



HAL
open science

Enhancing Privacy in Federated Learning: Secure Aggregation for Real-World Healthcare Applications

Riccardo Taiello, Cansiz Sergen, Vesin Marc, Cremonesi Francesco, Innocenti Lucia, Önen Melek, Lorenzi Marco

► **To cite this version:**

Riccardo Taiello, Cansiz Sergen, Vesin Marc, Cremonesi Francesco, Innocenti Lucia, et al.. Enhancing Privacy in Federated Learning: Secure Aggregation for Real-World Healthcare Applications. 5-th MICCAI Workshop on Distributed, Collaborative and Federated Learning in Conjunction with MICCAI 2024, Oct 2024, Marrachech, Morocco. hal-04855481

HAL Id: hal-04855481

<https://inria.hal.science/hal-04855481v1>

Submitted on 25 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Enhancing Privacy in Federated Learning: Secure Aggregation for Real-World Healthcare Applications

Riccardo Taiello^{1,2,3}[0000-0002-9890-9639], Sergen Cansiz¹, Marc Vesin¹,
Francesco Cremonesi¹, Lucia Innocenti¹, Melek Önen²[0000-0003-0269-9495], and
Marco Lorenzi^{1,3}[0000-0003-0521-2881]

¹ Epione Research Project, Inria, Sophia Antipolis, France

² EURECOM, Sophia Antipolis, France

³ Université Côte d’Azur, Nice, France

{riccardo.taiello,marco.lorenzi}@inria.fr
melek.onen@eurecom.fr

Abstract. Deploying federated learning (FL) in real-world scenarios, particularly in healthcare, poses challenges in communication and security. In particular, with respect to the federated aggregation procedure, researchers have been focusing on the study of secure aggregation (SA) schemes to provide privacy guarantees over the model’s parameters transmitted by the clients. Nevertheless, the practical availability of SA in currently available FL frameworks is currently limited, due to computational and communication bottlenecks. To fill this gap, this study explores the implementation of SA within the open-source Fed-BioMed framework. We implement and compare two SA protocols, Joye-Libert (JL) and Low Overhead Masking (LOM), by providing extensive benchmarks in a panel of healthcare data analysis problems. Our theoretical and experimental evaluations on four datasets demonstrate that SA protocols effectively protect privacy while maintaining task accuracy. Computational overhead during training is less than 1% on a CPU and less than 50% on a GPU for large models, with protection phases taking less than 10 seconds. Incorporating SA into Fed-BioMed impacts task accuracy by no more than 2% compared to non-SA scenarios. Overall this study demonstrates the feasibility of SA in real-world healthcare applications and contributes in reducing the gap towards the adoption of privacy-preserving technologies in sensitive applications.

Keywords: Federated Learning · Secure Aggregation · Healthcare Applications.

1 Introduction

Federated Learning (FL) is a distributed machine learning paradigm that enables multiple clients to collaboratively train a global model without sharing their local datasets. While researchers have largely focused in developing FL theories and

methods in a variety of applications, the deployment of FL in real-world scenarios is still challenging, particularly in terms of communication protocols, security, and customization bottlenecks.

A critical requirement for real-world applications of FL concerns the protection of the model’s parameters shared by the clients during model aggregation. To this end, privacy-preserving methodologies such as Secure Aggregation (SA) [17] are currently under study, to guarantee that aggregated data shared among participants do not reveal individual contributions. Contrarily to other privacy-enhancing technologies like Differential Privacy (DP) [11], the privacy guarantees of SA rely on the security proofs of established cryptographic primitives [13,15].

On the practical side, while DP requires only minor adjustments to the federated aggregation process through the injection of noise to the model’s parameters, implementing SA in production is more complex as it requires changes to the standard operational flow of the FL framework by incorporating new communication phases. As a result, the adoption of SA in currently available FL software frameworks is lagging behind. Existing SA solutions primarily target settings with a large number of clients, where hardware limitations can lead to protocol execution failures. Some preliminary solutions have been proposed in the framework FLOWER [4], which however introduce a non-negligible overhead. The approach provided by NVFLARE is simpler but suffers from a weak security model [21]. Finally SA in OPENFL [20] requires dedicated hardware solutions. Overall, the applications of these SA protocols in the cross-silo healthcare setting is suboptimal, due to the limited number of clients, and their general availability as compared to the cross-device setting.

To address these limitations, in this work we explore the implementation of SA schemes optimally customized for cross-silo healthcare applications. In particular, we study the two suitable categories of SA based on masking and additively homomorphic encryption [17]. We identify respectively LOW OVERHEAD MASKING [15] and JOYE-LIBERT [13] as the most relevant solutions for our application. These protocols are designed to protect individual updates from being exposed during the aggregation process.

This work is based on theoretical and experimental evaluation of these SA protocols within the Fed-BioMed framework [9]. In particular, we conducted a comprehensive comparison on four distinct medical datasets including medical images and tabular data: Fed-IXI [24], Fed-Heart [24], REPLACE-BG [1], and FedProstate [12]. We measured the computational resources required for training, encryption, and overall execution time. When training was performed on a CPU, we achieved a total computation overhead of less than 1%, while on a GPU, for larger machine learning models ($> 5M$ parameters), the overhead was less than 50%, with a protection phase that took less than 10 seconds. Furthermore, we analyzed the impact of SA on task accuracy, demonstrating that incorporating SA into Fed-BioMed affects accuracy by no more than 2% compared to non-SA scenarios. Overall this study demonstrates the feasibility of SA in real-world healthcare applications and contributes in reducing the gap towards the adoption of privacy-preserving technologies in sensitive applications.

2 Background

Federated Learning. As introduced by McMahan et al. [18], FL consists of a distributed machine learning paradigm where a group of clients, denoted as \mathcal{U} , collaboratively trains a global model with parameters $\theta \in \mathbb{R}^d$, under the guidance of a FL server. One of the first and popular methods used to train a FL model is the FedAvg scheme [18]. With FedAvg, at each FL round denoted by τ , each client $u \in \mathcal{U}$ trains the model $\theta_{u,\tau}$ on the private local data \mathcal{D}_u , for example through Stochastic Gradient Descent (SGD) [6]. Upon completion of the local training, each client forwards its updated model $\theta_{u,\tau}$ to the server and the local dataset size $w_u = |\mathcal{D}_u|$. When the server receives the updated models from all participating clients, it proceeds to the weighted aggregation step:

$$\theta_{\tau+1} \leftarrow \frac{\sum_{\forall u \in \mathcal{U}} w_u \theta_{u,\tau}}{\sum_{\forall u \in \mathcal{U}} w_u}.$$

This iterative process continues until the global model θ reaches some desired level of accuracy. The presence of a large number of FL clients significantly impacts the communication overhead. To mitigate this, instead of involving all clients in the training, at each FL round, the server selects a subset of clients (*client selection* [18]), denoted as $\mathcal{U}^{(\tau)} \subseteq \mathcal{U}$, with $|\mathcal{U}^{(\tau)}| = n$, and collects their parameters only for aggregation.

Secure Aggregation. SA [17] typically involves multiple *users* and a single *aggregator*. Each user possesses a private input, and the role of the aggregator is to calculate the sum of these inputs. A property of SA is that the aggregator learns nothing more than the aggregated sum, thereby preserving the privacy of individual user inputs.

SA has found significant applications in Federated Learning (FL), where it is used to securely aggregate the updated model parameters received from FL clients (aligned with the *user's* concept in SA) during each FL round, by instantiating an FL server (*aggregator* in the context of SA). The adoption of SA is motivated by the potential threats posed by adversaries having access to the client's updated model $\theta_{u,\tau}$ which may infer information about its private dataset \mathcal{D}_u [19,23]. Hence, the local models should remain confidential even against the FL server. SA in FL was first developed by Bonawitz et al. [5]. The protocol considered in that study faced two different challenges:

- *Threat models* defining the potential risks and behaviors that the security protocol is designed to protect against. The primary threat scenarios in SA include the honest-but-curious model where parties (server and clients) follow the protocol without tampering with the data but may attempt to infer additional information.
- *Client dropouts*, caused by factors such as connectivity issues or voluntary withdrawal, are common in real-world federated learning environments. Dropouts can significantly impact the computation and number of communication rounds of the SA protocol, as they often require the participation of all

selected clients within a training round. With communication rounds, we refer to the number of interactions required between the clients and the server to complete a particular phase of the protocol.

3 Related Works

In real-world deployments, only a few FL frameworks implement some form of SA: OPENFL[20], NVFLARE[21], and FLOWER [4].

FLOWER implements SECAGG+ [2], a masking-based protocol that ensures security in the honest-but-curious model. This protocol requires four communication rounds and uses Shamir’s Secret Sharing to recover missing masks in case of client dropout, ensuring the server can complete the aggregation. Compared to the SA schemes here introduced in Fed-BioMed, Flower’s approach is more costly in terms of communications, albeit accommodating for client dropout.

NVFLARE introduces an SA method that leverages the CKKS asymmetric homomorphic encryption scheme [8]. This threat model is considered weaker than typical state-of-the-art protocols because it requires clients to share a common secret key and assumes clients are honest. Clients protect their inputs using a public key, while the server, operating under the honest-but-curious model, aggregates these inputs and returns the aggregate to each client for decryption using the same secret key. This approach requires one communication round and allows client dropout.

OPENFL’s use of Trusted Execution Environments (TEEs) represents a further step in sandboxing and securing local computations, but requires specific hardware which may not be available in typical FL studies involving hospitals.

4 Methods

In this section, we detail the implementation in Fed-BioMed of the two SA protocols, JOYE-LIBERT (JL) and LOW OVERHEAD MASKING (LOM). From this point on, we adopt the terminology of Fed-BioMed, where a client is referred to as a node.

A general overview of SA is depicted in Figure 1, and a more detailed scheme is provided in Supplementary Figure 4. An SA protocol comprises two phases: **setup** and **online**. The setup phase, illustrated in Figure 1.1, is executed among all participating nodes in \mathcal{U} before the FL training. This step ensures that all parties have the appropriate cryptographic material necessary to run the specific SA protocol.

The online phase, Figure 1.2 is repeated during each FL round τ and consists of two steps: (i) *protect* and (ii) *aggregate*. In the *protection* step, each node protects its private local model using specific SA primitives and then sends the protected model to the server. In the *aggregation* step, the server receives the protected local models, computes the aggregate, and then decrypts it. To ensure the correct functioning of the cryptographic primitives, the locally-trained model vector of each node must be quantized beforehand.

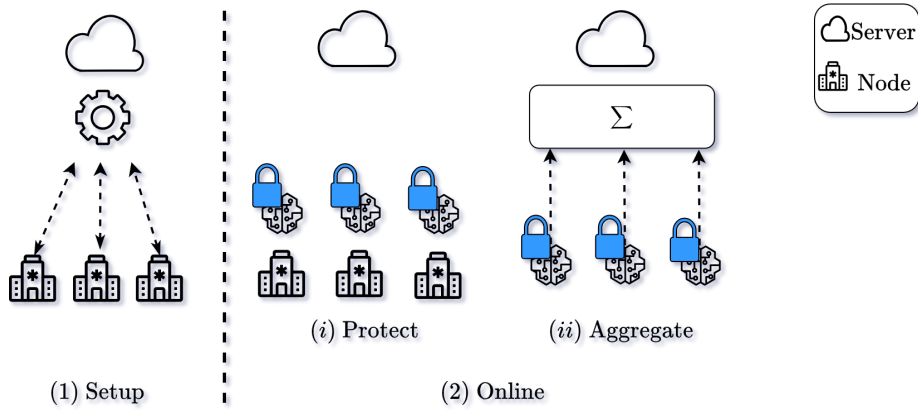


Fig. 1: Overview of Secure Aggregation phases.

Prerequisites. To perform FedAvg with SA, we first convert the node’s local parameters $\theta_{u,\tau} \in \mathbb{R}^d$ into integers $\mathbb{Z}_{2^L}^d$, where L represents the maximum number of bits of the plaintext. This conversion is achieved by applying uniform quantization, defined as: $Q(\theta_{u,\tau}) = \left\lfloor \frac{2^L \cdot (\theta_{u,\tau} - \theta_{\min})}{(\theta_{\max} - \theta_{\min})} \right\rfloor$. Here, $\lfloor \cdot \rfloor$ denotes the standard rounding function. To ensure that real values are within a desired range, we apply a clipping function, $\text{clip}(x, \theta_{\min}, \theta_{\max}) = \min(\max(x, \theta_{\min}), \theta_{\max})$, where θ_{\min} and θ_{\max} are the lower and upper bounds, respectively.

To apply weighted averaging over the integers, we assume that $w_u \in \mathbb{Z}_{2^{W_u}}$, where W_u is the number of bits to represent the node’s dataset size, and we define $W = \max(\{W_u\}_{\forall u \in \mathcal{U}})$.

The weighted local model is computed as $\mathbf{x}_{u,\tau} = Q(\theta_{u,\tau}) \cdot w_u$, resulting in $\mathbf{x}_{u,\tau} \in \mathbb{Z}_{2^{L+W}}^d$. To avoid overflow, we define $M = L+W+\log_2(n)$ as the maximum number of bits for sum computation. The aggregate $\mathbf{x}_{\tau+1} = (\sum_{u \in \mathcal{U}(\tau)} \mathbf{x}_{u,\tau}) \in \mathbb{Z}_{2^M}^d$ is then divided by $s = \sum_{u \in \mathcal{U}(\tau)} w_u$ and dequantized using the following formula: $\theta_{\tau+1} = Q^{-1}(\mathbf{x}_{\tau+1}) = \mathbf{x}_{\tau+1} \cdot \frac{(\theta_{\max} - \theta_{\min})}{2^L} + \theta_{\min}$.

In this context, we assume that quantization has been performed and omit the details of the dequantization process in the protocol explanation.

Joye-Libert

In Supplementary Figure 4a, we illustrate the Joye-Libert (JL) implementation in Fed-BioMed. During the **setup** phase, the participating nodes \mathcal{U} generate their private keys sk_u , and the server creates its server key sk_0 — which is the sum of the node keys — using Shamir Secret Sharing (SS) [22].

During the **online** phase, the protection and aggregation are applied as described in JL (Section 4 [13]). In the protection step, each node uses a private secret key sk_u at FL round τ with a one-time mask derived from sk_u and τ ,

to obtain a protected local model through modular exponentiation over a large modulus N . Using the server key sk_0 , the server can recover the aggregate of the nodes’ private local models in clear.

Our JL implementation works with vectors; the protection and aggregation algorithms are applied element-wise. We use the element’s index i to generate a unique FL round (need to guarantee a one-time mask) for each element in the vector. For instance, to protect $\mathbf{x}_{u,\tau}$, we execute protect and aggregate over the FL round $\tau||i$ and input $\mathbf{x}_{u,\tau}[i]$, where $\mathbf{x}_{u,\tau}[i]$ represents the i -th element of the vector $\mathbf{x}_{u,\tau}$.

The computation and communication of the protected local model is optimized by using vector encoding [16].

Software details. SS is integrated into Fed-BioMed using MP-SPDZ library [14]. The modulus N is provided by the server, and the modular operations are performed using the GMPY2⁴ Python library.

Low Overhead Masking

The second implementation, Low Overhead Masking (LOM) [15], which supports client selection, is depicted in Supplementary Figure 4b. During the **setup** phase, all participating nodes \mathcal{U} establish a pairwise secret $s_{u,v}$, such that $s_{u,v} = s_{v,u}$, with all nodes through the Diffie-Hellman Key Agreement (KA) [10], which will be used in the protect step.

In the **online** phase, during the *protection* step, a selected node $u \in \mathcal{U}^{(\tau)}$ runs the protect algorithm (Section 3.4 [15]). This algorithm protects the local model with a one-time mask derived through a Pseudo-Random Function (PRF) which uses the pairwise secret with the selected nodes $\mathcal{U}^{(\tau)}$ and the current FL round τ , and sends the protected local model to the server. The server then sums the protected local models and collects the final aggregate $\mathbf{x}_{\tau+1}$.

Software details. Diffie-Hellman KA and PRF are implemented in the CRYPTOGRAPHY⁵ Python library, with ECDH and the ChaCha20[3] stream cipher, respectively. Distribution of the DH public key is assumed outside of Fed-BioMed, offline, or through a Public Key Infrastructure.

5 Evaluation

In this Section we provide our theoretical and experimental evaluation of the two implemented SA protocols.

Complexity Analysis: JL’s node computation is $O(d)$, independent of the number of selected nodes, but requires modular exponentiation, and node communication for vector encoding is $O(d \cdot 2 \cdot M)$ [16]. The server’s computation is $O(n + d)$, involving n multiplications and d exponentiation [13].

⁴ GMPY2: <https://gmpy2.readthedocs.io/>

⁵ CRYPTOGRAPHY: <https://github.com/pyca/cryptography>

LOM’s node computation is $O(nd)$, dependent on the number of selected nodes, using faster modular addition and PRF evaluation. The server’s computation involves nd modular additions, and node communication is $O(d \cdot M)$ [15].

Experimental evaluation: The experimental evaluation consists of tracking the computation time between JL and the LOM. We carried out the experiments by considering varying FL hyper-parameters represented by the number of total nodes n_{tot} , the number of selected nodes n , the number of FL rounds T , the number of local SGD steps e , the batch size b and the learning rate η . For SA, the hyper-parameters we explored were the number of bits input L , the number of bits weight W . Moreover, we fixed the aggregation number of bits $M = 32$ and the clipping range min and max. Finally, we report the hardware used to train ML model. We report all this information for each experiments in Table 1.

We use four medical datasets to evaluate the task accuracy of our SA implementations over the aggregated global model at each FL round, using a dedicated tasks-specific test set, and tracking the required computational resources for the nodes.

SA	Time Train (s)	Time Enc. (s)	Time Tot. (s)
FedIXI ($d = 246K; n_{tot} = n = 3$)			
JL	68.10 ± 2.17	52.21 ± 0.85	121.48 ± 2.50
LOM	46.51 ± 1.39	0.62 ± 0.14	48.22 ± 1.03
FedHeart ($d = 258; n_{tot} = n = 4$)			
JL	0.24 ± 0.08	0.08 ± 0.01	0.68 ± 0.09
LOM	0.20 ± 0.09	> 0.01	0.59 ± 0.08
REPLACE-BG ($d = 256K; n_{tot} = 180; n = 18$)			
JL	N/A	N/A	N/A
LOM	53.72 ± 8.61	0.39 ± 0.06	57.42 ± 6.95
FedProstate ($d = 7.4M; n_{tot} = n = 4$)			
JL	N/A	> 300	> 300
LOM	7.65 ± 1.6	9.22 ± 0.38	23.86 ± 2.1

Table 1




Fig. 2

Table 1 Comparison of node average computation time across different SA protocols using four medical datasets. Each dataset is characterized by the total number of nodes (n_{tot}), the number of selected nodes (n), and the size of the local model (d). Fig. 2 Compare the task accuracy of the global model at each FedAvg aggregation with and without applying SA for FedIXI, FedProstate, FedHeart and REPLACE-BG. The SA is characterized by L bits for representing the local model, W bits for representing the maximum dataset size, and the specified maximum and minimum clipping range.

The four datasets are:

- Fed-Heart [24], providing patients’ demography and clinical history from four hospitals. The task is to predict the clinical status of a patient (binary

classification from tabular data). For FL training and testing we follow [24], and the target evaluation metric is the balanced accuracy.

- Fed-IXI [24], is composed by T1 and T2 brain magnetic resonance images (MRIs) from three hospitals. The task is supervised brain segmentation, and ground truth segmentations are provided. For FL training and testing we follow [24], and the target evaluation metric is the dice score.

- REPLACE-BG dataset [1] was obtained from a cohort of 202 participants. The task is prediction of blood glucose levels for the subsequent hour based on data from the last three hours, including glucose levels, insulin boluses, and CHO content.

- FedProstate dataset [12] provides T2 MRIs of the whole prostate from three publicly available datasets, and the task is supervised prostate segmentation. We defined the splitting criteria into different clients, the pre-processing methods, and the FL training and testing parameters coherently with [12].

Supplementary Table 2 reports the dataset details, the FL and SA hyper-parameters, and the hardware specific for model training across experiments. The code is publicly available⁶.

In Table 1, we report the required node’s computational resources comparing the two SA solutions. We present the average training time, encryption time, and total time. LOM consistently outperforms JL due to its faster underlying primitive (modular addition vs. modular exponentiation). Specifically, in all experiments where training runs on a CPU, LOM accounts for less than 1% of the total time. When a GPU (e.g., in FedProstate) is available, the overall encryption time is around 40% of the total time, considering a large input parameter dimension of $d = 7.4M$.

In Figure 2, we display the task accuracy comparison with and without SA for FedIXI, FedProstate, FedHeart and REPLACE-BG. This figure demonstrates that incorporating SA in Fed-BioMed affects the accuracy by no more than 2% compared to the case without SA.

6 Conclusion and Future Works

We have demonstrated that SA can be effectively implemented within the Fed-BioMed framework to enhance privacy in federated learning. Our evaluations using four medical datasets show that both Joye-Libert and Low Overhead Masking protocols protect privacy while maintaining task accuracy. The computational overhead is minimal, making SA a viable option for real-world deployments. As part of future work, we plan to replace MP-SPDZ with a direct implementation of additive secret sharing within Fed-BioMed. We also aim to replace JL with a quantum-resistant SA [7] using the SHELL C++ library ⁷.

⁶ [GitHub code](#)

⁷ [SHELL library: https://github.com/google/shell-encryption/](https://github.com/google/shell-encryption/)

Acknowledgements This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA0002, by the TRAIN project ANR-22-FAI1-0003-02, and by the ANR JCJC project Fed-BioMed 19-CE45-0006-01.

References

1. Aleppo, G., Ruedy, K.J., Riddlesworth, T.D., Kruger, D.F., Peters, A.L., Hirsch, I., Bergenstal, R.M., Toschi, E., Ahmann, A.J., Shah, V.N., et al.: Replace-bg: a randomized trial comparing continuous glucose monitoring with and without routine blood glucose monitoring in adults with well-controlled type 1 diabetes. *Diabetes care* **40**(4), 538–545 (2017)
2. Bell, J.H., Bonawitz, K.A., Gascón, A., Lepoint, T., Raykova, M.: Secure single-server aggregation with (poly) logarithmic overhead. In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. pp. 1253–1269 (2020)
3. Bernstein, D.J., et al.: Chacha, a variant of salsa20. In: *Workshop record of SASC*. vol. 8, pp. 3–5. Citeseer (2008)
4. Beutel, D.J., Topal, T., Mathur, A., Qiu, X., Fernandez-Marques, J., Gao, Y., Sani, L., Li, K.H., Parcollet, T., de Gusmão, P.P.B., et al.: Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390* (2020)
5. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K.: Practical secure aggregation for privacy-preserving machine learning. p. 1175–1191. *CCS ’17, Association for Computing Machinery, New York, NY, USA* (2017)
6. Bottou, L.: Stochastic learning. In: Bousquet, O., von Luxburg, U. (eds.) *Advanced Lectures on Machine Learning*, pp. 146–168. *Lecture Notes in Artificial Intelligence, LNAI 3176, Springer Verlag, Berlin* (2004), <http://leon.bottou.org/papers/bottou-mlss-2004>
7. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-lwe and security for key dependent messages. In: *Annual cryptology conference*. pp. 505–524. Springer (2011)
8. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: *Advances in Cryptology—ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I 23*. pp. 409–437. Springer (2017)
9. Cremonesi, F., Vesin, M., Cansiz, S., Bouillard, Y., Balelli, I., Innocenti, L., Silva, S., Ayed, S.S., Taiello, R., Kameni, L., et al.: Fed-biomed: open, transparent and trusted federated learning for real-world healthcare applications. *arXiv preprint arXiv:2304.12012* (2023)
10. Diffie, W., Hellman, M.E.: New directions in cryptography. In: *Democratizing Cryptography: The Work of Whitfield Diffie and Martin Hellman*, pp. 365–390 (2022)
11. Dwork, C.: Differential privacy. In: *International colloquium on automata, languages, and programming*. pp. 1–12. Springer (2006)
12. Innocenti, L., Antonelli, M., Cremonesi, F., Sarhan, K., Granados, A., Goh, V., Ourselin, S., Lorenzi, M.: Benchmarking collaborative learning methods cost-effectiveness for prostate segmentation. *arXiv preprint arXiv:2309.17097* (2023)

13. Joye, M., Libert, B.: A scalable scheme for privacy-preserving aggregation of time-series data. In: Sadeghi, A.R. (ed.) *Financial Cryptography and Data Security*. Springer Berlin Heidelberg (2013)
14. Keller, M.: Mp-spdz: A versatile framework for multi-party computation. In: *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*. pp. 1575–1590 (2020)
15. Kursawe, K., Danezis, G., Kohlweiss, M.: Privacy-friendly aggregation for the smart-grid. In: *Privacy Enhancing Technologies: 11th International Symposium, PETS 2011, Waterloo, ON, Canada, July 27-29, 2011. Proceedings 11*. pp. 175–191. Springer (2011)
16. Mansouri, M., Önen, M., Ben Jaballah, W.: Learning from failures: Secure and fault-tolerant aggregation for federated learning. In: *Proceedings of the 38th Annual Computer Security Applications Conference*. pp. 146–158 (2022)
17. Mansouri, M., Onen, M., Jaballah, W.B., Conti, M.: Sok: Secure aggregation based on cryptographic schemes for federated learning. *Proc. Priv. Enhancing Technol* (1), 140–157 (2023)
18. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-Efficient Learning of Deep Networks from Decentralized Data. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 54. PMLR (2017)
19. Nasr, M., Shokri, R., Houmansadr, A.: Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In: *2019 IEEE Symposium on Security and Privacy (SP)* (2019)
20. Reina, G.A., Gruzdev, A., Foley, P., Perepelkina, O., Sharma, M., Davidyuk, I., Trushkin, I., Radionov, M., Mokrov, A., Agapov, D., et al.: Openfl: An open-source framework for federated learning. *arXiv preprint arXiv:2105.06413* (2021)
21. Roth, H.R., Cheng, Y., Wen, Y., Yang, I., Xu, Z., Hsieh, Y.T., Kersten, K., Harouni, A., Zhao, C., Lu, K., et al.: Nvidia flare: Federated learning from simulation to real-world. *arXiv preprint arXiv:2210.13291* (2022)
22. Shamir, A.: How to share a secret. *Commun. ACM* (1979)
23. Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership inference attacks against machine learning models. In: *2017 IEEE Symposium on Security and Privacy (SP)* (2017)
24. Ogier du Terrail, J., Ayed, S.S., Cyffers, E., Grimberg, F., He, C., Loeb, R., Mangold, P., Marchand, T., Marfoq, O., Mushtaq, E., et al.: Flamby: Datasets and benchmarks for cross-silo federated learning in realistic healthcare settings. *Advances in Neural Information Processing Systems* **35**, 5315–5334 (2022)

A Appendix

Prerequisites Security paramter λ .

Parties: Server, nodes \mathcal{U} and selected nodes $\mathcal{U}^{(\tau)}$, s.t $|\mathcal{U}| = n_{tot}$ and $|\mathcal{U}^{(\tau)}| = n$.

<p>Public Parameters:</p> <ul style="list-style-type: none"> - $(\perp, pp^{JL}) \leftarrow \mathbf{JL.Setup}(\lambda)$ <p>Setup - Key Setup:</p> <p>Node u:</p> <ol style="list-style-type: none"> 1. $sk_u \xleftarrow{R} \mathbb{Z}_{N^2}$. 2. $\{(v, [sk_u]_v)\}_{v \in \mathcal{U}} \leftarrow \mathbf{SS.Share}(sk_u, t, \mathcal{U})$ 3. Send $\forall v \in \mathcal{U} \setminus \{u\}, [sk_u]_v$ 4. Receive $\{[sk_v]_u\}_{v \in \mathcal{U} \setminus \{u\}}$ 5. $[sk_0]_u \leftarrow \sum_{v \in \mathcal{U}} [sk_v]_u$ 6. Send $[sk_0]_u$ to Server <p>Server:</p> <ol style="list-style-type: none"> 1. Collect $\{[sk_0]_u\}_{v \in \mathcal{U}}$. 2. If $\mathcal{U} < t$, abort; otherwise, proceed. 3. $sk_0 \leftarrow \mathbf{SS.Recon}(\{[sk_0]_v\}_{v \in \mathcal{U}}, t)$ <p>Online - Protection (τ):</p> <p>Node $u \in \mathcal{U}$:</p> <ol style="list-style-type: none"> 1. $\mathbf{y}_{u,\tau} \leftarrow \mathbf{JL.Protect}(pp^{JL}, sk_u, \tau, \mathbf{x}_{u,\tau})$ 2. Send $\mathbf{y}_{u,\tau}$ to Server. <p>Online - Aggregation (τ):</p> <p>Server:</p> <ol style="list-style-type: none"> 1. Collect $\{\mathbf{y}_{u,\tau}\}_{v \in \mathcal{U}}$. 2. $\mathbf{x}_{\tau+1} \leftarrow \mathbf{JL.Agg}(pp, -sk_0, \tau, \{\mathbf{y}_{u,\tau}\}_{v \in \mathcal{U}})$ 	<p>Public Parameters:</p> <ul style="list-style-type: none"> - $(\perp, pp^{LOM}) \leftarrow \mathbf{LOM.Setup}(\lambda)$ - $(\perp, pp^{KA}) \leftarrow \mathbf{KA.Param}(\lambda)$ <p>Setup - Key Setup:</p> <p>Node $u \in \mathcal{U}$:</p> <ol style="list-style-type: none"> 1. $(c_u^{SK}, c_u^{PK}) \leftarrow \mathbf{KA.Gen}(pp^{KA})$ 2. Broadcast c_u^{PK} 3. Receive $\forall v \in \mathcal{U} \setminus \{u\}, c_v^{PK}$ 4. $\forall v \in \mathcal{U} \setminus \{u\},$ $s_{u,v} \leftarrow \mathbf{KA.Agree}(pp^{KA}, c_u^{SK}, c_v^{PK})$ <p>Online - Protection ($\tau, \mathcal{U}^{(\tau)}$):</p> <p>Node $u \in \mathcal{U}^{(\tau)}$:</p> <ol style="list-style-type: none"> 1. $\mathbf{y}_{u,\tau} \leftarrow \mathbf{LOM.Protect}(pp^{LOM}, \{s_{u,v}\}_{v \in \mathcal{U}^{(\tau)} \setminus u}, \tau, \mathbf{x}_{u,\tau})$ 2. Send $\mathbf{y}_{u,\tau}$ to Server. <p>Online - Aggregation (τ):</p> <p>Server:</p> <ol style="list-style-type: none"> 1. Collect $\{\mathbf{y}_{u,\tau}\}_{v \in \mathcal{U}^{(\tau)}}$. 2. $\mathbf{x}_{\tau+1} \leftarrow \mathbf{LOM.Agg}(pp^{LOM}, \{\mathbf{y}_{u,\tau}\}_{v \in \mathcal{U}^{(\tau)}}$ <p style="text-align: center;">(b) LOM</p>
--	---

(a) JL

Fig. 4: SA protocols implemented in Fed-BioMed

Dataset	FL Hyper-params						SA Hyper-params			Hardware spec.
	n_{tot}	n	T	e	b	η	L	W	max/min	
FedHIXI	3	3	75	10	2	1×10^{-3}	22	8	+20/-20	CPU
FedHeart	4	4	75	10	8	5×10^{-4}	15	17	+3/-3	CPU
REPLACE-BG	180	18	400	10	64	1×10^{-3}	13	15	+3/-3	CPU
FedProstate	4	4	75	6	8	1×10^{-3}	22	8	+2/-2	GPU

Table 2: FL hyper-params: number of total nodes n_{tot} , the number of selected nodes n , the number of FL rounds T , the number of local SGD steps e , the batch size b , and the learning rate η . SA hyper-params: number of bits input L , number of bits weight W and clipping range max/min.