

# Eluding Secure Aggregation in Federated Learning via Model Inconsistency\*

Dario Pasquini  
dario.pasquini@epfl.ch  
SPRING Lab, EPFL  
Lausanne, Switzerland

Danilo Francati  
dfrancati@cs.au.dk  
Aarhus University  
Aarhus, Denmark

Giuseppe Ateniese  
ateniese@gmu.edu  
George Mason University  
Fairfax, Virginia, USA

## ABSTRACT

Secure aggregation is a cryptographic protocol that securely computes the aggregation of its inputs. It is pivotal in keeping model updates private in federated learning. Indeed, the use of secure aggregation prevents the server from learning the value and the source of the individual model updates provided by the users, hampering inference and data attribution attacks.

In this work, we show that a malicious server can easily elude secure aggregation as if the latter were not in place. We devise two different attacks capable of inferring information on individual private training datasets, independently of the number of users participating in the secure aggregation. This makes them concrete threats in large-scale, real-world federated learning applications.

The attacks are generic and equally effective regardless of the secure aggregation protocol used. They exploit a vulnerability of the federated learning protocol caused by incorrect usage of secure aggregation and lack of parameter validation. Our work demonstrates that current implementations of federated learning with secure aggregation offer only a “false sense of security”.

## KEYWORDS

Federated Learning, Secure Aggregation, Model Inconsistency

## 1 INTRODUCTION

Deep learning is evolving rapidly but often at the expense of privacy and security. Neural networks may misbehave, hide backdoors, or be reverse-engineered to reveal sensitive information about the training datasets [5, 23, 54]. Data holders are thus reluctant to provide and share their datasets unless some level of protection is in place.

Cryptographic primitives, such as multi-party computation (MPC) and fully homomorphic encryption (FHE), offer only a partial solution to this problem: They enable learning while protecting sensitive information but at the expense of efficiency and scalability. Even state-of-the-art implementations of these primitives are highly inefficient and add a significant overhead to the learning process, making them unusable and inapplicable in practice.

Accordingly, researchers have looked at alternative solutions that rely on decentralization, where data remain local with the participants while the neural network evolves during the distributed learning process. Along this line of research, **federated learning** (FL) [11, 35, 36], along with its main implementations federated stochastic gradient descent (FedSGD) and federated averaging (FedAVG), has been proposed. At a high level, FL allows a set of users

to train a shared neural network without outsourcing their local datasets. To this end, they are only required to locally train the neural network and send model updates (e.g., gradients, model parameters) to a central server. The updates will be aggregated by the server, completing a round of the training. The informal security guarantee offered by FL is that sharing the (possibly scrambled [2]) updates does not leak any information about the actual training instances used by the users. Unfortunately, it has been shown that an adversary can invert an individual model update of a target user in order to leak a large amount of information about its dataset [30, 43, 46, 70].

For this reason, Bonawitz et al. [10] have proposed to combine **secure aggregation** (SA) protocols with FL as a first step to increase the security of FL, preventing the server from accessing individual model updates. Informally, SA is a specialized MPC protocol that allows a set of users to compute the sum of their private inputs securely. The security guarantee is the same as standard MPC protocols, i.e., nothing is leaked about the inputs except what can be inferred from the output (the sum of the values).

SA is believed to be one of the most robust defenses against gradient inversion and related inference attacks [32]. In particular, the application of SA in FL has two main objectives: (1) “*Privacy by aggregation*”: Aggregating together a suitable number of model updates smooths out the information carried out by individual contributions. In turn, this makes it unfeasible to assert or recover meaningful information on individual training instances that produced the aggregated value. (2) “*Privacy by shuffling*”: SA “hides” the source of the aggregated information; even if sensitive data is recovered from the aggregated model updates, this cannot be attributed to the user (the individual model update) who provided it. Thus, although the privacy of the set of users may be violated, the privacy of individuals is preserved.

Our work shows that a motivated and malicious server can easily violate both of these fundamental properties of current SA defenses. This vulnerability emerges from the federated learning protocol, not the SA protocol; it is caused by incorrect usage of the SA protocol. For example, even if we abstract the SA protocol with an ideal aggregation functionality, the protocol is still exploitable. In this case, the failure to validate SA inputs by the user is one of the protocol’s weaknesses, and we are not targeting a specific implementation of SA. The main intuition is that model updates (i.e., the inputs of SA) are under the indirect control of the malicious server since model updates are computed starting from the parameters sent by the server. A malicious server can leverage this control to tamper with the updates (that are the inputs of SA) so that their aggregation will leak information about the update of a target user.

\*In the proceedings of ACM Conference on Computer and Communications Security 2022 (CCS ’22).

In order to achieve this, the server exploits a new attack vector that we call **model inconsistency**. Here, the server distributes different views of the same model to different users within the same round. In this work, we show that model inconsistency can introduce new vulnerabilities in FL algorithms. The intuition is that a malicious server, providing different parameters to different users, can exploit behavioral differences in the model updates provided by the different models to infer information on users’ datasets, even when those model updates are securely aggregated before reaching the server.

To make this inherent vulnerability evident, we implement two attacks that give a representative view of the threat induced by the model inconsistency attack vector. Eventually, these attacks demonstrate how a malicious server can nullify the security offered by current SA-based defenses proposed for FL. That is: (1) individual model updates can be perfectly recovered from the final aggregated value, independently of the number of users participating in the aggregation, and, (2) the source of the recovered data can be attributed to individuals in the pool of active users.

Finally, we introduce multiple strategies for preventing the vulnerability that was brought up. The proposed solutions seamlessly integrate with current state-of-the-art SA protocols without impacting performance or utility.

## 1.1 Contributions

Our contributions can be summarized as follows:

*Model inconsistency.* We demonstrate that FL incorrectly leverages SA [10], and as a result, it is as secure as the original FL protocol (without SA). We prove this by introducing a new adversarial strategy, named **model inconsistency**, that leverages the following two observations: (i) in each protocol round, the SA’s input of user  $u_i \in \mathcal{U}$  is its model update  $\Delta_{\mathcal{D}_i}^\Theta$  and, (ii) the value of  $\Delta_{\mathcal{D}_i}^\Theta$  of each  $u_i \in \mathcal{U}$  depends on the parameters  $\Theta$  sent by the server  $S$  (i.e., different parameters produce different model updates). The combination of the above two observations implies that  $S$  could act maliciously and craft different parameters for different users in order to tamper with the inputs of SA. As our two attacks will demonstrate, at a different scale,  $S$  can exclude from the aggregation the updates of some non-target users  $\mathcal{U} \setminus \{u_{\text{trgt}}\}$ , forcing the SA to leak part of the model update  $\Delta_{\mathcal{D}_{\text{trgt}}}^\Theta$  of the target  $u_{\text{trgt}}$ .

*Gradient suppression attack.* In Section 5, we present a first attack, named **gradient suppression**. It shows that a malicious server can force the local training of the deep model  $f_\Theta$  executed by  $u_i$  to unconditionally produce a zeroed gradient  $\Delta_{\mathcal{D}_i}^\Theta = [0]$ . By combining both gradient suppression and model inconsistency, the server  $S$  can send the honest parameters  $\Theta$  to the target user  $u_{\text{trgt}}$  and the malicious parameters  $\tilde{\Theta}$  to the remaining non-target ones  $\mathcal{U} \setminus \{u_{\text{trgt}}\}$ . In turn, this will leak the honest  $u_{\text{trgt}}$ ’s gradient  $\Delta_{\mathcal{D}_{\text{trgt}}}^\Theta$  even if SA is in place. This is because SA will sum up the zeroed gradients of  $\mathcal{U} \setminus \{u_{\text{trgt}}\}$  and the honest gradient of  $u_{\text{trgt}}$ . The output will be equal to the gradient  $\Delta_{\mathcal{D}_{\text{trgt}}}^\Theta$  of the latter user  $u_{\text{trgt}}$ . This is the first practical attack that demonstrates that a malicious server can completely nullify SA. More importantly, the attack does not

require auxiliary information on the targets, e.g., the distribution of users’ datasets, or unrealistic architecture alterations.

*Canary-gradient attack.* We extend the first attack and devise a second approach called **canary-gradient**. Here, we show that a malicious server  $S$  can modify the target’s model to induce specific behavior in the derivative of a tiny subset  $\xi$  of its parameters (e.g., two out of millions of parameters). In particular,  $S$  can forge malicious parameters that force the model to produce non-zero gradients for  $\xi$  only when a specific adversarially-chosen property is present in the input batch used to compute the update. Then, the server can preserve the target’s gradient for  $\xi$  in the final aggregated value by forcing the non-target users to unconditionally produce zero gradients only for the parameters  $\xi$ . This allows  $S$  to recover the target’s gradient for  $\xi$  in “plaintext” and ascertain the presence of the queried property in the user’s private data (e.g., membership inference). Eventually, this demonstrates that a malicious server can cast extremely effective property inference attacks on individual users under SA, while ensuring the stealthiness of the attack.

*The (in)correct usage of SA in FL.* In cryptographic terms, SA protocols are specialized multi-party computation (MPC) protocols that implement the ideal functionality  $f^{\text{sa}}(v_1, \dots, v_n) = \sum_{u_i \in \mathcal{U}} v_i = v$ . They are built assuming that the inputs of honest users are untamperable, i.e., an adversary has no control over the input  $v_i$  of an honest user  $u_i \in \mathcal{U}$ . This holds in both the semi-honest and malicious security models.

Unfortunately, the above assumption does not hold in FL. As discussed earlier, a malicious server  $S$  can tamper with the inputs  $(v_1, \dots, v_n)$  of the honest users  $\mathcal{U}$ . Unlike other attacks, ours is the first that does not contradict the security of the underlying SA protocol and does not require auxiliary information about user data. Our attacks will succeed regardless of the number of FL protocol users, which is additional evidence that SA and FL cannot protect against adversarial servers.

*Preventing model inconsistency.* We discuss ways to help mitigate model inconsistency by integrating consistency checks during the FL protocol. In particular, we show that the two most influential SA protocols of Bonawitz et al. [10] (CCS’17) and Bell et al. [8] (CCS’20) can be modified to incorporate consistency checks without affecting the efficiency of the original FL protocol. This is achieved by linking SA’s masking values (generated by a particular user to hide its input) to the parameters received from the server. By doing it this way, when two or more users receive different parameters, their masks will look random. As a result, the malicious server cannot tell whose input is whose, reducing the attack efficacy.

Finally, we will discuss how DP techniques can be used in conjunction with SA, which remains a suitable solution to prevent a malicious server from breaching users’ privacy. We refer the reader to Section 7 for more details.

To make our results reproducible, we made our code available.<sup>1</sup>

<sup>1</sup><https://github.com/pasquini-dario/EludingSecureAggregation>.

## 2 RELATED WORK

### 2.1 Federated learning and secure aggregation

The distributed architecture of FL protocol [40, 53] provides a fertile ground for attackers [20, 30, 43, 46]. This is because a malicious party, mainly the server, has access to sensitive information such as model updates that can be exploited to violate users' privacy. Accordingly, SA has been proposed by Bonawitz et al. [10] as a fundamental step to increase the security of FL without modifying the original structure of the protocol. Subsequent works focus on the development of new SA protocols for FL with reduced communication/computation overhead [8, 16, 26, 33, 56], multiple servers [7], increased robustness against malicious updates [13, 49], or with verifiable aggregation [26, 62].<sup>2</sup>

Separately, several other works focused on building new protocols (or propose significant modifications of FL, e.g., protocol, architecture, etc.) to train a deep neural network without leaking unnecessary information about the datasets of the users. We refer the reader to Appendix A for more details.

### 2.2 Gradient Inversion

A core privacy concern in FL is the role of the server. In this direction, it has been shown that, without SA, even a semi-honest server can invert users' gradients (sent as a model update in a particular round of FL) and compute a close-enough approximation of users' local training datasets. In a nutshell, by leveraging  $f_\Theta$  (where  $\Theta$  are the parameters of the current round of FL) and the gradient locally computed by the user  $u$  using a subset  $\mathcal{D}$  of its local data, the server can recover  $\mathcal{D}$  by searching a set of instances  $\widehat{\mathcal{D}}$  that generates the gradient similar to the one sent by the user. Thanks to the inherent smoothness of the neural network  $f_\Theta$ , this searching problem can be defined as a second-order optimization, i.e.,

$$\arg\min_{\widehat{\mathcal{D}}} [d(\nabla_{\widehat{\mathcal{D}}}^\Theta, \nabla_{\mathcal{D}}^\Theta) \cdot \alpha r(\widehat{\mathcal{D}})] \quad (1)$$

where  $\widehat{\mathcal{D}}$  is the candidate solution of the malicious server  $S$ ,  $d$  is a distance function to measure the discrepancy between the gradient signals  $\nabla_{\widehat{\mathcal{D}}}^\Theta$  and  $\nabla_{\mathcal{D}}^\Theta$ ,  $r$  is a regularizer defined on the input domain, and  $\alpha$  is the weight associated to the regularization term in the optimization. In the work of Zhu et al. [70],  $d$  is set to be the Euclidean distance, and the L-BFGS solver is used to solve the optimization problem. The follow-up work of Geiping et al. [25] improves their approach/results by noting that the gradient signal is scale-invariant and accounts for that in defining the optimization objective in Equation (1). Their work improves the effectiveness of the inversion attack, drastically increasing its applicability on real-world architectures such as ResNets [29] and more realistic batch sizes and a number of FedAVG local iterations. These results have been further improved in the work of Yin et al. [65] by relying on additional regularization terms and tailored optimization techniques.

A more recent line of research dispensed with optimization-based approaches to focus on closed-form procedures to recover

<sup>2</sup>In FL with verifiable aggregations [26, 62], users are supposed to verify the integrity of the aggregated model updates and compute the new parameters (if the parameters are updated solely by the server, then the verifiability of aggregated gradients becomes meaningless). This differs from the original FL protocol [40, 53] in which the server updates the model for improved scalability.

data from gradients in deep neural networks [48, 69]. In this vein, a recent work of Fowl et al. [22] (which is concurrent to our work) improves over previous approaches by considering a malicious server that modifies the FL architecture and crafts the network parameters to artificially create a neural layer that retains information on the input batch. In particular, this is a linear layer followed by a ReLU activation whose parameters must be chosen considering the private sets' CDF for a given property, i.e., the attacker must have some auxiliary information on users' training datasets. Unfortunately, this extra knowledge may not be acquired in realistic scenarios (when users' data distributions are unknown) and weakens the applicability of the attack. In addition, the server needs to manipulate the model architecture and place a linear layer at the start of the network to maximize the attack effectiveness. However, this modification is unworkable for typical deep learning applications (e.g., in computer vision).

Lam et al. [38] showed that, if SA is enabled, a malicious server  $S$  can still try to reconstruct individual contributions by observing multiple training rounds of FL. This disaggregation process can be reduced to a matrix factorization problem. However, their attack is effective only in a particular restricted setting in which the malicious server  $S$  (i) alters the protocol execution by providing *always* the same parameters  $\widetilde{\Theta}$  at each round of FL (i.e., users compute their gradient updates always on the same model  $f_{\widetilde{\Theta}}$ ), (ii) leverages additional side-channel information about users' participation in the training rounds, and (iii) users are required to use the same local training dataset at each round of FL. Their approach enables the gradient inversion attack of [25, 65, 70] to scale and be effective in more realistic scenarios where there is a significant number of users participating in the protocol execution of FL. Nevertheless, its feasibility still depends on the number of active users, network parameters, and other factors such as the number of rounds that the server monitors to recover individual gradients accurately. Additionally, while the authors showed that this approach could handle noise, its applicability is inherently limited in FedSGD, where local training is performed on randomly selected batches rather than the entire, static, local dataset. Finally, this attack is not applicable in large-scale, real-world deployments of FL [27, 64], where users participate in the protocol only once with a high probability.

## 3 PRELIMINARIES

We use small letters (such as  $x$ ) to denote concrete values, calligraphic letters (such as  $\mathcal{X}$ ) to denote sets. For a string  $x \in \{0, 1\}^*$ , we let  $|x|$  be its length; if  $\mathcal{X}$  is a set,  $|\mathcal{X}|$  represents the cardinality of  $\mathcal{X}$ . In the setting of deep learning, we use the notation  $[\cdot]$  to express a tensor (i.e., vector) of arbitrary dimension. We write  $[x]$  for a tensor filled with the value  $x$ . When a set is included between square brackets (e.g.,  $[\mathbb{R}^+]$ ), the tensor is filled with arbitrary elements from that set.

### 3.1 Neural networks

We abstract a neural layer with respect to some parameters  $\Theta \in \mathbb{R}^m$  using the following notation:

$$\ell(x) = \phi(x \otimes \theta + b), \quad (2)$$

where the symbols  $\theta \in \Theta$  and  $b \in \Theta$  are arbitrarily shaped real tensors<sup>3</sup> that represent the learnable parameters of the layer  $\ell$ . Hereafter, we refer to  $\theta$  and  $b$  as the *kernel* and the *bias*, respectively. The operation  $\otimes$  abstracts the application of the kernel on the input tensor  $x$ . As an example,  $\otimes$  can be a matrix multiplication operator ( $\ell$  is a fully connected layer), a convolution operator ( $\ell$  is a convolutional layer), or a more complex parametric transformation such as the multi-head attention mechanism used to build transformer networks [60]. This definition captures other core building blocks such as normalization layers. The function  $\phi$ , instead, is the activation function of  $\ell$ , which makes  $\ell$  non-linear.

A deep neural network is a function  $f_\Theta : \mathcal{X} \rightarrow \mathcal{Y}$  defined by the composition of many layers, i.e.,  $f_\Theta(x) = \ell_{n-1}(\dots(\ell_1(\ell_0(x))))$  where  $\ell_i(x) = \phi(x \otimes \theta_i + b_i)$  (as defined in Equation (2)),  $\theta_i \in \Theta$ ,  $b_i \in \Theta$ , and  $x \in \mathcal{X}$ .

### 3.2 Federated learning

Federated learning (FL) allows a set of users  $\mathcal{U} = \{u_1, \dots, u_n\}$ , where each  $u_i$  holds a local training dataset  $\mathcal{D}_i$ , to train the deep neural network  $f_\Theta : \mathcal{X} \rightarrow \mathcal{Y}$  on a global dataset that is distributed among  $\mathcal{U}$ . A centralized server  $S$  coordinates the communications between the users to train a deep neural network on  $\mathcal{D} = \bigcup_{u_i \in \mathcal{U}} \mathcal{D}_i$  in such a way that each  $\mathcal{D}_i$  does not leave the  $u_i$ 's device. The learning phase of FL is an interactive process that is divided into rounds. In the setting of FL, we denote with  $t$  the current round of FL and we use the superscript “ $(t)$ ” (sometimes in combination with the subscript “ $i$ ”) to denote values used or generated (such as batches, model updates, model parameters) during the current round  $t$  (by a particular user  $u_i$ ). At each round  $t \in \mathbb{N}$ , the server  $S$  (holding the parameters  $\Theta^{(t)}$  of the deep neural network  $f$ ) sends to the subset of available users  $\mathcal{U}^{(t)} \subseteq \mathcal{U}$  the parameters  $\Theta^{(t)} \in \mathbb{R}^m$ .

The set  $\mathcal{U}^{(t)}$  is composed of the users entitled to participate in the learning phase during the current round  $t$ . Each  $u_i$  samples a random subset  $\mathcal{D}_i^{(t)}$  (known as batch) from its dataset  $\mathcal{D}_i$  and locally trains  $f_{\Theta^{(t)}}$  on  $\mathcal{D}_i$ . The final result of the local training is a model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  that is then forwarded to the server  $S$ . The latter will be responsible of computing the new configuration  $\Theta^{(t+1)}$  of the parameters with respect to average of the model updates  $\{\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}\}_{u_i \in \mathcal{U}}$  received by the server  $S$ . This process is iterated until the parameters  $\Theta$  converge.

*FedSGD and FedAVG.* The computation of the users' model updates  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  and model parameters  $\Theta^{(t)}$  vary according to the type of FL that is in place. The two main approaches are known as *federated stochastic gradient descent* (FedSGD) and *federated averaging* (FedAVG). In FedSGD, a single step of gradient descent is performed per round  $t \in \mathbb{N}$ . In other words, the model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  (computed by a user  $u_i$ ) corresponds to the gradients  $\nabla_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  computed with respect to the randomly chosen batch  $\mathcal{D}_i^{(t)} \subseteq \mathcal{D}_i$ . This gradient

$\nabla_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  is set to be the model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  of user  $u_i$ . On the server side, the single step of gradient descent is executed in order to compute the new model parameter  $\Theta^{(t+1)}$  from  $\Theta^{(t)}$  and  $\{\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}\}_{u_i \in \mathcal{U}}$ . More formally, the parameters are updated as follows:

$$\Theta^{(t+1)} = \Theta^{(t)} - \eta \frac{\sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}}{|\mathcal{U}^{(t)}|} \quad (3)$$

where  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}} = \nabla_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  and  $\eta$  is the learning parameter.

On the other hand, in FedAVG, users locally perform  $k \in \mathbb{N}$  iterations of stochastic gradient descent, producing new model parameters at each round. More formally, let  $\Theta^{(t)} = \Theta_i^{(t,1)}$  be the parameters received by  $u_i$  from the server at the beginning of round  $t$ . For every  $j \in \{1, \dots, k\}$ , a user  $u_i$  samples a random batch  $\mathcal{D}_i^{(t,j)} \subset \mathcal{D}_i$  and computes a gradient  $\nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta^{(t,j)}}$ . Then, it locally updates the model parameters as defined in Equation (3), i.e.,  $\Theta_i^{(t,j+1)} = \Theta_i^{(t,j)} - \eta \cdot \nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta^{(t,j)}}$ . Then, after  $k$  iterations of gradient descent, the final model parameters  $\Theta_i^{(t,k)}$  will be the model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  that user  $u_i$  will send to the server. Finally, the server  $S$  only needs to compute the new parameters  $\Theta^{(t+1)}$  that is the average of the parameters received from  $\mathcal{U}^{(t)}$ :

$$\Theta^{(t+1)} = \frac{\sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}}{\sum_{u_i \in \mathcal{U}^{(t)}} b_i} = \frac{\sum_{u_i \in \mathcal{U}^{(t)}} \Theta_i^{(t,k)}}{\sum_{u_i \in \mathcal{U}^{(t)}} b_i}, \quad (4)$$

where  $b_i = \sum_{j=1}^k |\mathcal{D}_i^{(t,j)}|$ .<sup>4</sup>

The complete description of the FL protocol (with FedSGD or FedAVG) is depicted in Figure 8 of the appendix.

### 3.3 Secure aggregation

A secure aggregation (SA) protocol is a specialized multi-party computation (MPC) protocol that allows a set of users to compute the summation (a.k.a. aggregation) of their inputs. Let  $\mathcal{U} = \{u_1, \dots, u_n\}$  be a set of users, each holding a secret input  $v_i$  (e.g., integer, group element, vector). A protocol  $\Pi$  is a secure SA protocol if it securely implements the following ideal functionality:

$$f^{\text{sa}}(v_1, \dots, v_n) = (v, \dots, v) \quad \text{for } v = \sum_{u_i \in \mathcal{U}} v_i, \quad (5)$$

i.e., a trusted third party executes  $f^{\text{sa}}$  computes and returns to all users  $\mathcal{U}$  the aggregation  $v$  of the users' inputs  $(v_1, \dots, v_n)$ . Informally, a SA protocol  $\Pi$  is considered secure if it is at least as secure as invoking the ideal functionality  $f^{\text{sa}}$ . This is formalized using the standard ideal and real-world paradigm of MPC [18] (see Appendix B). The security of  $\Pi$  does not guarantee that nothing is leaked about other users' inputs. Instead, it implies that nothing is leaked except what can be inferred from the final aggregation. The information that can be inferred is highly correlated to the inputs provided to  $f^{\text{sa}}$  (e.g., entropy).

<sup>4</sup>Observe that, in order to compute the new model parameters  $\Theta^{(t+1)}$ , the server needs to receive either  $b_i = \sum_{j=1}^k |\mathcal{D}_i^{(t,j)}|$  from each user  $u_i$  or the size of each  $\mathcal{D}_i^{(t,j)}$  needs to be fixed.

<sup>3</sup>The shape of these tensors depend on the operator  $\otimes$ .



To increase the security of FL and mitigate attacks such as gradient inversion (see Section 2), Bonawitz et al. [10] propose a communication-efficient, dropout resilient SA for FL.<sup>5</sup> When applied to FL, the protocol guarantees a secure aggregation of users’ model updates – sensitive values that carry important information about the users’ datasets. Hence, an adversary (e.g., malicious user or server) can not observe the individual model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  of a target user  $u_i$ . This decreases the amount of information about the local training dataset  $\mathcal{D}_i$  of the target user  $u_i$  that an adversary can leak (e.g., by performing gradient inversion attacks). We stress that SA protocols [8, 10] have the property of being robust: The final aggregation can always be computed when at most  $\delta$  users drop out during the execution of the SA protocol (i.e., more than  $n - \delta$  users are online). When this condition is failed (i.e., the online users are less than  $n - \delta$ ), the server cannot recover the final aggregated value. Hence, SA also guarantees that the server can only see the final aggregation when the latter contains at least  $n - \delta$  model updates (ensuring the desired level of “privacy by aggregation”).

**Both Bonawitz et al. [10] and Bell et al. [8] (the most influential SA protocols in the literature) demonstrate the security of their protocols in both the semi-honest and malicious models** (including collusion between server and users) [10, Theorems 6.2 and 6.6] and [8, Theorems 3.6 and 4.9]. However, we emphasize that the results in [8, 10] cover only the security of the SA protocol. Nothing is claimed about the security of the overall FL protocol with SA enabled.

## 4 THREAT MODEL

In this section we formalize the threat model in which our attacks are defined (Section 5 and Section 6). We adopt the *exact* same threat model for which the SA protocols of Bonawitz et al. [10, Section 6.2] (CCS ’17) and Bell et al. [8, Section 4] (CCS ’20) have been demonstrated to be secure: A malicious parameter server (aggregator) that can corrupt at most a fixed number  $m$  (out of  $n$ ) of users.<sup>6</sup> Observe that these SA protocols [8, 10] are the most influential and practical-oriented solutions in the field.

More formally, in each round  $t \in \mathbb{N}$ , the active users  $\mathcal{U}^{(t)}$  do not send their model updates  $\{\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}\}_{u_i \in \mathcal{U}^{(t)}}$  in the clear. Instead, they execute the SA protocol  $\Pi$  to securely compute the aggregation  $v = \sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  (see Section 3.3) for at least  $q$  users. Only the final aggregation  $v$  is revealed to the server  $S$ . More importantly, we do not target any specific implementation of SA. In fact, our attack exploits a vulnerability caused by the incorrect usage of SA in the FL protocol. In particular, the FL protocol does not validate the inputs of SA (i.e., model updates). For this reason, to keep the attack general, we replace the execution of the underlying SA protocol with the invocation of its ideal functionality  $f^{\text{sa}}$  (Section 3.3).

We model an adversarial server  $S$  whose objective is to learn information about the local dataset  $\mathcal{D}_{\text{trgt}}$  of one (or more) target user  $u_{\text{trgt}}$  that participates in the execution of the protocol, outside

what can be learned from the aggregated model updates.  $S$  can deviate from a honest execution (e.g., it sends arbitrary messages) as defined in [10, Section 6.2] and [8, Section 4]. Specifically, in our attacks, the malicious server  $S$  exploits the model inconsistency attack vector; that is,  $S$  provides arbitrary malicious parameters to arbitrary users even within the same training round  $t \in \mathbb{N}$ . Although the adversarial server  $S$  is allowed to collude with at most  $m$  users (as considered by [8, 10]), our attacks demonstrate that SA is ineffective in FL even when  $m=0$ , i.e., the server is malicious, but it does not collude with any user.<sup>7</sup> This ensures the effectiveness of the proposed attacks also outside the cross-device FL setting, where the parameter server does not perform user sub-sampling.

Lastly, we assume the standard, centralized, communication topology of FL as in real-world applications [27, 64], where each user is authenticated by a PKI and shares an encrypted channel with the server  $S$ . The SA protocols in [8, 10] are designed for this communication topology. Moreover, note that a PKI is a necessary assumption. Indeed, as described in [10, Section 6.2] and [8, Section 4.2], without a PKI, a server could break the privacy of users by simply launching a sybil attack. Assuming the existence of a PKI is enough to rule out such trivial sybil attacks [10, Section 6.2] and [8, Section 4.2].

## 5 GRADIENT SUPPRESSION ATTACK

In this section, we present our first attack, dubbed *gradient suppression*, in which a malicious server exploits the model inconsistency attack vector to bypass SA and leak the model update of a chosen target user.

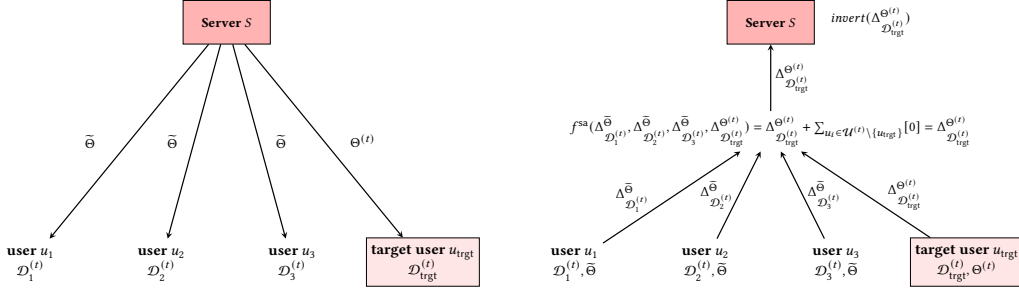
In a nutshell, the malicious server  $S$  selects a **target user**  $u_{\text{trgt}}$  among the set of active users  $\mathcal{U}^{(t)}$  of the current round  $t \in \mathbb{N}$ . The aim of  $S$  is then to preserve the target’s model update during the SA process by tampering with the parameters for the other **non-target users**. Here,  $S$  creates a set of malicious parameters  $\tilde{\Theta}$  that is sent to the non-target users  $\mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}$  whereas  $u_{\text{trgt}}$  receives the real parameters vector  $\Theta^{(t)}$ . The malicious  $\tilde{\Theta}$  is crafted in such a way that the local application of gradient descent performed by a non-target users produces a tampered model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\tilde{\Theta}}$ . The tampered model updates, when aggregated through SA, have the property of preserving the target’s model update  $\Delta_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}}$ , allowing the server to recover it. Once that the target’s model update is on the server-side,  $S$  can leak sensitive information about the batch  $\mathcal{D}_{\text{trgt}}^{(t)}$ , used during the current round  $t$  of FL by executing an arbitrary gradient inversion attack (see Section 2.2) or related inference attacks.

Given the fact that the SA performs the sum among the users’ model updates, the simplest way to achieve the isolation of the target’s signal is to **force the tampered model updates to have a negligible magnitude, or, more strictly, to be zero everywhere**. In this section, we first study the extreme case  $\Delta_{\mathcal{D}_i^{(t)}}^{\tilde{\Theta}} = [0]$ . That is, the aggregation of the model updates (i.e., gradients) is

<sup>5</sup>In [10] the aggregation  $v$  is obtained by the server  $S$  only. This can be represented as an ideal functionality  $f^{\text{sa}}(v_1, \dots, v_n, \perp) = (\perp, \dots, \perp, v)$  where the  $(n+1)$ -th input/output is associated to the server  $S$  and  $\perp$  represents the empty string.

<sup>6</sup>The number of corrupted users  $m$  depends on the implementation of the SA protocol. For instance, this is  $\lceil \frac{n}{3} \rceil - 1$  for Bonawitz et al. [10].

<sup>7</sup>However, collusion with users improves the effectiveness of the attacks when additional mechanisms such as distributed differential privacy are in place.



**Figure 1: Graphical representation of the gradient suppression attack (against FedSGD) with  $\mathcal{U}^{(t)} = \{u_1, u_2, u_3, u_4 = u_{trgt}\}$ . The left figure depicts the malicious parameters distribution (model inconsistency). The right figure depicts the secure aggregation of model updates, the collection of the target’s model update, and the inversion. The malicious parameters  $\tilde{\Theta}$  produce a tampered model update  $\Delta_{\mathcal{D}_i^{(t)}}^{(t)}$  (i.e., gradient) equal to  $[0]$  for each non-target user  $u_i \in \mathcal{U}^{(t)} \setminus \{u_{trgt}\}$ . The function  $invert(\cdot)$  denotes the technique used by the server  $S$  (e.g., gradient inversion) to extract sensitive information of the original dataset  $\mathcal{D}_{trgt}^{(t)}$  from the target gradient  $\Delta_{\mathcal{D}_{trgt}^{(t)}, \Theta^{(t)}}^{(t)}$ .**

equal to  $u_{trgt}$ ’s model update, i.e.,

$$\begin{aligned} f^{sa}(\Delta_{\mathcal{D}_1^{(t)}, \tilde{\Theta}}^{(t)}, \dots, \Delta_{\mathcal{D}_{i-1}^{(t)}, \tilde{\Theta}}^{(t)}, \Delta_{\mathcal{D}_{trgt}^{(t)}, \Theta^{(t)}}^{(t)}, \Delta_{\mathcal{D}_{i+1}^{(t)}, \tilde{\Theta}}^{(t)}, \dots, \Delta_{\mathcal{D}_n^{(t)}, \tilde{\Theta}}^{(t)}) \\ = f^{sa}([0], \dots, [0], \Delta_{\mathcal{D}_{trgt}^{(t)}, \Theta^{(t)}}^{(t)}, [0], \dots, [0]) = \Delta_{\mathcal{D}_{trgt}^{(t)}, \Theta^{(t)}}^{(t)} \end{aligned}$$

allowing  $S$  to exactly recover  $\Delta_{\mathcal{D}_{trgt}^{(t)}, \Theta^{(t)}}^{(t)}$ . Figure 1 depicts the gradient suppression attack against FedSGD. We stress that the attack applies to FedAVG as we will discuss in Section 5.2.

Next, we show how to compute the malicious parameters  $\tilde{\Theta}$  required to perform the gradient suppression attack. We focus on the most widely adopted class of deep learning models—the one based on the ReLU activation function. **However, in Appendix C, we show that our approach extends to arbitrarily composed architectures.**

## 5.1 Gradient suppression for ReLU layers: The dead-layer trick

The *Rectified Linear Unit* (ReLU) activation function:

$$ReLU(x) = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{if } x \leq 0 \end{cases} \quad (6)$$

is one of the core technical improvements that led to deep learning [45]. Nowadays, this function is ubiquitous in computer vision architectures, representing the core building block of highly successful and standardized models such as ResNet [29], DenseNet [31] and many others. Outside the computer vision domain, the ReLU activation function is currently finding its place in Natural Language Processing (NLP) applications thanks to the success of transformer networks [12, 50, 60].

The *dying-ReLU* problem [39] is a phenomenon that naturally occurs during the training of deep neural networks that rely on the ReLU activation function. When a layer  $\ell$  “dies”, it enters a state where it can only produce a constant output. More importantly, the dead layer  $\ell$  cannot produce any gradient during the gradient

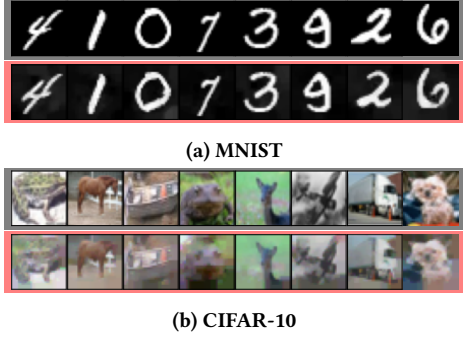
descent iterations, i.e., the derivatives of its trainable parameters are zero regardless of the given input and loss function.

Despite the *dying-ReLU* phenomenon can naturally occur, we show that it can also be intentionally induced by a malicious server to prevent a network from producing a gradient for one or more sets of parameters. Next, we describe how this can be achieved and exploited to perform our gradient suppression attack.

**5.1.1 Triggering the Dying-ReLU phenomenon with malicious parameters.** The *Dying-ReLU* phenomenon is due to the piece-wise non-differentiability of the ReLU activation function. Consider a neural layer  $\ell$  with a ReLU activation function, i.e.,  $\ell(x) = ReLU(x \otimes \theta + b)$ . From Equation (6), we can see that the ReLU function behaves as a constant function  $ReLU(x) = 0$ , whenever the input  $x$  is equal or less than zero. Since the derivative of a constant function is always 0, we can easily conclude that, for any loss function  $\mathcal{L}$ , then we have  $\frac{\partial \mathcal{L}}{\partial \theta} = [0]$  and  $\frac{\partial \mathcal{L}}{\partial b} = [0]$  when  $x \otimes \theta + b \leq 0$ . In other words, the layer  $\ell$  receives zero gradient for its trainable parameters  $\theta, b \in \Theta$  every time its pre-activation (i.e.,  $x \otimes \theta + b$ ) is less or equal to zero.

A malicious server  $S$  can exploit the above behavior of the ReLU function to kill a layer  $\ell$  of a neural network  $f$ , i.e., by forcing the pre-activation  $x \otimes \theta + b$  of  $\ell$  to be less or equal to zero. This can be accomplished (without control over the input  $x$  of an user) by computing some malicious trainable parameters  $\tilde{\theta}, \tilde{b} \in \tilde{\Theta}$  of the layer  $\ell$ .

In more detail, the operator  $\otimes$  (see Section 3.1) is generally based on a multiplication-like operation between the input  $x$  and the kernel  $\theta$ . Therefore, we can easily force the pre-activation to be  $[0]$  for any input  $x$  by just choosing  $\tilde{\theta} = [0]$  and  $\tilde{b} = [\mathbb{R}_{\leq 0}]$ . Alternatively, having some knowledge on the input  $x$ , we can rely on different setups for  $\tilde{\theta}$  and  $\tilde{b}$ . For instance, if  $x$  is strictly positive (e.g., because  $x$  is the output produced by a previous ReLU-layer or because of the adopted input normalization process), it is enough to produce a malicious  $\tilde{\theta}$  with negative numbers. Instead, if a bound on  $x$  is known (e.g.,  $x \in [-1, 1]$ ), we can just set the malicious



**Figure 2: Examples of reconstruction (red panels) obtained via gradient inversion attacks using [25] for two datasets.**

bias vector  $\tilde{b}$  to a large enough negative number (e.g., fix  $\tilde{\theta}$  and set  $\tilde{b} = -\max(\tilde{\theta} \otimes x)$ ).

Now, an attacker can exploit the *dead-layer trick* to force a ReLU-based network to produce zero gradient for every layer. In this direction, it is important to note that, for plain, ReLU-based feedforward architectures, the server  $S$  can just kill the kernels in the very first layer to suppress the gradient flow for the rest of the network.<sup>8</sup> Hence, killing all the kernels  $\{\theta_i\}$  in the original parameters  $\Theta$  is very often not necessary. Similarly, in modern architectures, neural layers tend to be arranged in the form:

neural layer  $\rightarrow$  normalization layer  $\rightarrow$  activation.

Therefore, to suppress the gradient for the network, it is enough to zero only the parameters of the normalization layers as these are often defined as  $\gamma\hat{x} + \beta$ , where  $\hat{x}$  is the normalized input. For instance, in the case of batch normalization, the attacker can kill the neural layer by setting the vectors  $\gamma$  and  $\beta$  of the batch normalization to  $[0]$ . Nevertheless, the strategy would work for every architectural configuration.

Moreover, even if ReLU is the most common activation function in deep learning, a server  $S$  can always maliciously choose a neural network architecture  $f$  that presents a ReLU activation function in the “right spots” of the model without requiring unrealistic architecture modifications. The only gradient signal that cannot be recovered by the server using the dead-ReLU trick is the one of the bias term of the last layer (details are given in Appendix C). This follows from the fact that the terminal layer of a network only rarely exhibits a ReLU activation function. A trivial solution for the malicious server is to avoid the bias term of the last layer when defining the architecture of the model. Alternatively, the malicious server can ignore the gradient of the bias term during the gradient inversion. Indeed, this represents only a tiny portion of the total number of trainable parameters of the network. For instance, in the case of a ResNet50 trained on ImageNet [29], the bias vector in the final layer counts for only  $4 \cdot 10^{-5}\%$  of the total number of parameters. In Section 5.2.2, we show that gradient inversion is unaffected from this missing gradient.

<sup>8</sup>However, if present, all the bias terms of the network should be set to values  $\leq 0$ .

## 5.2 Attack execution

Turning back to the attack described at the beginning of this section, we have now an effective and efficient approach to isolate the model update  $\Delta_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}}$  of the target user  $u_{\text{trgt}}$ . The malicious server  $S$  can just exploit the techniques discussed in Section 5.1 and Appendix C to generate the malicious parameters  $\tilde{\Theta}$  and suppress the model updates of non-target users, completely nullifying their contributions in the aggregated signal produced by SA. As previously described (Figure 1), the attack is composed of two phases.

**5.2.1 Distribution of the (malicious) parameters.** In the first phase of the attack,  $S$  creates the malicious parameters  $\tilde{\Theta}$  for the non-target users. Right after the choice of  $\tilde{\Theta}$ ,  $S$  must choose the target user  $u_{\text{trgt}}$  for the current round  $t \in \mathbb{N}$  of FL.  $S$  can either select the target at random (a trawling attack) from  $\mathcal{U}^{(t)}$  or target a specific (e.g., exploiting the IP address used to query the model by the user, if available). Then,  $S$  can enforce the model inconsistency by distributing the parameters (Figure 1). In more detail, upon receiving a request for the parameters from a user  $u_i$ , the parameter server answers by sending  $\Theta_i^{(t)}$  defined as follows:

$$\Theta_i^{(t)} = \begin{cases} \Theta^{(t)} & \text{if } i = \text{trgt} \\ \tilde{\Theta} & \text{otherwise} \end{cases}, \quad (7)$$

where  $\Theta^{(t)}$  are the honest parameters of the current round  $t \in \mathbb{N}$ . Optionally,  $S$  can send a maliciously crafted model to the target user to increase the information recovered from the inversion attack [22, 25].

**5.2.2 Aggregation, collection, and inversion.** After the distribution of the parameters, the malicious server  $S$  waits until it receives the output  $v$  of  $f^{\text{sa}}$ , i.e.,  $f^{\text{sa}}(\Delta_{\mathcal{D}_1^{(t)}}^{\tilde{\Theta}}, \dots, \Delta_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}}, \dots, \Delta_{\mathcal{D}_n^{(t)}}^{\tilde{\Theta}}) = v$  where

$$v = \Delta_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}} + \sum_{u_i \in \mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}} \Delta_{\mathcal{D}_i^{(t)}}^{\tilde{\Theta}}. \quad (8)$$

Then, it proceeds differently according to which algorithm (between FedSGD or FedAVG) is active.

In FedSGD, the output  $v$  of  $f^{\text{sa}}$  is the  $u_{\text{trgt}}$ ’s gradient, i.e.,  $v = \Delta_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}} = \nabla_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\Theta^{(t)}}$ . This is because, as discussed earlier, the malicious parameters  $\tilde{\Theta}$  produces  $\Delta_{\mathcal{D}_i^{(t)}}^{\tilde{\Theta}} = [0]$  for each non-target user  $u_i \in \mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}$  (Section 5.1 and Appendix C). After recovering the plaintext gradient, the server can reconstruct the target input by performing standard inversion attacks (Section 2.2) as done in the protocol without SA. Figure 2 reports examples of gradient inversion on the federated system with gradient recovered via the gradient suppression attack. In the examples, we keep the bias term of the last layer in the architecture and ignore it during the optimization [25]. Note that the attack is independent of the number of users participating in the aggregation, and this can be arbitrarily large. Similarly, the server can perform previously proposed inference attacks [43, 46] on the individual user.

On the other hand, in FedAVG, a model update is composed of the parameters of the local model rather than a gradient (see Section 3.2).

More formally,

$$\Delta_{\mathcal{D}_i^{(t)}}^{\Theta_i^{(t)}} = \Theta_i^{(t,k)} \quad \text{for } u_i \in \mathcal{U}^{(t)}, \quad (9)$$

where the local model  $\Theta_i^{(t,k)}$  of  $u_i$  is obtained by applying  $k$  iterations of SGD using the local dataset  $\mathcal{D}_i$  and  $\Theta_i^{(t)}$  (as defined Equation (7)) sent by server (recall that  $\Theta_i^{(t)} = \tilde{\Theta}$  for  $\mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}$ ). Now, when a non-target user performs the local training procedure using the malicious parameters  $\Theta_i^{(t)} = \tilde{\Theta}$ , we have that

$$\Theta_i^{(t,j+1)} = \Theta_i^{(t,j)} - \eta \cdot \nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta_i^{(t,j)}} \quad (10)$$

where  $\Theta^{(t,1)} = \tilde{\Theta}$ ,  $\mathcal{D}_i^{(t,j)} \subseteq \mathcal{D}_i$ , and  $j \in \{1, \dots, k\}$ . As in the case of FedSGD, we have that  $\nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta_i^{(t,j)}} = [0]$  for every non-target user

$u_i \in \mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}$  and we conclude that  $\Theta_i^{(t,j)} = \tilde{\Theta}$ .

By combining Equations (8) to (10), we obtain the equality  $v = (n-1) \cdot \tilde{\Theta} + \Theta_{\text{trgt}}^{(t,k)}$ . This equation can be solved with respect to the indeterminant  $\Theta_{\text{trgt}}^{(t,k)}$ . Once the malicious server  $S$  recovered the updated local model  $\Theta_{\text{trgt}}^{(t,k)}$  of the target  $u_{\text{trgt}}$ , it can determine the (pseudo) gradient signal by removing the honest parameters  $\Theta^{(t)}$  from  $\Theta_{\text{trgt}}^{(t,k)}$  and proceed with the gradient inversion/inference attack.

### 5.3 Impact

Current FL+SA implementations do not prevent the gradient suppression attack, making users actively susceptible to this simple yet powerful exploit. To a certain extent, this attack can be interpreted as an **invalid input validation vulnerability** present in the users' FL client software. Here, the latter permits users to perform computation on "semantically malformed inputs" sent by a non-trusted party, i.e., the server. This allows the server to control SA's inputs of users and eventually affect the aggregation. Furthermore, in contrast to most of the previous attacks introduced in FL, the disclosed vulnerability has the practical advantage of **being completely independent of the number of users participating in the current round**. Therefore, this procedure scales to millions of active users, making it applicable to real-world scenarios such as cross-device FL, which is currently being deployed in-the-wild [27, 64]. In the same direction, its effectiveness is independent of the size of the model or other nuisance factors such as the stillness of users' training datasets during the attack [38]. Additionally, unlike [22], this attack neither hinges on auxiliary information on the users' private sets nor requires unrealistic modifications of the model architecture; indeed, it can be applied to arbitrary architectures and loss functions. It is important to note that the gradient suppression attack can be iterated several times and arbitrarily alternated with honest training iterations. If the server wants to recover information on all the users, it has to iterate the attack several times by targeting a user at a time. Assuming no dropouts among users participating at the FL protocol, recovering the gradient of all users requires  $n$  iterations where  $n$  is the number of active users.

We stress that the gradient suppression attack shows the incorrect application of SA in FL, yielding a "false sense of security". As discussed in Section 1.1, the core motivation is that SA guarantees that nothing is leaked about the model updates of the users except what can be inferred from their aggregation. **This claim assumes that the inputs (i.e., model updates) of SA are fixed and are not under the control of an adversary. However, this does not hold in FL since, in this case, the value of inputs that needs to be aggregated depends on the parameters  $\Theta$  distributed by the server.** Hence, a malicious server that executes the gradient suppression attack can indirectly tamper with the SA's inputs to maximize the information leaked (e.g., leak the model update of a target user).

Although the gradient suppression attack is highly effective, it can be easily detected by non-target users (we delve into this topic in Section 7). Nevertheless, in the next Section 6 we introduce an extension of the gradient suppression attack that adds stealthiness and it is harder to detect. More generally, one can always trade effectiveness for stealthiness also in the gradient suppression framework. For instance, if recovering noisy model updates is acceptable, the server can send highly optimized models (e.g., obtained after some rounds of honest execution) to non-targets and an unoptimized model to the target. Intuitively, the gradient from the unoptimized model should dominate the aggregation given the low magnitude and sparsity of the one produced by the optimized models (see Figure 4).

## 6 CANARY-GRADIENT ATTACK FOR PROPERTY INFERENCE

Section 5 demonstrates that a malicious server  $S$  can force non-target users to produce a zero gradient during a round of FL. This allows  $S$  to bypass SA and, at the same time, maximize the leakage regarding the dataset of a target user. While the gradient suppression attack can be seen as the most extreme exploitation of the model inconsistency attack vector, more stealthy attacks can be created harnessing the same underlying intuition.

This section shows a general procedure that allows a malicious server to perform highly accurate property inference attacks on individual users, even if SA is enabled. The idea behind this approach is that the server can maliciously modify the parameters of the model in order to inject specific detectors in one or more subsets of the network. These detectors are specifically crafted to react to attacker-chosen trigger conditions that can be present in the users' training instances. Whether the detector is triggered during the local training procedure, the network produces a clear footprint in the model update. Then, upon receiving the latter from a user, the server can determine if the trigger condition has been met by looking for the footprint in the model update. This allows the server to infer information on the content of the user's training set; that is, the presence or absence of data with the specific property. For instance, using this approach, the server can perform an extremely accurate membership inference on a chosen target user. Hereafter, we refer to this general procedure as the **canary-gradient attack**.



As for the gradient suppression, the canary-gradient attack does not target any specific SA protocol but instead leverages a vulnerability of FL caused by its incorrect usage of SA. For this reason, we abstract the SA protocol with its ideal functionality  $f^{\text{sa}}$ .

### 6.1 The Conditional Dead-Layer trick

The main building block to construct the attack is a conditioned version of the dead-layer trick of Section 5.1. Informally, we want to “kill” a layer only if the instance  $x \in \mathcal{D}_{\text{trgt}}^{(t)}$  of the target  $u_{\text{trgt}}$  satisfies a particular condition. In other words, we would like a programmed death through the backdooring of the layer. For the sake of presentation, we introduce the conditional dead-layer trick assuming that SA is disabled. Then, in Section 6.2 we extend the discussion to the case of SA enabled.

Formally, given a layer  $\ell$ , we want to find some malicious parameters  $\tilde{\Theta}$  to enforce the following behavior

$$\frac{\partial \mathcal{L}(\mathcal{D}_{\text{trgt}}^{(t)}, \tilde{\Theta})}{\partial \xi} \neq 0 \iff \exists x \in \mathcal{D}_{\text{trgt}}^{(t)} : P(x) = \text{True}, \quad (11)$$

where  $\mathcal{D}_{\text{trgt}}^{(t)}$  is the batch (of the current round  $t$ ) used by the target user  $u_{\text{trgt}}$ ,  $\xi \in \tilde{\Theta}$  is a subset parameters of the network, and  $P$  is a predicate that defines the property the malicious server  $S$  wants to detect in the batch  $\mathcal{D}_{\text{trgt}}^{(t)}$  of  $u_{\text{trgt}}$ . In particular,  $\xi$  can be composed of the parameters of any logic partition in the neural network, such as a specific filter in a convolution layer or an element in the scale and shift vectors in a normalization layer.

As discussed in Section 5, suppressing the gradient for a set of parameters in a ReLU-based layer is about controlling the value of its pre-activation. Therefore, given a neural layer  $\ell$  with ReLU activation, we can substitute Equation (11) with:


$$\ell_{\xi} = (x_{\xi} \otimes \theta_{\xi} + b_{\xi}) > 0 \iff \exists x \in \mathcal{D}_{\text{trgt}}^{(t)} : P(x) = \text{True}, \quad (12)$$

where  $\xi = \{\theta_{\xi}, b_{\xi}\}$ , and  $x_{\xi}$  is the subset of the input of  $\ell$  that interacts with the parameters  $\xi$  and  $\ell_{\xi}$  refers to the subset of the output of the layer  $\ell$  computed using the parameters  $\xi$ .


The simplest and most natural way to find  $\xi$  that correctly induce Equation (12) is to explicitly train the layers preceding  $\ell$  and the parameters  $\xi$  to force  $\ell_{\xi}$  to produce a positive value only when the input of the network satisfies  $P$ . In other words, we train part of the network in a classification task, using the output  $\ell_{\xi}$  such as the output layer, where the classification threshold is centered in zero. Observe that, if the behavior of Equation (12) is correctly embedded in the network  $\tilde{f}_{\tilde{\Theta}}$ , a malicious server  $S$  will be able to determine the event  $\exists x \in \mathcal{D}_{\text{trgt}}^{(t)} : P(x) = \text{true}$  by only collecting gradient  $\nabla_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\xi}$  of  $u_{\text{trgt}}$  and check that the derivatives of  $\xi$  are different from zero. In Section 6.3 we show how this can be done in practice.

### 6.2 Targeted property inference attacks via model inconsistency

To perform the membership inference attack discussed in the previous section, the malicious server  $S$  needs to have access to the gradient  $\nabla_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\xi}$  of  $u_{\text{trgt}}$ . This is possible when SA is disabled. On the other hand, when SA is enabled, the  $S$  can inject the canary-gradient



**Figure 3: Graphical representation of the SA execution when the canary-gradient is applied.** On the left, each square represents a gradient update produced by a different user. The green square represents the target’s gradient and the inner small blue square represents the gradient for  $\xi$ . On the other hand, each red square represents the gradient produced by non-target users, with zero gradient for  $\xi$ . The square on the right represents the aggregation where the target’s gradient  $\nabla_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\xi}$  for  $\xi$  is preserved.



**Figure 4: Comparison between the test-set accuracy of a ResNet model trained on CIFAR10 and the sparsity of its gradient (i.e., percentage of parameters that receive zero gradient) during the training.**

functionality in the network and then perform the inference attack on the whole pool of active users. In this scenario,  $S$  would be able to infer that one of the users triggered the canary-gradient by observing that  $f_{\xi}^{\text{sa}}(\nabla_{\mathcal{D}_1^{(t)}}^{\xi}, \dots, \nabla_{\mathcal{D}_n^{(t)}}^{\xi}) = \sum_{u_i \in \mathcal{U}^{(t)}} \nabla_{\mathcal{D}_i^{(t)}}^{\xi} \neq [0]$  where  $f_{\xi}^{\text{sa}}$  denotes the inner idealized functionality of  $f^{\text{sa}}$  that performs the aggregation of the gradients of  $\xi$ . However, in this case, the privacy of users would be partially preserved as  $S$  would not be able to attribute the result of the inference attack to a specific user (“privacy by shuffling”).

Still, we show that model inconsistency can be exploited even in this case, allowing the malicious server  $S$  to bypass SA and target the specific target  $u_{\text{trgt}}$ . Analogously to the gradient suppression attack (Section 5),  $S$  needs to tamper with the honest parameters  $\Theta^{(t)}$  in order to produce two different malicious  $\tilde{\Theta}_1$  and  $\tilde{\Theta}_2$ . The target user  $u_{\text{trgt}}$  will receive  $\tilde{\Theta}_1$  that is the original model  $\Theta^{(t)}$  injected with a canary-gradient for the parameters  $\xi$  as discussed in Section 6.1. On the other hand, the non-target users  $\mathcal{U}^{(t)} \setminus \{u_{\text{trgt}}\}$  will receive  $\tilde{\Theta}_2$  that is a slight perturbation of the original model  $\Theta^{(t)}$  that has the additional property of unconditionally produce zero-gradient only for the parameters  $\xi$ , i.e.,  $\nabla_{\mathcal{D}_i^{(t)}}^{\xi} = [0]$ . This can be achieved by exploiting the dead-layer trick in a localized way. Instead of killing the gradient for the whole layer, it intentionally inhibits only the

gradient produced by the parameters  $\xi$ ; for instance, for just one filter in a convolution layer.

Now, when the target and non-target gradients are aggregated, the target’s gradient  $\nabla_{\mathcal{D}_{\text{trgt}}^{(t)}}^{\xi}$  for  $\xi$  will be preserved, allowing the server to state the activation or non-activation of the canary-gradient, and so, with high probability, the presence of the property  $P$  in the batch  $\mathcal{D}_{\text{trgt}}^{(t)}$  of  $u_{\text{trgt}}$ . This intuition is captured by Figure 3.

Finally, given that the number of parameters in  $\xi$  can be arbitrarily small, the attack will leave only a minimal footprint, making the detection non-trivial (in our experiments, we show that two parameters are enough). This is also supported by the fact that the gradient becomes more sparse as the training proceeds (see Figure 4), making it difficult to distinguish between natural “holes” and artificial ones in the gradient vector.

While we gave an abstract view on the attack strategy, next, we show how the attack can be carried out on a realistic architecture such as ResNet in a membership inference attack scenario.

### 6.3 Injecting canary-gradient for membership inference

We show a practical example of how to model a membership attack on the training dataset of a specific user. Specifically, we target training instance  $x_t$ , and we want to infer if  $x_t$  is contained in the batch  $\mathcal{D}_{\text{trgt}}^{(t)}$  used by the target user  $u_{\text{trgt}}$  to compute the gradient update in the current round  $t \in \mathbb{N}$  of FL. Following previous notation, we want to infer the following property:

$$P_{x_t}(x) = \text{True} \iff x = x_t.$$

We start by considering the case of FedSGD and carry out the attack on a ResNet20 network. However, for this network, we do not consider the batch normalization layers as those would make the attack trivial. Indeed, if batch normalization is used, we could detect the activation of  $\ell_{\xi}$  by checking the average computed and sent to the server to update the running mean. For this reason, we keep the attack general by substituting every batch normalization with layer normalization [6] that does not present this issue and has an overlapping role. We then extend the attack to FedAVG in Appendix D.

In our experiment, we inject the canary gradient in the last residual block of the network. This is because the terminal layers are usually the ones that receive the sparsest gradient during the training. Since the normalization layer precedes the ReLU activation function, we chose a subset of the parameters of the latter as our  $\xi$ . In particular, we can pick any pair  $(\gamma_i, \beta_i)$  in the scale and shift vectors  $\gamma$  and  $\beta$ . Thus, in this case,  $\xi$  is composed of only two parameters, that is about  $7 \cdot 10^{-4}\%$  of the total number of parameters in the network. Hereafter, we always choose  $i = 0$ ; however, choosing a different channel would not affect the attack.

To inject the canary gradient, we use a learning-based approach. In this direction, we assume that the adversary (i.e., malicious server) knows a shadow dataset  $\mathcal{D}_s$  defined in the same domain of the target point  $x_t$ . For instance, if  $x_t$  is a face image,  $\mathcal{D}_s$  contains face images as well. We stress that, as we will show later, the distribution of  $\mathcal{D}_s$  and one of the users’ datasets can be different. Intuitively, the role of  $\mathcal{D}_s$  is providing negative samples while training  $\ell_{\xi}$  to fire




Figure 5: Graphical representation of the feature-space ( $x$ -axis and  $y$ -axis) for the pre-activation  $\ell_{\xi}$  ( $z$ -axis). The gray plane represents  $z = 0$ , i.e., the threshold for the activation of the ReLU function.




Figure 6: Average accuracy of the canary-gradient attacks for six different setups with increasing batch size for a ResNet20.


on  $x_t$ . It is important to note that the canary-gradient attack is agnostic, and the simple training procedure discussed can be substituted with other techniques. For instance, approaches such as [22] and learning-based approaches that do not require negative samples (i.e., one-class classification) may be used to reach the same result. In this direction, our main contribution is introducing the idea that the server can manipulate the derivative of the current model to encode arbitrary messages (and that these messages can be retrieved under SA).

We can now inject the canary gradient by reducing the malicious injection to a classification problem. We train the split of the network up to  $\xi$  in producing positive  $\ell_{\xi}$  when the network’s input contains  $x_t$ . In our case, with  $\xi = \{\gamma_i, \beta_i\}$ , we have  $\ell_{\xi} = \gamma_i \bar{x}_i + \beta_i$ , where  $\bar{x}$  is the normalized input of the layer. Our loss function  $\mathcal{L}$  for a batch of size  $n$  is simply defined as the binary cross entropy:

$$-\frac{1}{n} \sum_{i=1}^n \begin{cases} \alpha_1 \cdot \log(\text{sigmoid}(\ell_{\xi})) & \text{if } x = x_t \\ \alpha_0 \cdot \log(1 - \text{sigmoid}(\ell_{\xi})) & \text{otherwise} \end{cases}, \quad (13)$$

where instances different from  $x_t$  are sampled from  $\mathcal{D}_s$ , and  $\alpha_1$  and  $\alpha_0$  weight the loss for the two events. In other words, we want to “overfit” the network to produce positive  $\ell_{\xi}$  only for the point  $x_t$ , while squashing under 0 the feature-space around it. This intuition is depicted in Figure 5. We stop the training when we reach a training loss lower than a given threshold.

Finally, the gradient for  $\xi$  in the non-target users is unconditionally suppressed by just setting both  $\gamma_i = \beta_i = 0$ , but other configurations are possible.



**Figure 7: Four different examples of false positive for the canary-gradient-based MIA. The first element in each sequence is the target, whereas the four following images are false triggers.**

**6.3.1 Results.** We evaluated the effectiveness of our attack. We test three different image datasets, namely, CIFAR10, CIFAR100 [37], and TinyImagenet [17]. We use all the possible permutations of those datasets to represent the private and shadow distribution for the canary-gradient attack, obtaining 6 different configurations.

To run the experiments, we pick  $x_t$  (i.e., the target of the membership inference) at random from the validation set of the private dataset. Then, we trained the model by injecting the canary-gradient in the channel 0 of the last normalization layer in the last residual block. After the injection, we evaluated the canary’s effectiveness by testing that the canary-gradient is non-zero when  $x_t$  is in the training batch and zero otherwise. To this end, we iterate over all the training data of the private dataset by selecting a batch  $X$  of size  $n$  at a time. Given  $X$ , we compute the gradient according to the original loss function and test that  $\xi$  has gradient zero (precision). Then, we insert  $x_t$  in  $X$ , and we test if the gradient of  $\xi$  is different from zero (recall). We perform the test on the three private datasets with different batch sizes and repeat the test for 50 different  $x_t$  for every case. We report our results in Figure 6, where we use the notation “*private-dataset*←*shadow-dataset*”.

**The canary gradient has perfect recall** (i.e., if  $x_t$  is in the batch, the canary is always triggered), but it can be subject to false-positive errors with low probability. The global accuracy of the method is about 99% for a batch size up to 128. The precision of the attacks slowly decreases when the batch size becomes larger. This follows from the fact that larger batches have more probability of including at least one “false trigger example” that causes a false positive. Intuitively, false positives are instances that induce a feature representation that is similar to the target one; a few examples are given in Figure 7. Therefore, while false triggers reduce the accuracy of the attack, these can still inform the attacker that instances *similar* to  $x_t$  are present in the victim’s local training set. The attacker can intentionally tune the definition of similarity by introducing the desired inductive bias in the canary injection process.<sup>9</sup> Regardless of the presence of false triggers, the attack is appreciably precise; the accuracy remains higher than 96% even in the worst configuration. **The canary-gradient attack also applies to FedAVG** by just requiring minimal effort to the malicious parameter server. In Appendix D, we empirically demonstrate the effectiveness of this approach.

<sup>9</sup>For instance, the server can train the canary to intentionally react only to “red cars”, by training it accordingly.

## 6.4 Impact

Although we presented the canary-gradient attack to execute a membership inference, the same approach can be applied to any type of property inference. Ideally, it is sufficient to define a different trigger condition and train the network accordingly. Moreover, the server can inject multiple canary-gradient with different triggers in the same network and infer non-binary properties. In the same direction, the server can simultaneously perform inference on multiple users without losing accuracy by carefully managing the allocation of conditional and unconditional dead-layers.

More importantly, the canary-gradient attack maintains the same practical properties as the gradient-suppression attack; mainly, it requires only a training round to perform, its effectiveness is independent of the number of users participating in the round, and it is loss agnostic (i.e., it works for any learning task). However, unlike the gradient suppression approach, this leaves only a minimal footprint in the model updates. Note that the canary gradient can be injected while training the model on the original task, and so minimizing the utility loss of the target model. On the other hand, suppressing a limited number of parameters in the non-target models (e.g., two in our example) has only a negligible impact on the utility.<sup>10</sup>

This attack demonstrates that a malicious server can perform highly accurate property inference attacks on individual users even if SA is adopted from the latter. Again, the introduced training-based approach is just an example, and more sophisticated techniques may be devised to reach the same (or better) results by relying on the same general construction.

For completeness, we emphasize that, even in the case of canary-gradient attacks, the discussion about the insecurity of the combination of SA and FL applies (discussed in Sections 1.1 and 5.3). As for the case of gradient suppression, the canary-gradient works if the SA protocol is perfectly secure (i.e., it behaves as the ideal functionality  $f^{sa}$ ).

## 7 MITIGATIONS

Next, we will discuss and introduce some mitigation approaches.

The first part focuses on heuristic approaches that strive to prevent model inconsistency and attack vectors similar to those seen in this work. While they do not stop all possible attacks, their practical relevance is still significant.

Then, we analyze how combining DP with SA can lead to a more general mitigation strategy.

### 7.1 Heuristic mitigations against model inconsistency

**SA dropout.** Although the gradient suppression attack (Section 5) is effective, it can be detected easily by non-target users. Indeed, non-target users can quickly discover an ongoing attack as their models would have zero gradients throughout.<sup>11</sup> However, the target (who is the victim of the attack) cannot detect or prevent the

<sup>10</sup>This is equivalent to perform dropout [57] on a single channel on a single layer in the network.

<sup>11</sup>A phenomenon that is very unlikely to be observed in reality, excluding numerical errors or pathologic overfitting.

attack without communicating with non-target users.<sup>12</sup> Therefore, a solution could be instructing users not to execute the SA protocol when they detect a null gradient after a local training iteration. In this case, only the attack victims would participate in the SA protocol, which is enough to prevent the server from recovering the target’s model update. Indeed, the server would not receive enough information to unmask the target model’s update and complete the execution of SA. This is because the dropout threshold  $\delta$  of the underlying SA protocol (i.e., the number of user dropout that the SA protocol can handle) must be set to  $\delta < n - 1$  where  $n$  is the number of users participating in the aggregation (otherwise, if  $\delta = n - 1$ , a malicious user can unmask the gradients of other users). See [10] for more details. Nevertheless, this detection strategy does not work in the case of the canary-gradient attack (Section 6). The “negligible” footprint of the model update does not allow for a reliable detection since sparse model updates are common in an honest execution of FL (see Figure 4).

*Parameter validation using signatures.* Another approach consists of checking that all users have received the same parameters  $\Theta^{(t)}$  in the current round  $t \in \mathbb{N}$  of FL. For example, each  $u_i \in \mathcal{U}^{(t)}$  sends the parameters  $\Theta_i^{(t)}$  (or its hash) sent by  $S$  to all other users. Then, each  $u_i \in \mathcal{U}^{(t)}$  checks that

$$\Theta_i^{(t)} = \Theta_j^{(t)} \text{ for every } u_i, u_j \in \mathcal{U}^{(t)}. \quad (14)$$

If the check fails, the user aborts. It is easy to see that this strategy does not allow  $S$  to execute model inconsistency attacks (note that it can still send a malicious parameters  $\Theta$  to all). Unfortunately, in the standard communication topology of FL (see Section 4), users do not have a direct communication channel. Each message needs to go through the server that, in turn, will forward the message to the intended receiver. Hence, a malicious server can perform a man-in-the-middle attack and substitute each  $\Theta_i^{(t)}$  sent by  $u_i$  with an arbitrary honest-looking parameters  $\Theta$ . If we want to preserve the communication topology of FL, digital signatures must be involved. For example, assuming that each  $u_i$  holds a key pair  $(sk_i, pk_i)$  then  $u_i$  needs to (i) sign its message before sending it to the server  $S$  and, (ii) verify the signature (using the public key  $pk_j$  of the signee) of each message received. If the security of the signature holds, then the server can not change the messages of users. As a consequence, users are guaranteed that the message is honest, and they can perform the check defined in Equation (14). We stress that this approach works under the assumption that users  $\mathcal{U}$  can trust and know (in advance) the public keys of all other users (e.g., there is a trusted certification authority). Note that this assumption is at the root of the SA protocols of Bonawitz et al. [10] and Bell et al. [8].

This approach would add one round of communication without modifying the implementation of the underlying SA protocol. However, this round could be merged with the ones of SA, e.g., the third round of the SA protocol of Bonawitz et al. [10, Figure 2]. Regarding the communication complexity, exchanging the signatures of the received parameters increases communication by  $n \cdot \ell$  where  $n$  is the number of active users and  $\ell$  is the size of one signature (e.g., 256 bits).

*Conditional secure aggregation.* Another mitigation consists of building a modified SA version for FL that performs the aggregation only if a particular condition  $C$  is satisfied. Otherwise, it outputs a random value (or fixed value  $\perp$  denoting that the aggregation did not occur). Intuitively, by setting the condition  $C$  to Equation (14), the FL protocol executes the aggregation (and continue its execution) only if all the users have received the same parameters  $\Theta^{(t)}$  in the current round  $t \in \mathbb{N}$ . This would hinder a malicious server  $S$  from exploiting the model inconsistency attack vector. Naturally, such a protocol can be built leveraging general MPC techniques. However, this would yield an inefficient aggregation that will not be deployed in practice. A candidate practical implementation of this SA for the specialized condition of Equation (14) can be easily obtained by modifying the SA protocol of Bonawitz et al. [10]. Still, we stress that this approach can also be applied to the SA protocol of Bell et al. [8]. In a nutshell, in [10] (and [8]) there is an ordering over the users  $\mathcal{U}$  and each pair  $(u_i, u_j)$  such that  $u_i \neq u_j$  share a random secret  $s_{i,j}$ . During the aggregation, each user  $u_i$  masks its input  $v_i$  in the following way:

$$y_i = v_i + \sum_{u_j \in \mathcal{U}: u_i < u_j} G(s_{i,j}) - \sum_{u_j \in \mathcal{U}: u_i > u_j} G(s_{j,i})$$

where  $G(\cdot)$  is a secure pseudorandom generator (PRG). Assuming no dropouts, the server can compute the aggregation as  $v = \sum_{u_i \in \mathcal{U}} v_i = \sum_{u_i \in \mathcal{U}} y_i$ . To enforce the condition  $C$  of Equation (14), we can simply substitute the PRG  $G(\cdot)$  with the evaluation of a pseudorandom function (PRF)  $F(s_{i,j}, \Theta_i^{(t)})$  where  $\Theta_i^{(t)}$  are the parameters received by  $u_i$  from  $S$ . It is easy to see that the aggregation remains hidden if two (or more) honest users receive two different parameters. Hence, the server  $S$  can not execute a model inconsistency attack. We stress that the presented solution (as discussed in [10]) is not resilient to dropouts. Still, the same technique can be applied seamlessly to both SA protocols (of Bonawitz et al. [10] and Bell et al. [8]) that handle users dropouts. This second approach does not need any additional round of communication. Moreover, unlike the previous approach (i.e., parameter validation using signatures), it preserves the communication complexity of the original SA protocols [8, 10] since we do not need to exchange signatures between active users.

Note that SA protocols such as [8, 10] (secure in the malicious setting) already require a PKI. Hence, the proposed solutions (parameter validation using signatures and conditional secure aggregation) can be implemented by using the existing PKI of the SA protocol. Nevertheless, patching model inconsistency becomes consistently harder for variations of the vanilla FL protocol such as Asynchronous Federated Learning [47, 61] which is gaining substantial interest thanks to its practical advantages. For instance, in the asynchronous SA protocol proposed in [55], solving model inconsistency would be a difficult task as aggregating model updates produced by different models is allowed by design. In these directions, solving model inconsistency efficiently remains an open problem.

<sup>12</sup>Note that the standard implementation of FL does not allow users to communicate.



## 7.2 Differential privacy and Secure Aggregation

One way to protect against the proposed attacks would be to mix SA with Differential Private-SGD algorithms [2]. However, unlike previous solutions (Section 7.1), differential privacy (DP) comes with a high utility cost, especially in the context of FL [66]. Regardless, standard central-DP approaches [42, 51] are ineffective when the parameter server is malicious. Here, the DP-noise is applied only after the model updates have been aggregated by the server. Trivially, if the server is malicious, it can just skip this step and obtain the target’s gradient in clear. Therefore, the proposed attacks remain unaffected. Peculiarly, pure central-DP is the approach used in state-of-the-art implementations [1] and employed in real-world deployments of FL [41, 51]. Nonetheless, user-level differential privacy can still be efficiently obtained in the presence of a dishonest parameter server by relying on the distributed-DP model [3, 14, 15, 34] that combines (partial) local noise application and SA to securely simulate central-DP.<sup>13</sup> More generally, as previous works have exhaustively shown it [9, 22, 32, 70], the application of local noise (i.e., local / distributed-DP) is sufficient to prevent gradient inversion and inference attacks on users’ model updates, and, therefore, the introduced attacks. In Appendix E, we offer a more detailed analysis of the privacy provided by the combination of SA and local-DP against the proposed attacks.

Nevertheless, the combination of SA and DP [3, 14, 15, 34, 58] is still partially susceptible to model inconsistency attacks since the adversary can isolate the gradient/parameters of the target user in the aggregation. That is, the server can still force the final aggregated value to be a function of the sole target’s training set. Thus, while a suitable amount of local noise still ensures the “privacy by aggregation” property of SA<sup>14</sup>, the “privacy by shuffling” property remains violated. Indeed, the information leaked from the aggregated model updates can still be traced back to the target—which is known to be the only source of information in the final aggregated value.

## 8 CONCLUSION

Our research found that federated learning implementations are susceptible to a critical vulnerability caused by incorrect usage of secure aggregation. As a result, the latter does not provide any additional security to users against a malicious server (even if a trusted PKI is assumed). The primary reason for the security issue is the lack of parameter validation, which would have prevented the server from providing inconsistent views of the global parameters to users.

We emphasize that the proposed attacks are just representative examples of threats induced by model inconsistency, and that other attacks may be devised by exploiting the same general intuition. In order to protect users’ privacy from current and future attacks, we argue that federated learning implementations must account for model inconsistency and prevent it at its source.

<sup>13</sup>Note that the server can still exploit the  $m$  compromised users allowed by the threat model and force them to participate in SA with zero gradient without applying the required noise. In the distributed DP model, this lets the server weaken the privacy guarantee for the target model update proportionally to  $m$ .

<sup>14</sup>To be precise, it implies a consistently stronger form of privacy than “privacy by aggregation” alone as the model update is now also differentially private.

## ACKNOWLEDGMENTS

The first author was supported by *Fondation Botnar*. The second author was supported by the Carlsberg Foundation under the Semper Ardens Research Project CF18-112 (BCM). The third author was supported by a grant from the Commonwealth Cyber Initiative (CCI). We acknowledge the generous support of Accenture and the collaboration with their Labs in Sophia Antipolis.

## REFERENCES

- [1] 2021. TensorFlow Federated. <https://www.tensorflow.org/federated> Accessed: 2021-12-09.
- [2] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) (CCS ’16). Association for Computing Machinery, New York, NY, USA, 308–318. <https://doi.org/10.1145/2976749.2978318>
- [3] Naman Agarwal, Peter Kairouz, and Ziyu Liu. 2021. The Skellam Mechanism for Differentially Private Federated Learning. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 5052–5064. <https://proceedings.neurips.cc/paper/2021/file/285baacbd8fda1de94b19282acd23e2-Paper.pdf>
- [4] Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shihō Moriai, et al. 2017. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security* 13, 5 (2017), 1333–1345.
- [5] Giuseppe Ateniese, Luigi V Mancini, Angelo Spognardi, Antonio Villani, Domenico Vitali, and Giovanni Felici. 2015. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *International Journal of Security and Networks* 10, 3 (2015), 137–150.
- [6] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *arXiv:1607.06450* [stat.ML]
- [7] Constance Beguier and Eric W Tramel. 2020. SAFER: Sparse Secure Aggregation for Federated Learning. *arXiv preprint arXiv:2007.14861* (2020).
- [8] James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. 2020. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 1253–1269.
- [9] Franziska Boenisch, Adam Dziedzic, Roei Schuster, Ali Shahin Shamsabadi, Ilia Shumailov, and Nicolas Papernot. 2021. When the Curious Abandon Honesty: Federated Learning Is Not Private. *CoRR abs/2112.02918* (2021). *arXiv:2112.02918* <https://arxiv.org/abs/2112.02918>
- [10] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1175–1191.
- [11] K. A. Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmityr Huba, Alex Ingberman, Vladimir Ivanov, Chloé M Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. 2019. Towards Federated Learning at Scale: System Design. In *SysML 2019*. <https://arxiv.org/abs/1902.01046> To appear.
- [12] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, J. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. Henighan, R. Child, A. Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. (2020).
- [13] Lukas Burkhhalter, Hidde Lycklama à Nijeholt, Alexander Viand, Nicolas Küchler, and Anwar Hithnawi. 2021. RoFL: Attestable Robustness for Secure Federated Learning. *arXiv:2107.03311* [cs.CR]
- [14] Wei-Ning Chen, Christopher A Choquette-Choo, Peter Kairouz, and Ananda Theertha Suresh. 2022. The Fundamental Price of Secure Aggregation in Differentially Private Federated Learning. *arXiv preprint arXiv:2203.03761* (2022).
- [15] Wei-Ning Chen, Ayfer Ozgur, and Peter Kairouz. 2022. The Poisson Binomial Mechanism for Unbiased Federated Learning with Secure Aggregation. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (Eds.). PMLR, 3490–3506. <https://proceedings.mlr.press/v162/chen22s.html>
- [16] Beongjun Choi, Jy-yong Sohn, Dong-Jun Han, and Jaekyun Moon. 2020. Communication-Computation Efficient Secure Aggregation for Federated Learning. *arXiv preprint arXiv:2012.05433* (2020).

- [17] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. 2017. A Downsampled Variant of ImageNet as an Alternative to the CIFAR datasets. *CoRR abs/1707.08819* (2017). arXiv:1707.08819 <http://arxiv.org/abs/1707.08819>
- [18] Ronald Cramer and Ivan Damgård. 2005. Multiparty computation, an introduction. In *Contemporary cryptography*. Springer, 41–87.
- [19] Ye Dong, Xiaojun Chen, Liyan Shen, and Dakui Wang. 2020. EaSTFLy: Efficient and secure ternary federated learning. *Computers & Security* 94 (2020), 101824.
- [20] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2020. *Local Model Poisoning Attacks to Byzantine-Robust Federated Learning*. USENIX Association, USA.
- [21] Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Helen Möllering, Thien Duc Nguyen, Phillip Rieger, Ahmad-Reza Sadeghi, Thomas Schneider, Hossein Yalame, et al. 2021. SAFELearn: Secure Aggregation for private Federated Learning. *IACR Cryptol. ePrint Arch.* 2021 (2021), 386.
- [22] Liam Fowl, Jonas Geiping, Wojtek Czaja, Micah Goldblum, and Tom Goldstein. 2021. Robbing the Fed: Directly Obtaining Private Data in Federated Learning with Modified Models. *arXiv preprint arXiv:2110.13057* (2021).
- [23] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 1322–1333.
- [24] David Froelicher, Juan R Troncoso-Pastoriza, Apostolos Pyrgelis, Sinem Sav, Joao Sa Sousa, Jean-Philippe Bossuat, and Jean-Pierre Hubaux. 2021. Scalable Privacy-Preserving Distributed Learning. *Proceedings on Privacy Enhancing Technologies* 2 (2021), 323–347.
- [25] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting Gradients - How easy is it to break privacy in federated learning?. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 16937–16947. <https://proceedings.neurips.cc/paper/2020/file/c4ede56bbd98819ae6112b20ac6bf145-Paper.pdf>
- [26] Xiaojie Guo, Zheli Liu, Jin Li, Jiqiang Gao, Boyu Hou, Changyu Dong, and Thar Baker. 2020. V eri FL: Communication-Efficient and Fast Verifiable Aggregation for Federated Learning. *IEEE Transactions on Information Forensics and Security* 16 (2020), 1736–1751.
- [27] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. 2019. Federated Learning for Mobile Keyboard Prediction. arXiv:1811.03604 [cs.CL]
- [28] Hanieh Hashemi, Yongqin Wang, Chuan Guo, and Murali Annamavaram. 2021. Byzantine-robust and privacy-preserving framework for fedml. *arXiv preprint arXiv:2105.02295* (2021).
- [29] K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [30] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. 2017. Deep models under the GAN: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 603–618.
- [31] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 2261–2269. <https://doi.org/10.1109/CVPR.2017.243>
- [32] Yangsibo Huang, Samyak Gupta, Zhao Song, Kai Li, and Sanjeev Arora. 2021. Evaluating Gradient Inversion Attacks and Defenses in Federated Learning. In *Thirty-Fifth Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=0CDKGyYaxC8>
- [33] Swanand Kadhe, Nived Rajaraman, O Ozan Koyluoglu, and Kannan Ramchandran. 2020. Fastsecagg: Scalable secure aggregation for privacy-preserving federated learning. *arXiv preprint arXiv:2009.11248* (2020).
- [34] Peter Kairouz, Ziyu Liu, and Thomas Steinke. 2021. The distributed discrete gaussian mechanism for federated learning with secure aggregation. In *International Conference on Machine Learning*. PMLR, 5201–5212.
- [35] Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. 2016. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. arXiv:1610.02527 [cs.LG]
- [36] Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2017. Federated Learning: Strategies for Improving Communication Efficiency. arXiv:1610.05492 [cs.LG]
- [37] Alex Krizhevsky. 2009. Learning Multiple Layers of Features from Tiny Images.
- [38] Maximilian Lam, Gu-Yeon Wei, David Brooks, Vijay Janapa Reddi, and Michael Mitzenmacher. 2021. Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix. In *Proceedings of the 38th International Conference on Machine Learning*.
- [39] Lu Lu. 2020. Dying ReLU and Initialization: Theory and Numerical Examples. *Communications in Computational Physics* 28, 5 (Jun 2020), 1671–1706. <https://doi.org/10.4208/cicp.0a-2020-0165>
- [40] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [41] H. Brendan McMahan and Thakurta Abhradeep. 2022. Federated Learning with Formal Differential Privacy Guarantees. <https://ai.googleblog.com/2022/02/federated-learning-with-formal.html>
- [42] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BJ0hF1Z0b>
- [43] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. 2019. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 691–706.
- [44] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. 2021. PPFL: Privacy-Preserving Federated Learning with Trusted Execution Environments. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services* (Virtual Event, Wisconsin) (MobiSys '21). Association for Computing Machinery, New York, NY, USA, 94–108. <https://doi.org/10.1145/3458864.3466628>
- [45] Vinod Nair and Geoffrey E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning* (Haifa, Israel) (ICML '10). Omnipress, Madison, WI, USA, 807–814.
- [46] Milad Nasr, Reza Shokri, and Amir Houmansadr. 2019. Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning. In *2019 IEEE Symposium on Security and Privacy (SP)*. 739–753. <https://doi.org/10.1109/SP.2019.00065>
- [47] John Nguyen, Kshitiz Malik, Hongyuan Zhan, Ashkan Yousefpour, Michael Rabbat, Mani Malek, and Dzmityr Huba. 2021. Federated learning with buffered asynchronous aggregation. *arXiv preprint arXiv:2106.06639* (2021).
- [48] Xudong Pan, Mi Zhang, Yifan Yan, Jiaming Zhu, and Min Yang. 2020. Theory-Oriented Deep Leakage from Gradients via Linear Equation Solver. arXiv:2010.13356 [cs.CR]
- [49] Krishna Pillutla, Sham M. Kakade, and Zaid Harchaoui. 2019. Robust Aggregation for Federated Learning. arXiv:1912.13445 [stat.ML]
- [50] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
- [51] Swaroop Ramaswamy, Om Thakkar, Rajiv Mathews, Galen Andrew, H Brendan McMahan, and Françoise Beaufays. 2020. Training production language models without memorizing user data. *arXiv preprint arXiv:2009.10031* (2020).
- [52] Sinem Sav, Apostolos Pyrgelis, Juan R Troncoso-Pastoriza, David Froelicher, Jean-Philippe Bossuat, Joao Sa Sousa, and Jean-Pierre Hubaux. 2020. POSEIDON: Privacy-preserving federated neural network learning. *arXiv preprint arXiv:2009.00349* (2020).
- [53] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. 1310–1321.
- [54] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.
- [55] Jinhyun So, Ramy E. Ali, Basak Güler, and Amir Salman Avestimehr. 2021. Secure Aggregation for Buffered Asynchronous Federated Learning. *CoRR abs/2110.02177* (2021). arXiv:2110.02177 <https://arxiv.org/abs/2110.02177>
- [56] Jinhyun So, Başak Güler, and A Salman Avestimehr. 2021. Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning. *IEEE Journal on Selected Areas in Information Theory* 2, 1 (2021), 479–489.
- [57] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15, 56 (2014), 1929–1958. <http://jmlr.org/papers/v15/srivastava14a.html>
- [58] Timothy Stevens, Christian Skalka, Christelle Vincent, John Ring, Samuel Clark, and Joseph Near. 2022. Efficient Differentially Private Secure Aggregation for Federated Learning via Hardness of Learning with Errors. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA. <https://www.usenix.org/conference/usenixsecurity22/presentation/stevens>
- [59] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. 2019. A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*. 1–11.
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [61] Chenhao Xu, Youyang Qu, Yong Xiang, and Longxiang Gao. 2021. Asynchronous Federated Learning on Heterogeneous Devices: A Survey. *CoRR abs/2109.04269*

- (2021). arXiv:2109.04269 <https://arxiv.org/abs/2109.04269>
- [62] Guowen Xu, Hongwei Li, Sen Liu, Kan Yang, and Xiaodong Lin. 2019. Verifynet: Secure and verifiable federated learning. *IEEE Transactions on Information Forensics and Security* 15 (2019), 911–926.
- [63] Runhua Xu, Nathalie Baracaldo, Yi Zhou, Ali Anwar, and Heiko Ludwig. 2019. Hybridalpha: An efficient approach for privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*. 13–23.
- [64] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. 2018. APPLIED FEDERATED LEARNING: IMPROVING GOOGLE KEYBOARD QUERY SUGGESTIONS. *ArXiv abs/1812.02903* (2018).
- [65] H. Yin, A. Mallya, A. Vahdat, J. M. Alvarez, J. Kautz, and P. 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>
- [66] Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. 2020. Salvaging Federated Learning by Local Adaptation. *CoRR abs/2002.04758* (2020). arXiv:2002.04758 <https://arxiv.org/abs/2002.04758>
- [67] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. 2020. Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning. In *2020 {USENIX} Annual Technical Conference ({USENIX} {ATC} 20)*. 493–506.
- [68] Wenting Zheng, Ryan Deng, Weikeng Chen, Raluca Ada Popa, Aurojit Panda, and Ion Stoica. 2021. Cerebro: A platform for multi-party cryptographic collaborative learning. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*.
- [69] Junyi Zhu and Matthew B. Blaschko. 2020. R-GAP: Recursive Gradient Attack on Privacy. *CoRR abs/2010.07733* (2020). arXiv:2010.07733 <https://arxiv.org/abs/2010.07733>
- [70] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep Leakage from Gradients. 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. <https://proceedings.neurips.cc/paper/2019/file/60a6c4002cc7b29142def8871531281a-Paper.pdf>

## A ALTERNATIVES TO FEDERATED LEARNING

Aono et al. [4] leverage a homomorphic encryption scheme to permit users to share encrypted model updates (e.g., gradients) with the server. The latter will compute the new model parameters in an encrypted fashion, using the additive homomorphic properties of the underlying encryption scheme. The new parameters remain encrypted on the server-side.

Truex et al. [59] use a similar approach by leveraging the homomorphic properties of Pailler’s cryptosystem. Here, the Pailler encryption scheme is necessary to allow the server to aggregate the model updates in an encrypted form. In contrast to [4], the aggregation will be decrypted by the users and made public to the server.

An analogous approach is used by EastFly [19] and BatchCrypt [67] with the main difference of keeping the aggregated updates secret to the server: The aggregation is locally decrypted by the users that will also compute the new model parameters (this differs from the original FL protocol [40, 53] in which the parameters are stored and updated by the server). In addition, [67] proposes a batching encoding scheme (that preserves the homomorphic properties of the underlying homomorphic encryption schemes) to reduce the number of encryption operations and speed up the efficiency of the protocol.

HybridAlpha [63] uses a multi-input functional encryption scheme to compute the aggregation. Informally, each user  $u_i$  computes the encryption  $Enc_{pk_i}(\Delta_{\mathcal{D}_i}^\Theta)$  of its update  $\Delta_{\mathcal{D}_i}^\Theta$ . Once the server has received all the ciphertexts, it computes the aggregation by decrypting them using the decryption key  $sk_f$  for the functionality  $f(\Delta_{\mathcal{D}_1}^\Theta, \dots, \Delta_{\mathcal{D}_n}^\Theta) = \sum_{i=1}^n \Delta_{\mathcal{D}_i}^\Theta$ . Analogously, SAFElearn [21]

leverages either multi-party computation or fully homomorphic encryption to protect the individual’s updates of users.

Poseidon [52] significantly deviates from the original FL architecture. In such a system, users are organized according to a tree hierarchy, and a multi-key fully homomorphic encryption scheme is used to protect users’ updates and the parameters of the neural network. At each round of the training phase, the root user sends the encrypted parameters to all users that, in turn, compute their updates according to their local training data. Then, each encrypted user’s update is sent to the parent that will then aggregate all children’s updates. At the end of this recursive process, the root receives the encrypted aggregation (that contains the updates of all users) and will update the parameters. The entire computation of the training process is executed inside the multi-key fully homomorphic encryption scheme. This permits them to keep the updates and the parameters encrypted. The latter remains encrypted even after the training process, and the model evaluation requires further computation inside the multi-key fully homomorphic encryption scheme. Note that Poseidon [52] is an extension of SPINDLE [24]. The former handles complex machine learning models (such as neural networks), while the latter supports only generalized linear models.

Lastly, we mention Cerebro [68] that proposes a compiler to automatically transform Python-like domain-specific language into an optimized MPC protocol for collaborative learning allowing users to keep their plaintext data secret. We stress that Cerebro [68] does not relate in any way with FL [40, 53] and it must be interpreted as an alternative to FL.

*Attack applicability.* As discussed in Section 4, our attacks are general and equally effective independently from the SA protocol used. This because they exploit a vulnerability of FL caused by the incorrect usage of SA. More precisely, our main attack applies to all the FL/SA protocols that do not prevent the parameter server from deviating from the honest execution. This class also includes most of the schemes that rely on fully homomorphic encryption (FHE) e.g., [21]. The reason is that FHE would still allow the server to multiply the individual encrypted parameters by the constant 0 and produce 0 gradient everywhere when used by the non-target users (see Appendix C). The only requirement is that server is required to collude with a least one user in order to access the result of the attack (the target model update) once decrypted by the pool (e.g.,  $m > 0$ ).

Protocols based on Trusted Execution Environment (TEE) can stop the server from executing malicious code [28, 44]. Thus, a properly deployed TEE environment with ideal/perfect hardware and secure and authenticated communication would prevent all the active attacks currently in the literature [9, 22, 38], including ours. However, the reality is that trusted hardware is vulnerable to side-channel attacks, and there is significant performance degradation when extending side-channel protections to arbitrary computations. Therefore, it remains unclear whether side-channel attacks can be entirely eliminated.



## Federated Learning

*Inputs.* For every  $u_i \in \mathcal{U}$ , user  $u_i$  holds a local training dataset  $\mathcal{D}_i = \{(x_j, y_j)\}_{j \in \{1, \dots, \ell\}}$ . The server  $S$  fixes the architecture  $f$  of the deep neural network, the learning parameter  $\eta \in [0, 1]$ , and the initial parameters  $\Theta^{(1)}$ .

*Goal.* The server  $S$  obtains the final model parameters  $\Theta$  of the deep neural network  $f$ .

*The protocol:*

- **For  $t \in \mathbb{N}$ , until  $\Theta^{(t)}$  converges.**
  - (1)  $S$  samples a random subset of users  $\mathcal{U}^{(t)} \subseteq \mathcal{U}$  that will participate in training of  $f$  in the current round  $t$ .  $S$  sends the parameters  $\Theta^{(t)}$  to each  $u_i \in \mathcal{U}^{(t)}$ .
  - (2) Each user  $u_i \in \mathcal{U}^{(t)}$  receives  $\Theta^{(t)}$  and proceeds as follows:
    - (a) **If FedSGD.**  $u_i$  samples a random training batch  $\mathcal{D}_i^{(t)} \subseteq \mathcal{D}_i$  and computes the gradient  $\nabla_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$ . Finally, it sets its model update  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}} = \nabla_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$ .
    - (b) **If FedAVG.** Let  $\Theta_i^{(t,1)} = \Theta^{(t)}$ . For  $j \in \{1, \dots, k\}$ ,  $u_i$  samples a random training batch  $\mathcal{D}_i^{(t,j)} \subseteq \mathcal{D}_i$  and computes the gradient  $\nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta^{(t,j)}}$ . Then, it updates the model parameters by computing  $\Theta_i^{(t,j+1)} = \Theta_i^{(t,j)} - \eta \cdot \nabla_{\mathcal{D}_i^{(t,j)}}^{\Theta^{(t,j)}}$ . Finally, it sets  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}} = \Theta_i^{(t,k)}$  and  $b_i = \sum_{j=1}^k |\mathcal{D}_i^{(t,j)}|$ .
  - (3) Each user  $u_i \in \mathcal{U}^{(t)}$  sends its model update as follows:
    - (a) **If SA disabled.** It sends  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  to  $S$ . In addition, **if FedAVG**,  $u_i$  sends  $b_i$  to  $S$ .
    - (b) **If SA enabled.** It provides the input  $\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  to  $f^{\text{sa}}$ . In addition, **if FedAVG**,  $u_i$  sends  $b_i$  to  $S$ .
  - (4)  $S$  computes the new model parameters  $\Theta^{(t+1)}$  in the following way:
    - (a) **If SA disabled.**
      - (i) **If FedSGD.**  $S$  receives  $\{\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}\}_{u_i \in \mathcal{U}}$  from users and computes  $\Theta^{(t+1)} = \Theta^{(t)} - \eta \cdot v/b$  where  $v = \sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  and  $b = |\mathcal{U}^{(t)}|$ .
      - (ii) **If FedAVG.**  $S$  receives  $\{(\Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}, b_i)\}_{u_i \in \mathcal{U}}$  from users and computes  $\Theta^{(t+1)} = v/b$  where  $v = \sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  and  $b = \sum_{u_i \in \mathcal{U}^{(t)}} b_i = \sum_{u_i \in \mathcal{U}^{(t)}} \sum_{j=1}^k |\mathcal{D}_i^{(t,j)}|$ .
    - (b) **If SA enabled.**  $S$  receives  $v = \sum_{u_i \in \mathcal{U}^{(t)}} \Delta_{\mathcal{D}_i^{(t)}}^{\Theta^{(t)}}$  from  $f^{\text{sa}}$ . Then, it proceeds as follows:
      - (i) **If FedSGD.**  $S$  computes  $\Theta^{(t+1)}$  as in Item 4(a)i using  $v$  outputted by  $f^{\text{sa}}$ .
      - (ii) **If FedAVG.**  $S$  receives  $\{b_i\}_{u_i \in \mathcal{U}}$  from users and computes  $\Theta^{(t+1)}$  as in Item 4(a)ii using  $v$  outputted by  $f^{\text{sa}}$ .

**Figure 8: Description of FedSGD-based and FedAVG-based FL with SA either enabled or disabled. The execution of SA is represented by the ideal functionality  $f^{\text{sa}}$ .**

## B SECURITY OF SECURE AGGREGATION

*Additional notation.* We model cryptographic algorithms (e.g., adversary) as (possibly interactive) Turing machines. If  $A$  is a deterministic (resp. randomized) algorithm, we write  $y = A(x)$  to denote a run of  $A$  on input  $x$  and output  $y$ ; if  $A$  is randomized,  $y$  is a random variable. An algorithm  $A$  is probabilistic polynomial-time (PPT) if  $A$  is randomized and for any input  $x \in \{0, 1\}^*$  the computation of  $A(x)$  terminates in a polynomial number of steps (in the input size). We denote by  $\lambda \in \mathbb{N}$  the security parameter of cryptographic primitives, and we implicitly assume that every algorithm takes as input the security parameter (written in unary). A function  $v : \mathbb{N} \rightarrow [0, 1]$  is called negligible in the security parameter  $\lambda$  if it vanishes faster than the inverse of any polynomial in  $\lambda$ , i.e.  $v(\lambda) \in O(1/p(\lambda))$  for all positive polynomials  $p(\lambda)$ . We write  $\text{negl}(\lambda)$  to denote an unspecified negligible function in the security parameter. Similarly, we write  $\text{poly}(\lambda)$  to denote all possible polynomials  $p(\lambda)$ . Let  $X$  and  $Y$  be two random variables. We say that  $X$  and  $Y$  are computationally indistinguishable, denoted  $X \approx_c Y$ , if for all PPT distinguishers  $D$

we have

$$\left| \Pr[D(1^\lambda, X) = 1] - \Pr[D(1^\lambda, Y) = 1] \right| \leq \text{negl}(\lambda).$$

We now describe the ideal and real-world paradigm of MPC [18] that defines the security of MPC protocols (including SA).

*Real world.* The real world refers to the scenario in which the real protocol  $\pi$  is executed between the users  $\mathcal{U}$ . During its execution the parties  $\mathcal{U}$  exchange messages between themselves in order to compute the  $f(x_1, \dots, x_n) = (y_1, \dots, y_n)$ . In this setting, we assume the presence of an adversary  $A$  that can be either semi-honest or malicious. A semi-honest  $A$  can take control of a subset of users  $\tilde{\mathcal{U}} \subset \mathcal{U}$ . This will allow  $A$  to have access to their inputs  $\{x_i\}_{u_i \in \tilde{\mathcal{U}}}$ , the messages received, and the final output  $f_i(x_1, \dots, x_n) = y_i$ . In addition, if  $A$  is malicious, then it can also program a corrupted user  $u_i \in \tilde{\mathcal{U}}$  to deviate from the original protocol specification (e.g., send malicious messages). We use the notation  $V_i^\pi(x^*)$  to denote the view of the  $i$ -th party, i.e.,

$$V_i^\pi(x^*) = (x_i, r_i, m_1, \dots, m_k)$$



where  $x^* = (x_1, \dots, x_n)$ ,  $r_i$  are the (private) random coins of  $u_i$ , and  $m_j$  is the  $j$ -th message received by  $u_i$ .

*Ideal world.* In the ideal world, the protocol execution is replaced by a trusted third party that honestly computes  $f$ . In more detail, each user  $u_i$  sends its input  $x_i$  to the third party. The latter computes  $f(x_1, \dots, x_n) = (y_1, \dots, y_n)$  and returns  $y_i$  to each user  $x_i$ . Even in this setting, the (either semi-honest or malicious) adversary  $A$  can corrupt a set of parties  $\tilde{\mathcal{U}}$  and, in turn, obtains their private inputs  $\{x_i\}_{u_i \in \tilde{\mathcal{U}}}$  and outputs  $\{y_i\}_{u_i \in \tilde{\mathcal{U}}}$ .

*Security.* At a high level, the security of  $\pi$  is defined by comparing the ideal and real world. In particular,  $\pi$  is secure if there exists a simulator  $S$  that simulates the view of a real-world adversary  $A$  by leveraging the interactions performed in the ideal world. This implies that any attack in the real world “corresponds” to an attack in the ideal world that, in turn, provides the maximum security guarantees that can be achieved. In other words, a secure protocol  $\pi$  provides at least the same level of security as if we had an honest, trusted party computing  $f$  without exposing to an adversary the inputs of uncorrupted users.

*Definition B.1.* Let  $\pi$  and  $\mathcal{U} = \{u_1, \dots, u_n\}$  be a  $n$ -party protocol that correctly computes an  $n$ -inputs  $n$ -outputs function  $f$  and the set of users that participate in the protocol execution, respectively. Let  $x_i$  be the input of user  $u_i \in \mathcal{U}$  and let  $x^* = (x_1, \dots, x_n)$ . For  $\tilde{\mathcal{U}} = \{u_{i_1}, \dots, u_{i_k}\} \subset \mathcal{U}$ , we let  $V_{\tilde{\mathcal{U}}}^\pi(x^*) = (V_{i_1}^\pi(x^*), \dots, V_{i_k}^\pi(x^*))$  and  $f_{\tilde{\mathcal{U}}}(x^*) = (y_{i_1}, \dots, y_{i_k})$ . We say that  $\pi$  securely computes  $f$  if there exists a PPT simulator  $S$  such that, for every  $\tilde{\mathcal{U}} \subset \mathcal{U}$ , for every input  $x^* = (x_1, \dots, x_n)$ , we have:

$$\{S(\tilde{\mathcal{U}}, \tilde{x}, f_{\tilde{\mathcal{U}}}(x^*))\}_{x^* \in \{0,1\}^n} \approx_c \{V_{\tilde{\mathcal{U}}}^\pi(x^*)\}_{x^* \in \{0,1\}^n}$$

where  $\tilde{x} = (x_{i_1}, \dots, x_{i_k})$  for  $u_{i_j} \in \tilde{\mathcal{U}}$ .

## C GRADIENT SUPPRESSION FOR ARBITRARY ARCHITECTURES

In Section 5 we focused on ReLU-based layers as these are instrumental for our second attack (Section 6). Nevertheless, gradient suppression can be achieved for networks composed of arbitrary activation functions at the cost of flexibility and granularity. Intuitively, we can force any differentiable function to unconditionally produce a zero gradient by constraining it to become constant with respect to the differentiated terms. In the context of neural networks, this means making the loss function constant with respect to the trainable parameters of the network. Formally, given a neural network  $f_{\Theta^{(t)}}$  and a loss function  $\mathcal{L}$ , this can be achieved by bringing  $\Theta^{(t)}$  in a state such that:

$$\forall x, y, \mathcal{L}(y, f_{\Theta^{(t+1)}}(x)) = c, \quad (15)$$

where  $c$  is an arbitrary constant and  $\Theta^{(t+1)}$  represents any possible alteration of the parameters of the network achievable starting from  $\Theta^{(t)}$ . Intuitively, when in this state, the parameters of the network have zero gradient since their alteration does not affect the loss function.

Bringing an arbitrary architecture in the state described in Equation (15) is trivial, and, for standard feedforward networks, this requires to act only on the kernels of the *last and penultimate layer*.

Hereafter, we refer to the composition of these last two layers as  $\phi_l(\phi_p(a \otimes \theta_p + b_p) \otimes \theta_l + b_l)$  where  $\phi_l$  and  $\phi_p$  are the two arbitrary activation functions for the last and penultimate layer respectively and  $a$  represents the intermediate state of the network up to the penultimate layer.

To suppress the gradient for most of the network, the kernel of the last layer  $\theta_l$  must be set to  $[0]$ . In this way, all transformations applied to the network’s input by the layers up to the last layer are nullified. Intuitively, this cuts off all parameters of the layers up to the last one from the loss computation, making their derivative 0 for every input. However, in this state, the kernel  $\theta_l$  still receives the gradient since its alteration still affects the network’s output (and so, the loss function.)

Now, to suppress the gradient for the kernel  $\theta_l$  of the last layer, the attacker needs to act on its input; that is, the output of the penultimate layer. In particular, in order to cut off the contribute of the kernel  $\theta_l$  from the loss computation, the attacker needs the output of the penultimate layer to be  $[0]$ , i.e.,  $\forall a, \phi_p(a \otimes \theta_p + b_p) = [0]$ . Indeed, given the multiplicative nature of the operation  $\otimes$  between the kernel  $\theta_l$  and the  $\phi_p(a \otimes \theta_p + b_p)$ , the assignment  $\phi_p(a \otimes \theta_p + b_p) = [0]$  completely nullifies the contribute of  $\theta_l$  in the layer output, making its derivative zero.

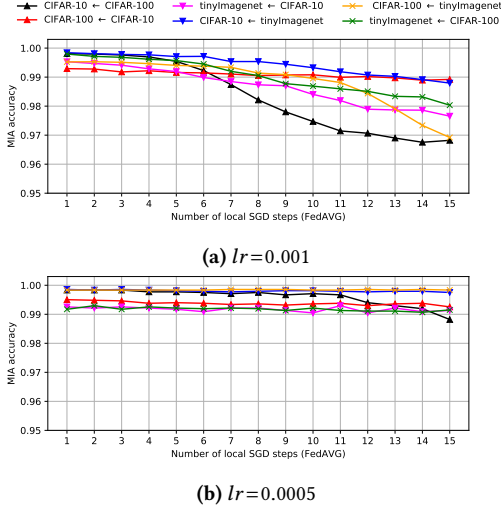
This can be achieved by choosing appropriate values for the  $\theta_p$  and  $b_p$  in the penultimate layer. In particular, the attacker can set  $\theta_p = [0]$  and play with the bias term  $b_p$  to force the activation  $\phi_p$  to output zero. Indeed, when  $\theta_p = [0]$ , we have that  $\phi_p(x \otimes \theta_p + b_p) = \phi_p(b_p)$  and the attacker just needs to set  $b_p$  in such a way that  $\phi_p(b_p) = [0]$ . This is possible for almost every activation function. In practice, this is even possible for the *Sigmoid* function that is zero only when its input is  $-\infty$ .<sup>15</sup>

After having nullified the contribution of all kernels in the network, the model is now the constant function  $f_\Theta(x) = \phi_l(b_l)$ . Therefore, the bias term  $b_l$  (if any) can still receive a non-zero gradient during the application of SGD. Also in this case, the server can either remove the last bias term from the architecture or ignore it during the attack phase (e.g., gradient inversion).

## D CANARY-GRADIENT ON FEDAVG

At least theoretically, extending the canary-gradient attack presented in Section 6 to the FedAVG protocol can be problematic. In FedAVG, a user locally applies multiple iterations of SGD, modifying the parameters of the model. Thus, there is the possibility that the user overwrites the canary functionality during the process. However, this can be easily prevented by the server. Indeed, given the stateless nature of users in FL, the parameter server must distribute the hyper-parameters needed for the training process. In this direction, the server takes care of the learning rate. Therefore, the server can simply select a low learning rate to preserve the canary’s functionality. The rationale here is that the perturbation of the local parameters induced by the local training steps is proportional to the learning rate. A low learning rate ensures to the server a bounded modification of the local parameters and, so, a limited perturbation of the hidden canary functionality originally injected in the model. Instead, the partial dead-kernel in the non-target’s

<sup>15</sup>It is sufficient to set  $b_p$  to a large negative number.



**Figure 9: Average accuracy of canary-gradient attacks for six different setups with increasing number of FedAVG local training steps for a ResNet20.**

model cannot be revived regardless of the number of iterations and the given learning rate.


In practice, even setting a standard learning rate such as 0.001 (which is quite common, especially towards the end of the training) seems enough to preserve the canary functionality during the local learning iterations. This is shown in Figure 9a, where the canary functionality is tested after an increasing number of local training steps. We rely on the same setup of Section 6.3.1, and we use a batch size of 64. As it can be observed, the accuracy of the attack does naturally decrease with the increase in the number of iterations. However, this performance loss is limited, and the accuracy always remains in the 97%-range even in the worst case. Nevertheless, there is a trade-off between effectiveness and stealthiness of the attack, and such a performance loss can be further reduced by considering a lower learning rate. This is shown in Figure 9b, for a learning rate of 0.0005. In this case, we do not need model inconsistency, and the server can send the same hyper-parameters (learning rate) to all the users in the pool without distinction.

Once computed the output of the secure aggregation, the server can state if the canary has been triggered by checking if  $\xi_t \neq \xi_{t+1}$ . Intuitively, the parameters  $\xi$  are modified if and only if the target user produced a non-zero gradient for  $\xi$  during the training at least once (i.e., the trigger condition has been met).

## E A NOTE ON THE COMBINATION OF SA WITH LOCAL DIFFERENTIAL PRIVACY AS A DEFENSE

When SA is combined with local-DP, the amount of noise applied to the model update recovered by our attacks is proportional to the number of users participating in the round.

In the local-DP setting without SA, the model update of the user  $u_{\text{trgt}}$  accessible by the parameter server is:



**Figure 10: Noise amplification effect on the target's model update when SA is combined with local-DP.**

$$\hat{\Delta}_{\mathcal{D}_{\text{trgt}}}^{\Theta^{(t)}} + \mathfrak{N}_{\epsilon},$$

where  $\mathfrak{N}_{\epsilon}$  is the local noise applied by  $u_{\text{trgt}}$  to make the model updated  $(\epsilon)$ -differentially-private and  $\hat{\Delta}_{\mathcal{D}_{\text{trgt}}}^{\Theta^{(t)}}$  is the clipped version of the model update. However, the situation is different when SA is combined with local-DP. In this case, assuming that all the active users  $\mathcal{U}^{(t)}$  apply an equal amount of noise, then the attacks recover a model update defined as:

$$\hat{\Delta}_{\mathcal{D}_{\text{trgt}}}^{\Theta^{(t)}} + \mathfrak{N}_{\epsilon} + (|\mathcal{U}^{(t)}| - 1) \cdot \mathfrak{N}_{\epsilon}. \quad (16)$$

As the target model update inherits the input noise added by non-target users, the amount of noise and, therefore, the degree of protection increases with the number of active users.

In practice, considering DP-SGD based on the Gaussian mechanism [2] and assuming an ideal functionality of SA working on real vectors<sup>16</sup>, we can rewrite Equation (16) as:

$$\hat{\Delta}_{\mathcal{D}_{\text{trgt}}}^{\Theta^{(t)}} + \mathcal{N}(0, |\mathcal{U}^{(t)}| \cdot \sigma_{(\epsilon, \delta)}^2),$$

where  $\sigma_{(\epsilon, \delta)}^2$  is the variance of the Gaussian distribution required to achieve  $(\epsilon, \delta)$ -differential-privacy. Figure 10 shows the “privacy amplification effect” of the combination of SA and local-DP with respect to an increase in active users for three different settings of local-DP:  $(\epsilon=4.58, \delta)$ ,  $(\epsilon=2.88, \delta)$ , and  $(\epsilon=1.63, \delta)$  with  $\delta=5 \cdot 10^{-5}$ . The y-axis reports the final privacy budget of the model update recovered by the server after the application of SA. In the plot, we consider FedSGD with a batch size of 64 and a local training set of 64 instances per user. Noise multipliers are 1, 1.5, and 2.5 respectively, with  $\ell_2$ -norm-clip of 1. The privacy budget is computed with [2] via its tensorflow-privacy implementation.

To summarize, when local-DP is combined with SA, it is possible to obtain a privacy amplification effect that is proportional to the number of non-target users. Compared to having no SA, users need to add less noise in the local-DP+SA regime to achieve the same level of protection when using a suitable trust-model.<sup>17</sup>

<sup>16</sup>Current SA protocols work in the discrete domain. However, the sum of discrete (independent) Gaussians is not a discrete Gaussian (see [34] for details).

<sup>17</sup>Other users must be honest and add the expected amount of noise.