Ticao Zhang and Shiwen Mao *Dept. Electrical & Computer Engineering, Auburn University, Auburn, AL, USA*

Editors: Shiwen Mao and Michelle Gong

AN INTRODUCTION TO THE FEDERATED LEARNING STANDARD

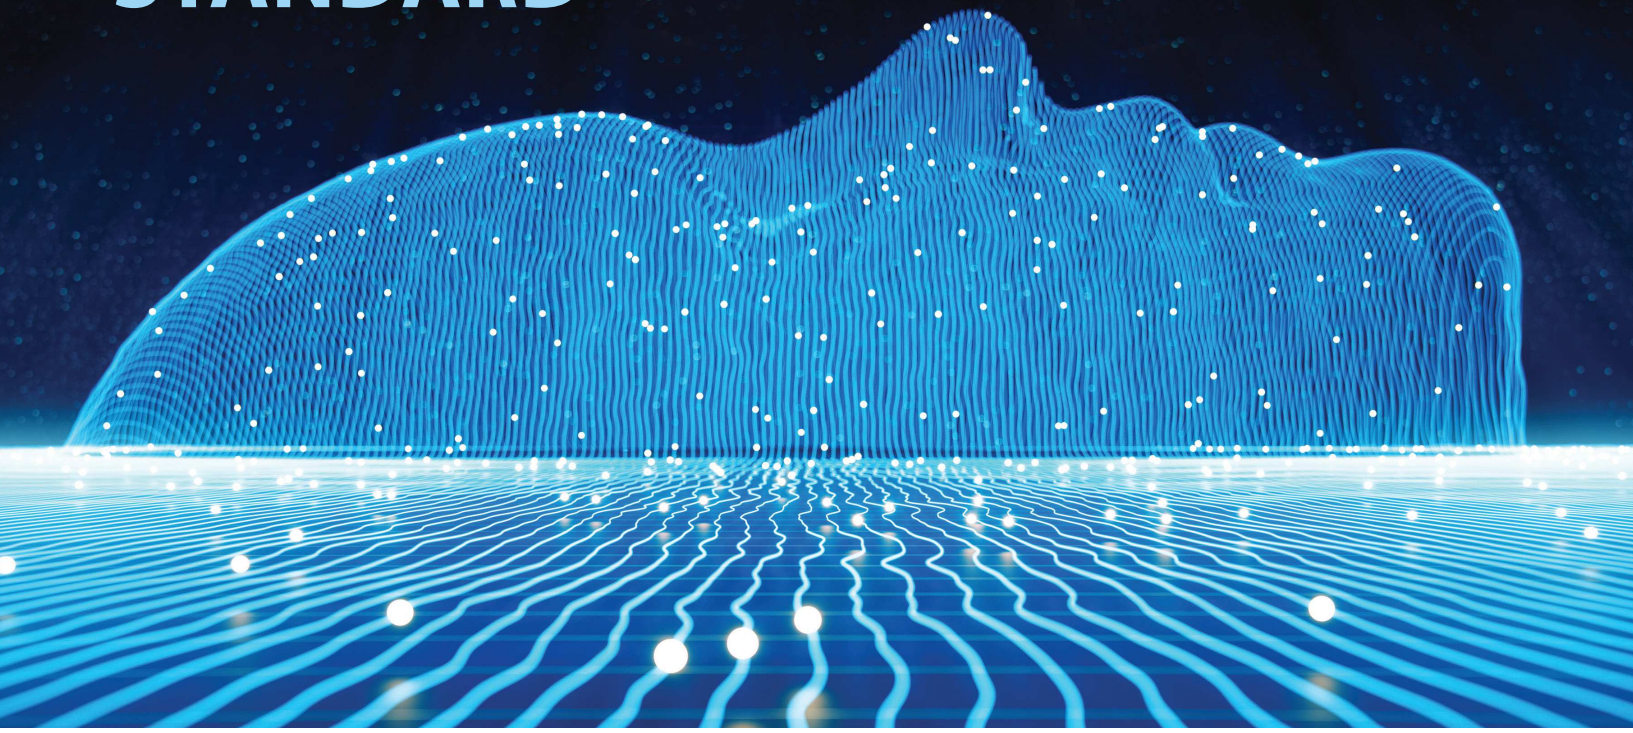

With the growing concern on data privacy and security, it is undesirable to collect data from all users to perform machine learning tasks. Federated learning, a decentralized learning framework, was proposed to construct a shared prediction model while keeping owners' data on their own devices. This paper presents an introduction to the emerging federated learning standard and discusses its various aspects, including i) an overview of federated learning, ii) types of federated learning, iii) major concerns and the performance evaluation criteria of federated learning, and iv) associated regulatory requirements. The purpose of this paper is to provide an understanding of the standard and facilitate its usage in model building across organizations while meeting privacy and security concerns.

Artificial intelligence (AI), driven by big data, has already been applied to various aspects of our daily life, such as transportation, agriculture, insurance, healthcare, and others. Companies and organizations are collecting increasingly more detailed data about their users. For example, some technology companies develop health apps by analyzing the data generated at users' wearable watches; banks evaluate customers' financial risks by analyzing their credit card usage and loan history; and retail companies deploy automatic recommendation systems based on customers' shopping data. The conventional AI approach requires an integration of data from multiple sources to build the AI model [1]. However, collecting such data may be costly and time consuming. Meanwhile, with the increasing concern on data privacy and security, some regulations forbid data sharing among different organizations. For example,

Illustration, istockphoto.com

the European Union passed the General Data Protection Regulation (GDPR) in May 2018 [2], which states that any institutions or organizations do not have authority to use the users' private data without an agreement. The establishment of the regulation helps protect data leakage and promote security. Similarly, the General Security Law of China also states that network operators should not destroy or disclose the personal information they collect [3]. These regulations help protect users' information and prevent its leakage, but it also brings challenges to AI model training.

To address this challenge, the concept of a decentralized learning framework, termed federated learning, was proposed by Google [4,5]. In this framework, training is performed over a set of federation of distributed learners with data stored and used in model training at each individual learner locally. Each learner can improve his local model without explicitly accessing other learners' private data. The term federated learning was initially introduced for mobile and edge computing applications and was later extended by researchers to cover secure distributed learning across multiple organizations, such as health centers or banks, using their local private data [6,7]. It is believed that federated learning will bring about the opportunity for different data owners to collaborate and share data to build AI models [8].

To reduce the cost and risks of business collaboration on AI projects with data from different sources, the federated machine learning group (C/AISC/FML) starts work on a new IEEE standard (IEEE Std 3652.1-2020) on federated learning [9,10]. The Project Authorization Request (PAR) for this standard started within the IEEE Standards Association (IEEE SA) on Oct. 14, 2018. The projected completion date for submittal to the Review Committee (RevCom) was in Oct. 2021. In this paper, we shall present an introduction to this ongoing standard. The reminder of this paper is organized as follows. "Federated Learning Overview" provides an overview of federated learning principles and basics. "Federated Learning Standard" introduces the federated learning standard in detail, including the types of federated learning, the major concern and performance evaluation of federated learning, the associated regulatory requirements, and applications. This article ends with "Conclusions."

FEDERATED LEARNING OVERVIEW

Federated machine learning is a distributed machine learning framework where dataset owners collaboratively train a global model for a given task, such as classification, prediction, or regression. In federated learning, the dataset owners collaboratively train a model without exchanging their raw data and no dataset owner can infer the private information of other dataset owners. The main goal of federated learning is to make sure the performance of the federated learning model is close to that of the desired model trained with a centralized approach while also preserving the privacy of each data owner.

To illustrate how federated learning works, consider the practical example in [11]. As shown in Figure 1, the user devices communicate with the cloud server periodically to train an emoji predictor in a distributed fashion. In each communication round, a subset of mobile users is selected to perform local training using their own data. Instead of sending the raw data to the cloud server, the data owner devices upload their trained model parameters to the cloud. After aggregating the local updates, the central server distributes the updated global model to another subset of model users. This process continues in an iterative manner until a stopping criterion is met. The experiment in [11] shows that this federated learning approach can train production-quality models for emoji prediction while keeping users' data locally.

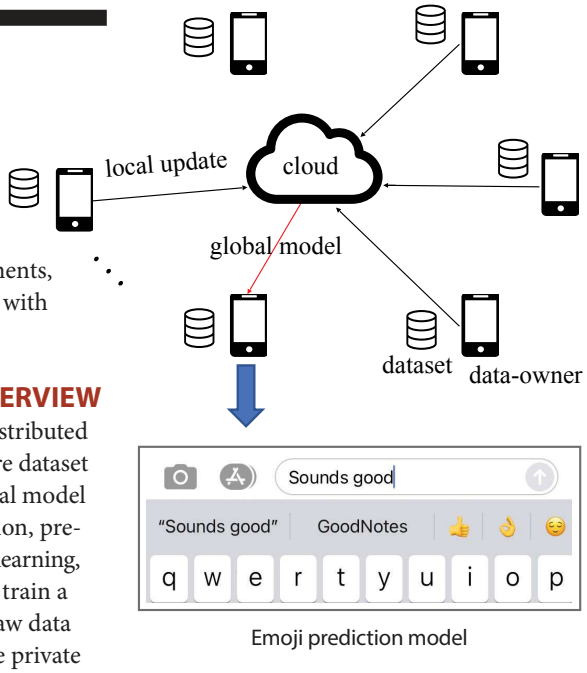

FIGURE 1. An example application of federated learning for emoji prediction [11].

Design Challenges

Distributed learning is a combination of distributed computing and machine learning. The main goal is to perform a global estimation by collecting and aggregating results from distributed computing units. Federated learning adopts the distributed approach. However, it is different from the conventional distributed learning because it requires the global model feedback and is more concerned with the security and privacy issues. These differences bring new challenges to federated learning.

1) Privacy Concern

In federated learning, the information might be traced back to the source via the uploaded, locally trained model, which could still raise a privacy concern. For example, the gradients can be used to infer the data used for computing the gradients [12,13]. Moreover, malicious users can upload unreliable or tempered model parameters to the central server to attack the training process [14].

2) Heterogeneity in Datasets

Distributed learning and parallel computing assume that the data are independent and identically distributed (i.i.d.) among different devices. In federated learning, the data at different devices could be unbalanced and non-i.i.d., which poses great statistical challenges to build a unified global model [15]. In addition to the heterogeneity in data distribution, the computation, storage, and network connectivity capability of the participating

devices may also differ considerably. The device hardware is heterogeneous, and some devices can be unreliable. These pose great challenges of fault tolerance [16].

3) Communication Cost

The repeated exchange of updated training models between the device and the central server incurs a massive data transmission cost. The large communication loads across devices over wireless communication links can limit the scalability of federated learning [17]. Indeed, a federated learning system can potentially comprise a massive number of devices, e.g., thousands of mobile phones. The limited communication resources at the devices, such as power, bandwidth, and energy, could be a bottleneck of the entire system.

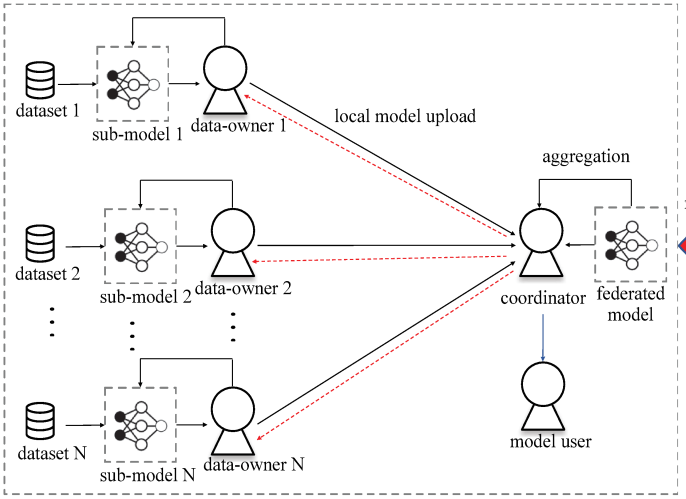

FIGURE 2. A schematic diagram of the federated learning framework.

FEDERATED LEARNING STANDARD

In this section, we introduce the federated learning standard (IEEE Std 3652.1-2020) [9,10], which specifies the following aspects of federated learning: the architecture of federated learning, the dataset and user role description, the application scenarios to which each category applies, a set of evaluation criterion, and the associated regulatory requirements.

Architecture Description

A federated learning architecture is presented in Figure 2. It consists of data, user, and functional modules. In this architecture, the data is distributed across different data owners to collaboratively train a federated learning model with secure and privacy-preserving techniques. The users are of different types to play different roles. Finally, multiple functional modules are provided to jointly support the federated learning services management.

Dataset Description

Federated machine learning data is generally stored in a standard database format, e.g., table. Assume a data owner has a dataset, which consists of several data samples. Each data sample consists of both features and labels. Moreover, a unique sample ID is associated with each data sample. In federated learning, data from different dataset owners may overlap in sample IDs and/or feature attributes. As shown in Figure 3, depending on the extent of overlap, federated learning can be categorized as horizontal federated learning, vertical federated learning, and federated transfer learning.

1) Horizontal Federated Learning

Horizontal federated learning splits the datasets horizontally. The user features of the two datasets overlap considerably while the user IDs overlap a bit. Intuitively, different data owners own different samples that share similar features. As a result, horizontal federated learning can increase the user sample size. For example, two banks from two different regions want to collaboratively train a model for their business. The user groups (i.e., the sample space) for the two banks are quite different. However, the business types (i.e., the feature space) are quite similar. The example we provide in Figure 1 belongs to horizontal federated learning, where each user device performs local training and uploads their parameters to the cloud server.

2) Vertical Federated Learning

Vertical federated learning splits the datasets vertically. The user features of two datasets overlap to certain extent, while the user IDs overlap considerably. Intuitively, different data owners own similar samples with different features. Thus, vertical federated learning can increase the feature space. For example, there is a bank and an E-commerce company in the same city. Their user groups are basically the same, which include the residents of the city. However, their businesses are quite different. The bank records users' credit and salary

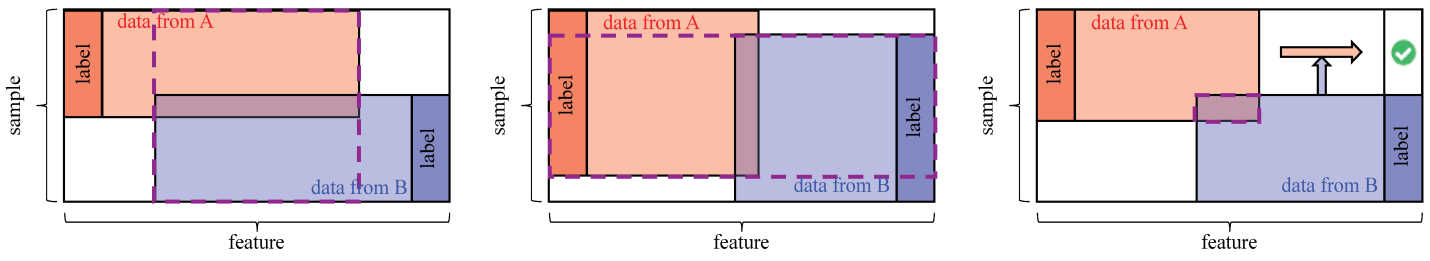

FIGURE 3. Categories of federated learning.

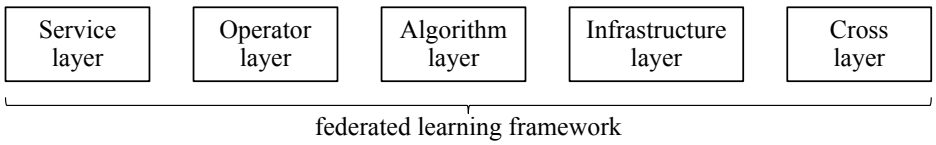

FIGURE 4. Federated learning framework.

while the E-commerce company tracks users' shopping history. Vertical federated learning can help aggregate the features to build a comprehensive model. A novel privacy-preserving tree-boosting system known as *SecureBoost* was proposed in [18], which allows multiple parties with common user groups but different feature sets to perform the training process together. This framework achieves the same level of accuracy as the centralized training approach while also protecting data privacy.

3) Federated Transfer Learning

If the overlaps of both the sample space and feature space are small, it will be hard to carry out effective federated learning. Federated transfer learning was proposed in [12] to provide a common representation based on limited common sample sets and common features. In this framework, federated learning is used to build a machine learning model while transfer learning takes advantage of the reusable knowledge across different domains and overcomes the limitations of conventional machine learning techniques. For example, there is a bank in city A and an e-commerce company in city B. Their user groups and business features overlap very little. Federated transfer learning can help to solve the small sample size problem to achieve an enhanced performance.

User Description

According to the federated learning standard, the architecture in Figure 2 consists of four types of users, i.e., data owners, data coordinator, model users, and auditors. The main role of the data owner is to collect and maintain its private data locally, and send their local model parameters to the coordinator. Meanwhile, they can ask for payments and receive rewards for participating in the training process of federated learning. The main role of the coordinator user is to develop algorithms and services for all the participating users, aggregate the local training results of all dataset owners, and feed back the global model parameter to each data owner. In addition to training and testing the model, the coordinator also designs appropriate incentive schemes consisting of calculating the payments to data owners and model users. The main role of the model user is to put requests to the coordinator and pay for the use of the federated learning model. Finally, the auditors are responsible for checking the correctness of the federated learning process and ensuring that the entire process complies with regulations and performance requirements.

Functional Module Description

A federated learning framework is shown in Figure 4, which comprises five layers, including the service layer, operator layer, algorithm layer, infrastructure layer, and cross layer. Each layer consists of different functional modules, such as data service module and user service module. These modules implement different elementary activities and may be included or omitted from a specific federated learning system.

The service layer provides services to help the model users access the federated learning model (i.e., the user service module), supports the management of local data repository (data service module), supports the management of federated learning tasks (task management service module), and so on. The operator layer aggregates the sub-models of the respective data owners (i.e., the aggregator module), provides activation operation support (i.e., the activation module), optimization implementation support (i.e., the operation module), and so on. The algorithm layer implements the federated machine learning algorithms (i.e., the algorithm module), evaluates the performance of federated learning based on various criterions (i.e., algorithm evaluation module), calculates the payments to participants (i.e., the economic incentive calculation module), and so on. The infrastructure layer supports all the functions and interfaces needed by traditional machine learning by providing computation, storage, and communication components supports. The cross-layer interacts with the other layers to jointly support service management (i.e., the operating functional module), ensure security (i.e., the system security functional module), and regulate the entire process (i.e., the regulation and audit module).

Evaluation Criterion

There are a series of metrics to confirm and evaluate the performance of federated learning. These performance metrics are widely adopted by both academic and industrial researchers to validate the performance of federated learning systems. The federated learning standard defines the following four evaluation measures.

1) Privacy and Security

For the sake of privacy-preserving, the federated learning system requires that both the extent of leakage and the probability of privacy disclosing be kept at a low level. The severity of leakage attack is characterized by the amount of data being disclosed (i.e., the extent of leakage) and the probability of private information being inferred from the disclosed data (i.e., the influence of leakage).

For the sake of security, the federated learning system is considered attack-proof if it can effectively defend against attacks that aim to tamper with the federated learning system. The extent of attacks is evaluated by the amount of data being altered (the extent of alternation) and the degradation of model performance caused by the altered data (the influence of alternation).

2) Model Performance

The federated learning model should achieve a performance that is equivalent or more competitive to that of the centralized training model. The metric to evaluate the differences between the two models is called model performance discrepancy. It can vary in different applications in terms of accuracy, prediction, image quality, and other measures.

3) Computation Efficiency

The federated learning training process is considered efficient if it takes a reasonable amount of time in training and testing or

consumes a reasonable amount of memory. The training (testing) time is measured by the ratio between the training (testing) time and the number of training (testing) samples. The memory usage is measured by the amount of memory needed by the data and code.

4) Economic Viability

In order to maintain the economic sustainability, the coordinator should design proper economic mechanism for all users. The economic viability is evaluated based on the individual rationality index (IRI), the budget surplus margin (BSM), the efficiency index (EI), the data offering rate (DOR), and the fairness index (FI).

IRI is a number between 0 and 1, which shows the weighted percentage of the rational users that are willing to stay in the federation. BSM, varying from $-\infty$ to 1, denotes the revenue difference between the model users' payment and revenue received from using the federated learning model. EI, varying from $-\infty$ to 10, indicates whether the Pareto efficiency is achieved. Pareto efficiency is achieved when no further changes can be made to make any data owners better off without making at least one data owner worse off. DOR is a number between 0 and 1, which measures the willingness of data owners to offer their data for training. FI is a number between 0 and 1, which measures the variance of the payment of a unit of effective data across all data owners.

Federated Learning Standard Applications

The application area of federated learning is divided into three categories: business-to-consumer (B2C), business-to-business (B2B), and business-to-government (B2G). B2C applications include telecommunications, education, internet-of-things (IoT), and others. B2B applications include finance, health, and marketing. B2G applications include urban computing and government services. For each type of application, the federated learning standard IEEE Std 3652.1-2020 specifies the role design and the main activities of each role. Moreover, the federated learning standard specifies the requirements for each application type. For example, the standard shows that a health application (B2B) should satisfy the level 4 security requirement, which means that the system should successfully defend the model control against three types of attacks defined in the standard. The health application also needs to satisfy the level 2 privacy requirement, which means the system should prevent privacy leakage and data inference from malicious participants.

CONCLUSIONS

With the development of big data and AI, people are more concerned about their data privacy. Federated learning is regarded as an effective solution to allow a distributed learning scheme without explicit data sharing and private data leakage. This paper presents an introduction to the emerging federated learning standard proposed by the federated machine learning working group (C/AISC/FML). We hope this paper will help reveal insights of the framework and application guidelines for federated machine learning. Different organizations in the field of healthcare, education, marketing, telecommunications, etc., can reduce the cost and risks of business collaboration on AI projects with the federated learning approach defined in this standard. ■

Ticao Zhang received his B.E. and M.S. degrees from the School of Electronic Information and Communications, Huazhong University of Science and Technology, Wuhan, China, in 2014 and 2017, respectively. He is currently pursuing a Ph.D. in electrical and computer engineering at Auburn University. His research interests include wireless networks, machine learning, and optimization.

Shiwen Mao received his Ph.D. in electrical and computer engineering from Polytechnic University, Brooklyn, NY in 2004. He a professor and Earle C. Williams Eminent Scholar, and Director of Wireless Engineering Research and Education Center at Auburn University, Auburn, AL. His research interests include wireless networks, multimedia communications, and smart grid. He is an IEEE Fellow.

This work is supported in part by the NSF (CNS-2107190).

REFERENCES

[1] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Pearson, 2002.

[2] J.P. Albrecht. 2016. How the GDPR will change the world, *Eur. Data Prot. L. Rev.*, vol. 2, no. 3, 287—289.

[3] M. Parasol. 2018. The impact of China's 2016 cyber security law on foreign technology firms, and on China's big data and smart city dreams, *Computer Law & Security Review*, vol. 34, no. 1, 67–98, Feb. 2018.

[4] J. Konečný, H.B. McMahan, F.X. Yu, P. Richtarik, A.T. Suresh, and D. Bacon. 2017. Federated learning: Strategies for improving communication efficiency, *arXiv preprint arXiv:1610.05492*, Oct. 2017. https://arxiv.org/abs/1610.05492.

[5] J. Konečný, H.B. McMahan, D. Ramage, and P. Richtárik. 2016. Federated optimization: Distributed machine learning for on-device intelligence, *arXiv preprint arXiv:1610.02527*, Oct. 2016. https://arxiv.org/abs/1610.02527.

[6] L. Li, Y. Fan, M. Tse, and K.-Y. Lin. 2020. A review of applications in federated learning, *Computers & Industrial Engineering*, 106854, Nov. 2020.

[7] P. Kairouz, et al. 2019. Advances and open problems in federated learning, *arXiv preprint arXiv:1912.04977*, Dec. 2019. https://arxiv.org/abs/1912.04977.

[8] T. Li, A.K. Sahu, A. Talwalkar, and V. Smith. 2020. Federated learning: Challenges, methods, and future directions, *IEEE Signal Processing Magazine*, vol. 37, no. 3, 50–60, May 2020.

[9] Q. Yang, L. Fan, R. Tong, and A. Lv. 2021. IEEE Federated Machine Learning, in *IEEE Federated Machine Learning – White Paper*, 1–18, June 2021.

[10] IEEE. IEEE 3652.1-2020 - IEEE guide for architectural framework and application of federated machine learning." https://standards.ieee.org/standard/3652_1-2020.html

[11] A. Hard, et al., Federated learning for mobile keyboard prediction, *arXiv preprint arXiv:1811.03604*, https://arxiv.org/abs/1811.03604.

[12] Q. Yang, Y. Liu, T. Chen, and Y. Tong. 2019. Federated machine learning: Concept and applications, *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, 12, 1-19. Feb. 2019.

[13] Y. Aono, T. Hayashi, L. Wang, and S. Moriai. 2018. Privacy-preserving deep learning via additively homomorphic encryption, *IEEE Transactions on Information Forensics and Security*, vol.13, no. 5, 1333—1345, May 2018.

[14] L. Lyu, H. Yu, and Q. Yang. 2020. Threats to federated learning: A survey, *arXiv preprint arXiv:2003.02133*, Mar. 2020. https://arxiv.org/abs/2003.02133.

[15] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar. 2017. Federated multi-task learning, in *Proc. 31st Conf. Neural Inf. Process. Syst. (NIPS'17)*, Long Beach, CA, Dec. 2017, 4424–4434.

[16] K. Bonawitz, et al. 2019. Towards federated learning at scale: System design, arXiv preprint arXiv:1902.01046v2, Mar. 2019. https://arxiv.org/abs/1902.01046.

[17] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B.A.Y. Arcas. 2017. Communication-efficient learning of deep networks from decentralized data, in *Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS'17)*, Fort Lauderdale, FL, Apr. 2017, 1273–1282.

[18] K. Cheng, T. Fan, Y. Jin, Y. Liu, T. Chen, and Q. Yang. 2019. Secureboost: A lossless federated learning framework, *arXiv preprint arXiv:1901.08755*, Jan. 2019. https://arxiv.org/abs/1901.08755.