEMNIST dataset and (b) the VGG-11 model on the CIFAR-10 dataset.




Fig. A5. Comparison of the training speed of PruneFL with and without *PriPrune* during the training process with (a) representing FL training on the FEMNIST dataset and (b) on the CIFAR-10 dataset.

model parameters. Figure A4 shows a significant reduction in the number of model parameters communicated between clients and the server with the application of *PriPrune*.

Figure A5 displays the training speed access through the total time of clients' local training across 1000 iterations of FL. *PriPrune* involves additional steps of updating the defense pruning rate, which is computationally lightweight. Figure A5 demonstrates that *PriPrune* introduces only a minor overhead and the additional steps brought by *PriPrune* has a minimal impact on training speed.

C.3 Hyper-parameter Selection

Table A1 has searched λ_{acc} in the list of [1, 5, 10], λ_{pri} in the list of [1, 10, 15] and λ_{sha} in the list of [2e-06, 2e-05, 2e-04], which showing the comparisons of privacy levels during the hyper-parameter search. As shown in Table A1, when we increase the λ_{pri} , the NMI value becomes lower, indicating better privacy protection. We explore different combinations of hyper-parameters and finally utilize the best set of values that could achieve the best privacy-accuracy tradeoff.

For different pruning methods and different datasets, we are adopting different combinations of hyper-parameters so as to achieve better privacy-utility tradeoff, as detailed in Table A2.

Parameter	$\lambda_{pri} = 1$			$\lambda_{pri} = 10$			$\lambda_{pri} = 15$		
	$\lambda_{sha} = 2e-04$	$\lambda_{sha} = 2e-05$	$\lambda_{sha} = 2e-06$	$\lambda_{sha} = 2e-04$	$\lambda_{sha} = 2e-05$	$\lambda_{sha} = 2e-06$	$\lambda_{sha} = 2e-04$	$\lambda_{sha} = 2e-05$	$\lambda_{sha} = 2e-06$
$\lambda_{acc} = 1$	0.166	0.171	0.156	0.138	0.153	0.154	0.138	0.159	0.151
$\lambda_{acc} = 5$	0.154	0.158	0.136	0.125	0.121	0.123	0.135	0.123	0.130
$\lambda_{acc} = 10$	0.148	0.142	0.140	0.132	0.126	0.118	0.135	0.132	0.140

Table A1. Comparison of privacy levels (NMI) during hyper-parameter selection: with regards to our hyperparameters in loss functions, we search λ_{acc} in the list of [1, 5, 10], λ_{pri} in the list of [1, 10, 15] and λ_{sha} in the list of [2e-06, 2e-05, 2e-04].

	Dataset	Method	λ_{pri}	λ_{sha}	λ_{acc}
FEMNIST	PruneFL+ <i>PriPrune</i>	15	2×10^{-5}	5	
	Synflow+ <i>PriPrune</i>	1	2×10^{-6}	10	
	SNIP+ <i>PriPrune</i>	1	2×10^{-6}	10	
	FedDST+ <i>PriPrune</i>	8	2×10^{-5}	10	
CIFAR-10	PruneFL+ <i>PriPrune</i>	10	2×10^{-5}	5	
	Synflow+ <i>PriPrune</i>	1	2×10^{-6}	5	
	SNIP+ <i>PriPrune</i>	1	2×10^{-6}	10	
	FedDST+ <i>PriPrune</i>	10	2×10^{-5}	10	

Table A2. Parameter settings for λ_{pri} , λ_{sha} , and λ_{acc} in FEMNIST and CIFAR-10 datasets.

Received 29 April 2024; revised 26 August 2024; accepted 7 October 2024