

# Comments on “Privacy-Enhanced Federated Learning Against Poisoning Adversaries”\*

Thomas Schneider<sup>✉</sup>, Ajith Suresh<sup>✉</sup>, and Hossein Yalame<sup>✉</sup>  
 Technical University of Darmstadt, Germany

**Abstract**—In August 2021, Liu et al. (IEEE TIFS’21) proposed a privacy-enhanced framework named PEFL to efficiently detect poisoning behaviours in Federated Learning (FL) using homomorphic encryption. In this article, we show that PEFL does not preserve privacy. In particular, we illustrate that PEFL reveals the entire gradient vector of all users in clear to one of the participating entities, thereby violating privacy. Furthermore, we clearly show that an immediate fix for this issue is still insufficient to achieve privacy by pointing out multiple flaws in the proposed system.

**Index Terms**—Federated Learning (FL), Homomorphic Encryption, Poisoning and Inference Attacks, Data Privacy.

## I. INTRODUCTION

Federated Learning (FL) is a new distributed machine learning approach that allows multiple entities to jointly train a model without sharing their private and sensitive local datasets with others. In FL, clients locally train models using their local training data, then send model updates to a central aggregator, which merges them into a global model. FL is used in a variety of applications such as word prediction for mobile keyboards in GBoard [7] and medical imaging [10]. Despite its benefits, FL has been shown to be susceptible to model poisoning [8] and inference attacks [2]. In model poisoning attacks, an adversary injects poisoned model updates by corrupting a subset of clients, with which the adversary can compromise the user’s data privacy as well as the FL model’s integrity [4], [5]. The recent work of Liu et al. [6] proposed a privacy-enhanced framework called PEFL to detect poisoning behaviors in FL. PEFL aims to prevent malicious users from inferring memberships by uploading malicious gradients and semi-honest servers from invading users’ privacy. Furthermore, PEFL claims to be the first effort to detect poisoning behaviors in FL while using ciphertext and uses homomorphic encryption (HE) as the underlying technology.

In this article, we have a closer look at the PEFL system of [6] and identify multiple privacy vulnerabilities. In particular, we show that each of the three main protocols in PEFL – SecMed, SecPear, and SecAgg, reveals significant information about the user’s gradients to one of the computing servers thereby compromising privacy. Furthermore, we demonstrate that combining information from the protocols enables a computation server to learn the gradient vectors of all users in clear, thereby breaking the PEFL system.

## II. LIU ET AL.’S PROTOCOLS ARE NOT PRIVATE

In this section, we revisit the Privacy Enhanced Federated Learning (PEFL) system in [6], but with our notations for clarity. We begin with an overview of PEFL’s four entities:

- *Key Generation Center (KGC)*: Trusted entity managing public and private keys  $(pk, sk)$  for HE.
- *Data Owners ( $U_x$ )*: Each data owner  $U_x$ , for  $x \in [m]$ , locally trains the local model on their data and computes the gradient vector  $\vec{g}_x = \{g_x^1, \dots, g_x^n\}$ . Here,  $m$  denotes the total number of users in the system and  $n$  denotes the dimension of the gradient vector.
- *Service Provider (SP)*: SP receives all gradients submitted by data owners and aggregates them (usually by averaging) to produce an optimized global model.
- *Cloud Platform (CP)*: CP assists SP in the computations and operates on a pay-per-use basis.

The threat model assumes that SP and CP are both semi-honest, whereas data owners can be maliciously corrupt. Furthermore, the four entities mentioned above are non-colluding.

We now examine the PEFL system in-depth, focusing on the amount of information visible to each entity. More specifically, we are interested in how much information the cloud platform (CP) learns from the protocol execution.

### A. Calculation of Gradients

The protocol begins with each user  $U_x$  training the model locally and obtaining the corresponding gradient vector  $\vec{g}_x = \{g_x^1, \dots, g_x^n\}$ . User  $U_x$  then encrypts and sends the gradient vector to SP using CP’s public key  $pk_c$ . As shown in (1), the gradient vectors corresponding to all users can be viewed as a matrix  $\vec{G}_{m \times n}$ .

$$\vec{G}_{m \times n} = \begin{matrix} & c_1 & c_2 & \dots & c_i & \dots & c_{n-1} & c_n \\ \begin{matrix} U_1 \\ U_2 \\ \vdots \\ U_m \end{matrix} & \begin{pmatrix} g_1^1 & g_1^2 & \dots & g_1^i & \dots & g_1^{n-1} & g_1^n \\ g_2^1 & g_2^2 & \dots & g_2^i & \dots & g_2^{n-1} & g_2^n \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots & \vdots \\ g_m^1 & g_m^2 & \dots & g_m^i & \dots & g_m^{n-1} & g_m^n \end{pmatrix} \end{matrix}. \quad (1)$$

### B. Median Computation using SecMed

The SP and CP use the SecMed algorithm (cf. Figure 4 in [6]) to compute the median value for each of the  $n$  coordinates. SP sends  $g_j^i + r_i$  to CP for each user  $U_j$  corresponding to a coordinate  $c_i$ , where  $r_i$  denotes a random pad sampled by

\*Please cite the journal version of this work published at IEEE TIFS’23 [9].

SP for each coordinate  $c_i$  (but the same for all users). The CP decrypts and computes the medians based on these padded values. The CP then encrypts the medians before sending them to the SP. Finally, the SP removes the pad  $r_i$  to achieve the desired results  $\vec{g}_y$  by utilizing the underlying encryption scheme's homomorphic property.

**Leakage:** The view of CP while executing SecMed is consolidated in the matrix  $\vec{V}_{\text{SecMed}}$ :

$$\vec{V}_{\text{SecMed}} = \begin{matrix} & c_1 & \dots & c_i & \dots & c_n \\ \begin{matrix} U_1 \\ \vdots \\ U_m \end{matrix} & \begin{pmatrix} g_1^1 + r_1 & \dots & g_1^i + r_i & \dots & g_1^n + r_n \\ \vdots & & \vdots & & \vdots \\ g_m^1 + r_1 & \dots & g_m^i + r_i & \dots & g_m^n + r_n \end{pmatrix} \end{matrix}. \quad (2)$$

We observe that for each coordinate  $c_i$ , CP learns a ‘‘shifted’’ distribution of gradients in clear across all users. This is clearly a violation of privacy [4], [5] because it leaks a lot more information to CP and thus does not meet the design goal of ‘Privacy’ claimed in [6]. The main source of the leakage is that SP uses *same random pad*  $r_i$  for all users with respect to a coordinate  $c_i$ . While the aforementioned leakage could be prevented by using different random pads across users, we emphasize that the use of the same pad is unavoidable for the SecMed algorithm to remain correct. In detail, the median of  $g_j^i$  values is calculated by first computing the median of  $g_j^i + r_i$ , then removing  $r_i$  from the result. This requires that the same  $r_i$  value be associated with each  $g_j^i$  value, or the computation’s correctness will be violated.

### C. Computing Pearson correlation coefficient using SecPear

Once the coordinate-wise medians are computed, the next step in PEFL is to calculate the Pearson correlation coefficient  $\rho_{x,y}$  between the coordinate-wise medians  $\vec{g}_y$  and the gradient of the user  $U_x$ . This is achieved via the SecPear protocol (cf. Figure 5 in [6]) where SP communicates  $\vec{g}_x \cdot p_x$  and  $\vec{g}_y \cdot p_y$  to CP. The view of CP in SecPear with respect to  $\vec{G}_{m \times n}$  is  $\vec{V}_{\text{SecPear}}$ :

$$\vec{V}_{\text{SecPear}} = \begin{matrix} & c_1 & \dots & c_i & \dots & c_n \\ \begin{matrix} U_1 \\ \vdots \\ U_m \end{matrix} & \begin{pmatrix} g_1^1 \cdot p_1 & \dots & g_1^i \cdot p_1 & \dots & g_1^n \cdot p_1 \\ \vdots & & \vdots & & \vdots \\ g_m^1 \cdot p_m & \dots & g_m^i \cdot p_m & \dots & g_m^n \cdot p_m \end{pmatrix} \end{matrix}. \quad (3)$$

**Leakage:** Similar to the problem with SecMed above, here CP learns the correlation between each coordinate in the gradient vector  $\vec{g}_x$ . SP uses the *same random pad*  $p_x$  for all coordinates, which causes leakage. However, using different pads for the coordinates does not address the issue since the use of the same pad is required for the SecPear algorithm to remain correct (cf. Proposition 1 in [6]). More elaborately, for  $d_x = \vec{g}_x \cdot p_x$  and  $d_y = \vec{g}_y \cdot p_y$ , computation of  $\rho_{x,y}$  involves computing the covariance  $Cov(d_x, d_y)$  and the standard deviations  $\sigma(d_x)$  and  $\sigma(d_y)$ . As shown in Proposition 1 in [6], the correctness of  $\rho_{x,y}$  relies on the following observations:

$$\begin{aligned} Cov(d_x, d_y) &= p_x p_y \cdot Cov(\vec{g}_x, \vec{g}_y), \\ \sigma(d_x) &= p_x \cdot \sigma(\vec{g}_x), \quad \sigma(d_y) = p_y \cdot \sigma(\vec{g}_y). \end{aligned}$$

If different pads are used for the coordinates, the above relations do not hold, and hence  $\rho_{d_x, d_y} = \frac{Cov(d_x, d_y)}{\sigma(d_x)\sigma(d_y)} \neq \rho_{x,y}$ .

### D. Aggregating the gradients using SecAgg

SecAgg, the final stage in PEFL, aggregates the gradients after scaling them with a factor based on the Pearson correlation coefficient calculated in SecPear. SP communicates  $\vec{g}_x + s_x$  to CP for this purpose, as shown in  $\vec{V}_{\text{SecAgg}}$ :

$$\vec{V}_{\text{SecAgg}} = \begin{matrix} & c_1 & \dots & c_i & \dots & c_n \\ \begin{matrix} U_1 \\ \vdots \\ U_m \end{matrix} & \begin{pmatrix} g_1^1 + s_1 & \dots & g_1^i + s_1 & \dots & g_1^n + s_1 \\ \vdots & & \vdots & & \vdots \\ g_m^1 + s_m & \dots & g_m^i + s_m & \dots & g_m^n + s_m \end{pmatrix} \end{matrix}. \quad (4)$$

**Leakage:** Again, similar to SecMed,  $\vec{V}_{\text{SecAgg}}$  reveals a ‘‘shifted’’ distribution of each user’s gradient values across all the coordinates to CP. When combining information from  $\vec{V}_{\text{SecPear}}$  (3) and  $\vec{V}_{\text{SecAgg}}$  (4), a more significant leakage occurs. Consider the gradient at coordinates  $i, j$  for user  $U_x$ . From  $\vec{V}_{\text{SecPear}}$ , CP learns  $a_1 = g_x^i \cdot p_x$  and  $a_2 = g_x^j \cdot p_x = (g_x^i \cdot \delta_x^{ij}) \cdot p_x$  where  $\delta_x^{ij} = g_x^j / g_x^i$ . Similarly, CP learns  $b_1 = g_x^i + s_x$  and  $b_2 = g_x^j + s_x = (g_x^i + \Delta_{ij}^x) + s_x$  from  $\vec{V}_{\text{SecAgg}}$ , where  $\Delta_{ij}^x = g_x^j - g_x^i$ . Given that CP can compute both  $\delta_x^{ij}$  and  $\Delta_{ij}^x$  in clear, CP learns  $g_x^i$  and  $g_x^j$  by solving the equations. For instance,  $a_2 = (g_x^i + \Delta_{ij}^x) \cdot p_x = a_1 + \Delta_{ij}^x \cdot p_x$  reveals  $p_x$ . Using this method, CP learns the entire gradient matrix  $\vec{G}_{m \times n}$ , thereby breaching the PEFL system’s privacy.

### E. Practical and probabilistic attacks

Another practical attack on PEFL would be to allow CP to register as an honest user  $U_{m+1}$  in the PEFL system and submit its gradients. This action does not violate the semi-honest assumption of CP in the PEFL threat model and may represent scenarios in which CP has some side channel information about some user gradients. Knowing  $\vec{g}_{m+1}$ , CP learns  $r_i$  for all  $i \in [n]$  and hence the gradient matrix  $\vec{G}_{m \times n}$  given in (1) in clear from (2) and thereby breaking the privacy of the entire PEFL system.

Considering the matrices  $\vec{V}_{\text{SecMed}}$  and  $\vec{V}_{\text{SecPear}}$  together, we note that it is sufficient for CP to be aware of just one gradient, say  $g_x^i$ , to compromise the system’s privacy. In particular,  $g_x^i$  will allow CP to learn the random pad  $r_i$  corresponding to the  $i$ -th coordinate  $c_i$  in  $\vec{V}_{\text{SecMed}}$ , revealing the gradients of all users at that coordinate. Similarly, knowing  $g_x^i$  allows CP to learn the random pad  $p_x$  corresponding to user  $U_x$  in  $\vec{V}_{\text{SecPear}}$  and reveals the gradient vector  $\vec{g}_x$  to CP in clear. CP will now learn the entire gradient matrix  $\vec{G}_{m \times n}$  by combining the information from  $\vec{V}_{\text{SecMed}}$  and  $\vec{V}_{\text{SecPear}}$ .

Finally, we note that CP can launch probabilistic attacks by looking for similar values in the matrices  $\vec{V}_{\text{SecMed}}$  and  $\vec{V}_{\text{SecPear}}$  and attempting to correlate the random pads. This is possible in PEFL because CP is aware of the correlation between different rows of  $\vec{V}_{\text{SecMed}}$  as well as columns of  $\vec{V}_{\text{SecPear}}$ .

#### ADDITIONAL REMARKS

Although our privacy issues mentioned in Section II have been published in January 2023 [9], several subsequent papers continued to reference [6] as a potential solution for private federated learning.<sup>1</sup> While a few works, e.g. [3], [12], [14],

<sup>1</sup>As of September 25, 2024, [6] has been cited 206 times according to Google Scholar.

have acknowledged the privacy concerns we raised, several of subsequent works either propagate these errors or adopt the constructions from [6], thereby unintentionally inheriting the same privacy vulnerabilities [1], [11], [13], [15]. We believe this oversight is partly due to the limited visibility of our comments paper at TIFS'23 [9]. Consequently, to prevent the continued propagation of the flawed algorithms in [6] into future research, we also put this article to an ePrint.

## ACKNOWLEDGEMENTS

This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No.850990 PSOTI) and was co-funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB1119 CROSSING/236615297 and GRK2050 Privacy & Trust/251805230.

## REFERENCES

- [1] J. Chen, H. Yan, Z. Liu, M. Zhang, H. Xiong, and S. Yu, "When Federated Learning Meets Privacy-Preserving Computation," *ACM Comput. Surv.*, 2024. [Online]. Available: <https://doi.org/10.1145/3679013>
- [2] H. Fereidooni, S. Marchal, M. Miettinen, A. Mirhoseini, H. Möllering, T. D. Nguyen, P. Rieger, A. Sadeghi, T. Schneider, H. Yalame, and S. Zeitouni, "SAFElearn: Secure Aggregation for private FEderated Learning," in *IEEE S&P Workshops (DLSP)*, 2021. [Online]. Available: <https://doi.org/10.1109/SPW53761.2021.00017>
- [3] T. Gehlhar, F. Marx, T. Schneider, A. Suresh, T. Wehrle, and H. Yalame, "SafeFL: MPC-friendly Framework for Private and Robust Federated Learning," in *IEEE S&P Workshops (DLSP)*, 2023. [Online]. Available: <https://doi.org/10.1109/SPW59333.2023.00012>
- [4] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller, "Inverting Gradients - How easy is it to break privacy in federated learning?" in *NeurIPS*, 2020. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/c4ede56bbd98819ae6112b20ac6bf145-Abstract.html>
- [5] B. Hitaj, G. Ateniese, and F. Pérez-Cruz, "Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning," in *ACM CCS*, 2017. [Online]. Available: <https://doi.org/10.1145/3133956.3134012>
- [6] X. Liu, H. Li, G. Xu, Z. Chen, X. Huang, and R. Lu, "Privacy-Enhanced Federated Learning Against Poisoning Adversaries," *IEEE Trans. Inf. Forensics Secur.*, 2021. [Online]. Available: <https://doi.org/10.1109/TIFS.2021.3108434>
- [7] B. McMahan and D. Ramage, "Federated Learning: Collaborative Machine Learning without Centralized Training Data," <https://research.google/blog/federated-learning-collaborative-machine-learning-without-centralized-training-data/>, 2017, accessed: September 24, 2024.
- [8] T. D. Nguyen, P. Rieger, H. Chen, H. Yalame, H. Möllering, H. Fereidooni, S. Marchal, M. Miettinen, A. Mirhoseini, S. Zeitouni, F. Koushanfar, A. Sadeghi, and T. Schneider, "FLAME: Taming Backdoors in Federated Learning," in *USENIX Security*, 2022. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/nguyen>
- [9] T. Schneider, A. Suresh, and H. Yalame, "Comments on "Privacy-Enhanced Federated Learning Against Poisoning Adversaries,"" *IEEE Trans. Inf. Forensics Secur.*, vol. 18, pp. 1407–1409, 2023. [Online]. Available: <https://doi.org/10.1109/TIFS.2023.3238544>
- [10] M. Sheller, A. Reina, B. Edwards, and J. Martin, "Federated Learning for Medical Imaging," in *Intel AI*, 2018.
- [11] Y. Wan, Y. Qu, W. Ni, Y. Xiang, L. Gao, and E. Hossain, "Data and Model Poisoning Backdoor Attacks on Wireless Federated Learning, and the Defense Mechanisms: A Comprehensive Survey," *IEEE Commun. Surv. Tutorials*, 2024. [Online]. Available: <https://doi.org/10.1109/COMST.2024.3361451>
- [12] J. Wu, W. Zhang, and F. Luo, "On the Security of "LSFL: A Lightweight and Secure Federated Learning Scheme for Edge Computing,"" *IEEE Trans. Inf. Forensics Secur.*, 2024. [Online]. Available: <https://doi.org/10.1109/TIFS.2023.3331274>
- [13] G. Xia, J. Chen, C. Yu, and J. Ma, "Poisoning Attacks in Federated Learning: A Survey," *IEEE Access*, 2023. [Online]. Available: <https://doi.org/10.1109/ACCESS.2023.3238823>
- [14] M. Xu and X. Li, "FedG2L: a privacy-preserving federated learning scheme base on "G2L" against poisoning attack," *Connect. Sci.*, 2023. [Online]. Available: <https://doi.org/10.1080/09540091.2023.2197173>
- [15] J. Zhao, H. Zhu, F. Wang, Y. Zheng, R. Lu, and H. Li, "Efficient and Privacy-Preserving Federated Learning against Poisoning Adversaries," *IEEE Trans. Serv. Comput.*, 2024. [Online]. Available: <https://doi.org/10.1109/TSC.2024.3377931>