

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2023.1120000

A Security-Oriented Overview of Federated Learning Utilizing Layered Reference Model

JIAxing LU, NORIHIRO FUKUMOTO, and Akihiro Nakao, (Member, IEEE)

Room A105, Daiwa Ubiquitous Academic Research Building, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan (e-mail: lujiaxing@g.ecc.u-tokyo.ac.jp, fukumotonorihiro@g.ecc.u-tokyo.ac.jp, nakao@nakao-lab.org)

Corresponding author: Akihiro Nakao (e-mail: nakao@nakao-lab.org).

ABSTRACT With the continuous development of Artificial Intelligence (AI), AI services are becoming increasingly influential in society, affecting both individual lives and enterprise production. However, the field of AI model training grapples with a conflict between the desire to maximize the utilization of private and external data, and the necessity to limit collaborative data sharing for privacy protection. Strict regulations on sensitive data give rise to data silos, impeding the smooth flow of information as well. In response to these challenges, Federated Learning (FL) emerges as a promising solution, enabling collaborative machine learning model training across isolated data silos. Despite its potential, securing FL systems remains a challenge, primarily due to the absence of a canonical reference model that hierarchically summarizes existing works in this field. This lack complicates users' understanding of federated learning in the context of data flow and impedes their ability to pinpoint specific security issues and corresponding solutions when utilizing an FL system. To address this gap, we propose a layered reference model for federated learning through a comprehensive survey. The model encompasses five layers: the data interaction layer, client management layer, local model layer, network transmission layer, and remote management layer. Prioritizing the promotion of system security awareness, we conduct a threat analysis for each layer and explore corresponding defense strategies drawn from existing techniques. As a result, readers can gain insights into the federated learning system from the perspective of data flow, comprehend the security risks their private sensitive data might encounter at each step, and explore relevant solutions to safeguard their information.

INDEX TERMS Federated Learning, Reference Model, Information Security, Privacy Protection

I. INTRODUCTION

The rapid growth of mobile communication technology has resulted in a surge in data generation and traffic, with predictions indicating that data traffic will constitute over 70% of global Internet connections, with 45% attributed to mobile connectivity (3G, 4G, or 5G) by 2023, according to statistics from Cisco [1]. The advent of 5G enables the provision of wired and mobile broadband services anytime and anywhere for users [2]. Besides the improvement of network transmission, end devices like smartphones and Internet of Things (IoT) sensors are continuously improving in performance, allowing for more complex crowdsensing tasks, such as medical applications [3]. However, the substantial amounts of data generated pose challenges for both network transmission and data processing within the system.

In parallel, Artificial Intelligence (AI) has demonstrated significant capabilities in handling large volumes of data, leading to the increasing importance of cloud computing-

based AI in modern society [4]. However, the mandatory requirement to transfer dispersed original data to a central cloud server for model training presents challenges for the further application, encompassing issues like poor private data protection [5] [6], weak data source robustness [7], and high transmission bandwidth requirements [8] [9], etc.

In contrast, Federated Learning (FL), introduced by Google in 2016, addresses these challenges by enabling collaborative AI model training on distributed data without transmitting the data to a central cloud server [10] [11] [12]. In FL, models are trained locally on each end device, and only model update information is transmitted to a central server for aggregation. This process ensures user privacy, reduces the burden of transmitting vast data volumes over public networks. Furthermore, FL could integrate data from different silos, expanding the data sources and enabling the extraction of deeper insights.

While security threats in distributed learning like FL have

been identified by Baruch et al. [13], existing surveys often follow tree-like structures, as illustrated in Figure 1 [14]. It poses challenges for users to understand FL in the context of data flow, making it difficult to narrow down the scope to specific security issues and corresponding solutions when utilizing a federated learning system. In general, current work has the following issues:

- Lacking a step-by-step and clearly defined layered summary based on the data flow throughout the entire federated learning system.
- An explicit system information security threats analysis along the data flow is also missing.
- Simply listing some security methods rather than delving into specific solutions for the various security threats faced by users.

To fill this gap, we propose a 5-layer reference model from a security-oriented perspective based on the data flow, from data generation to model training, in the federated learning system and conduct a security analysis of each layer according to it. We also introduce some defense strategies to secure the system for users. We aspire for this survey to serve as a valuable resource for readers to enhance the information security of their FL system, with the following specific contributions:

- We propose a security-oriented layered reference model for the federated learning system based on the flow of data. This model divides the entire FL system into five layers: the data interaction layer, client management layer, local model layer, network transmission layer, and remote management layer.
- We analyze the security threats that FL encounters at each layer and introduce appropriate defense strategies based on existing techniques.
- Readers can gain insights into the federated learning system from the perspective of data flow. They can understand the types of security risks their private sensitive data might encounter at each step and explore corresponding solutions to safeguard their private information.

The rest of this paper is organized as follows: Section II introduces related work. Section III defines and outlines the constitution of federated learning after showing a definition by us. Our proposed layered reference model is presented in Section IV. In Section V, we discuss common threats, and Section VI analyzes the system security risks layer by layer. Corresponding defense strategies are outlined in Section VII, and challenges and conclusion are discussed in Section VIII.

II. RELATED WORK

Federated Learning is evolving rapidly, emerging as a promising approach for addressing critical data privacy and confidentiality. While there are many existing survey works on FL, and some of them specialize in privacy and security, they merely list common security issues of FL rather than providing a systematic analysis based on the data flow utilizing a

reference model. This approach limits the depth of knowledge conveyed to readers, as a comprehensive exploration of security concerns and vulnerabilities in federated learning necessitates a more nuanced and detailed examination within the data flow architecture. Table 1 presents a comprehensive summary of some recent surveys, offering a comparative analysis of their distinctive characteristics.

Distinguished from other prevailing surveys, our survey delves into the intricate realms of security and privacy in federated learning, extending along the trajectory of data flow. What sets our survey apart is the proposal of a security-oriented layered reference model, which serves as a scaffold for unraveling its underlying security issues.

Furthermore, we introduce corresponding countermeasures aligned with each layer of the reference model. This strategic approach equips readers not only with theoretical knowledge but also with pragmatic solutions. Consequently, when confronted with security issues, readers can judiciously narrow down their focus and select the most pertinent countermeasures.

III. OVERVIEW OF FEDERATED LEARNING

In this section, we delve into a comprehensive overview of FL, offering an exploration that spans its definition, the pivotal components constituting its architecture, and the common federated training process.

A. DEFINITION

Based on the concept put forth by Google [10] [11] [12], and taking into account the opinions of Yang et al. [26] and Kairouz et al. [27], we define federated learning as follows:

Federated learning is a decentralized machine learning setting where original data is distributed across multiple client compute nodes, such as enterprise data centers or individual smart devices [26]. In this collaborative framework, client compute nodes collectively train a machine learning model, either globally shared or personalized [14], under the coordination of an aggregation control server, either synchronously or asynchronously [27] [28].

Compared to traditional centralized machine learning, this approach prioritizes privacy protection by keeping raw data on client compute nodes, minimizing data leakage risks [27] [26]. However, it is crucial to recognize some inherent challenges in practical federated learning models, especially in terms of model evaluation, where accuracy is a notable concern. While federated learning models are anticipated to surpass models trained individually on local devices, they might not reach the same performance level as those trained in a centralized approach [14]. This trade-off is essential to maintain privacy while effectively leveraging distributed data sources.

B. CONSTITUTION

Referring to the study by Li et al. [14], we define three major components in federated learning, containing two hardware

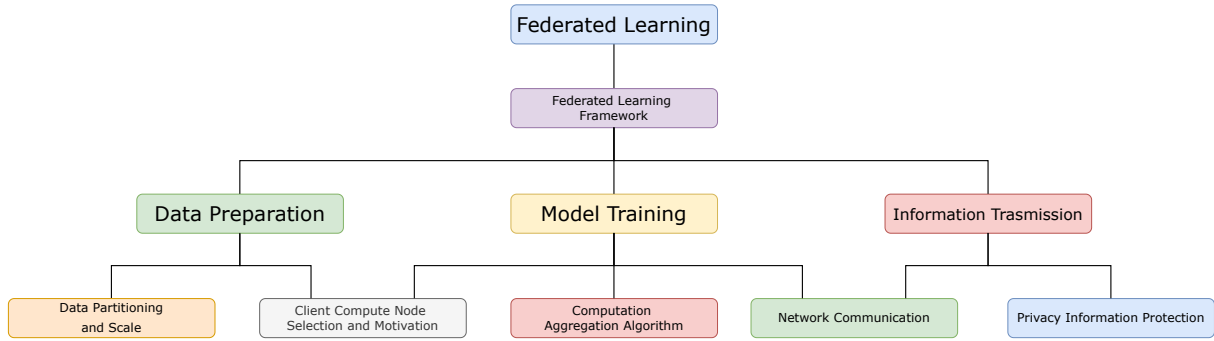


FIGURE 1. Tree-like Structure in Current Surveys of Federated Learning [14]

and one software: the aggregation control server, client compute node, and the computation communication framework that connects the aggregation control server and client compute nodes together, as illustrated in Figure 2.

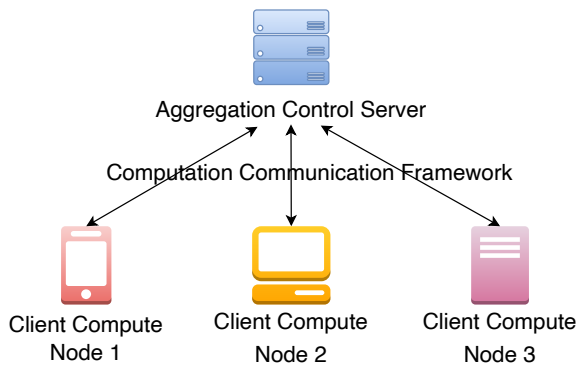


FIGURE 2. Three Major Components in Federated Learning

- **Aggregation Control Server:** The aggregation control server orchestrates the federated learning process, acting as a central controller in a manner similar to centralized schemes. It manages the entire model training process and communications between two client compute nodes or between client compute nodes and itself. Alternatively, in decentralized schemes [29], like SimFL [30], client compute nodes update their model over the model parameter space by aggregating information through communication with their one-hop neighbors, acting in both client compute node and aggregation control server roles at the same time.
- **Client Compute Node:** Client compute nodes, spanning from smartphones to data centers, function as both the owners of original data and users of the trained model. They possess non-Independent Identically Distributed (non-IID) private data generated from daily use, irrespective of the scale of data distribution [14] [27]. This implies that the local private data stored on a specific client compute node cannot represent the entire data distribution [31], posing challenges for analysis due to differing data distributions among nodes.

- **Computation Communication Framework:** The computation communication framework facilitates communication and parameters transmission between client compute nodes and the aggregation control server, supporting local model training on client compute nodes and gradient aggregation on the aggregation control server. It plays a crucial role in the entire model training process, incorporating tasks like client selection mechanism, for example, FedCS [32], incentive mechanism, aggregation mechanism, and transmission mechanism.

C. COMMON FEDERATED TRAINING PROCESS

We summarize the common training processes of federated learning models according to the research by Kairouz et al. [27], encompassing both centralized and decentralized scenarios, as a complement to our proposed layered reference model

- 1) **Nodes Selection:** In centralized federated learning, the server, which serves as the aggregation control server, selects a group of client compute nodes who meet eligibility requirements. While in decentralized federated learning, the aggregation control server can be a normal client compute node that initiates the training process. The selection mechanism of client compute nodes is based on various criteria, such as battery status, computational resources, and other relevant factors.
- 2) **Initial Model Broadcast:** The aggregation control server disseminates the initial global model weights to the selected client compute nodes through computation communication framework.
- 3) **Model Training:** Each selected client compute node independently trains and computes model updates using its sensitive private data.
- 4) **Update Aggregate:** After several local iterations, the aggregation control server collects and aggregates updates from the selected client compute nodes via computation communication framework. This step often incorporates security and efficiency algorithms to ensure data privacy.
- 5) **Model Update:** The aggregation control server updates the global model based on the aggregated update information, under the instruction of aggregation mech-

TABLE 1. Comparisons among Some Recent Surveys on Federated Learning

References	Published Date	Brief Content	Security Related	Providing Model
Lim W. et al. [15]	April 2020	Introduce background and fundamentals of FL, highlight the challenges of implementation and review existing solutions, present the applications, challenges, and future research directions	Yes	No
Li T. et al. [16]	May 2020	Discuss the characteristics and challenges of FL, provide a overview of current approaches, and outline several directions	Briefly introduce three protection mechanisms	No
Aledhari M. et al. [17]	July 2020	Software and hardware platforms, protocols, real-life applications and use-cases of FL	No	No
Mothukuri V. et al. [18]	October 2020	Description of approaches and various implementation styles with an examination of the current challenges in FL	Yes	No
Xu J. et al. [19]	December 2020	Solutions to the statistical challenges, system challenges, and privacy issues in FL, and the implications and potentials in healthcare	Partially related	No
Zhang C. et al. [20]	January 2021	Introduce FL from: data partitioning, privacy mechanism, machine learning model, communication architecture and systems heterogeneity	Only some privacy protection means	No
Li Q. et al. [14]	November 2021	Introduce definition of FL, and provide a categorization to six aspects	Briefly mention privacy mechanisms	No
Gosselin R. et al. [21]	September 2022	Privacy and security risks of FL, state-of-the-art approaches to counteract problems	Yes	No
Coelho. K. K. et al. [22]	May 2023	Data security and privacy applications for FL in Internet of Healthcare Things networks	Yes	No
Hasan J. [23]	July 2023	Comprehensive taxonomy of security and privacy challenges in FL	Yes	No
Gabrielli E. et al [24]	August 2023	Summarize existing decentralized FL approaches, identify emerging challenges and promising research directions	Partially related	No
Kandati D.R. et al. [25]	August 2023	Examine FL's privacy and security concerns, and deal with several issues	Yes	No
Ours		Security-oriented survey of security risks for FL utilizing layered reference model	Yes	Yes

anism offered by computation communication framework.

- 6) **Evaluate and Repeat:** Engineers evaluate the global model's performance and determine whether to continue training or terminate the process and deploy the model. If training continues, the above steps are repeated. Some high-quality client compute nodes would be incentivized if an incentive mechanism is introduced in the computation communication framework.

IV. LAYERED REFERENCE MODEL

A. MOTIVATIONS AND METHODOLOGIES IN THE PROPOSAL

It is well-accepted that federated learning is a secure variant of the distributed system that addresses data silos issue while keeping private data decentralized, involving different kinds of participants with diverse requirements and constraints [33] [27] [34]. However, Lo et al. emphasize the lack of references for an end-to-end federated learning framework, attributing it to the absence of a unified reference model [34]. We agree that such absence impedes the further development of FL and may cause some repetitive work. Motivated by the willingness to fill this gap, we decide to design a general federated learning

system reference model.

Based on the empirical summarization collected through our systematic literature review about federated learning from different fields, especially survey papers mentioned but not limited to those in II, and experiences from conducting FL-related experiments, we arrive at a preliminary requirement, which is to separate the problems as much as possible to reduce the coupling of the system. Second, for the sake of addressing diverse issues and requirements of different participants, standardization work should be carried out. Third, the reference model must be able to provide an easier and clearer view for checking data privacy protection, since privacy protection is the most important principle in FL [21] [35].

Hoping to alleviate the work about deploying FL and inspired by the OSI 7-layer network model [36] that is dominant in the current networking field, we consider a layered model as the optimal solution for us to satisfy these three requirements summarized above. With the assistance of the layered model, engineers could focus on only some separated problems at one time while using standardized interfaces between two adjacent layers. Moreover, scalability and flexibility are improved as different issues are isolated in different layers. Lastly, the data flow becomes clearer under the layered model, which will benefit analyzing security threats.

B. OVERVIEW OF THE MODEL

In response to the gap mentioned in IV-A and aiming to simplify the work while hoping to standardize the development of FL, we propose a security-oriented 5-layer federated learning system reference model, as shown in Figure 3. This model organizes federated learning system functions into distinct layers along the data flow and control flow, enhancing the standardization level, scalability, flexibility and comprehensibility of federated learning system research and implementation.

The proposed reference model comprises the following layers, organized from the orientation of security along the data flow: data interaction layer, client management layer, local model layer, network transmission layer, and remote management layer. Each lower layer provides foundational services and information to the upper layer, while the upper layer offers management and reply functions to the lower layer.

C. DEFINITION OF EACH LAYER

In this subsection, we will define and elucidate the roles of each layer in the layered reference model, aiming to provide readers with a clear understanding.

- **Data Interaction Layer:** This layer functions as the physical device layer, encompassing client compute nodes where users deploy applications to satisfy their business. It serves as the source of original data, presenting model inference results to users, and collecting node status information for the remote management layer to determine training eligibility. Essentially, it acts as the

interactive interface between users/applications and the federated learning system.

- **Client Management Layer:** This layer focuses on data processing, addressing variations in data formats and structures through data normalization and standardization to enhance model accuracy and reliability [37] [38]. Additionally, it executes client selection decisions made by the remote management layer, as it is impractical to include every client compute node that requests to join in the training process. It also handles client incentive decisions and evaluates client applications based on factors such as throughput, computation capability, and wireless resources [32] [39].
- **Local Model Layer:** This core layer is pivotal in the federated learning system. It employs approaches like gradient descent algorithms to train a local model on selected client compute nodes using local processed data. After local model computes, gradient information is forwarded to the network transmission layer. The layer also deploys models issued from the remote management layer and handles model inference tasks.
- **Network Transmission Layer:** The network transmission layer establishes communication links between peers, the client compute node and the aggregation control server in centralized scheme, or client compute node's one-hop neighbor in decentralized scheme. It encapsulates and decapsulates information, transmitting updated model gradient information from selected client compute nodes to the remote aggregation control server for global model updates. The communication link may vary from traditional Ethernet to advanced 5G networks. After global model aggregation, it issues the updated model to selected client compute nodes and transmits control information between the client and the server. However, this transmission step is susceptible to private information leakage, as highlighted by Zhang et al. [40]. As the result, the network transmission layer plays a crucial role in safeguarding transmitted data. It could apply security mechanisms such as encryption, differential privacy [41] [42], and homomorphic encryption [43] [44] to enhance data security before transmission to address this issue. On the receiver side, the network transmission layer would restore the data for utilization. It is important to note that adopting these protection mechanisms may introduce additional complexities in communication and computation [45].
- **Remote Management Layer:** The remote management layer is another crucial functional layer in federated learning. Its main role is to aggregate the updated gradient information from selected client compute nodes, using it as input to update the global model on the aggregation control server. Additionally, this layer addresses fairness concerns arising from the diverse distribution of source data among participating client compute nodes. It acts as the decision-making hub for client selection and incentive allocation, while also plays a key role in

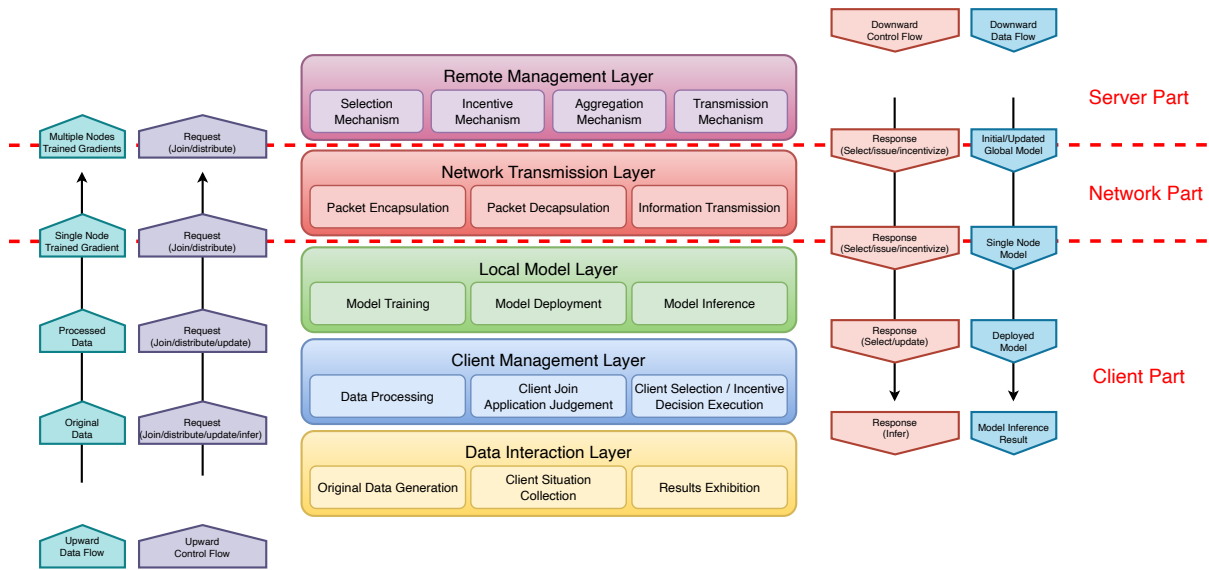


FIGURE 3. Proposal of 5-layer Federated Learning Reference Model

enhancing the security of the transmission phase.

- **Data Flow and Control Flow:** Data flow and control flow are two essential types of information movement within the layered reference model. Upward data flow involves the processing of data, while downward data flow deals with the deployment and utilization of the trained model. Upward control flow focuses on the request flow initiated by end applications, while downward control flow involves responses from the remote management layer. The low-level Data Flow Diagram (DFD) based on our model, from the perspective of DeMarco et al. [46], is depicted in Figure 4.

We believe the 5-layer reference model we propose, although in its early stages, holds the potential to offer a standardized and open reference interface for researchers in the federated learning community. However, we acknowledge the need for refinement and expansion in future endeavors.

D. APPLICATION CASE EXAMPLE

To illustrate the practical relevance and applicability of our proposed security-oriented 5-layered federated learning reference model, in the subsection, we will present an application case to readers as an example and explain it using our model.

Assume the background of the application case is as follows: several financial enterprises in a city hope to establish a cross-enterprise personal credit AI model to better serve the citizens. These enterprises include insurance companies, commercial banks, and government tax departments. While their customers belong to the same group, sharing customer data is impossible due to individual privacy protection requirements. Fortunately, federated learning could assist them in achieving this mission. However, the deployment is complex since there are no unified design standards in place.

This task would become simpler if they adopt our proposed

layered model layer by layer. According to our model, two flow paths must be followed to complete the design of FL model training system: bottom-up data flow and top-down control flow.

- **Bottom-up Data Flow:**

- Data Interaction Layer: These three participants need to separately select the original data for model training, negotiate information collected for client situation monitoring, and determine the desired output.
- Client Management Layer: The only objective is to set the unified output standards for data processing.
- Local Model Layer: Each enterprise needs to deploy some servers for local model training, model deployment and model inference as specified by the local model layer. Until this step, participants still maintain autonomy to make decisions following unified standards and interfaces. This could protect their privacy data to the greatest extent, and comply with their respective internal rules. No matter rules and regulations in the insurance companies, commercial banks, or government tax departments.
- Network Transmission Layer: Participants upload the model update information following the packet encapsulation algorithm specified by the network transmission layer.
- Remote Management Layer: The final step is to aggregate this information to update the global model on a shared server that is trusted by all parties.

- **Top-down Control Flow:**

- Remote Management Layer: Participants have to negotiate client compute node selection mechanisms, incentive mechanisms, global model aggregation mechanisms, and data transmission mecha-

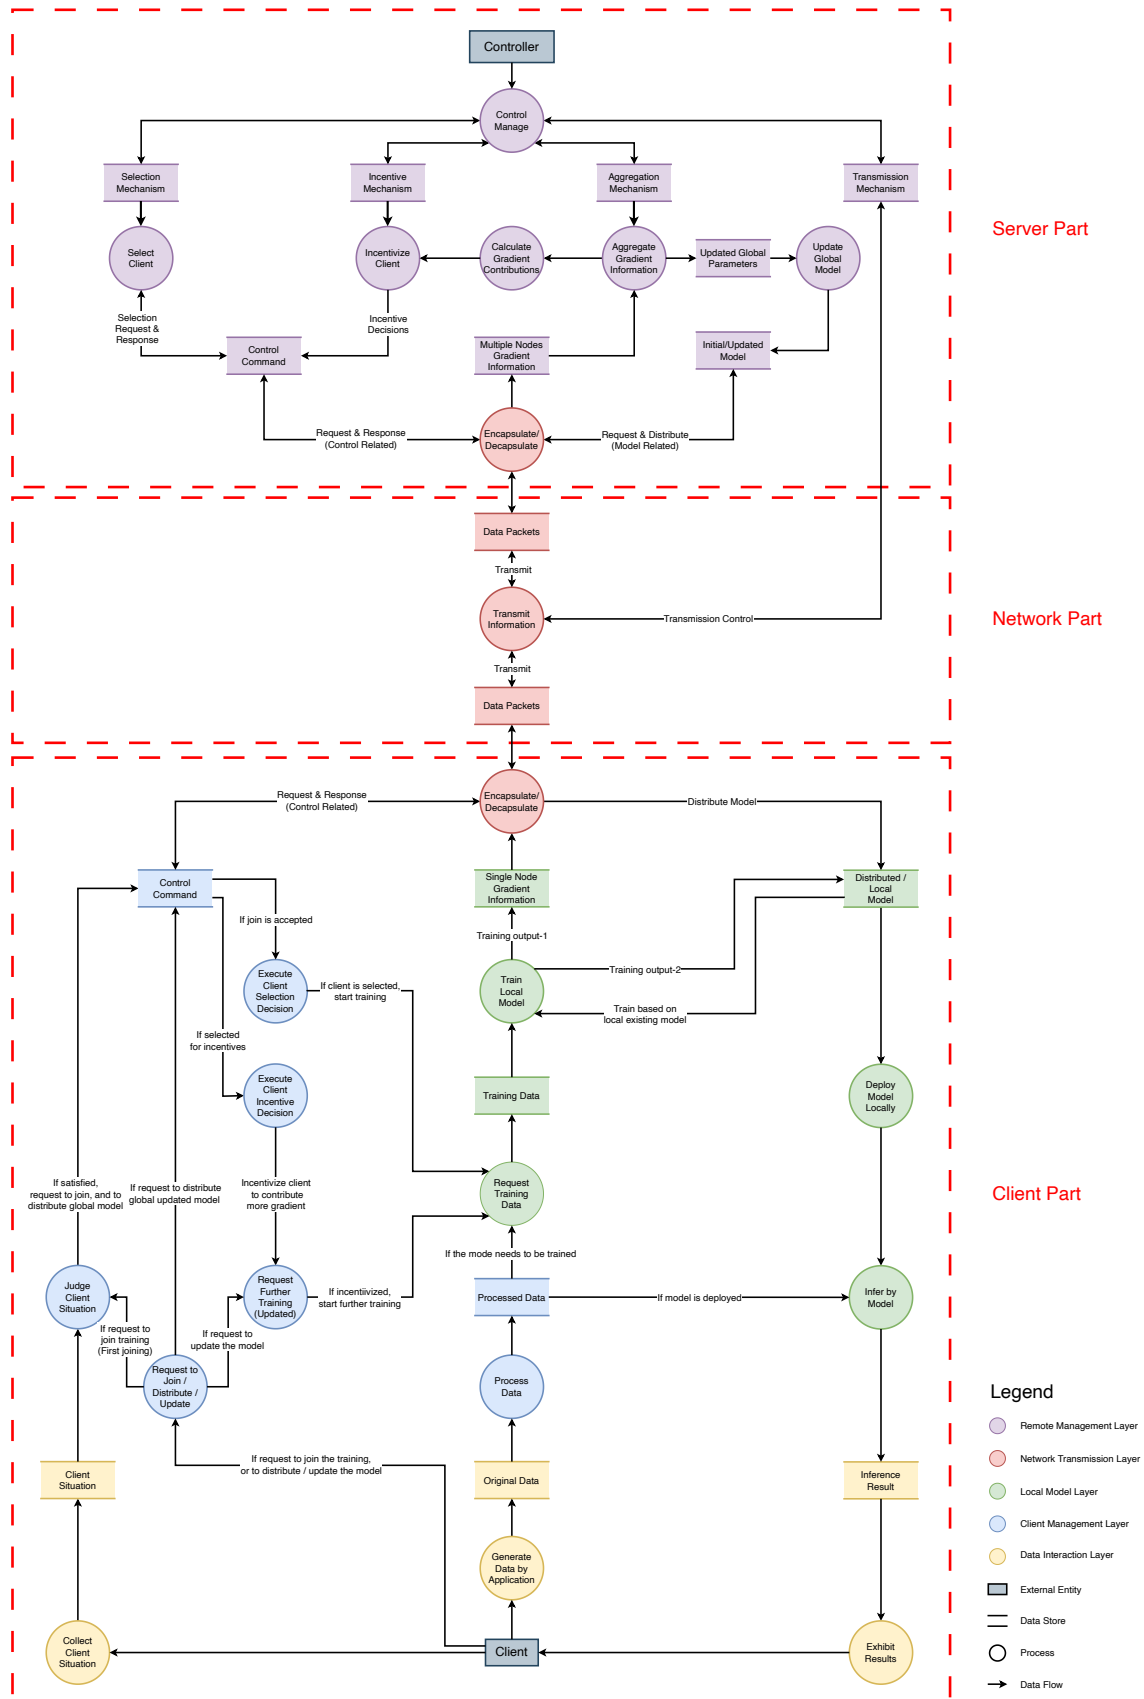


FIGURE 4. Low-level Data Flow of Federated Learning

nisms as the control flow input based on the perspective of remote management layer.

- Network Transmission Layer: For the sake of successful transmission among each participant, a unified packet encapsulation and decapsulation algorithm is required in the network transmission layer. Besides, control information transmission standards including client compute node selection, global model issue, participated node incentivize, etc. are needed to be set up.
- Local Model Layer: The execution of client compute node selection results, updated model deployment, and inference API for common use should be considered at this layer.
- Client Management Layer: The problem becomes simpler as we come to this lower layer. How to process the original inference output from local model layer is the main objective. Since in most instances, it is difficult to directly leverage such original output results. Note that starting from this layer, the control flow becomes an internal process for the participants, which means that, based on following unified standards, participants have room to freely develop according to their actual conditions. For example, the government tax department may have higher requirements than either the insurance company or commercial bank in protecting sensitive information.
- Data Interaction Layer: The only thing the engineers have to do is exhibit the final results to the end user through the data interaction layer.
- **Security Enhancement:** After designing both bottom-up data flow and top-down control flow, a prototype of the federated learning model training system among different participants is created. Then, engineers can follow the low-level data flow diagram proposed in Figure 4 to analyze flows and further enhance the protection of sensitive private information.

With the assistance of the application case example, we could find that our security-oriented 5-layered federated learning reference model simplifies the complex design of cross-enterprise FL model training system into five separate layers, focusing on two flows in opposite directions: bottom-up data flow and top-down control flow. Our proposal could assist engineers to address isolated and simple problems at a time, simplify their work and save system design and deployment time. Additionally, system scalability is enhanced with the improvement in standardization as well. Moreover, our model provides a clear approach to analyze security threats for systematically enhancing data protection.

V. THREATS INTRODUCTION

While one of the primary objectives of federated learning is to protect the privacy of training participants by avoiding the uploading of private and sensitive data to the centralized server, recent research indicates that privacy and security in

federated learning can still be vulnerable under certain circumstances, as highlighted in a study by Lim et al. [15]. This section conducts a brief introduction to the threats inherent in the federated learning system.

A. THREATS MODELS

Before delving into attacks on federated learning, let's introduce the STRIDE model, a simple threat model developed by Microsoft [47] for classifying threats. The STRIDE model is commonly used in threat modeling and risk assessment to identify potential threats and vulnerabilities in a system or application. By categorizing threats into the following six distinct categories, users can gain a better understanding of the types of attacks that could impact their systems and take appropriate measures to mitigate these risks. Table 2 provides an overview of the threats typically associated with each element utilizing STRIDE model, based on Adam's work [47].

- **Spoofing:** Attackers impersonate someone else to gain unauthorized access to the system.
- **Tampering:** Unauthorized modification or alteration of data, such as changing data contents or inserting new data.
- **Repudiation:** Denying that an action or data alteration occurred.
- **Information Disclosure:** Unauthorized exposure, release, or sharing of sensitive or confidential data with individuals or systems that should not have access to it.
- **Denial of Service:** Disrupting the availability of a system or service due to attacks.
- **Elevation of Privilege:** Attacker gains higher-level access or permissions to system or data than they should have.

B. CLASSIFICATION BASED ON ATTACKER'S KNOWLEDGE

Before delving into specific threats in federated learning, another classification is based on the level of knowledge attackers possess about the target machine learning system. This classification comprises three main branches: white-box attacks, black-box attacks, and gray-box attacks [48].

- **White-Box Attack:** In a white-box attack, adversaries have comprehensive knowledge and unrestricted access to the machine learning system, encompassing the model network architecture, model parameters, and even the training data [49]. However, this type of attack is primarily conducted to assess the vulnerability of the system [48].
- **Black-Box Attack:** In a black-box attack, attackers possess little to no relevant knowledge about the model or the original training data. The primary approach involves interacting with the model through limited query access, such as Application Programming Interfaces (APIs), to obtain complete or partial information [50]. Black-box attacks align more closely with real-world scenarios compared with white-box attack, resembling normal us-

TABLE 2. STRIDE Threats Per Element

Element	Interaction	Spoofing	Tampering	Repudiation	Information Disclose	Denial of Service	Elevation of Privilege
External Entity	External interacting entities pass input into the process	✓		✓	✓		
	External interacting entities get input from the process	✓					
Process	Process has external data transferred to the data store	✓		✓			
	External interacting entities get input from the process	✓		✓	✓	✓	✓
	Process sends output to an external interacting entity (Application)	✓		✓	✓	✓	
	Process sends output to an external interacting entity (User)			✓			
	Process has an input data stream from the data store	✓	✓			✓	✓
	Process has an input data stream from another process	✓		✓		✓	✓
	Process has input data streams from external interacting entities	✓				✓	✓
Data Flow (Request/response)	Crossing machine boundaries		✓		✓	✓	
Data Store (Database)	Process has output data flowing to the data store		✓	✓	✓	✓	
	Process has an input data stream from the data store	✓		✓	✓		

age conditions. For example, Shukla et al. successfully launch such an attack with fewer than 1000 queries [51].

- **Gray-Box Attack:** The attacker may know partial training data or some knowledge about the model, encompassing internal states or decision logic of the model, model architecture, or hyperparameters before they launch an attack by approaches like side-channel attacks, model reverse engineering, etc. Compared to white-box attack and black-box attack, gray-box attack is often more challenging and requires more sophisticated attack strategies. It is also close to real-world scenarios.

However, using this taxonomy to analyze security threats in FL systems is challenging, as it just describes the knowledge level of attackers rather than illustrating vulnerabilities. There are other taxonomies as well, for example, Suci et al. propose a framework to classify attackers' knowledge and control from four dimensions [52].

C. FEDERATED LEARNING SPECIFIC THREATS

Similar to other machine learning systems, federated learning systems are susceptible to various forms of attacks. Lyu et al. [53] specifically identify two prominent threats in federated learning: poisoning attacks and inference attacks. In addition to these, we will also discuss two other important threats: backdoor attacks and free-riding attacks. Besides, the NIST (National Institute of Standards and Technology) points out that evasion attack is generating increased interest in machine learning research space [48]. However, it's important to note that each type of attack has unique characteristics.

1) Poisoning Attack

The overarching intent of this assault is to perturb the compartment of the target model's inferences, thereby inducing it to deviate and render erroneous predictions. According to Huang et al.'s work [54], it can be classified into two categories: random attack and targeted attack. The former seeks to curtail model accuracy, while the latter endeavors to

manipulate the model into producing predefined adversarial label information. The classification of poisoning attacks can be based on their source, resulting in two categories: data poisoning attack and model poisoning attack. [53]

- **Data Poisoning Attack:** Data poisoning attacks are recognized as the most common type of attack in the realm of machine learning, and this holds true not only within the domain of federated learning [55]. This prevalent attack primarily occurs during local data collection. Adversarial agents deliberately introduce corruption into training data on a subset of participant devices during the learning phase, thereby compromising the final global model's accuracy. Both definite and indefinite targeted devices could potentially serve as targets for the injection of poisoned data [56]. Tolpegin et al. point out data poisoning attacks can be either targeted or non-targeted [57]. In the targeted variant, the attacker aims to exert influence over a specific subset of data classes, consequently undermining the overall model accuracy. In contrast, non-targeted variants lack the presence of any specific subset of data classes to be targeted for attack [56].

Two forms of data poisoning attacks are recognized: clean-label attack [58] and dirty-label attack [53] [59].

- **Clean-label Attack:** In this scenario, adversaries inject malicious information into the data while keeping the data labels unchanged. Attackers attempt to introduce subtle perturbations to the dataset, causing the model's predictions for this data to be misleading, all the while preserving the true data labels, due to the existing validation process. Detecting such an attack can prove challenging, as the tampered data retains identical labels to the original data, thereby circumventing conventional detection methods.
- **Dirty-label Attack:** Adversaries inject malicious data with disguised labels into the dataset to disrupt the training and performance of the model, leading it to generate misleading predictions. The focus of this attack is on altering the labels of the data, rather than solely the data content itself. Such an attack can yield a more enduring influence on the model, as the malevolent label will continue to impact the model's performance during both training and testing phases. A notable illustration of a dirty-label attack is the label-flipping attack [60].

The extent of impact stemming from a data poisoning attack within the context of federated learning is contingent upon factors such as the number of participating client compute nodes and the proportion of poisoned data within the comprehensive training dataset.

- **Model Poisoning Attack:** This attack occurs during the training of models. In contrast to data poisoning attacks, model poisoning attacks have been demonstrated to be considerably more potent by Bhagoji et al. [61]. In this

scenario, adversaries taint local model updates rather than the underlying local data. These poisoned updates, when introduced, aim to perturb the global model's classification on a chosen set of inputs characterized by high confidence levels [62]. It is important to underscore that the manipulation is intended to influence the model training process, resulting in misclassifications [63].

According to various studies, model updates into which a poisoning attack is injected can be generated through various methods [64] [61], such as the single-shot attack as exemplified by Bagdasaryan et al [62]. Lyu et al. [53] indicate that a model poisoning attack could exploit an alternating minimization strategy to enhance attack stealth and elude detection.

However, due to the fact that a data poisoning attack ultimately affects a subset of model updates, a model poisoning attack is sometimes considered to encompass a data poisoning attack [60].

2) Inference Attack

The phenomenon of inference attacks arises from the utilization of gradients derived from the private data features of participants within various layers, encompassing sequential fully connected layers as well as convolutional layers. These gradients can inadvertently leak sensitive private information, such as class or membership details. This phenomenon is referred to as an inference attack. Notably, Zhu et al. [65] demonstrate the startling revelation that original samples can be reconstructed from uploaded gradients, obviating the need for direct access to the training dataset. Unlike localized poisoning attacks, inference attacks manifest more prominently during the transmission of updated gradient information between participants and the central server. The taxonomies proposed by Lyu et al. [53] and Liu et al. [49] classify inference attacks into seven sub-classes: representative inference, membership inference, property inference, training inputs and labels inference, attribute inference, model extraction, and model inversion.

- **Representative Inference:** In this category, a malicious participant exploits the real-time nature of the training process to train a Generative Adversarial Network (GAN) capable of generating prototypical samples representative of the private targeted training data [66]. Notably, this form of attack exclusively targets class representatives, as the generated samples convincingly mirror the distribution of the original dataset. This strategy relies on the premise that a single participant contributes the entire training corpus, and only when the data of all members are similar, can the reconstructed representatives resemble the original training data. Analogous to a model inversion attack [67], this assumption is often challenging to meet.
- **Membership Inference:** This attack, highlighted by researchers from Cornell University, aims to deduce whether a specific sample originates from private data belonging to a single participant or a group [68]. This

category bifurcates into two variants [53]: passive membership inference attack and active membership inference attack.

- Passive Membership Inference Attack: This variant hinges upon the passive observation of transmitted parameters within updated models, thereby enabling inference without manipulation of any data. This inference can occur during both local and global training procedures [69].
- Active Membership Inference Attack: In this manifestation, attackers manipulate model training protocols to induce the provision of supplementary information regarding the private local data of interest by introducing malicious updates. Nasr et al. [70] term this approach a gradient ascent attack.
- **Property Inference:** Adversaries execute property inference to deduce global information, such as whether the training dataset possesses a specific characteristic [71] [69]. However, the property here is not explicitly stated as the attribute in the training set [72]. Property inference is premised on the assumption that auxiliary training data are labeled with the relevant property. Two variations emerge according to Lyu et al [53]:
 - Passive Property Inference Attack: Adversaries solely observe or eavesdrop on updated information, leveraging a binary property classifier for the purposes of inference.
 - Active Property Inference Attack: This variant employs multi-task learning to induce better data separation, thereby facilitating subsequent extraction of private information.

Notably, the performance of property inference attacks remains unaffected by property presence during model training.

- **Training Inputs and Labels Inference:** An exemplary instance within this category is the concept of Deep Leakage from Gradient (DLG), introduced by MIT [65]. This framework facilitates the restoration of pixel-wise accurate original images and token-wise matching text contents from a limited number of iterations. Zhao et al. [73] extend this concept with the notion of Improved Deep Leakage from Gradient (iDLG), harnessing connections between shared gradient features and labels to extract labels. Differentiable models trained with cross-entropy loss over one-hot labels are susceptible to this approach [73].
- **Attribute Inference:** The attribute inference attack aims to uncover concealed sensitive attributes within a given sample [74] [75]. Fredrikson et al. suggest that the input information encompassed in the confidence output can serve as a measure, as the model prediction results typically incorporate reasoning information about the sample [74]. Liu et al. define the difference between attribute inference attack and property inference attack as follows: an attribute inference attack extracts features

related to the primary task, whereas a property inference attack extracts features unrelated to the primary task [49].

- **Model Extraction:** The goal of a model extraction attack is to extract information about the model architecture and parameters, or other sensitive information, by making numerous queries to the model. As illustrated in NIST's work, there are predominantly three approaches to achieve it: direct extraction, using learning methods, and using side-channel information [48]. Each of these methods has been demonstrated by several literature works.

There are many subclasses of model extraction attacks according to existing works. We introduce four common subclasses here based on Gong et al.'s work [76]:

- Model Parameter Extraction: The objective of model parameter extraction is to recover the model parameters through black-box access.
- Hyperparameter Extraction: Hyperparameter extraction aims to restore the potential hyperparameters that may bring models with significantly different performances with different values [77].
- Model Architecture Extraction: Taking Deep Neural Network model as the target example, attackers attempt to infer the number of neural network layers utilizing a pre-trained model.
- Decision Boundary: Adversaries leverage the approach of training a substitute model based on the labels received from a large number of queries to manipulate the decision boundary of the model. Finally, the model would generate some incorrect predictions under specific inputs.

Note that a model extraction attack often serves as a step towards other powerful attacks, and it could also support transferable adversarial attacks [78] [79], since the attackers have already obtained detailed information about the target model.

- **Model Inversion:** Model inversion attack (also known as attribute inference attack) [80], predominantly leverages the machine learning system's APIs, and non-sensitive attributes of the test input data to obtain hidden sensitive information about the model [49]. This allows attackers to reconstruct class representatives of the original training data [67]. Although it is unable to recover the training data completely [48], the data reconstructed by model inversion still have higher accuracy than characteristics solely inferred from data not present in the training dataset [81].

Dibbo et al. categorize model inversion into optimization-based approach and surrogate model training approach depending on the incorporated techniques [72].

- Optimization-based Approach: Attackers attempt to reconstruct the training instances without the need for training any additional model.
- Surrogate Model Training Approach: Adversaries

seek to invert the sensitive attributes or training samples by exploiting auxiliary information and a surrogate model that utilizes input-output correlation in the target model.

However, it is important to note that both model inversion attacks and membership inference attacks can be executed in either black-box attacks or white-box attacks [81] [82].

3) Backdoor Attack

The backdoor attack, alternatively referred to as a targeted attack [55], involves a malicious participant or a group of such participants intentionally introducing a backdoor function into the global model [59] [62]. In backdoor attacks, the adversary manipulates individual features or small regions of the original training datasets, with the intention of embedding backdoors [53] that afford the attacker control over the behavioral patterns of the global model. Notably, this control is achieved without compromising the overall accuracy of the model. Sun et al.'s research [83] demonstrates that the severity of such attacks is correlated with the proportion of compromised participants and the model's capacity. Remarkably, the insidious nature of backdoor attacks renders them elusive to detection, as they remain inconspicuous within the model's performance on clean data.

4) Free-riding Attack

The free-riding attack occurs when a malicious participant, contributing little or nothing to the federated learning process, seeks to exploit the benefits derived from the global model. The attackers pretend to possess an extremely small training dataset, inducing other participants to allocate additional resources to facilitate the model training process. This attack poses a significant threat, especially in sensitive applications of federated learning [84].

According to research, there are two primary incentives for client compute nodes to engage in free-riding by submitting fake parameters:

- 1) The client compute node is concerned about data privacy, or even have no available local data for model training. In the latter case, to evade security detection, the client compute node has to submit fake gradient information [85].
- 2) The client compute node aims to save its local computing resources such as CPU or GPU cycles, or local storage space [85] [86].

Depending on whether the nodes possess local data and computing resources, Wang et al. classify free-riders in the FL system into two types [87]: anonymous free-riders and selfish free-riders.

- **Anonymous Free-riders:** In this case, there is no available data or computing resources possessed by the free-riders, the attackers submit random Gaussian noise as the training updates [88] [89].

- **Selfish Free-riders:** Selfish free-riders possess local data and computing resources, but are loath to contribute them for model training. Instead, they upload fake updates to the aggregation control server [88] [89].

5) Evasion Attack

In evasion attacks, the primary objective is to construct specific adversarial samples whose classification can be arbitrarily manipulated to a specific class controlled by the attacker at the time of model deployment, while introducing minimal perturbations that are imperceptible to humans [48] [90]. In other words, an evasion attack aims to mislead the model into producing an incorrect prediction result rather than altering the model itself [49]. Various works, such as those conducted by Kearns and Li [91] and Szegedy et al. [63], have successfully demonstrated the feasibility of evasion attacks, and current research continues to focus on designing of adversarial examples.

Although Liu et al. categorize evasion attacks into targeted attacks with class-specific errors and untargeted attacks that do not consider class-specific errors based on the optimization objective [49], Vassilev et al. categorize them according to the perspective of the adversary's knowledge into white-box evasion attack and black-box evasion attack [48].

6) Summary

After introducing some federated learning specific threats above, some comparisons are illustrated in this summary part. Table 3 presents comparisons of the common attack methods that appear in federated learning, and Table 4 emphasizes some other differences among those attack methods. And Table 5 shows the differences among variations in inference attacks.

While these empirical comparisons may not offer a complete assessment, they highlight the intricate security challenges faced by federated learning systems. This prompts us to delve deeper into security threat analysis to bolster the security measures of FL systems.

VI. THREATS ANALYSIS OF EACH LAYER

In order to improve the security level of federated learning while assisting users in narrowing down the scope of the issue, we conduct a layer by layer security analysis along the data flow in the federated learning system, with the help of DFD in this section. Figure 4 provides the detailed information, while Figure 5 illustrates the high-level data flow in federated learning. Note that we only analyze security threats in traffic that terminates at the target layer and ignore traffic originating from the target layer when traffic is bidirectional.

A. DATA INTERACTION LAYER

The utilization of sensors assists researchers in gaining a better understanding of environmental features in their projects. For example, Lu et al. develop a multilevel regression model based on the data they collected [95]. AI techniques like federated learning provide a promising approach to processing

TABLE 3. Comparisons of Various Common Attack Methods in Federated Learning 1

Threats	Category	Variations	Attack Objective	Attack Timing	Attack Approach
Poisoning Attack	Data Poisoning	Clean-label	Mislead model's prediction	Local data collection, training phase	Introduce subtle perturbations to the dataset
		Dirty-label			Alter the labels of the data, rather than the data content
	Model Poisoning	None	Perturb the global model's classification on specific inputs	Training of global model	Local model updates
Inference Attack	Representative Inference	None	Infer training data distribution	Transmission of updated gradient information, during or after the model training phase, inference phase	Train GAN Model, analyze model output
	Membership Inference	Passive Membership Inference	Deduce whether a specific sample originates from private data [68]		Observe transmitted parameters
		Active Membership Inference			Introduce malicious update
	Property Inference	Passive Property Inference	Infer global information about the training data distribution [48]		Eavesdrop on updated information
		Active Property Inference			Employ multi-task learning
	Training Inputs and Labels Inference	None	Restore the training data		Harness connections between shared gradient features and labels
	Attribute Inference	None	Uncover concealed sensitive attributes [92]		Infer based on partial information from select training data and a model trained on that data [93]
	Model Extraction [76]	Model Parameter Extraction	Recover the model parameters		Perform large queries on the model
		Hyperparameter Extraction	Restore the potential hyperparameter		
		Model Architecture Extraction	Infer the architecture of the model		
Decision Boundary		Extract the decision boundary of the targetted model			
Model Inversion	Optimization-based Approach	Reconstruct sensitive features of training data using auxiliary information [94]	Reconstruct of training instances [72]		
	Surrogate Model Training Approach		Train a surrogate model [72]		
Backdoor Attack	None	None	Embed backdoors	Local data preparation phase, before uploading the local model	Manipulate parts of the original training datasets
Free-riding Attack	Anonymous Free-riders	None	Capitalize on the benefits yielded by the global model	Malicious participate joining, training phase	Submit random value noise as model updates
	Selfish Free-riders				Upload fake updates
Evasion Attack [49]	Targeted attack	None	Decline the performance of the model	During the model is training, or inference phase	Modify input data
	Untargeted attacks				

such a huge amount of data. As a result, the data interaction layer serves as a bridge between the federated learning system and user applications. Consequently, it becomes the entry point vulnerable to external threats.

As illustrated in Figure 6, whether the client compute node could join the training process depends on the assessment of the collected client system environment situation. Consequently, the attackers could tamper with or repudiate the sent information, or infer operation environment information according to the client situation information after invading

this system-level process by injecting a backdoor or elevating themselves to higher privileges without being noticed by the users. These attack approaches would disturb the further client application judgment. Furthermore, since the original source data stem from users' daily interactions, attackers could tamper with the raw data generated by applications or launch further data poisoning attacks, including clean-label attacks and dirty-label attacks. Property inference attack is easy to be achieved at this step as well, which is another kind of information disclosure. Alternatively, it is possible

TABLE 4. Comparisons of Various Common Attack Methods in Federated Learning 2

Threats	Attack Targets		Attack Point		Auxiliary Knowledge Required
	Model	Private Information	Aggregation Control Server	Client Compute Node	
Poisoning Attack	✓			✓	Training data distribution, model architecture, data injection
Inference Attack		✓	✓		Model architecture, data distribution, training data
Backdoor Attack	✓			✓	Training data, backdoor injection, system control
Free-riding Attack	✓			✓	Only little about fake information generation
Evasion Attack	✓		✓	✓	Model decision boundaries, mode architecture, output analysis

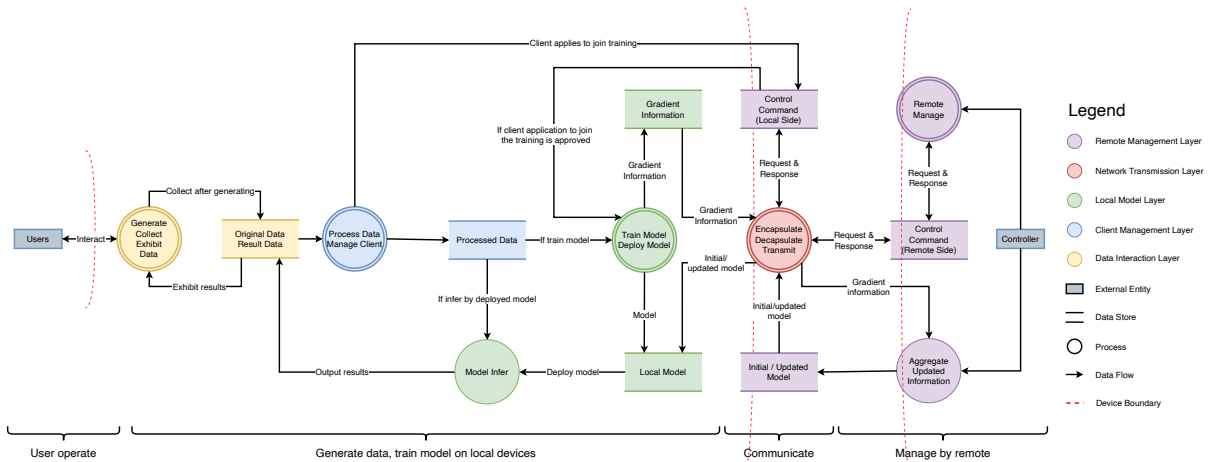


FIGURE 5. High-level Data Flow of Federated Learning Training

for adversaries to disturb the raw data generation process by utilizing repudiation attacks or even denial of service after injecting backdoors in the applications.

Besides, not all users can be guaranteed to act in good faith. As a result, the inference result exhibition part in the data interaction layer is another area with high risk. Once the inferred results are provided by the model, malicious users may launch inference attacks such as representative inference, membership inference, property inference, etc. These could be regarded as private information disclosures. Moreover, they could tamper with the original results to achieve evasion attack. What's more, there remains a possibility of spoofing or repudiation attacks, which are the consequences of results leakage at the results exhibition stage.

B. CLIENT MANAGEMENT LAYER

As discussed in Section IV, the primary role of the client management layer is data processing and client-side client selection, serving as the starting point of FL model training. However, since the decision-making process of client selection occurs at another layer, we will only focus on the security threats analysis of the data processing and join request/response. The execution part of client selection is considered as a reliable security zone.

There are four potential paths for attackers, as illustrated in Figure 7. The first one originates from the request of joining the training or distributing the model. Besides spoofing fake requests or repudiating real requests sent from the client, attackers could infer client system information according to the leaked data. Moreover, it is possible for attackers to overwhelm the client resources by launching a large number of denial of service attacks. On the other hand, the privilege of attackers could be elevated if the adversaries invade, since either join request or distribute request is generated based on the client system's current situation, which requires high privilege to collect. Besides, malicious users who possess little data can launch a free-riding attack to access the training data. While this attack may not always compromise private data, it can significantly degrade the overall training performance due to its zero contribution. Moreover, such free-riding nodes can potentially become vectors for more malicious attacks, as distinguishing between genuine free-riding nodes and malicious ones pretending to be so is impractical. In other words, the attackers are hidden among a large number of normal users.

The second path is related to the collected client environment situation. Apart from the attack methods analyzed

TABLE 5. Comparison of Several Variations of Inference Attacks

Attack Type	Possible Information Leakage	Attack Success Conditions	Impact Scope	Attack Complexity	Attack Detectability	Defense Strategy Example
Representative Inference	Sensitive data distribution	Sufficient data diversity	Training data	Only analyze the statistical characteristics of training data	No frequent queries, few traces left	Improve training data diversity and privacy protection
Membership Inference	Private data	Member data	Data samples	Train an attack model, get shadow member data	Large number of queries may raise suspicion	Differential privacy
Property Inference	Sensitive data distribution	Leaked attribute	Partial attribute	Analyze leaked attribute information	No frequent queries, few traces left	Prevent attribute information leakage
Training Inputs and Labels Inference	Part of training data	Leaked data	Partial data	Independently reconstruct the training data	Require many queries, but uneasy to be detected	Fundamentally prevent data leakage
Attribute Inference	Model information	High model transparency	Model information	Train an auxiliary model based on leaked attribute information because of transparency	Large number of queries	Limit model transparency
Model Extraction	Model information	Enough queries	Model information	Reconstruct the model	Large number of queries	Indistinguishable from common queries
Model Inversion	Part of training data	Private data and leaked model	Partial data	Need knowledge of training data and model parameters	Many queries	Enhance model protection, prevent training data leakage

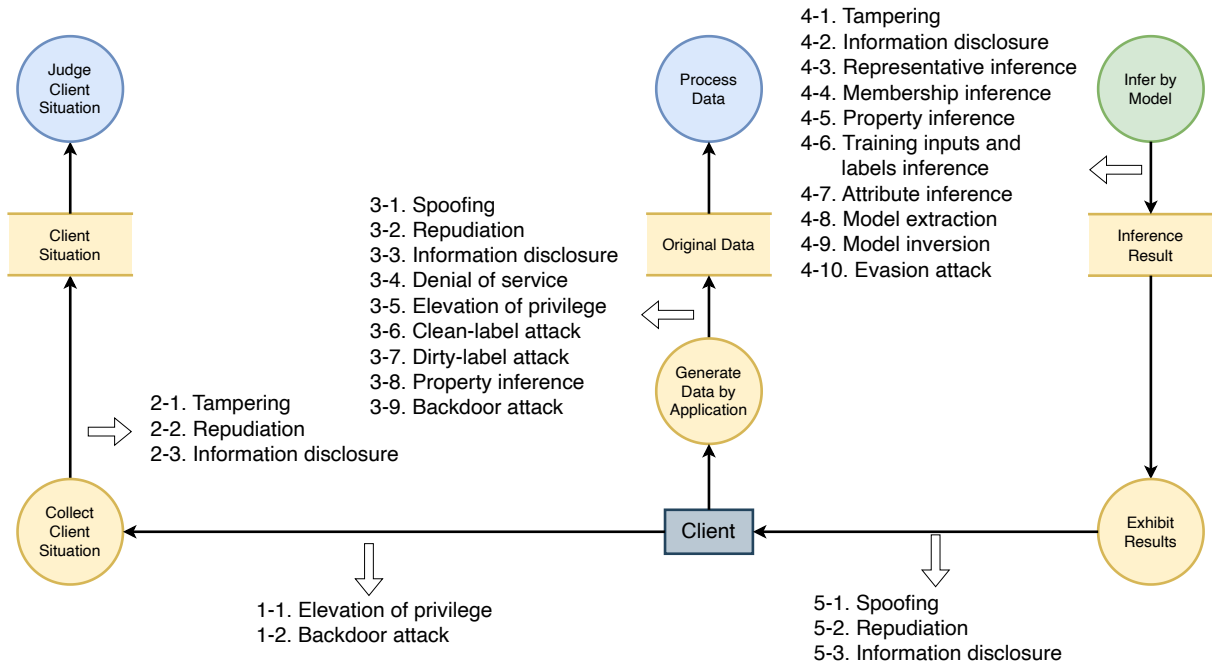


FIGURE 6. Threats in Data Interaction Layer

in the data interaction layer, a malicious participant could initiate attacks from stored data. Spoofing or tampering is possible at this step. There is also the possibility of information disclosure if client compute nodes that don't meet the conditions are allowed to join. However, preventing this risk is challenging, as it may require detecting physical hardware allocation situations while checking the data at the point of entry.

The third attack path comes from the decapsulation of response data sent from the remote aggregation control server. Since the aggregation control server and client compute node are often connected through the network in many scenarios, attacks in this path are a subset of network security threats. In addition to tampering with or acquiring private decapsulated data, attackers could also paralyze the entire training system by repudiating requests or making a large number of repeated requests. Such threats would be discussed in the network transmission layer again.

The last invasion approach is related to data flow. Security threats on this path are similar to those in the previous step in the data interaction layer, as these potential threats all stem from the original data generated by applications. Excluding the free-riding attack, poisoning attacks including clean-label attack and dirty-label attack, and property inference attacks become available since original data are exposed to the adversaries. This could become worth if they inject backdoors into the original data generated by applications, or into the applications themselves.

C. LOCAL MODEL LAYER

Wu et al. [96] state that the federated learning model training paradigm comprises two key components:

- 1) **The client-side:** which trains models on local individual user private data based on a distributed global model and subsequently uploads the updated model information to the central server.
- 2) **The server-side:** responsible for aggregating the uploaded model update information to train the global model.

In our layered referenced model, the client-side model training corresponds to the local model layer that leverages processed data to train the local model. And the server-side model training corresponds to the remote management layer, which would be discussed later.

There are three sources from which security threats originate: the request for training data before the model training phase, the execution of client selection/incentive decisions, and the distribution of the global model to local participants, as depicted in Figure 8.

Initially, as the model is trained on the client-side by requesting training data from client compute nodes, attackers could spoof the data as being generated by the authorized application to disrupt the training dataset. Moreover, there is a significant possibility of private data being disclosed to malicious attackers who breach this phase, as the processed data are often not encrypted. This vulnerability could lead to further clean-label attacks and dirty-label attacks. Attackers may passively or actively infer property information of the training data by utilizing the disclosed data. Furthermore, the

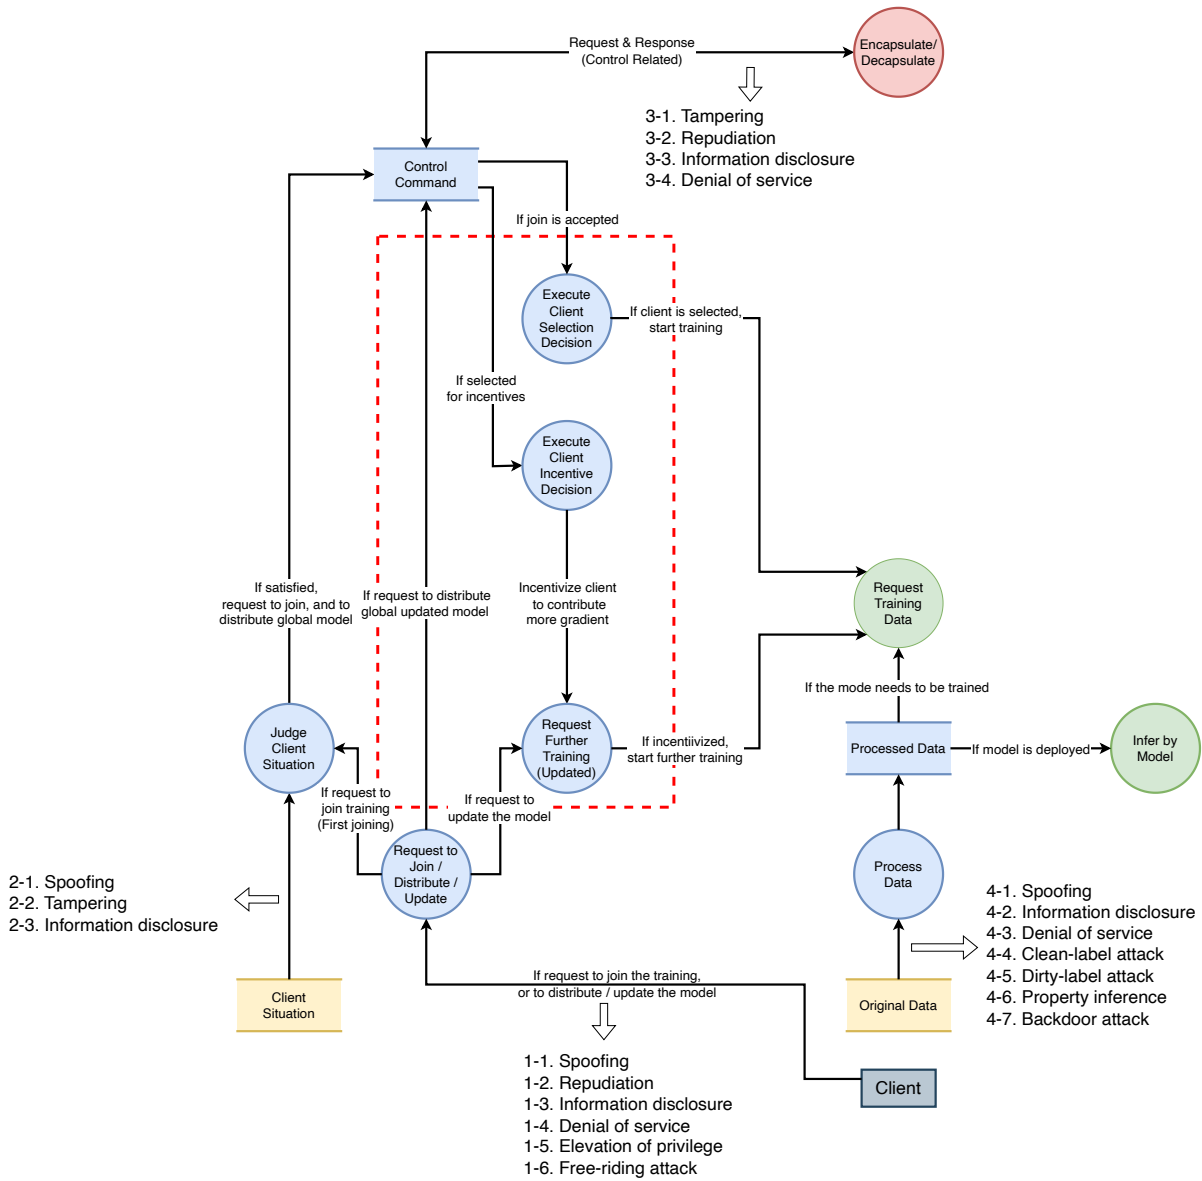


FIGURE 7. Threats in Client Management Layer

system might be overwhelmed if adversaries send a large number of data requests in a short time. The local model layer would become even more vulnerable if backdoors are injected.

Moreover, attackers can easily disrupt client selection/incentive by spoofing decisions sent from the aggregation control server. System situations may become transparent as a significant amount of information is disclosed. Additionally, clients might become non-functional if adversaries initiate many repudiations or a large number of requests in a short time, due to limited resources being exhausted. Moreover, attackers will possess higher privileges through previously injected backdoors, while free-riding attacks are another approach to steal trained model information. On the other hand, attackers could launch attacks on the model infer-

ence process from the same data source. During the inference stage, processed data becomes a vulnerable point of attack before being sent into the model for subsequent inference. Apart from spoofing the real data source, attackers could manipulate the data during transmission between processes or utilize the real data to infer user private training data. Furthermore, if the attackers gain the privilege of sending processed data to the model for inference, the risk of launching denial of service attacks also increases.

The third source of risk arises from the decapsulation of data distributed from the remote aggregation control server via the network. The transmission phase faces the same security threats as the network transmission layer, which will be explained in the next subsection. Since the trained model is issued to the client compute node at this step, malicious

attackers could easily poison or even tamper with the model to mislead the subsequent inference results. Moreover, model-related information such as model parameters and model architecture is also exposed to them, which could lead to private information disclosure. Additionally, the model's performance would also be affected by the evasion attack initiated from adversaries. However, verifying the received information and detecting abnormal activities could effectively decrease the security risk originating from this source.

D. NETWORK TRANSMISSION LAYER

In contrast to conventional machine learning model training paradigms, federated learning highlights communication as a pivotal bottleneck that must be paid attention to, since communication costs often dominate computation costs in federated learning, according to the research by McMahan et al. [97]. Within the domain of federated learning, several rounds of communication between the aggregation control server and client compute nodes are required for the model's accuracy. Furthermore, given that many participants are often operating under public network environments such as Ethernet, the presence of unreliable network conditions and asymmetric public links could exacerbate the security situation.

Figure 9 illustrates the data flow of the network transmission layer, which involves the transmission of various data types, including gradients, models, and control commands. In the gradient transmission part after training on client compute node (Threats Group 1), it is very easy for attackers to infer the membership and property of training data through observing the gradient information. If this malicious observation continues, the possibility of more complex attack approaches like model extraction attack and model inversion attack increases, since deeper information could be revealed by comparing the differences among gradient information transmitted in each rounds. This would cause more serious information leakage. Moreover, malicious attackers could also simply disturb the gradient computing results through spoofing, i.e., sending fake gradient information, or make the gradient information unavailable by repudiation and denial of service attacks. The situation of model-related request and distribution flow is similar (Threats Group 5). The only difference is that adversaries could launch model poisoning attacks instead of membership inference and property inference as the entire models are available during the model issuance process.

As for the control flow, including request commands sent by client compute nodes (Threats Group 2), control responses (Threats Group 4), and transmission controls (Threats Group 6), the security risks faced by them are almost the same due to the same data type in their traffic. Attackers could spoof unauthorized control command sources to affect the control decision-making, or exploit control information to infer user system-level situations, which are information disclosure. Moreover, adversaries may repudiate effective commands or even launch denial of service attacks to interfere with the normal operation of the FL system. However, since the trans-

mission control is often exposed to insecure public network environments in many scenarios, attackers may initiate network attacks for higher privileges to further invade the entire training system.

The information transmission step (Threats Group 3) is the core part of the network transmission layer. This highlights the fact that the network transmission layer is susceptible to all six common threats: spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege, since the entire system is exposed to the public through this layer. Compared with invasion from the data interaction layer, invading via this layer is much easier due to the complexity and anonymity offered by the public network. Moreover, as gradient information is transmitted, there's a risk of revealing sensitive private information [98] [67] [69], and it can even result in deep leakage, as indicated by Zhu et al. [65]. Notably, even partially updated gradients could be exploited to unveil local data, as shown by Aono et al. [99]. Such leakage could occur either to a third party or even to the malicious central server [100] [101] [53]. Once the attackers gain access to the transmitted information, serious attacks including model poisoning attack, model extraction attack, and model inversion attack become possible. Except for the deep leakage issue mentioned above, attackers could utilize the model update information to poison the subsequent global model. In general, the network transmission layer is the most vulnerable part of the entire federated learning system, so it may serve as another potential entry point for malicious attackers from outside the system.

E. REMOTE MANAGEMENT LAYER

Serving as the orchestrator, the remote management layer not only aggregates the gradient information uploaded from the participating client compute nodes to train an updated global model, but also performs various management functions. Such functions include client compute node selection, incentive mechanisms, aggregation mechanisms, and transmission management.

As Figure 10 shows, there is another security zone within it. The entrance of data from the network transmission layer is the main threat, while the invasion from the controller is the secondary threat. Regarding the data received from the network transmission layer, it is the most vulnerable point as discussed in the previous subsection. The situation becomes worse if the data type is a command about client selection, client incentive, model distribution, and transmission control (Threats Group 1, 4, and 5). Because malicious invaders can inject a backdoor during data transmission at the network transmission layer, through which they can control either the entire system or the private data without being noticed. If attackers deliberately tamper with the control request command, they would seriously interfere with the decision-making mechanism. Moreover, there is a possibility for them to infer the private data including system status and training data quantity by analyzing the control request commands. This represents a threat of information disclosure. The con-

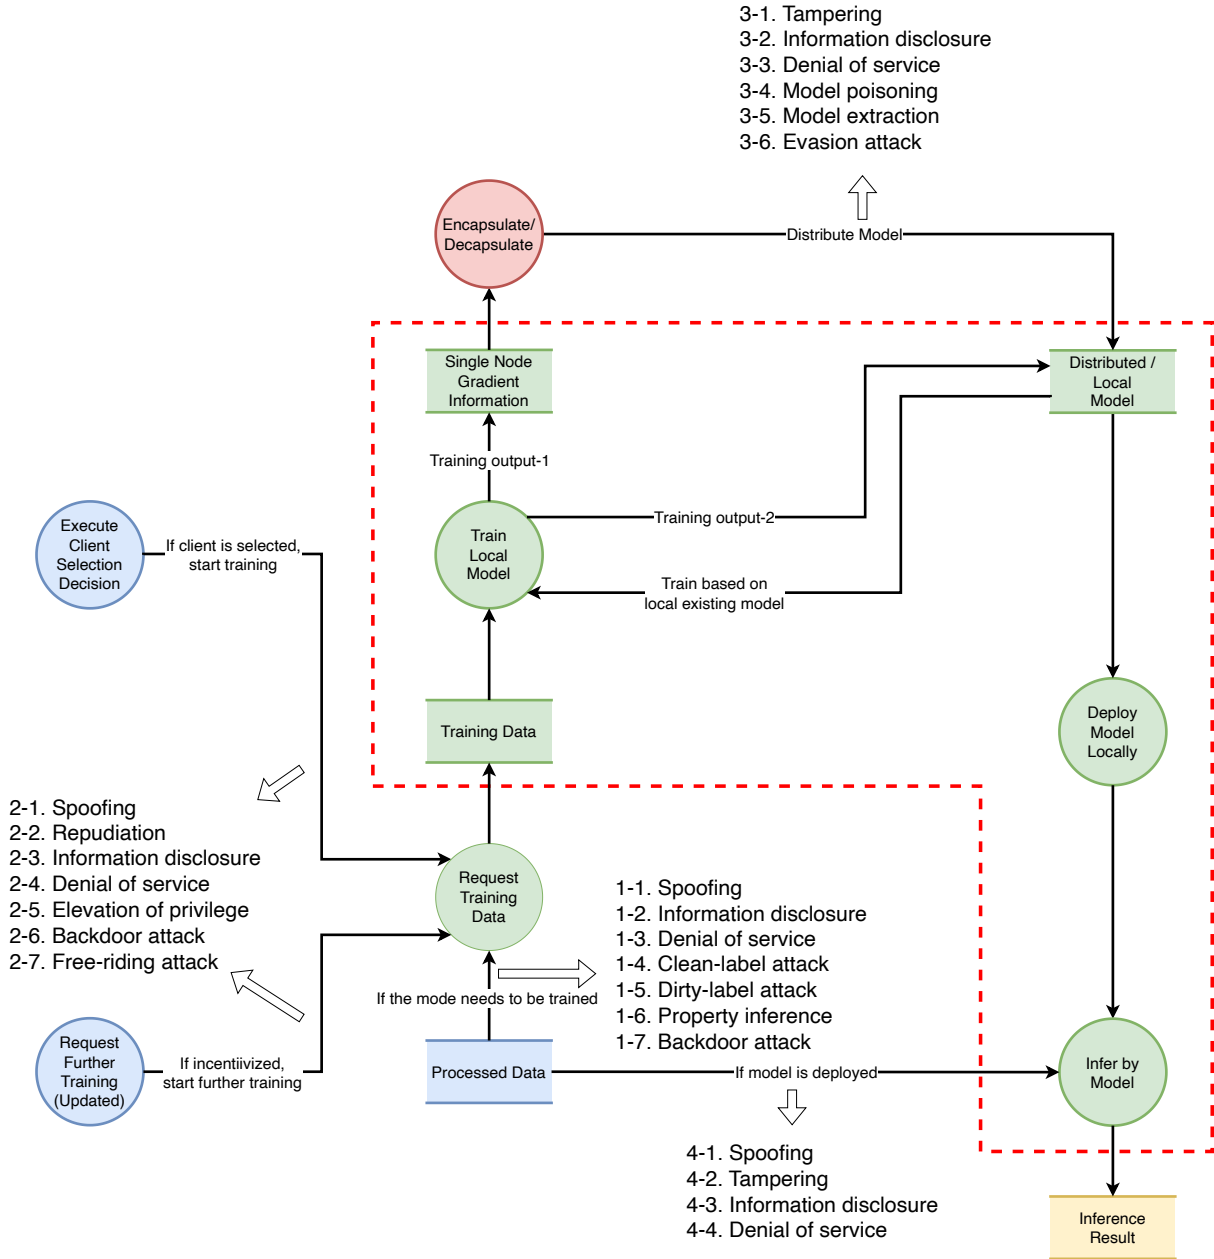


FIGURE 8. Threats in Local Model Layer

troller might run out of service if they successfully exhaust the controller's resources deliberately or frequently repudiate normal control requests. However, issuing a command while identifying the actual role of the requester is a difficult task.

Moreover, attackers could spoof a legal client compute node to disturb the updated gradient information aggregation of the controller if they invade the decapsulation process through which data is transmitted into the remote management layer. Adversaries could poison the model by tampering with updated gradient information as well. Once some private gradient information is leaked at this step, model inference attacks including membership inference, property inference,

etc., could be utilized for delving more private information.

As for the threat issues from the controller side, although it is within the entire system, meaning that the operators are often well-trained and the server is well-protected, since it plays the role of managing the whole federated learning system, small operation mistakes could evolve into attacks like spoofing, elevation of privilege, which could cause serious information disclosure. Furthermore, some backdoors may be unintentionally injected if software security is not given sufficient attention.

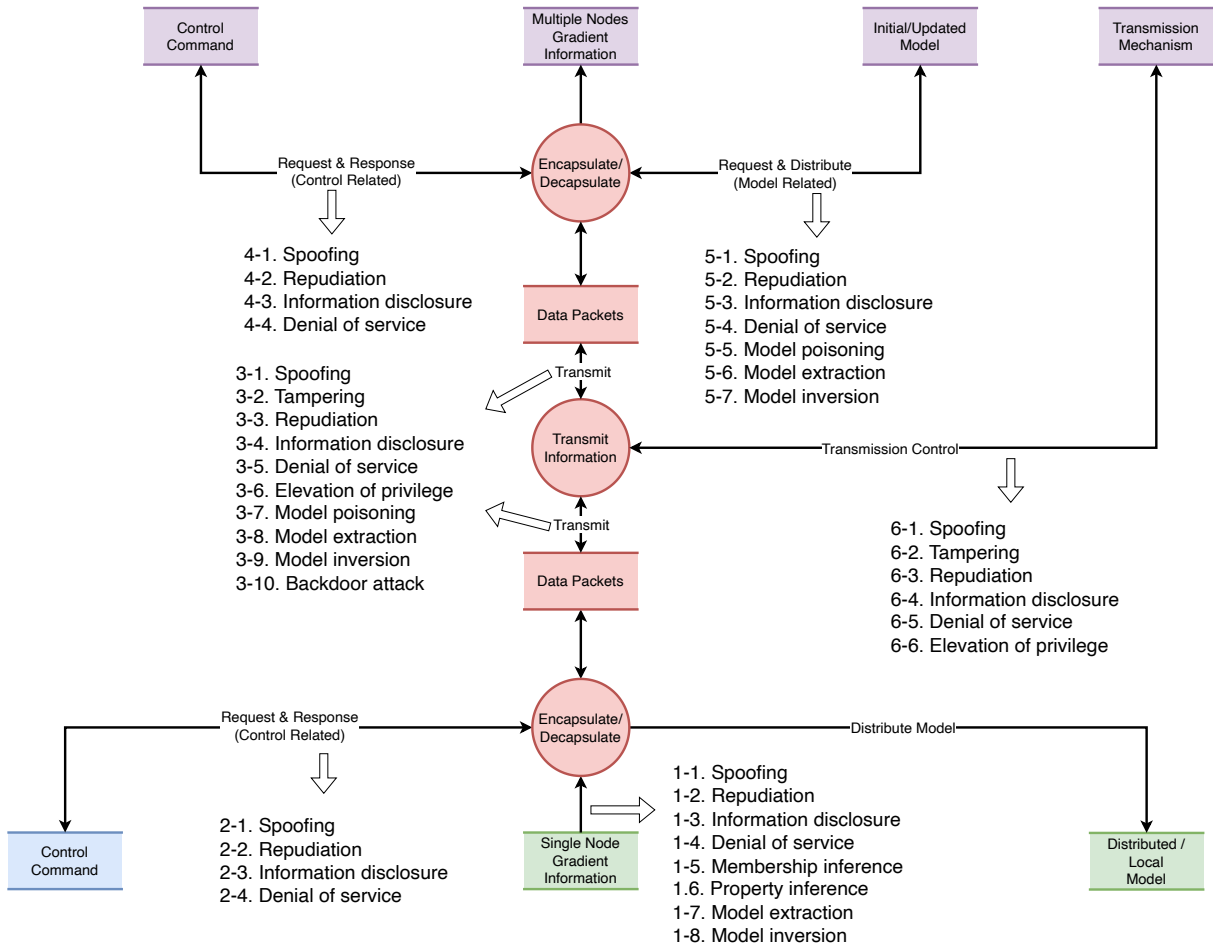


FIGURE 9. Threats in Network Transmission Layer

F. SUMMARY

In addition to the threats analysis layer by layer above, the following two critical vulnerabilities in the design of federated learning protocols identified by Lyu et al. [53] are useful for the security analysis as well:

- 1) The control server continuously monitoring updated information, could potentially be malicious, leading to tampering with the model training process and altering participants' view of the global parameter.
- 2) Participants might also be able to observe or control the global parameters and the upload process of those parameters.

According to these analysis, we could draw a conclusion that client compute node and the remote aggregation control server are the two primary vulnerable entities in the entire federated learning system from the perspective of constitution, rather than the computational communication framework itself. While data interaction layer, network transmission layer, and remote management layer are the three main attack entrances from outside, from the perspective of our proposed layered model.

However, for the sake of increasing the information secu-

urity level of the federated learning system, users can leverage some common strategies to defend against these issues, which will be introduced in Section VII.

VII. THREATS DEFENSE STRATEGIES

Federated learning, especially in the context of cross-silo federated learning, offers a collaborative approach for different organizations to jointly train machine learning models by aggregating gradient updates without exposing sensitive data. However, despite its merits, there remain certain security concerns within this framework as we have analyzed above. In order to bolster information security and privacy, several established security mechanisms can be integrated into the framework of federated learning.

In this section, we will briefly introduce some common defense strategies based on the STRIDE model mentioned in Section VI, followed by strategies against FL-specific threats. Finally, we will illustrate four specific approaches that can effectively improve the overall security level of the system.

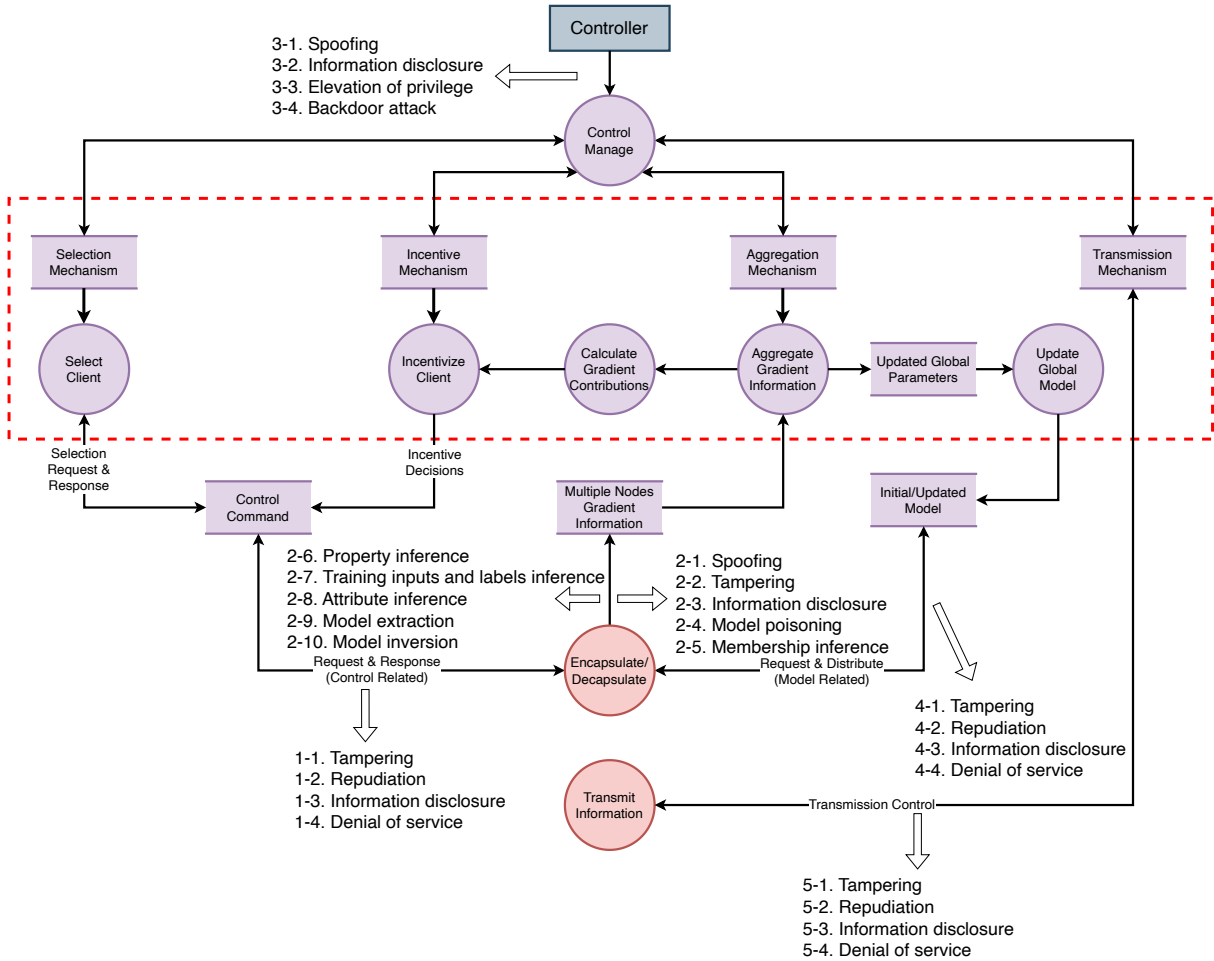


FIGURE 10. Threats in Remote Management Layer

A. DEFENSE STRATEGIES BASED ON STRIDE MODEL

As we discussed in Section V-A, spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege are the six common threats identified according to the DFD. Threats of different risk levels may require different mitigation strategies. In Table 6, we list some common strategies for addressing these threats. Note that these defense strategies are not specific to issues in the federated learning system, but rather general solution.

B. DEFENSE STRATEGIES AGAINST FL-SPECIFIC THREATS

In Section V-C, we have identified five specific threats unique to federated learning. These threats pose challenges to the security and privacy of the training and inference phase. However, it's essential to note that for each of these threats, there exist corresponding defense strategies tailored to address them effectively.

- **Poisoning Attack:** Employing an identification mechanism to scrutinize malicious participants through their model updates before integrating them into the round of model averaging in each learning iteration is an effective strategy against data poisoning attacks. Additionally, the

adoption of rejection based on error rate and loss function emerges as a prevalent approach to mitigate model poisoning attacks [102]. These defend strategies can be employed independently, or in conjunction with the identification mechanisms to defend against poisoning attacks.

- **Inference Attack:** In essence, the successful execution of inference attacks necessitates substantial computational resources and advanced technical capabilities on the part of adversaries. Furthermore, these attackers must sustain their efforts consistently over several training rounds. The employment of homomorphic encryption (HE) [26] could serve as an efficacious mechanism for safeguarding shared gradients from the vulnerabilities posed by inference attacks. Additionally, differential privacy (DP) offers another approach to defend against inference attacks by introducing random noise to the FL system, and the introducing location could vary from data to model parameters [103]. These two methods will be further discussed in Section VII-C. Apart from HE and DP, strengthening the security level of the network where the FL system is located, and promptly identify-

TABLE 6. Common Defense Strategies to STRIDE Threats

Threats	Attribution	Example Defense Strategy
Spoofing	Certification	Basic certification, digital signature, verification code, hash table
Tampering	Integrity	Integrity control, Access Control List (ACL)s, digital signature, verification code
Repudiation	Udeniability	Security logging and auditing, digital signature, Trusted third party
Information Disclosure	Confidentiality	Encryption technology, ACLs
Denial of Service	Availability	ACLs, filtering, authorization, high availability design
Elevation of Privilege	Authorization	ACLs, group or role affiliation, privilege ownership, permissions, input validation

ing potentially malicious nodes that have existed for a long time but rarely accept management by the remote management layer can help reduce the likelihood of inference attacks. Furthermore, Jia et al. propose a novel method that adds crafted noise to each confidence score vector to defend against membership inference attacks [104], falling under the sub-category of information obfuscation defense. In addition to information obfuscation defense, other approaches such as limiting query control and creating more robust model architectures [48] are also efficient in defending against model extraction attacks. The split FL framework called ResSFL proposed by Li et al. shows another novel approach to mitigate model inversion attack [105].

- **Backdoor Attack:** To counteract backdoor attacks, Mammen et al. advocate for the application of weak differential privacy or the imposition of a norm thresholding mechanism on updates [55]. They suggest that norm thresholding can eliminate models with artificially boosted parameters, while the adoption of differential privacy introduces a degree of protection, albeit potentially influencing the overall model's performance [106]. It is imperative to note that the task of distinguishing malicious participants from those exhibiting misbehavior presents a formidable challenge.
- **Free-riding Attack:** The mitigation of free-riding attacks can be advanced through the incorporation of blockchain technology [107] [108]. Such an approach entails the verification on exchanged model updates among devices, thereby engendering an environment that incentivizes participants to contribute actively to the training process. It is noteworthy that an increase in communication payload is found, primarily attributed to block size. However, these findings still prompt the potential for a fully decentralized federated learning architecture. Furthermore, Lin et al. propose STD-DAGMM, which is effective in detecting free-riders [85]. Moreover, the FRAD mechanism designed by Wang et al. outperforms their baselines in their experiments on defending against free-rider attacks [88]. In addition to defense methods, some researchers also focus on attack

approaches. For example, Zhu et al. conducted an advanced free-rider attack [86], contributing from another perspective.

- **Evasion Attack:** In general, there are two main approaches to mitigate evasion attacks: empirical defense and certified defense [49]. Although empirical defense may not guarantee foolproof defense, it is efficient in addressing attacks, especially when designed to target specific types of attacks. Hong et al. [109] identify four classes within empirical defense:

- 1) Gradient-Masking Defense [110] [111]: This defense involves modifying the model inference process.
- 2) Input-Transformation Defense [112] [113]: This defense preprocesses the input before transformation.
- 3) Adversarial Training [114] [115]: First proposed by Goodfellow et al. [78], this approach aims to enhance model robustness by introducing adversarial examples into the training data. However, adversarial training requires more time for training as it iteratively generates adversarial examples [48].
- 4) Adversarial Examples Detection [116] [117]: This defense focuses on the robustness of the model.

On the other hand, certified defense aims to guarantee consistent predictions for a classifier when encountering adversarial examples input. Randomized smoothing, introduced by Lecuyer et al. [118], is the first method in the field of certified defense that is able to certify arbitrary classifiers of any scale [109]. In addition to randomized smoothing, Katz et al. also propose a method based on mathematically rigorous techniques [119].

C. OTHER DEFENSE STRATEGIES

In addition to the problem-specific solutions mentioned above, there are some other approaches that could improve the security level of the entire system to enhance the protection of private data. For example: Homomorphic Encryption (HE), Differential Privacy (DP), Secure Multiparty Computation (SMC), and Trusted Execution Environment (TEE).

1) Homomorphic Encryption (HE)

Homomorphic Encryption (HE) is a cryptographic technique that allows mathematical operations to be performed on encrypted data without requiring prior decryption. In conventional encryption, a shared key pair (comprising a public key and a private key) is utilized for message encryption. However, this method raises privacy concerns in certain scenarios, such as cloud services, where the control rights over data can be a subject of contention. Additionally, third parties with no direct involvement may access the contents of shared encrypted data, even without access to the encryption keys [120]. Furthermore, the identification of user elements may endure, even long after the users terminate their engagements with the service. To address these challenges, the concept of HE is proposed as a means to enable secure data operations without revealing the original data. While the theoretical framework for Fully Homomorphic Encryption (FHE) was introduced by Gentry et al. in 2009 [121], practical realization of its generality across diverse platforms remains an ongoing challenge [120].

Accordingly, HE can be applied in an additive manner within industrial federated learning frameworks to ensure that no update is revealed during the aggregation process. However, this approach may introduce computational and communication bottlenecks, which have been mitigated by the development of the BatchCrypt solution [122].

2) Differential Privacy (DP)

Differential Privacy (DP) is a privacy-preserving technique that involves adding random noise to data or model parameters to prevent the inference of private data [42] [123]. DP, grounded in information-theoretic principles [124] [125], provides robust statistical privacy guarantees by making it difficult for adversaries to extract specific information from the data.

Within federated learning, DP can be classified into two main types: Central Differential Privacy (CDP) and Local Differential Privacy (LDP) [126].

- **Central Differential Privacy (CDP):** In CDP, a central server is entrusted with a high degree of trust. Model parameters generated in each update round are initially aggregated on this centralized server before being perturbed. This approach ensures that model parameters remain concealed from other participating client nodes.
- **Local Differential Privacy (LDP):** LDP introduces a higher level of privacy, as model parameters are protected not only from external entities but also from the central server. Differential private transformations are applied to data on each participant before being transmitted to the central server. Despite offering stronger security, LDP may engender challenges in reconciling security with utility preservation [127].

Research underscores that client-level DP could effectively thwart any malicious participant's attempt to reconstruct the private sensitive data of others through the utilization of

the global model. [126]. Accordingly, numerous researchers have endeavored to implement client-level DP in federated learning. For instance, Geyer et al. have integrated CDP into federated learning frameworks [106], while Pihur et al. have proposed a novel federated learning framework based on LDP [128].

Despite its capability to counter inference attacks and provide privacy guarantees through the introduction of noise into clipped model parameters before aggregation [106], a trade-off between privacy protection and model accuracy exists in both CDP and LDP, since the introduction of additional noise may compromise model accuracy [126].

3) Secure Multiparty Computation (SMC)

Secure Multiparty Computation (SMC), an alternative to homomorphic encryption, is a technique that enables participants to collaboratively compute results while solely revealing computation results to a specific subset of involved participants. It is a crucial branch of techniques in Secure Computation (SC) that has already been operational in federated learning [55].

SMC safeguards the original model accuracy while furnishing robust security assurances for participants. This is particularly useful for federated learning, because computation results are shared while maintaining data privacy with each participant. SMC ensures that even the central server could gain minimal information and only observe the aggregated results during each iteration.

Despite the additional computation and communication costs associated with SMC, it remains a significant research area. For instance, Bonawitz et al. have proposed a protocol based on Shamir's Secret Sharing, a widely-used approach in federated learning, to secure updates aggregation [129].

4) Trusted Execution Environment (TEE)

The Trusted Execution Environment (TEE) serves as a secure platform embedded within a device's primary processor, ensuring the confidentiality and integrity of both data and code [55]. By enforcing a dual-world implementation, TEE achieves robust isolation and attestation of secure compartments, effectively mitigating threats posed by compromised or malicious system software within the normal operational realm [130]. In contrast to SC, TEE introduces a lower computational overhead when integrated into the implementation of federated learning [131].

Significant academic and industrial endeavors have been directed towards the advancement of TEEs, aiming to provide secure programmability for various device types, including high-end and mobile devices. Notable examples include Intel Software Guard Extensions (Intel SGX) [132] and the open-source framework Open Portable Trusted Execution Environment (OP-TEE) [133]. An inherent advantage of TEE lies in its substantial reduction of the Trusted Computing Base (TCB), simultaneously maintaining execution speed while enhancing security. Trust is established through the execution of code within the secure environment. Nonetheless, careful

consideration must be given to designing the size of the TEE and mitigating potential overhead during the swapping of memory between secure and unprotected domains [130]. It is pertinent to note that Mo et al. emphasize the necessity of compact TEEs due to the associated heightened vulnerability to attacks [131]. The Privacy-Preserving Federated Learning (PPFL) approach introduced by them exemplifies the feasibility of robust privacy assurances coupled with comparable utility of machine learning models, all while exerting minimal impact on communication and system costs [130].

In order to enhance security and privacy in federated learning, a combination of these mechanisms can be employed [59] [134] [135] [136] by users to establish a robust security framework for the entire system.

VIII. CHALLENGES AND CONCLUSION

Federated learning, while promising and versatile, presents practical challenges as research progresses. According to Sone et al., privacy protection in federated learning focuses on minimizing privacy leakage and adhering to privacy compliance [126]. However, practical scenarios often overlook the diverse stipulations governing heterogeneous parties. Existing techniques offer uniform privacy protection across all parties [14], but in many cases, adherence to individual data owner privacy compliance is acceptable, tailoring treatment based on individual restrictions while maintaining equivalent model performance. Techniques like heterogeneous differential privacy [137] show their promise, but the increase in communication overhead and the decrease in model accuracy are a thorny problem for most solutions. Challenges like fairness among nodes, client incentives [138], fault tolerance, one-shot federated learning [139], etc., are also key research areas in the future.

This paper provides a security-oriented overview of federated learning, focusing on offering readers a systematic introduction to security and risk prevention methods from the perspective of data flow. We expand a comprehensive survey about FL using the proposed 5-layer reference model, which includes the data interaction layer, client management layer, local model layer, network transmission layer, and remote management layer. Additionally, we conduct a brief security threats analysis based on the STRIDE model and specific issues faced by each FL layer. Finally, we introduce some existing defense strategies for each security issue.

In summary, our proposed layered model aims to offer a concise compendium of federated learning along with data flow and control flow in the federated learning system for general readers, emphasizing the importance of system security and sensitive data protection. For researchers, we hope it could provide a standardized and layered reference model to simplify and accelerate their complex work in designing FL model training system, alleviate development pressure, improve FL system scalability, and systematically enhance the information security of federated learning by providing a clear threats analysis approach.

REFERENCES

- [1] U. Cisco, "Cisco annual internet report (2018–2023) white paper," *Cisco: San Jose, CA, USA*, 2020.
- [2] A. Dogra, R. K. Jha, and S. Jain, "A survey on beyond 5g network with the advent of 6g: Architecture and emerging technologies," *IEEE Access*, vol. 9, pp. 67 512–67 547, 2020.
- [3] R. Pryss, M. Reichert, J. Herrmann, B. Langguth, and W. Schlee, "Mobile crowd sensing in clinical and psychological trials—a case study," in *2015 IEEE 28th International Symposium on Computer-Based Medical Systems*. IEEE, 2015, pp. 23–24.
- [4] W. Li and M. Liewig, "A survey of ai accelerators for edge environment," in *World Conference on Information Systems and Technologies*. Springer, 2020, pp. 35–44.
- [5] W. He, G. Yan, and L. Da Xu, "Developing vehicular data cloud services in the iot environment," *IEEE transactions on industrial informatics*, vol. 10, no. 2, pp. 1587–1595, 2014.
- [6] D. Oletic and V. Bilas, "Design of sensor node for air quality crowdsensing," in *2015 IEEE Sensors Applications Symposium (SAS)*. IEEE, 2015, pp. 1–5.
- [7] R. Jhawar, V. Piuri, and M. Santambrogio, "Fault tolerance management in cloud computing: A system-level perspective," *IEEE Systems Journal*, vol. 7, no. 2, pp. 288–297, 2012.
- [8] H. Li, K. Ota, and M. Dong, "Learning iot in edge: Deep learning for the internet of things with edge computing," *IEEE network*, vol. 32, no. 1, pp. 96–101, 2018.
- [9] X. Yi, F. Liu, J. Liu, and H. Jin, "Building a network highway for big data: architecture and challenges," *Ieee Network*, vol. 28, no. 4, pp. 5–13, 2014.
- [10] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, "Federated optimization: Distributed machine learning for on-device intelligence," *arXiv preprint arXiv:1610.02527*, 2016.
- [11] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [12] H. B. McMahan, E. Moore, D. Ramage, and B. A. y Arcas, "Federated learning of deep networks using model averaging," *arXiv preprint arXiv:1602.05629*, 2016.
- [13] G. Baruch, M. Baruch, and Y. Goldberg, "A little is enough: Circumventing defenses for distributed learning," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [14] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He, "A survey on federated learning systems: vision, hype and reality for data privacy and protection," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [15] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y.-C. Liang, Q. Yang, D. Niyato, and C. Miao, "Federated learning in mobile edge networks: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 2031–2063, 2020.
- [16] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [17] M. Aledhari, R. Razzak, R. M. Parizi, and F. Saeed, "Federated learning: A survey on enabling technologies, protocols, and applications," *IEEE Access*, vol. 8, pp. 140 699–140 725, 2020.
- [18] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantanha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, pp. 619–640, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167739X20329848>
- [19] J. Xu, B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang, "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, vol. 5, pp. 1–19, 2021.
- [20] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021.
- [21] R. Gosselin, L. Vieu, F. Loukil, and A. Benoit, "Privacy and security in federated learning: A survey," *Applied Sciences*, vol. 12, no. 19, 2022. [Online]. Available: <https://www.mdpi.com/2076-3417/12/19/9901>
- [22] K. K. Coelho, M. Nogueira, A. B. Vieira, E. F. Silva, and J. A. M. Nacif, "A survey on federated learning for security and privacy in healthcare applications," *Computer Communications*, vol. 207, pp. 113–127, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S014036642300172X>
- [23] J. Hasan, "Security and privacy issues of federated learning," 2023.

- [24] E. Gabrielli, G. Pica, and G. Tolomei, "A survey on decentralized federated learning," 2023.
- [25] D. R. Kandati and S. Anusha, "Security and privacy in federated learning: A survey," *Trends in computer science and information technology*, p. 29–37, Aug 2023. [Online]. Available: <https://www.peertechzpublications.org/articles/TCSIT-8-166.php>
- [26] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.
- [27] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings et al., "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [28] K. Bonawitz, P. Kairouz, B. McMahan, and D. Ramage, "Federated learning and privacy: Building privacy-preserving systems for machine learning and data science on decentralized data," *Queue*, vol. 19, no. 5, pp. 87–114, 2021.
- [29] A. Lalitha, S. Shekhar, T. Javidi, and F. Koushanfar, "Fully decentralized federated learning," Dec 2018. [Online]. Available: <http://bayesiandeeplearning.org/2018>
- [30] Q. Li, Z. Wen, and B. He, "Practical federated gradient boosting decision trees," in *Proceedings of the AAAI conference on artificial intelligence*, 2020, pp. 4642–4649.
- [31] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-iid data," *arXiv preprint arXiv:1907.02189*, 2019.
- [32] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in *ICC 2019-2019 IEEE international conference on communications (ICC)*. IEEE, 2019, pp. 1–7.
- [33] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang, "A survey on federated learning: challenges and applications," *International Journal of Machine Learning and Cybernetics*, vol. 14, no. 2, pp. 513–535, 2023.
- [34] S. K. Lo, Q. Lu, H. Paik, and L. Zhu, "FLRA: A reference architecture for federated learning systems," *CoRR*, vol. abs/2106.11570, 2021.
- [35] N. Truong, K. Sun, S. Wang, F. Guitten, and Y. Guo, "Privacy preservation in federated learning: An insightful survey from the gdpr perspective," *Computers & Security*, vol. 110, p. 102402, 2021.
- [36] H. Zimmermann, "Osi reference model-the iso model of architecture for open systems interconnection," *IEEE Transactions on communications*, vol. 28, no. 4, pp. 425–432, 1980.
- [37] T. Jayalakshmi and A. Santhakumaran, "Statistical normalization and back propagation for classification," *International Journal of Computer Theory and Engineering*, vol. 3, no. 1, pp. 1793–8201, 2011.
- [38] D. Chicco, "Ten quick tips for machine learning in computational biology," *BioData mining*, vol. 10, no. 1, pp. 1–17, 2017.
- [39] J. Xu and H. Wang, "Client selection and bandwidth allocation in wireless federated learning networks: A long-term perspective," *IEEE Transactions on Wireless Communications*, vol. 20, no. 2, pp. 1188–1200, 2020.
- [40] K. Zhang, X. Song, C. Zhang, and S. Yu, "Challenges and future directions of secure federated learning: a survey," *Frontiers of Computer Science*, vol. 16, 2021.
- [41] K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, S. Jin, T. Q. Quek, and H. V. Poor, "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3454–3469, 2020.
- [42] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [43] C. Gentry, "Computing arbitrary functions of encrypted data," *Communications of the ACM*, vol. 53, no. 3, pp. 97–105, 2010.
- [44] X. Zhang, S. Ji, H. Wang, and T. Wang, "Private, yet practical, multi-party deep learning," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. Los Alamitos, CA, USA: IEEE Computer Society, jun 2017, pp. 1442–1452.
- [45] C. Briggs, Z. Fan, and P. Andras, "A review of privacy-preserving federated learning for the internet-of-things," *Federated Learning Systems*, pp. 21–50, 2021.
- [46] T. DeMarco, "Structure analysis and system specification," in *Pioneers and Their Contributions to Software Engineering: sd&m Conference on Software Pioneers, Bonn, June 28/29, 2001, Original Historic Contributions*. Springer, 1979, pp. 255–288.
- [47] A. Shostack, "Experiences threat modeling at microsoft." *MODSEC@MoDELS*, vol. 2008, p. 35, 2008.
- [48] A. Vassilev, A. Oprea, A. Fordyce, and H. Andersen, "Adversarial machine learning: A taxonomy and terminology of attacks and mitigations," *NIST Trustworthy and Responsible AI*, 2024.
- [49] P. Liu, X. Xu, and W. Wang, "Threats, attacks and defenses to federated learning: issues, taxonomy and perspectives," *Cybersecurity*, vol. 5, no. 1, pp. 1–19, 2022.
- [50] C. Guo, J. Gardner, Y. You, A. G. Wilson, and K. Weinberger, "Simple black-box adversarial attacks," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2484–2493.
- [51] S. N. Shukla, A. K. Sahu, D. Willmott, and Z. Kolter, "Simple and efficient hard label black-box adversarial attacks in low query budget regimes," in *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2021, pp. 1461–1469.
- [52] O. Suciu, R. Marginean, Y. Kaya, H. Daume III, and T. Dumitras, "When does machine learning {FAIL}? generalized transferability for evasion and poisoning attacks," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 1299–1316.
- [53] L. Lyu, H. Yu, and Q. Yang, "Threats to federated learning: A survey," *arXiv preprint arXiv:2003.02133*, 2020.
- [54] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinstein, and J. D. Tygar, "Adversarial machine learning," in *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, 2011, pp. 43–58.
- [55] P. M. Mammen, "Federated learning: Opportunities and challenges," *arXiv preprint arXiv:2101.05428*, 2021.
- [56] G. Sun, Y. Cong, J. Dong, Q. Wang, and J. Liu, "Data poisoning attacks on federated machine learning," *CoRR*, vol. abs/2004.10020, 2020.
- [57] V. Tolpegin, S. Truex, M. E. Gursoy, and L. Liu, "Data poisoning attacks against federated learning systems," *CoRR*, vol. abs/2007.08432, 2020.
- [58] A. Shafahi, W. R. Huang, M. Najibi, O. Suciu, C. Studer, T. Dumitras, and T. Goldstein, "Poison frogs! targeted clean-label poisoning attacks on neural networks," *Advances in neural information processing systems*, vol. 31, 2018.
- [59] T. Gu, B. Dolan-Gavitt, and S. Garg, "Badnets: Identifying vulnerabilities in the machine learning model supply chain," *arXiv preprint arXiv:1708.06733*, 2017.
- [60] C. Fung, C. J. Yoon, and I. Beschastnikh, "Mitigating sybils in federated learning poisoning," *arXiv preprint arXiv:1808.04866*, 2018.
- [61] A. N. Bhagoji, S. Chakraborty, P. Mittal, and S. Calo, "Analyzing federated learning through an adversarial lens," in *International Conference on Machine Learning*. PMLR, 2019, pp. 634–643.
- [62] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2938–2948.
- [63] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.
- [64] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [65] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [66] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the gan: information leakage from collaborative deep learning," in *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*, 2017, pp. 603–618.
- [67] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 1322–1333.
- [68] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 3–18.
- [69] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 691–706.
- [70] M. Nasr, R. Shokri, and A. Houmansadr, "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)*. IEEE, 2019, pp. 739–753.
- [71] C. Song and V. Shmatikov, "Overlearning reveals sensitive attributes," *arXiv preprint arXiv:1905.11742*, 2019.
- [72] S. V. Dibbo, "Sok: Model inversion attack landscape: Taxonomy, challenges, and future roadmap," in *2023 IEEE 36th Computer Security Foundations Symposium (CSF)*. IEEE, 2023, pp. 439–456.

- [73] B. Zhao, K. R. Mopuri, and H. Bilen, "idlg: Improved deep leakage from gradients," *arXiv preprint arXiv:2001.02610*, 2020.
- [74] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart, "Privacy in pharmacogenetics: An {End-to-End} case study of personalized warfarin dosing," in *23rd USENIX security symposium (USENIX Security 14)*, 2014, pp. 17–32.
- [75] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, "Privacy risk in machine learning: Analyzing the connection to overfitting," in *2018 IEEE 31st computer security foundations symposium (CSF)*. IEEE, 2018, pp. 268–282.
- [76] X. Gong, Q. Wang, Y. Chen, W. Yang, and X. Jiang, "Model extraction attacks and defenses on cloud-based machine learning models," *IEEE Communications Magazine*, vol. 58, no. 12, pp. 83–89, 2020.
- [77] B. Wang and N. Z. Gong, "Stealing hyperparameters in machine learning," in *2018 IEEE symposium on security and privacy (SP)*. IEEE, 2018, pp. 36–52.
- [78] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.
- [79] J. Li, A. S. Rakin, X. Chen, L. Yang, Z. He, D. Fan, and C. Chakrabarti, "Model extraction attacks on split federated learning," *arXiv preprint arXiv:2303.08581*, 2023.
- [80] B. Jayaraman and D. Evans, "Evaluating differentially private machine learning in practice," in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 1895–1912.
- [81] M. Veale, R. Binns, and L. Edwards, "Algorithms that remember: model inversion attacks and data protection law," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 376, no. 2133, p. 20180083, 2018.
- [82] M. Abadi, U. Erlingsson, I. Goodfellow, H. B. McMahan, I. Mironov, N. Papernot, K. Talwar, and L. Zhang, "On the protection of private information in machine learning systems: Two recent approaches," in *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*. IEEE, 2017, pp. 1–6.
- [83] Z. Sun, P. Kairouz, A. T. Suresh, and H. B. McMahan, "Can you really backdoor federated learning?" *arXiv preprint arXiv:1911.07963*, 2019.
- [84] Y. Fraboni, R. Vidal, and M. Lorenzi, "Free-rider attacks on model aggregation in federated learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 1846–1854.
- [85] J. Lin, M. Du, and J. Liu, "Free-riders in federated learning: Attacks and defenses," *arXiv preprint arXiv:1911.12560*, 2019.
- [86] Z. Zhu, J. Shu, X. Zou, and X. Jia, "Advanced free-rider attacks in federated learning," in *the 1st NeurIPS Workshop on New Frontiers in Federated Learning Privacy, Fairness, Robustness, Personalization and Data Ownership*, 2021.
- [87] J. Wang, X. Chang, R. J. Rodríguez, and Y. Wang, "Assessing anonymous and selfish free-rider attacks in federated learning," in *2022 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 2022, pp. 1–6.
- [88] B. Wang, H. Li, X. Liu, and Y. Guo, "Frad: Free-rider attacks detection mechanism for federated learning in aiot," *IEEE Internet of Things Journal*, 2023.
- [89] J. Wang, "Pass: Parameters audit-based secure and fair federated learning scheme against free rider," *arXiv preprint arXiv:2207.07292*, 2022.
- [90] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrnđić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*. Springer, 2013, pp. 387–402.
- [91] M. Kearns and M. Li, "Learning in the presence of malicious errors," in *Proceedings of the twentieth annual ACM symposium on Theory of computing*, 1988, pp. 267–280.
- [92] X. Gong, Y. Chen, Q. Wang, M. Wang, and S. Li, "Private data inference attacks against cloud: Model, technologies, and research directions," *IEEE Communications Magazine*, vol. 60, no. 9, pp. 46–52, 2022.
- [93] B. Jayaraman and D. Evans, "Are attribute inference attacks just imputation?" in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1569–1582.
- [94] Y. Zhang, R. Jia, H. Pei, W. Wang, B. Li, and D. Song, "The secret revealer: Generative model-inversion attacks against deep neural networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 253–261.
- [95] S. Lu, W. Oh, R. Ooka, and L. Wang, "Effects of environmental features in small public urban green spaces on older adults' mental restoration: Evidence from tokyo," *International Journal of Environmental Research and Public Health*, vol. 19, no. 9, p. 5477, Apr 2022. [Online]. Available: <http://dx.doi.org/10.3390/ijerph19095477>
- [96] J. Wu, Q. Liu, Z. Huang, Y. Ning, H. Wang, E. Chen, J. Yi, and B. Zhou, "Hierarchical personalized federated learning for user modeling," *Proceedings of the Web Conference 2021*, 2021.
- [97] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [98] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," *arXiv preprint arXiv:1812.00984*, 2018.
- [99] Y. Aono, T. Hayashi, L. Wang, S. Moriai et al., "Privacy-preserving deep learning via additively homomorphic encryption," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 5, pp. 1333–1345, 2017.
- [100] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang, "Learning differentially private language models without losing accuracy," *CoRR*, vol. abs/1710.06963, 2017.
- [101] N. Agarwal, A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan, "cpsgd: Communication-efficient and differentially-private distributed sgd," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [102] M. Fang, X. Cao, J. Jia, and N. Gong, "Local model poisoning attacks to {Byzantine-Robust} federated learning," in *29th USENIX Security Symposium (USENIX Security 20)*, 2020, pp. 1605–1622.
- [103] N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman, "Sok: Security and privacy in machine learning," in *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2018, pp. 399–414.
- [104] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, "Memguard: Defending against black-box membership inference attacks via adversarial examples," in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, 2019, pp. 259–274.
- [105] J. Li, A. S. Rakin, X. Chen, Z. He, D. Fan, and C. Chakrabarti, "Ressfl: A resistance transfer framework for defending model inversion attack in split federated learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 10 194–10 202.
- [106] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.
- [107] H. Kim, J. Park, M. Bennis, and S.-L. Kim, "On-device federated learning via blockchain and its latency analysis," *arXiv preprint arXiv:1808.03949*, 2018.
- [108] J. Weng, J. Weng, J. Zhang, M. Li, Y. Zhang, and W. Luo, "Deepchain: Auditable and privacy-preserving deep learning with blockchain-based incentive," *IEEE Transactions on Dependable and Secure Computing*, vol. 18, no. 5, pp. 2438–2455, 2019.
- [109] H. Hong and Y. Hong, "Certifiable black-box attack: Ensuring provably successful attack for adversarial examples," *arXiv preprint arXiv:2304.04343*, 2023.
- [110] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, "Mitigating adversarial effects through randomization," *arXiv preprint arXiv:1711.01991*, 2017.
- [111] G. S. Dhillon, K. Azizzadenesheli, Z. C. Lipton, J. Bernstein, J. Kossaifi, A. Khanna, and A. Anandkumar, "Stochastic activation pruning for robust adversarial defense," *arXiv preprint arXiv:1803.01442*, 2018.
- [112] J. Chen, M. I. Jordan, and M. J. Wainwright, "Hopskipjumpattack: A query-efficient decision-based attack," in *2020 IEEE symposium on security and privacy (sp)*. IEEE, 2020, pp. 1277–1294.
- [113] F. Liao, M. Liang, Y. Dong, T. Pang, X. Hu, and J. Zhu, "Defense against adversarial attacks using high-level representation guided denoiser," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1778–1787.
- [114] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *arXiv preprint arXiv:1706.06083*, 2017.
- [115] F. Tramer and D. Boneh, "Adversarial training and robustness for multiple perturbations," *Advances in neural information processing systems*, vol. 32, 2019.
- [116] K. Roth, Y. Kilcher, and T. Hofmann, "The odds are odd: A statistical test for detecting adversarial examples," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5498–5507.
- [117] S. Jain, A.-M. Cretu, and Y.-A. de Montjoye, "Adversarial detection avoidance attacks: Evaluating the robustness of perceptual hashing-based client-side scanning," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2317–2334.

- [118] M. Lecuyer, V. Atlidakis, R. Geambasu, D. Hsu, and S. Jana, "Certified robustness to adversarial examples with differential privacy," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 656–672.
- [119] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient smt solver for verifying deep neural networks," in *Computer Aided Verification: 29th International Conference, CAV 2017, Heidelberg, Germany, July 24–28, 2017, Proceedings, Part I 30*. Springer, 2017, pp. 97–117.
- [120] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys (Csur)*, vol. 51, no. 4, pp. 1–35, 2018.
- [121] C. Gentry, *A fully homomorphic encryption scheme*. Stanford university, 2009.
- [122] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "{BatchCrypt}: Efficient homomorphic encryption for {Cross-Silo} federated learning," in *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, 2020, pp. 493–506.
- [123] Q. Li, Z. Wu, Z. Wen, and B. He, "Privacy-preserving gradient boosting decision trees," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 784–791.
- [124] C. Dwork, "Differential privacy: A survey of results," in *International conference on theory and applications of models of computation*. Springer, 2008, pp. 1–19.
- [125] C. Dwork, A. Roth et al., "The algorithmic foundations of differential privacy." *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3-4, pp. 211–407, 2014.
- [126] L. Song, H. Wu, W. Ruan, and W. Han, "Sok: Training machine learning models over multiple sources with privacy preservation," *arXiv preprint arXiv:2012.03386*, 2020.
- [127] R. Bassily, K. Nissim, U. Stemmer, and A. Guha Thakurta, "Practical locally private heavy hitters," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [128] V. Pihur, A. Korolova, F. Liu, S. Sankuratripati, M. Yung, D. Huang, and R. Zeng, "Differentially-private draw and discard" machine learning," *arXiv preprint arXiv:1807.04369*, 2018.
- [129] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1175–1191.
- [130] F. Mo, H. Haddadi, K. Katevas, E. Marin, D. Perino, and N. Kourtellis, "Ppfl: privacy-preserving federated learning with trusted execution environments," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, pp. 94–108.
- [131] F. Mo and H. Haddadi, "Efficient and private federated learning using tee," Dresden, Germany, Mar 2019, poster presentation. [Online]. Available: <https://eurosys2019.org>
- [132] V. Costan and S. Devadas, "Intel sgx explained," *Cryptology ePrint Archive*, 2016.
- [133] OP-TEE Core Team, "Open portable trusted execution environment," 2022. [Online]. Available: <https://www.op-tee.org>
- [134] G. Kaissis, A. Ziller, J. Passerat-Palmbach, T. Ryffel, D. Usynin, A. Trask, I. Lima, J. Mancuso, F. Jungmann, M.-M. Steinborn et al., "End-to-end privacy preserving deep learning on multi-institutional medical imaging," *Nature Machine Intelligence*, vol. 3, no. 6, pp. 473–484, 2021.
- [135] R. Xu, N. Baracaldo, Y. Zhou, A. Anwar, and H. Ludwig, "Hybrid-pha: An efficient approach for privacy-preserving federated learning," in *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, 2019, pp. 13–23.
- [136] S. Goryczka and L. Xiong, "A comprehensive comparison of multiparty secure additions with differential privacy," *IEEE transactions on dependable and secure computing*, vol. 14, no. 5, pp. 463–477, 2015.
- [137] M. Alaggan, S. Gams, and A. Kermarrec, "Heterogeneous differential privacy," *CoRR*, vol. abs/1504.06998, 2015.
- [138] A. Richardson, A. Filos-Ratsikas, and B. Faltings, "Rewarding high-quality data via influence functions," *ArXiv*, vol. abs/1908.11598, 2019.
- [139] N. Guha, A. Talwalkar, and V. Smith, "One-shot federated learning," *arXiv preprint arXiv:1902.11175*, 2019.



JIAXING LU received his MS degree at the Graduate School of Interdisciplinary Information Studies (GSII), the University of Tokyo. Currently he is a Ph.D student. His research interests include federated learning and edge computing, right now he is focusing on the application of federated learning on network traffic optimization.



NORIIHIRO FUKUMOTO received the B.E., M.E., and D.E. degrees from Waseda University, Tokyo in 1999, 2001, and 2015 respectively. He joined KDDI R&D Laboratories, Inc. in 2001 and has been engaged in research and development on speech application services, voice packetization systems over IP networks, QoS/QoE Management, and 5G/B5G network. He is currently a Visiting Associate Professor at the University of Tokyo.



AKIHIRO NAKAO (M' 14) received B.S. (1991) in Physics, M.E. (1994) in Information Engineering from the University of Tokyo. He was at IBM Yamato Laboratory, Tokyo Research Laboratory, and IBM Texas Austin from 1994 till 2005. He received M.S. (2001) and Ph.D. (2005) in Computer Science from Princeton University. He taught as an associate professor (2005-2014) and as a professor (2014-2021) in Applied Computer Science, at Interfaculty Initiative in Information Studies, Graduate School of Interdisciplinary Information Studies, the University of Tokyo.

He has served as Vice Dean of the University of Tokyo's Interfaculty Initiative in Information Studies (2019-2021). In April 2021, he has moved to School of Engineering, the University of Tokyo (2021-present). Since April 2023, he has been serving as Head of Department of System Innovations, School of Engineering. He was appointed as an adviser to the President of the University of Tokyo (2019-2020) and has been a special adviser to the President of the University of Tokyo (2020-present). He is serving as Director, Collaborative Research Institute for NGCI, (Next-Generation Cyber Infrastructure), the University of Tokyo (2021-present).

For social services, he has been playing several important roles in Japanese government and also at research societies. He has also been appointed Chairman of the 5G Mobile Network Promotion Forum (5GMF) Network Architecture Committee by Japanese government. He has been appointed as Chairman of 5G/Beyond 5G committee, Space ICT Promotion Initiative Forum, International Committee, and Beyond 5G Promotion Consortium as well (2020-present). From 2020 to present, he is a chair and advisor of IEICE technical committee on network systems (NS) as well as a chair of IEICE technical committee on cross-field research association of super-intelligent networking (RISING). He will be the president of Communication Society, IEICE, in 2024.

AUTHORS

...