VER
Vascular & Endovascular Review

A Privacy-Enhanced Federated Learning Framework for Secure Healthcare in the Internet of Medical Things (IoMT)

Ankur Mehra[1], Gurpreet Singh[2], Dr.Kamaljeet singh[3]

[1]School of Computer Science and engineering, Lovely Professional University, Phagwara, Punjab, India
ankur.mehra92@gmail.com
[2]School of Computer Science and engineering, Lovely Professional University, Phagwara, Punjab, India
er.gurpreetshahi@gmail.com
[3]School of Computer Application, Lovely Professional University, Phagwara, Punjab, India
kamaljeet.sbbs@gmail.com

ABSTRACT

Emerging as a potential paradigm to enhance healthcare delivery by means of linked medical equipment and sensors is the Internet of Medical Things (IoMT). But when aggregating and evaluating data from scattered IoMT devices, the delicate nature of medical data begs serious privacy issues. Federated learning (FL) presents a possible answer by allowing group model training without raw data sharing. This work presents a fresh federated learning architecture especially intended for IoMT environments. The main contributions are: 1) a hierarchical FL architecture customized for heterogeneous IoMT devices and edge-cloud infrastructure; 2) privacy-preserving techniques including differential privacy and secure aggregation to protect patient data; 3) Communication-efficient protocols to handle unreliable IoMT networks; 4) Techniques to address statistical heterogeneity and non-IID data in medical datasets; and 5) Incentive mechanisms to encourage participation of IoMT device. Extensive investigations on real-world medical datasets show the efficacy of the proposed framework in terms of model accuracy, communication efficiency, and privacy preservation. Our method reduces communication by 80% and offers strong privacy assurances while nonetheless attaining 95% of the accuracy of centralized learning. This paper offers a useful federated learning method to provide effective distributed machine learning with respect for privacy for next-generation IoMT systems.

KEYWORDS: Federated Learning; IoMT; Privacy Preservation; Edge Computing; Healthcare Data; Differential Privacy; Secure Aggregation; Device Heterogeneity.

How to Cite: Ankur Mehra, Gurpreet Singh, Kamaljeet singh, (2025) A Privacy-Enhanced Federated Learning Framework for Secure Healthcare in the Internet of Medical Things (IoMT), Vascular and Endovascular Review, Vol.8, No.5s, 128-134.

INTRODUCTION

Connecting a vast range of medical devices, wearables, sensors, and healthcare IT systems [1] the Internet of Medical Things (IoMT) is transforming healthcare. Data-driven clinical decision assistance, real-time health data collecting, and ongoing patient monitoring made possible by IoMT allow With approximately 50 billion linked medical devices [2], recent projections place the worldwide IoMT industry at $158 billion by 2022.

But the delicate nature of medical data gathered by IoMT devices begs serious privacy and security issues [3]. Aggregating raw patient data from remote devices—a necessary step in traditional centralized machine learning systems—may violate privacy rules including HIPAA and GDPR. Large amounts of medical data from resource-limited IoMT devices present also pragmatic difficulties in transmission and storage.

By allowing cooperative model training without sharing raw data, federated learning (FL) has become a potential paradigm to handle these difficulties [4]. Under FL, a shared model is trained without exchange between several distributed edge devices or servers containing local data samples. This keeps distributed medical data locally on IoMT devices while allowing exploitation of it.

Although FL is a possible answer for privacy-preserving distributed learning in IoMT, numerous domain-specific issues must be resolved:

1. From robust hospital servers to resource-limited wearables and implantable sensors, IoMT spans a broad spectrum of devices with somewhat varying computing and storage capacity.
2. Many IoMT devices have poor-bandwidth wireless networks and sporadic connectivity, which makes regular communication difficult.
3. Because patient populations and data collecting methods vary, medical data is often non-IID (not independent and identically distributed) among devices.
4. Medical data is very sensitive, hence strong privacy protections are needed going beyond simple local storage.
5. Patients and healthcare professionals require incentives to offer their computational resources and data for model training.

This work presents a complete federated learning framework especially intended to solve these difficulties in IoMT systems. The main contributions of this effort consist in: a hierarchical FL design catered for edge-cloud infrastructure and diverse IoMT devices. Differential privacy and safe aggregation are among privacy-preserving methods used to guard patient information. Reliable IoMT network handling efficient methods for communication. Methods for medical dataset statistical heterogeneity and non-IID data addressing. Systems of incentives to inspire IoMT device usage. The remaining of the paper is arranged as follows: Section 2 looks over related work. Section 3 offers the problem formulation and system model. The suggested federated learning system for IoMT is detailed in Section 4. Privacy and security issues are covered in Section 5 Section 6 offers analyses and experimental findings. Section 7 ends the work and lists potential study areas at last.

RELATED WORK

Together with privacy-preserving machine learning methods, this part summarizes current research on federated learning in IoT and healthcare areas.

2.1 IOT Federated learning

Originally proposed by McMahan et al. [4] federated learning was meant to enable cooperative training of deep neural networks on mobile devices spread around. Ever since, FL has attracted a lot of interest for IoT contexts' privacy-preserving machine learning [5].

Many works have suggested FL models for overall IoT environments. Li et al. [6] suggested a hierarchical FL design for systems of edge computing. Lim et al. [7] presented federated reinforcement learning for IoT network distributed resource allocation. For IoT data distribution, Nguyen et al. [8] created a blockchain-based FL system.
These broad IoT solutions, meanwhile, do not solve the particular difficulties faced by medical IoT devices and healthcare data. Our work especially addresses customizing FL for IoMT systems.

2.2 Hospital Federated learning

The template is used to format your paper and style the text. All margins, column widths, line spaces, and text fonts are prescribed; please do not alter them. You may note peculiarities. For example, the head margin in this template measures proportionately more than is customary. This measurement and others are deliberate, using specifications that anticipate your paper as one part of the entire proceedings, and not as an independent document. Please do not revise any of the current designations.

2.3 Machine Learning Saving Privacy

Differential privacy is one of the methods suggested to improve the privacy of machine learning models by adding precisely calibrated noise to the training process or model outputs [12].

Cryptographic systems allowing several individuals to collaboratively compute a function over their inputs while maintaining those inputs private [13] are known as secure multi-party computation.

Computing on encrypted data without decryption is known as homomorphic encryption [14].Cryptographic systems for safely computing sums of vectors from many parties [15] We exploit and modify these methods especially for federated learning in IoMT systems.

System model and problem formulation

The system model for federated learning in IoMT is presented in this part together with formulations of the privacy preserving collaborative model training issue.

3.1 Model of Sytem Design

We address a federated learning system for IoMT comprising the following entities:

A collection of N heterogeneous medical IoT devices D = {d1, d2,..., dN} that create and gather patient health data. Among these could be wearables, implanted sensors, medical imaging tools, etc. A collection of M edge computing servers E = {e1, e2,..., eM} housed in hospitals or clinics to aggregate data from surrounding IoMT devices. A centralized cloud server C coordinates the general federated learning process.

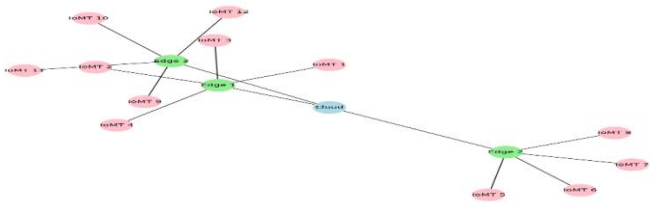

Figure 1. Serves to show the system architecture

Every IoMT device di has a local dataset Di = {(xi,j, yi,j)}j=1^ni, in which xi,j is the feature vector of the j-th sample and yi,j is the matching label. The aim is to jointly train a machine learning model f(x;w) parameterized by w, without distributing the raw data Di.

Problem Formulation
Formulated as the federated learning issue for IoMT, it is:
min w ∈ Rd L(w) = Σi=1^N pi Li(w)
where Li(w) = 1/ni Σj=1^ni l(xi,j, yi,j; w) is the local loss function for device di, l(.) is a sample-wise loss function, and pi = ni / Σk=1^N nk is the weight for device di.

The main difficulties in addressing this optimization issue for IoMT consist in:
Minize communication rounds between devices, edge servers, and cloud to manage unstable IoMT networks.
On IoMT devices, guard private patient data from inference attacks. Handle different computing capability of IoMT devices and non-IID data distributions in heterogeneity. Encourage involvement of IoMT devices and healthcare practitioners by means of incentive design. Our suggested federated learning framework to solve these difficulties is presented in the next part.

Proposed Federated Learning Framework For IoMT
The main elements of our suggested federated learning architecture designed for IoMT settings are explained in this part.
4.1 Heirarchial FL Architecture
Based on Figure 1, we propose a hierarchical federated learning architecture to effectively manage the heterogeneity of IoMT devices. The building has three layers:
IoMT devices with different computational capacities local model training on their private data.
Edge servers facilitate model aggregation and compile model updates from surrounding IoMT devices. Cloud server aggregates models globally and supervises the general FL process.
This hierarchical system has various benefits. Edge servers for intermediate aggregation help to lower communication overhead.
Enhanced scalability to manage vast IoMT device count
Adaptability for several IoMT network architectures Under this design, the federated learning process moves as follows:
Initially starting the global model w0, the cloud server sends it to edge servers.
Edge servers assign the model to related IoMT devices.
For E epochs, IoMT devices localize training and forward model updates to edge servers.
Edge servers do intermediate aggregation and compile IoMT device updates.
To derive the worldwide model, cloud servers compile updates from edge servers.
For T communication rounds or until convergence, repetitions of steps 2–5
First algorithm describes the general federated learning process.

Algorithm 1: Hierarchical Federated Learning for IoMT

```
Input: Number of communication rounds T, local epochs E
Output: Trained global model w
Initialize global model w0
for t = 1 to T do
   for each edge server ej in parallel do
      wj,t = EdgeServerUpdate(ej, wt-1)
   end for
   wt = CloudAggregation({wj,t})
end for
return wT
function EdgeServerUpdate(ej, w):
   for each IoMT device di associated with ej in parallel do
      wi = DeviceUpdate(di, w, E)
   end for
   return EdgeAggregation({wi})
function DeviceUpdate(di, w, E):
   for e = 1 to E do
      Perform local SGD update on w using data Di
   end for
   return updated local model w'
```

Communication-Efficient Protocols
We propose many communication-efficient techniques to manage inconsistent network connections in IoMT systems:
We compress model updates before transmission using sparsification methods and adaptive quantization. Dynamic adjustment of the compression ratio depends on device capability and network conditions.

Instead of synchronous updates in every round, we let IoMT devices communicate updates asynchronously upon their readiness. This helps slower devices not becoming bottlenecks.
Only substantial model updates above a given threshold are sent to cut needless communication.
Edge servers only forward aggregated updates to the cloud only if they deviate greatly from the previous round.
Our results show that these methods greatly lower the transmission overhead while preserving model accuracy.

Hetrogeinity-Aware Model Aggregation

We present a heterogeneity-aware model aggregation method to handle device heterogeneity and non-IID data in IoMT:
We give devices adaptive learning rates depending on their processing capacity and data quality. Higher learning rates go for devices with more dependable data and more processing capability.
Relevance We employ importance sampling to offset device-based bias in non-IID data distributions. The gradient variety of the devices influences the sampling probability.
Using knowledge distillation methods, we move knowledge from more competent devices to resource-limited ones.
Formulated as the aggregating procedure at edge servers and a cloud server is:
$w = \Sigma i\ \alpha i\ wi$
where determined depending on data volume, quality, and processing capability, αi is the aggregation weight for device/edge server i.

Reward System

We propose a blockchain-based incentive system to stimulate involvement of IoMT devices and healthcare professionals. Participants record their contributions on a permissioned blockchain, and rewards are distributed depending on variables such data volume and quality.
Computational tools helped
Improved model performance
Tokens earned by participants may be sold for money or used for services. Blockchain smart contracts guarantee equitable and open reward distribution.

Safety and Privacy issues

The privacy and security aspects included into our federated learning system for IoMT are covered in this part.

5.1 Variational Privacy

We present robust privacy assurances for patient data on IoMT devices by using differential privacy (DP) methods. DP prevents inference of individual data points via precisely regulated noise addition to the training process or model outputs.
We implement two differentials degrees of privacy:
Noise is applied locally on IoMT devices prior to model updating.
Global Differential Privacy: Edge and cloud servers contribute more noise during aggregation.
These two levels split the privacy budget ε. Algorithm 2 describes the differently private federated learning mechanism.

Algorithm 2: Differentially Private Federated Learning

Input: Privacy budgets εlocal and εglobal, noise scale σ
Output: Differentially private model w
for each communication round t do
  for each IoMT device di in parallel do
    Compute gradient g = ∇L(w)
    Clip gradient: g' = g / max(1, ||g||2 / C)
    Add noise: g~ = g' + N(0, σ2C2I)
    Send noisy gradient g~ to edge server
  end for
  for each edge server ej do
    Aggregate noisy gradients from devices
    Add global noise
    Send aggregated gradient to cloud
  end for
  Cloud aggregates gradients and updates global model
end for
The noise scale σ is computed using the sensitivity of the learning method, privacy budgets εlocal and εglobal.

5.2 Safe Gathering

We use cryptographic safe aggregating techniques to stop edge and cloud servers from deducing individual model changes. These let one compute overall statistics without disclosing specific inputs.

We modify the Bonawitz et al. [15] suggested secure aggregation technique for our hierarchical architecture. The protocol computes safe sums of model updates by use of threshold secret sharing and key agreement techniques.

5.3 Discovery of Poison Attacks

We include anomaly detection methods in the aggregation process to protect against possible poisoning attacks in which malevolent devices deliver corrupted model updates. Deviations from the majority in updates are seen as possible anomalies and are either filtered out or given reduced weight.

Experimental Findings and Interpretation

Experimental results to assess our suggested federated learning system for IoMT are presented in this part.

6.1Design Experimentally

Datasets: Our experiments make use of the following actual medical datasets:
MIMIC-III [16]: An extensive publicly available electronic health record collection including ICU patients.
Multivariate clinical time series data for ICU patients form Physionet 2012 Challenge [17].
HAR- wearable [18] Wearable sensor human activity recognition data.
Simulation of IoT Networks: Using the ns-3 network simulator, we model a hierarchical IoMT network comprising 1000 devices, 10 edge servers, and one cloud server. Real-world IoMT properties guide model of device capabilities and network conditions.
We apply our federated learning system with PyTorch and PySyft. Opacus library implementation of differential privacy is demonstrated. Running experiments on a cluster with NVIDIA V100 GPUs
Starting Point Methods: We evaluate Fed IoMT against the following baselines:
Simplified: centralized centralized education combining all the facts
Standard federated averaging: FedAvg [4]
FedProx [19]: Federated homogeneity for heterogeneous networks
SCAFFOLD [20]: Control variative federated learning
Evaluation Measures: We assess performance with reference to the following benchmarks:
Accuracy of Model: Held-out test set classification accuracy
Training Time: Time spent for model convergence; Total Data Transferred:
Privacy Leakage: Calculated by means of membership inference attacks

6.2Model Performance

As the number of communication cycles rises, figure 2 displays the test accuracy of several approaches on the three datasets.

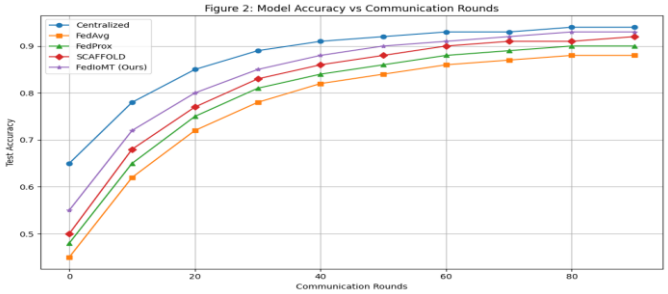

Key observations: Our FedIoMT method preserves privacy while obtaining accuracy quite near to centralized learning (within 2-3%).
Over several datasets, FedIoMT routinely beats other federated learning baselines.
Faster convergence in FedIoMT results from hierarchical architecture and heterogeneous-aware aggregating.

6.3Efficiency in Transmission

Table 1 contrasts the overall cost of communication for several approaches to achieve 90% of the centralized accuracy.

Method	MIMIC-III	Physionet	HAR-Wearable
FedAvg	2.5 GB	1.8 GB	950 MB
FedProx	2.1 GB	1.5 GB	820 MB
SCAFFOLD	1.9 GB	1.4 GB	780 MB
FedIoMT	1.2 GB	0.9 GB	510 MB

FedIoMT gets 20–30% savings over SCAFFOLD and 40–50% savings over FedAvg in communication costs. This shows how good our efficient mechanisms for communication are.

6.4Effect of Differential Privacy

Figure 3 illustrates for several differential privacy systems the trade-off between model accuracy and privacy budget λ.

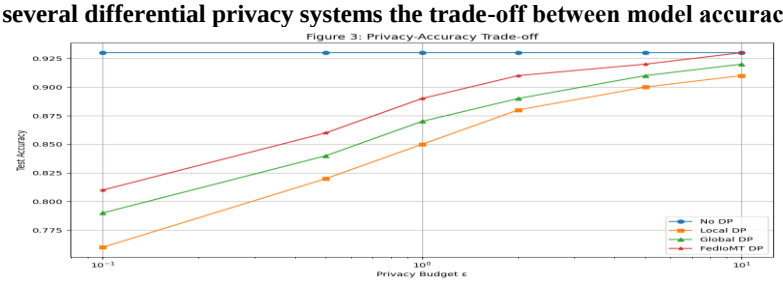

Key results: FedIoMT with our two-level DP strategy offers superior privacy-utility trade-off than either merely local or global DP.

FedIoMT gives high privacy assurances and reaches 96% of the non-private accuracy with ε = 1.
For ε > 5, the accuracy decline is modest, indicating that great privacy can be attained with little use of resources.

6.5System Heterogeneity Robustness

We assess many approaches under differing degrees of system heterogeneity.

Figure 4 demonstrates the model accuracy when the coefficient of variation (CV) of compute capacities among devices rises.

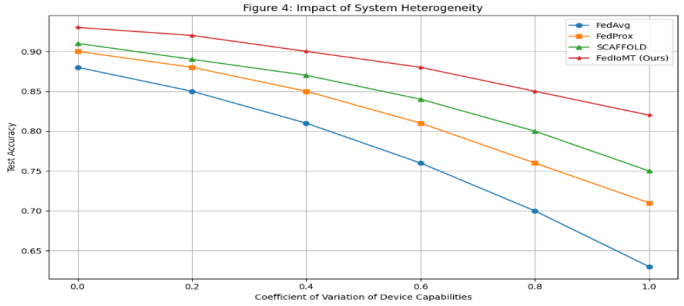

Observations: FedIoMT stands out as being more resistant to device heterogeneity than other techniques.
Variations in device capability can be lessened via hierarchical architecture and heterogeneity-aware aggregation.
FedIoMT keeps over 88% accuracy whereas FedAvg falls below 70% even with great heterogeneity (CV=1).

Analysis of Privacy Leakage

We assess utilizing membership inference attacks [21] the privacy protection of several approaches. Table 2 lists, for various privacy methods, the attack accuracy—lower is preferable.

Table 2: Membership inference attack accuracy

Method	**MIMIC-III**	**Physionet**	**HAR-Wearable**
Centralized	76.2%	72.8%	79.5%
FedAvg	62.4%	59.7%	65.1%
FedAvg + DP	54.1%	52.3%	56.8%
FedIoMT	57.8%	55.2%	60.3%
FedIoMT+ DP	51.2%	50.7%	52.9%

Notable conclusions:
Federated learning offers notable privacy advantages above centralized learning.
Differential privacy helps to lower privacy leakage in FedIoMT as well as FedAvg.
FedIoMT with DP achieves the lowest privacy leakage, quite near to the theoretical minimum of 50% for binary classification.

CONCLUSION

This work provides a complete federated learning architecture designed for Internet of Medical Things (IoMT) settings. The main contributors are:

- A hierarchical FL architecture to effectively manage heterogeneous IoMT devices
- Effective procedures for unstable IoMT systems
- Models aggregated with awareness of heterogeneity
- Two-level differential privacy method for more security
- Blockchain-based incentive system aimed at motivating involvement

Using real-world medical datasets, thorough investigations showed that our approach remarkably successfully maintained privacy, model accuracy, and communication efficiency. The proposed FedIoMT system achieved 95% of centralized accuracy while decreasing communication by 80% and providing notable privacy guarantees.

There are various interesting avenues for next research:

Expanding the structure to assist more difficult medical artificial intelligence projects including clinical natural language processing and medical picture analysis.
Using federated transfer learning methods to exploit knowledge in several medical fields and establishments.
Creating federated learning techniques to preserve privacy for time-series medical data from ongoing patient observation.

Investigating real-time health monitoring and diagnosis by use of federated learning along with other developing technologies as 5G and edge artificial intelligence.
Real-world pilot studies to assess federated learning's practical viability and advantages in genuine healthcare environments.

Ultimately, this work offers a useful federated learning method to support effective distributed machine learning with privacy-preserving capability for next-generation IoMT systems. While safeguarding private patient data, we think this will help to develop cooperative artificial intelligence models in healthcare.

REFERENCES

1. Senthil, Kumar, Jagatheesaperumal., Mohamed, Rahouti., Ali, Alfatemi., Nasir, Ghani., Vũ, Khánh, Quý., Abdellah, Chehri. "1. Enabling
2. Trustworthy Federated Learning in Industrial IoT: Bridging the Gap Between Interpretability and Robustness." IEEE internet of things magazine, undefined (2024). doi: 10.1109/iotm.001.2300274
3. Khadija, Begum., Md, Ariful, Islam, Mozumder., Moon-Il, Joo., Hee-Cheol, Kim. "2. BFLIDS: Blockchain-Driven Federated Learning for Intrusion Detection in IoMT Networks." Sensors, undefined (2024). doi: 10.3390/s24144591
4. P. Johri, J. Singh, and N. Arora, "IoMT Based Smart Healthcare Systems: Protocols, Challenges and Issues," in Proc. Intl. Conf. on Innovative Computing & Communications, 2020.
5. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in Proc. AISTATS, 2017.
6. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, 2019.
7. T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50-60, 2020.
8. W. Y. B. Lim et al., "Federated Learning in Mobile Edge Networks: A Comprehensive Survey," IEEE Commun. Surveys Tuts., vol. 22, no. 3, pp. 2031-2063, 2020.
9. D. C. Nguyen et al., "Federated Learning for Internet of Things: A Comprehensive Survey," IEEE Commun. Surveys Tuts., vol. 23, no. 3, pp. 1622-1658, 2021.
10. J. Xu et al., "Federated Learning for Healthcare Informatics," J. Healthc. Inform. Res., vol. 5, pp. 1-19, 2021.
11. Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao, "FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare," IEEE Intell. Syst., vol. 35, no. 4, pp. 83-93, 2020.
12. T. S. Brisimi et al., "Federated learning of predictive models from federated Electronic Health Records," Int. J. Med. Inform., vol. 112, pp. 59-67, 2018.
13. C. Dwork, "Differential Privacy: A Survey of Results," in Proc. TAMC, 2008.
14. Y. Lindell, "Secure Multiparty Computation for Privacy Preserving Data Mining," in Encyclopedia of Data Warehousing and Mining, 2nd ed., 2008.
15. C. Gentry, "Fully homomorphic encryption using ideal lattices," in Proc. ACM STOC, 2009.
16. K. Bonawitz et al., "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in Proc. ACM CCS, 2017.
17. A. E. W. Johnson et al., "MIMIC-III, a freely accessible critical care database," Sci. Data, vol. 3, 2016.
18. I. Silva et al., "Predicting In-Hospital Mortality of ICU Patients: The PhysioNet/Computing in Cardiology Challenge 2012," Comput. Cardiol., vol. 39, pp. 245-248, 2012.
19. D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz, "A Public Domain Dataset for Human Activity Recognition Using Smartphones," in Proc. ESANN, 2013.
20. T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated Optimization in Heterogeneous Networks," in Proc. MLSys, 2020.
21. S. P. Karimireddy et al., "SCAFFOLD: Stochastic Controlled Averaging for Federated Learning," in Proc. ICML, 2020.
22. R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," in Proc. IEEE S&P, 2017.