A Survey on Secure and Private Federated Learning Using Blockchain: Theory and Application in Resource-constrained Computing

Ervin Moore, Ahmed Imteaj, Shabnam Rezapour, and M. Hadi Amini, *Senior Member, IEEE*

Abstract—**Federated Learning (FL) has gained widespread popularity in recent years due to the fast booming of advanced machine learning and artificial intelligence along with emerging security and privacy threats. FL enables efficient model generation from local data storage of the edge devices without revealing the sensitive data to any entities. While this paradigm partly mitigates the privacy issues of users' sensitive data, the performance of the FL process can be threatened and reached a bottleneck due to the growing cyber threats and privacy violation techniques. To expedite the proliferation of FL process, the integration of blockchain for FL environments has drawn prolific attention from the people of academia and industry. Blockchain has the potential to prevent security and privacy threats with its decentralization, immutability, consensus, and transparency characteristic. However, if the blockchain mechanism requires costly computational resources, then the resource-constrained FL clients cannot be involved in the training. Considering that, this survey focuses on reviewing the challenges, solutions, and future directions for the successful deployment of blockchain in resource-constrained FL environments. We comprehensively review variant blockchain mechanisms that are suitable for FL process and discuss their trade-offs for a limited resource budget. Further, we extensively analyze the cyber threats that could be observed in a resource-constrained FL environment, and how blockchain can play a key role to block those cyber attacks. To this end, we highlight some potential solutions towards the coupling of blockchain and federated learning that can offer high levels of reliability, data privacy, and distributed computing performance.**

Index Terms—**Federated Learning, Blockchain, Security, Privacy, Resource Limitations, Optimization.**

I. INTRODUCTION

A. Motivation, Comparison, and Contributions

Data-driven technologies can be limited by factors such as limited computing resources or the need for large amounts of quality data. Earlier online methodologies transferred raw data which created potential data privacy risks. Data silos are formed in organizations that are inaccessible to other departments and often incompatible with datasets within

Corresponding Author: M. Hadi Amini, Florida International University, Miami, FL 33199, moamini@fiu.edu

Ervin Moore and M. Hadi Amini are with Knight Foundation School of Computing and Information Sciences, Florida International University, Miami, FL 33199, USA. They are also with the Sustainability, Optimization, and Learning for InterDependent networks laboratory (solid lab) at FIU.

Ahmed Imteaj is with Sustainability, Privacy and Edge intElligence for Distributed networks (SPEED) Laboratory, Southern Illinois University, Carbondale, USA.

Shabnam Rezapour is with Smart Decision-Making for Network-Centric Systems, Florida International University, Florida, USA.

the same organization. FL is introduced as a solution for privacy preservation due to FL applying differential privacy to communicated data. In FL machine learning parameters are exchanged instead of raw data. Centralized computing architectures can be directly targeted and denied because of the single-point-of-failure vulnerability. Blockchains being a decentralized approach can improve system resilience from adversaries. A combination of FL and blockchain have increased robustness in comparison to earlier methodologies that had diminishing performance, due to features such as lack of privacy preservation. This survey examines blockchain-based FL as a solution for privacy preservation and secure resource management. Recent blockchain and FL surveys did not comprehensively analyze combining the two to provide a secure learning setting for a resource-constrained environment.

B. Organization of the Survey

The rest of this paper is organized as follows. In Section **II**, we present an overview and taxonomy of FL with a comprehensive list of existing studies. In Section **III**, we review distributed optimization and ML approaches. Section **IV** presents a detailed analysis of the major challenges of FL while applying on resource-constrained devices, which is followed by Section **V**, where we discuss the potential solutions of those emerging challenges. In Section **VI**, we present the existing FL applications, and in Section **VII**, we highlight the future research direction in the FL-based IoT domain. Finally, Section **VIII** concludes the paper.

II. BACKGROUND STUDY

Online computing environments can put private information at risk when communicating online. Ensuring data privacy throughout communications is important and can be approached by FL. FL transmits user data through machine learning updates instead of raw data. Classical data-sharing techniques can be improved by FL and its privacy protection mechanisms. Blockchain, a decentralized approach, offers additional security and resources for applications such as FL. BCFL is an appropriate solution for resource-constrained computing environments due to: peer-to-peer computing, network participation incentives, validation protocols, and other resource-preserving mechanisms. The following section details mechanisms within FL and blockchain.

A. Brief Introduction to Federated Learning

FL proposed by Google in 2016, protects users' privacy by allowing devices to train a machine learning model using local data collaboratively. FL can process model updates synchronously or asynchronously. FL is a collaborative machine-learning architecture that stores local data on a device, this local data reduces the need for data to be stored online in a cloud. Sensitive user data is protected by adding noise to personally identifiable information. The goal is for data to be unidentifiable while preserving good qualities for machine learning model optimization. Finding the balance of data retention improves performance. FL, as a manner of distributed machine learning, can significantly preserve clients' private data from being exposed to external adversaries [1].

Data structures encourage distinct FL implementations for optimal data management. Each FL structure looks at feature spaces and datasets differently. Three different types of FL implementations include Horizontal FL, Vertical FL, and Federated Transfer Learning.

- Horizontal (sample-based) FL - Data distribution contains a consistent set of features and different samples. For example, a database containing predefined rows of relevant user data has different entries for each unique user. Some examples of Horizontal FL are predicting smartphone user behavior, personalized recommendations, identifying risk factors or predicting patient diseases, optimizing manufacturing product outcomes, and recognizing fraudulent transactions [2].
- Vertical (feature-based) FL - Data distribution contains recurring samples with different features. For instance, consider a bank and a superstore in the same area. Most of their customers may be the same, but their business structure, i.e., the feature space, is different, and thus the user-space intersection is quite large [3]. Some examples of vertical FL are predicting the likelihood of patient admission considering different types of health records, lab results in hospitals, customer shopping behavior, and likelihood to make purchases analyzing heterogeneous data types of multiple retail stores [4].
- Federated Transfer Learning (FTL) - Combination of horizontal and vertical FL that contains different features and samples. The knowledge of an existing machine learning model is transferred to another model for improved performance. Transfer learning reuses learned lessons and re-purposes information towards a related problem. FTL currently has applications in wearable healthcare [5], autonomous driving, classification of EEG signals [6], industrial fault diagnostics [7], and image steganalysis [8], [9].

Traditional FL has a single point of failure limitation. Figure 1 showcases an uninterrupted online learning environment, although a denial of service could halt the repeated process. Internet of Things (IoT) devices differ in performance and reliability. The FL training process could be computationally expensive depending on available resources, which may require devices to drop out from the learning process. Devices may need the motivation to participate in FL training, which can be encouraged by incentives. Many of the pitfalls of FL can be augmented by blockchain mechanisms.
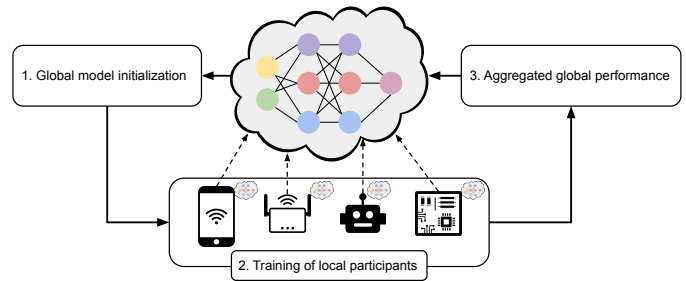

Fig. 1. Introduction to FL cycle.

B. Brief Introduction to Blockchain

Blockchain technology connects data blocks into a digital ledger similar to a database. The blockchain ledger is distributed throughout the network and considered decentralized; thus, third-party intermediaries are not required for blockchain processes. Each block in a blockchain contains relevant transaction history, such as unique identifiers, block computing costs, and other network-related information. Network transactions are directly recorded into the blockchain once authenticated. Peer-to-peer transactions allow workloads to share resources and information. Peers can participate as suppliers and consumers of resources. The authors in [10] proposed that the nodes or agents involved within a blockchain network are called participants and miners. The participants are the agents who perform any transaction, and the miners are responsible for validating or rejecting a block [10]. In a blockchain, network processing happens simultaneously, creating a large optimization surface area. Resources such as processing power, storage space, and available network bandwidth all have to be considered for resource-constrained computing environments.

Many processes operate at the same time within a blockchain. When blocks are being generated at a high rate, performance can be lower. Merge issues can appear when two or more blocks with the same hash information are successfully mined simultaneously. Merge issues can create forks in the blockchain, causing alternative chains to emerge until solved by consensus. Having a proper consensus protocol can discourage multiple versions of a blockchain. A rule of thumb is to reference the largest chain of blocks as the main blockchain to deter conflicts. Protecting the integrity of blockchain allows historical data to sync efficiently.

Blockchain architectures contain several layers:

- The Hardware layer is necessary for hosting a blockchain architecture. The core infrastructure requires computers, graphics processing units (for miners), and miscellaneous data storage. Computing resources can be emulated through virtual machines as a process.
- The Network layer is responsible for peer-to-peer communication. Block transaction data is publicly visible in a public ledger. As a result, public blockchain architectures
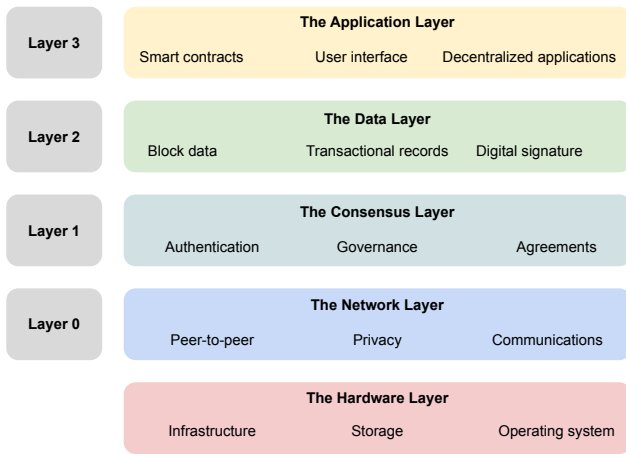
Fig. 2. Blockchain architecture layers.

can share processes with unknown participants that can directly interact with one another.

- The Consensus layer authenticates transactions based on agreed-upon protocols. This layer is responsible for validating transactions. If a fork in the blockchain emerges, the consensus layer handles disputes using protocol logic.
- The Data layer includes blocks and block transactional records. Blocks specifically contain information about: the previous block hash, the timestamp of the block creation, the generated block hash, and the transaction nonce, which is used for verifying transactions.
- The Application layer contains the front end of the blockchain and is considered the visible layer for the user. This layer contains an application programming interface (API), displays smart contracts, and includes basic configuration for quickly exchanging information with other blockchains. Blockchain interoperability allows similar blockchain structures to communicate efficiently.

1) Blockchain Fundamentals:
Blockchain architectures use cryptographic hash functions for transactions. Miners solve cryptographic puzzles to generate blocks and may form mining pools to boost performance. Nonce (numbers only used once) are numbers that miners must calculate before solving a cryptographic block. Cryptographic hash functions are encrypted, requiring efficient computing techniques to reasonably solve. The consensus layer authenticates that block conditions are met before integrating blocks into the blockchain. Blockchain agreements are transacted through smart contracts, rule-based digital agreements between two parties (i.e., miners and participants) that automatically execute when conditions are met. Once a block is authenticated, a new block is sought for mining. The authors in [10] proposed that typically in a blockchain environment, constant creation of new blocks, even with no transaction, is crucial for maintaining security as it prevents malevolent users from creating longer, tampered blockchain [10].

2) Blockchain Categories:
Generally, there are three types of blockchains: public, private, and consortium. Each type of blockchain involves different permissions and participation specifics. Different types of consensus and authority mechanisms are included in each blockchain type.

- Public blockchain - The most well-known blockchain is the public blockchain which is permissionless and open to all participants. Public blockchain ledgers are publicly available. Each participant has equal permissions in this peer-to-peer networking environment. Open participation in blockchains has been known to affect the number of resources in many ways. Cases, where open participation leads to a large number of dishonest participants can cause blockchain performance to diminish. Examples of popular public blockchains include cryptocurrency platforms such as Etheruem or Bitcoin.
- Private blockchain - Private blockchains are commonly referred to as permissioned blockchains. Only authorized participants are allowed to join and view private blockchain ledgers. The central entity of the private blockchain has permission to alter and manage protocols. The authors in [10] mention that the central entity has the power of validation and can change any rule of the blockchain (e.g., block consensus) [10].
- Consortium blockchain - Consortium blockchains include public and private blockchain qualities. Consortium blockchains are partially decentralized, permissioned and transparent. Selected authorities have increased permissions compared to the public. Increased permissions for consortium authorities include the ability to participate in block validation. The authors in [11] utilized a consortium blockchain for pre-selecting a limited number of trusted miners to maintain the distributed ledger of an intelligent transportation system.

TABLE I
SUMMARY OF INCLUDED PAPERS REGARDING FL, BLOCKCHAIN AND BLOCKCHAINED FL (BCFL).

Ref No.	Area	Year	Contribution	Resource-constrained considerate	Participation incentivized
[12]	FL	2022	Adaptive FL survey for minimizing resource consumption within communication constrained environments	Yes	No
[13]	FL	2021	Comprehensive FL survey focusing on security and privacy in FL systems, the authors examine FL limitations and countermeasures	No	No
[14]	BCFL	2022	Blockchain empowered FL model for blade icing estimation	No	Yes
[15]	BCFL	2022	Article that proposes a Blockchain-based communication-efficient federated learning framework that compresses communications	Yes	No
[16]	BCFL	2021	Blockchain-empowered decentralized horizontal FL framework that improves 5G-Enabled UAV privacy	Yes	No
[10]	Blockchain	2021	Comprehensive book on fundamental topics related to Blockchain	No	Yes
[3]	FL	2021	Comprehensive survey on FL challenges that occur when applied to resource-constrained Internet of Things devices	Yes	Yes
[17]	BCFL	2022	Comprehensive Blockchained FL survey regarding approaches for securing Internet of Things devices	Yes	Yes
[18]	Blockchain	2021	Comprehensive survey on public Blockchain for Internet of Things devices with bounded computing capabilities	Yes	Yes
[19]	BCFL	2019	Blockchained FL paper that examines end-to-end latency	No	No
[20]	BCFL	2022	In-depth survey on Blockchained FL as a framework	No	Yes
[21]	BCFL	2021	Blockchained FL paper that explores opportunities and challenges in multi-access edge computing	Yes	Yes
[22]	FL	2022	Survey on FL applications within smart healthcare	Yes	Yes
[23]	FL	2019	Hybrid FL system that protects against inference threats and produces models with high accuracy.	No	No
[24]	FL	2019	Adaptive FL paper focused on gradient-descent based FL for reducing computing resource budgets	Yes	No
[4]	FL	2019	Comprehensive survey on FL concepts and applications	No	Yes
[25]	BCFL	2020	Privacy-preserving FL paper that leverages customers data through IoT devices to assist home appliance manufacturers	Yes	Yes

Consortium or permissioned blockchains are suitable when the integrity of participants is in question. Dishonest participants may sabotage others to compete for incentives or resources. Table I summarizes various papers that consider par-

ticipation incentives and resource constraints. The following section examines potential attacks, vulnerabilities, and threats found in BCFL environments.

C. Threat Models and Integration Motivation

FL includes privacy protection mechanisms that deter unwanted access to data. Meticulous attack attempts can breach FL security in different ways. Attacks can occur during device training, parameter upload and download, central aggregation, and post-aggregation phases. A range of various attacks can exploit FL vulnerabilities during each stage:
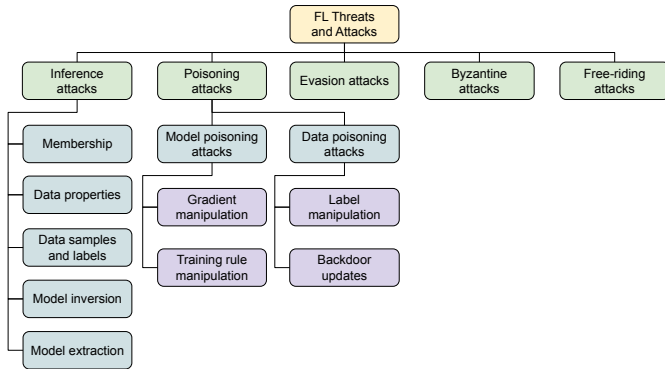

Fig. 3. Classifying FL threats and attacks.

1) Inference attacks:

During inference attacks, an adversarial agent attempts to capture sensitive information throughout training. Training data, participant data, and label data are recorded without permission in hopes of correlating encrypted information with real values. Generative Adversarial Networks (GANs) can create powerful inference attacks. Inference attacks fall into five general categories:

  (i) Membership inference attack - Attackers determine if an individual was present in the training dataset. The authors in [26] constructed a membership inference attack that exploits the observation that machine learning models often behave differently on the data they were trained on versus the data they see for the first time [26]. Adversaries attempt to infer from specific samples that can be captured from model outputs. The objective is for adversaries to determine if a specific record exists in the model training dataset.

 (ii) Data properties inference attack - Data properties of a machine learning model are inferred using the parameters shared during model training. The particular data property is fully investigated to learn the frequency of the property within training data. For example, the authors in [27] demonstrated that a classifier that recognizes smiling faces also leaks information about the relative attractiveness of the individuals in its training set [27]. For example, a data properties inference attack on a model trained to screen FL participants by threshold may mistakenly reveal hardware or host types within the FL architecture.

(iii) Data samples and labels inference attack - An adversary uses inference to capture targeted data samples and labels. FL model classes and participant training input labels can be reconstructed for inference. A malicious participant calculates distributed gradients to infer private information about other participants or model parameters. The authors in [28] demonstrated a label-only inference attack that could capture private information from machine learning models without access to confidence scores.

(iv) Model inversion attack - Model inversion attacks aim to reconstruct private information from training data. Attackers with access to model parameters can look at the model's confidence score of predicted classes for inference. Successful model inversion attacks infer realistic representations from training data. The authors in [29] demonstrated the applicability of model inversion attacks on decision trees for lifestyle surveys as used on machine-learning-as-a-service systems and neural networks for facial recognition [29].

 (v) Model extraction attacks - Adversary attack that aims to steal exact model parameters. Hu and Pang [30], mention how model extraction attacks aim to duplicate a machine learning model through query access to a target model. Hu and Pang studied two attacks on GANs: fidelity extraction attacks and accuracy extraction attacks [30]. A strategic position can be extracted from learned functionality. Transfer learning attacks can occur when attacks are trained on vulnerabilities of similar models with the same framework, then transferred to attack the target model. Extracting knowledge of the target models' training data and functionality increases future attack effectiveness.

2) Poisoning attacks:

Poisoning attacks include misleading data injected by adversaries as inputs. Poisoning attacks can be classified into two categories: (1) model poisoning and (2) data poisoning attacks.

  (i) Model poisoning attacks - In model poisoning attacks, the attacker reduces the model's performance on targeted sub-tasks (e.g., classifying planes as birds) by uploading "poisoned" updates [31]. Poisoning attacks change the model's weights and biases, leading to misclassifying data. The model's confidence decreases when interacting with poisoned data.

   • Gradient manipulation attack - Transmitted gradients and parameters are manipulated throughout the FL training process to reduce global model performance. As a result, global model accuracy is reduced by adversaries injecting malicious updates. In addition, misclassification can result from gradient manipulation, allowing adversaries to increase the success rate of their attacks.

   • Training rule manipulation attack - An adversary attempts to minimize the difference between correct and incorrect training updates during training. The strength of this attack increases when multiple participants become dishonest and poison their parameters. Additionally, adversaries may bribe or collude

with multiple participants to increase misclassification and lower detection mechanisms. The objective is for eventual harmful updates to go undetected based on the similarity of correct and incorrect training updates.

(ii) Data poisoning attacks - Training data is altered to impact the FL model negatively. During training, contaminated data is introduced to corrupt the central aggregator. Data poisoning attacks inject malicious data into the training dataset before the learning process starts [32]. Such an attack allows adversaries to disguise harmful data, and eventually, the large amounts of poisoned data increase training times.

- Label manipulation attack - Label manipulation attacks cause the model to mislabel training data. Label flipping includes training sample labels being flipped around to lower the judgment of the model. Classification errors can result from label flipping. Label flipping is an example of dirty-label attacks. In contrast, clean-label attacks are created when an adversarial makes poisoned training data indistinguishable from non-poisoned training data. Clean-label attacks are difficult but extremely powerful.

- Backdoor attack - A backdoor attack tricks the model into associating a backdoor pattern with a specific target label so that, whenever this pattern appears, the model predicts the target label, otherwise, behaves normally [33]. Backdoor attacks are also referred to as Trojan attacks. Adversaries inject clean or dirty backdoor updates into data samples in the training phase. When the global model performs aggregation, the model performance on the specific input is reduced depending on the number of undetected backdoors. When the backdoor is triggered, incorrect predictions occur on the targeted task while the main task performance appears untampered. Backdoor attacks compromise a subset of samples exchanged with adversaries during local training.

3) Evasion attacks:

Evasion attacks - An adversary may attempt to evade a deployed system at test time by carefully manipulating attack samples [34]. Evasion attacks can bypass detection during training time with carefully manipulated attack samples. The input looks unaltered to humans but deceiving to the machine learning model. Evasion attacks contaminate training data during aggregation with undetectable misleading data.

4) Byzantine attacks:

Byzantine attacks - In a Byzantine attack, a malicious device gets processed with honest devices. Byzantine attacks aim to harm consensus and decrease model performance. In FL, a malicious attacker may control multiple clients, known as Byzantine users [35]. Byzantine users can upload fake data due to unreliable communication channels, corrupted hardware, or malicious attacks. This leads to the global model being manipulated by attackers and cannot be converged [35]. Larger amounts of Byzantine users increase the effectiveness of Byzantine attacks. In the worst case, a 51% attack could occur when the majority of participants are dishonest.

5) Free-riding attacks:

Free-riding attacks - Within FL environments, the global model is distributed to all participants regardless of their contributions. "Free-riders" may emerge that do not contribute their fair share during the training process. Free-riders create free-rider attacks that lower model performance due to fake model updates being exchanged. A dishonest participant may become a free rider to withhold private information, reduce resource costs, or receive participation rewards while lacking participation requirements. The authors in [36] mention two types of free-riders. Plain free-riders, which do not update the local parameters during the iterative federated optimization, and disguised free-riders, which employ sophisticated disguising techniques relying on stochastic perturbations of the parameters [36].

TABLE II
SUMMARY OF FL ATTACKS AND RESOURCE COSTS.

Attack type	Ref.	Interop.	Privacy risks	Resource Consumption		Prevention	Impact
				Adversary	Server		
Inference	[26]	High	High	High	Medium	Medium	High
Poisoning	[31]	High	Medium	High	High	Medium	High
Evasion	[34]	High	Low	High	Low	Low	Medium
Byzantine	[37]	Medium	High	High	High	Medium	High
Free-riding	[36]	Low	Low	Low	Medium	High	Medium

D. Threats during FL phases

A range of various attacks can compromise FL security. FL attacks can occur during four distinct FL phases: the training phase, parameters exchange phase, parameters aggregation phase, and prediction phase. Each phase has a different attack surface, most susceptible to model poisoning and inference attacks. While FL performs continuous learning, each phase is repeated in multiple iterations. As a result, undetected attacks strengthen as resources are drained from the environment. Four distinct FL phases and potential threats:

(i) Training phase - During the training phase, training data is sent to participants for learning. Local training data is not validated during this phase, allowing adversarial attacks. Early attacks in the training phase occur when the participant's reputation is unavailable. Malicious participants can pollute the training phase by conducting poisoning attacks. Adversaries may target the model performance for misclassification or overall training times through misleading data. Increasing amounts of adversarial data negatively impact model effectiveness.

(ii) Parameters exchange phase - The parameter exchange phase consists of A) participants downloading the global model parameters or B) participants uploading the local model parameters. External adversaries may attempt to capture sensitive information during parameter exchange through inference attacks. Inference attacks during this phase can steal parameters (model extraction), combine data to reconstruct the dataset, and estimate the distribution of continuously exchanged data. In addition, eavesdropping attacks can occur when communications are intercepted and potentially compromised online.

(iii) Parameters aggregation phase - During the parameters aggregation phase, the central server performs a weighted

average of participant parameters. Misleading parameters during this phase affect system-wide performance. If the central model architecture is well understood, adversaries can target the central server directly. The central server coordinates the aggregation of online learning and may learn to become malicious due to negative influences such as poisoning attacks. Fig. 4 showcases the danger of adversarial updates being successfully processed into the parameter aggregation phase, thus lowering central server performance. Preserving central server effectiveness is important for the health of the FL architecture. Honest-but-curious central servers attempt to learn all possible information, overstepping boundaries through conducting inference or model poisoning attacks. Inefficient central servers incorrectly redistribute information to participants.

(iv) Prediction phase - The prediction phase occurs after the global model has been deployed to devices. Adversaries may continually obtain the deployed global model to perform a combination of evasion and inference attacks. Evasion attacks deceive end devices, causing predictions to be inaccurate. Inference attacks on deployed predictions attempt to extract private information such as participant data, model parameter data, and training data.
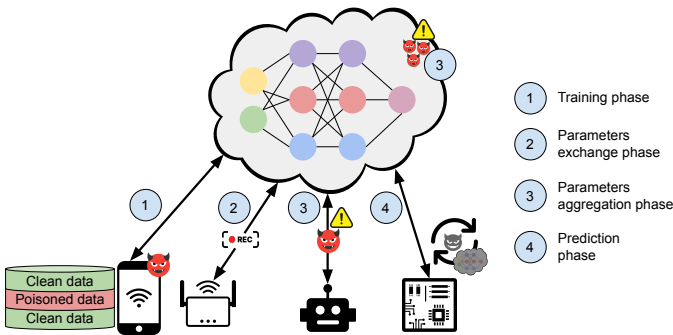

Fig. 4. Potential threat observations within the four phases of federated learning.

E. Frog-boiling attacks in online environments

Frog-boiling attacks reveal the limitations of anomaly detection in online environments. The authors in [38] proposed the frog-boiling attack, where an adversary disrupts the network while consistently operating within the threshold of rejection [38]. Frog-boiling refers to the phenomenon that a frog placed in hot water will instantly jump out. In contrast, a frog placed in slightly warm water that gradually heats up will remain in the water and eventually boil to death. Adversaries in online environments may continually attempt smaller attacks to learn the system's state and eventually disrupt the network. The authors in [38] mention three variants to the frog-boiling attack: the basic-targeted attack, the network-partition attack, and the closest-node attack [38]. Frog-boiling attacks can avoid outlier detection and increase network-wide latency.

1) Insider and outsider attacks:
FL environments can be threatened by the insider and outsider attacks. Insider attacks can be launched by either the FL server or the participants in the FL system [39]. Insider attacks affect global model aggregation performance and can infer real values from noisy data. On the inside, a dishonest participant may attempt to decode FL updates through malicious tactics. Honest participants on the inside may be impersonated or bribed by external adversaries. In comparison, outsider attacks do not have direct access to internal communications of the network. From the outside, adversaries may eavesdrop on communications between the client and the server. Transferring data wirelessly can be intercepted, leading to compromised FL updates.

2) Device latency:
IoT devices are heterogeneous and fluctuate in resource constraints. Each device may have different amounts of noise depending on the environment. An FL central server waiting for participants to complete the training phase can be delayed by slower devices. Devices may have unreliable memory bandwidth or unsatisfactory hardware. Delays at the device level can cause delays throughout a synchronous FL architecture. According to Issa et al. [17], the speed of rounds in synchronous FL is restricted to the speed of the slowest device. This causes a "straggler effect", which can cause inefficient processing [17].

3) Consensus and dishonesty:
Consensus performance is affected by the number of dishonest participants. The authors in [40] stated that the blockchain network security level is directly proportional to the amount of hash computing power that supports the blockchain. As the miners increase in the mining process, it becomes more difficult for an attacker to attack the blockchain [40]. The amount of honest and dishonest participants directly influences consensus mechanisms in FL and blockchain. Adversaries may employ Sybil attacks, where multiple fake identities are created to throw off consensus mechanisms. The weight of dishonest participants can cause consensus to malfunction. In the worst case, a 51% attack can control network consensus.

4) Additional noise against adversaries:
Additional noise creates a trade-off between levels of security and central server performance. The authors in [13] proposed that Federated learning intrinsically protects the data stored on each device by sharing model updates, e.g., gradient information, instead of the original data. However, model updates, which are based on original data, can reveal sensitive information [13]. Adding additional noise can increase the uncertainty of adversaries and reduce adversarial attack effectiveness. Jia et al. [41] proposed to add a carefully crafted noise vector to a confidence score vector to turn it into an adversarial example that misleads the attacker's classifier [41].

Table III contains details on how columns in the expanded table of FL attacks and costs were evaluated and scored. Interoperability was evaluated as the compatibility of attacks in an FL environment. In most cases, adversaries use high amounts of resources to conduct attacks. Some attacks explicitly drain the resources of the central server, while other attacks avoid server resource consumption to evade detection. The prevention column compares detection levels and proposed defenses for specified attacks. The impact column considers

TABLE III
EXPANDED TABLE OF FL ATTACKS AND RESOURCE COSTS [INTER. = INTEROPERABILITY, REF. = REFERENCE].

Attack type	Interop.	Privacy risks	Resource Consumption		Prevention	Impact
			Adversary	Server		
Inference	High	High	High	Medium	Medium	High
Membership	High	High	High	Low	Medium	High
Data properties	High	High	High	Low	Medium	High
Data samples and labels	High	High	High	Low	Medium	High
Model inversion	High	High	High	Low	Low	Medium
Model extraction	Medium	Low	High	Medium	Medium	High
Poisoning	High	Medium	High	High	Medium	High
Model poisoning	High	Low	High	High	Medium	High
Gradient manipulation	High	Low	High	Medium	Low	Medium
Training rule manipulation	High	Low	High	Medium	Low	Medium
Data poisoning	High	Low	High	High	Medium	High
Label manipulation	High	Medium	High	Medium	Medium	Medium
Backdoor	Medium	Medium	High	Medium	Medium	Medium
Evasion	High	Low	High	Low	Low	Medium
Byzantine	Medium	High	High	High	Medium	High
Free-riding	Low	Low	Low	Medium	High	Medium

Fig. 5. Various vulnerabilities found in FL and BCFL.

the importance of privacy preservation and architecture robustness. Well-planned combinations of attacks increase attack effectiveness.

F. Vulnerabilities of FL and BCFL

The attack surface of FL can be reduced by combining architectures with blockchains. FL has a large attack surface due to the single point of failure component of the FL system. An adversary could directly attack the central server to cause instability throughout the FL network. System-wide latency occurs when communications between devices and the central server are unreliable. Blockchain FL offers devices with distributed reliability for computing performance. When a device signals availability, the application layer of the blockchain automates a smart contract for a conditional agreement. Smart contracts on blockchains are immutable, so once created, the smart contract cannot be modified. Imperfect smart contracts include unwanted loopholes or vagueness.

Distributed machine learning architectures such as BCFL are vulnerable to outsider attacks while communicating over internet connections. Communication attacks such as eavesdropping attacks, man-in-the-middle attacks, and spoofing attacks are potential threats to the exchange of private information. A denial of service on network communications can be highly disruptive. Flooding a network with malicious data is resource-draining for distributed computing architectures.

Adversaries can commit poisoning attacks that reduce performance by sending deceptive inputs to mislead calculations. Dishonest participants redirect optimal resource allocations by providing inaccurate data. Byzantine users who upload fake data can purposely reduce global model performance. A 51% attack occurs when the number of dishonest participants outweighs the number of honest participants; thus, most dishonest participants can gain control of an unbalanced network.

The authors in [42] mention the cost of creating a 51% attack is surprisingly low if hash power is abundantly available [42]. In FL, a 51% attack can mislead FL training when the majority of participant data is untrustworthy. In blockchain, a 51% attack can control the consensus mechanisms that intervene with transaction authentication. Blockchain transaction authentication can be delayed when conflicting transactions are completed at the same time. A race attack creates a potential fork in the blockchain within the short window of
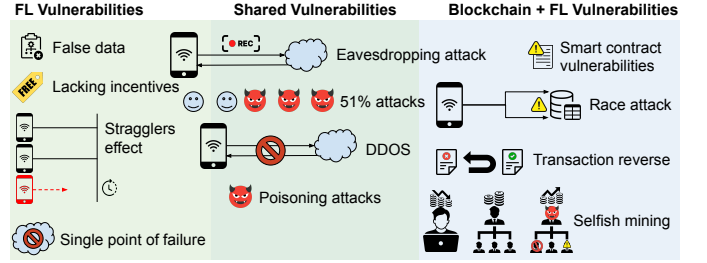
time required to authenticate transactions. Selfish mining pools may purposely generate or withhold blocks simultaneously to gain an advantage over other miners. Transaction reverse is a dishonest mining tactic where the processing of potentially generated blocks is reversed to cause Blockchain delays.

G. Integration Motivation of FL and Blockchain

Blockchains have inherent interoperability, allowing blockchains to communicate with other blockchains efficiently. Blockchain interoperability increases robustness when useful information is transferred between blockchains. For example, if a blockchain layer undergoes a denial of service attack, a similar blockchain could offer alternative services to supplement security. Similarly, participant data such as reputation can be communicated between blockchains to increase scalability. Blockchain can solve the problem of trust establishment among distributed systems through distributed node verification, and consensus mechanisms [43] [44].

Optimal device selection reduces the dangers of adversaries. A device selection phase can reduce resource costs associated with unfavorable behavior. Ideally, the behavioral patterns between an honest and dishonest participant are largely different. Behavioral auditing of malicious patterns from questionable participants can increase security. The authors in [45] proposed the Reject On Negative Impact (RONI) defense to determine whether a candidate training instance is malicious. [45] states if adding the candidate instance to a training set causes the resulting classifier to produce substantially more classification errors, reject the instance as detrimental in its effect.

III. LEARNING ON RESOURCE-CONSTRAINED DEVICES

Latency on IoT devices can increase based on the distance between data centers. Solutions to device latency include moving cloud resources closer to devices, which is seen in Mobile edge computing (MEC). The authors in [21] determined that MEC servers are now becoming a weak point due to data privacy concerns and high data communication overheads. Mobile devices may be subject to attack if not properly secured, and may require additional security when deploying important communications in MEC. The authors in [43] mention how blockchain can solve the security problem of edge computing.

As we know, IoT devices may have limited resources, and quantifying the expected resource usage can help reduce excess resource expenses [46]. If a device is showing

signs of being a straggler, device dropout can be applied to reduce communication costs. However, if we have a large number of stragglers, then simply dropping them can put a negative impact on the overall learning process. The authors in [47] developed FedAR algorithm that can select only the effective and trustworthy agents for learning process. Besides, FedPARL framework is specially designed for resource-constrained FL devices that can prune large model and allow feasible local task allocation for the edge devices [48]. To deal with heterogeneous resources of the agents, the authors in [49] proposed an FL model that can generate multiple global models considering the resource status of the agents and accelerate the learning process. Moving forward, increasing device participation can cause diminishing returns if devices
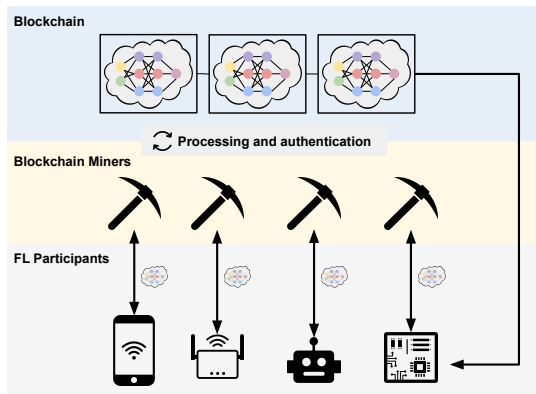

Fig. 6. Condensed workflow with integrated blockchain-based FL.

perform poorly during training. Adaptive methodologies, such as adaptive FL can help balance resource costs. The authors in [24] proposed an adaptive FL control algorithm based on gradient-descent for reducing computational resource budgets. The proposed control algorithm determined the best trade-off between local update and global parameter aggregation to minimize the loss function under a given resource budget [24]. The authors in [12] proposed AFAFed, an Adaptive FL solution for IoT devices affected by packet-loss communication. One of the key features of AFAFed is the implementation of distributed sets of local adaptive tolerance thresholds and global centralized adaptive fairness coefficients. AFAFed's features allow the algorithm to calculate the right personalization vs. fairness trade-off in various resource-constrained computing environments.

A. Verification of Local Model Updates

Verifying local model updates helps prevent resource-consuming[1] data from being aggregated with the global model. The authors in [1] found that blockchain-assisted decentralized FL frameworks can prevent malicious clients from poisoning the learning process and thus provides a self-motivated and reliable learning environment for clients [1]. Blockchain consensus mechanisms boost network agreement, thus deterring unreasonable data from being aggregated. The authors in [50]

[1]Resource-consuming can be considered negative, harmful, or malicious data.

mentions four different blockchain consensus protocols for confirming the correctness of a global model and executing global model aggregation. Blockchain consensus protocols include: Proof of Work (PoW), Proof of Stake (PoS), Raft, and practical Byzantine fault tolerance [50].

- Proof of Work (PoW) discourages harmful blocks from being added to the network by requiring the completion of a cryptographic puzzle by miners before verification. The authors in [51] proposed that PoW is the most famous algorithm for Bitcoin. Although it works effectively for protection from Sybil attacks and data manipulation, it is expensive because of the required hash power and long block interval [51].

- Proof of Stake (PoS) randomly selects data validators based on a total stake or overall impact on the network. Individuals with high amounts of holdings are given the opportunity to validate and approve new blocks being added. Validators in PoS create validator nodes to assist in network consensus. The validator's assets are locked while staking to help support network consensus. Validators receive incentives for vouching for the legitimacy of transactions.

- Raft supports consensus within decentralized blockchain environments. A raft can improve the understandability of the verification by breaking consensus problems down using decomposition. According to Ongaro and Ousterhout [52], Raft separates the key elements of consensus, such as leader election, log replication, and safety, and it enforces a stronger degree of coherency to reduce the number of states that must be considered [52].

- Practical Byzantine fault tolerance (PBFT) ensures a decentralized network continues operating even if a portion of nodes fail or act maliciously. The authors in [53] mention that PBFT is an optional consensus protocol for consortium blockchains scenarios, where strong consistency is required [53]. Consortium BCFL can prevent the straggler effect in FL, where the speed of the slowest device causes network-wide delays.

B. Global Model Aggregation

Global model aggregation continually occurs online. Blockchains can provide decentralized applications, such as smart contracts, to achieve global model aggregation without a central server. The distributed training of the global model can encourage dishonest participants to collude with others and launch a coordinated attack. The authors in [54] suggested that past clients can coordinate with current or future participants to participate in attacks against current or future updates to the global model [54]. Global model updates may not change much depending on context. While global model parameters are not changing much, parameters can be frozen to reduce communication costs. The authors in [55] found that it is unnecessary to always synchronize the full FL model in the entire training process because many parameters gradually stabilize prior to the ultimate model convergence [55]. Devices with low bandwidth benefit from frozen parameters to reduce device resource costs. Yang et al. [50] proposes a decentralized

blockchain-based FL architecture that can resist failures or attacks of servers and devices by building trustworthy global model aggregation with secure model aggregation based on blockchain consensus protocol among multiple servers.

C. Incentive Mechanism

Federated learning does not reward local device participation. Blockchain can offer an incentive mechanism to improve participation. The authors in [56] mention that there could be heterogeneous devices with different computational and data resources in an FL system. Therefore, participants with better resources must have extra benefits compared to participants with little contribution [56]. Honest contribution is advantageous to all participants and should be rewarded appropriately. Besides, the authors in [57] mention how the blockchain's incentive mechanism can track the contribution of each data provider towards the globally optimized model so that participants can be treated fairly, thereby attracting more data sharers. Blockchain-based federated learning architectures such as BlockFL [19], enable exchanging devices' local model updates while verifying and providing their corresponding rewards [19]. While the authors in [58] found that BCFL enables all clients to verify the learning results that are recorded on the blockchain, whereby distributed clients can be rewarded incentives to participate, and untrusted learning models can be detected.


Fig. 7. Example BCFL figure with example incentive mechanism for residential buildings.

D. Privacy and Security Protection

Adding noise to data can reduce overall efficiency. Too much noise and performance are diminished, while not enough noise and the state is potentially insecure. Levels of noise carry a trade-off that should be considered. FL can ensure data privacy by only exchanging model parameters with devices. Local device data is never communicated online, boosting the security of data. Multiple approaches examine noise trade-off. Truex et al. [23] proposed an FL approach that utilizes both differential privacy and secure multiparty computation (SMC). [23] combined differential privacy with secure multiparty computation to enable the growth of noise injection to be reduced as the number of parties increases without sacrificing privacy.

Security measures include anticipating anomalous behavior. Typical technologies for defending FL security include malicious participant detection, and malicious impact mitigation [32]. Irregular behavior can be indicative of an adversary or

potential machine learning attack. The authors in [59] mention six anomaly detection models for effectively detecting anomalous behavior: generative architectures, classification-based models, clustering-based models, nearest neighbor models, statistical & analytical models, and reinforcement learning-based models [59]. In addition, participant contribution can be measured to classify behavior between honest vs. dishonest participation.

IV. EXISTING RESEARCH CHALLENGES AND FUTURE RESEARCH DIRECTIONS

A. Potential Solutions of Existing Research Challenges

Prior works demanded high resource consumption for large scale data-driven environments. FL and its form of distributed computing can process participation from different IoT devices, thus allowing a variety of resourceful participation. Resourceful participation has three main challenges:

(i) High communication costs can be discouraged based on network reliability. Communication delays from service providers may require a device to drop out at slower speeds. The authors in [15] examine compressed communication to reduce communication overheads in BCFL. Cui et al., [15] created a communication-efficient framework that could reduce the training time by about 95% without compromising model accuracy. Approaches that reduce communication overheads are advantageous for lowering participation requirements.

(ii) Trade-offs between security and performance exist. Increasing noise against participants boosts security but also lowers performance. Choosing the best devices for participation can reduce the dependence on trade-offs between security and performance. Consortium blockchains are present solutions for the possible diminishing returns of open participation. Consortium blockchains screen participating devices to encourage optimal participation selection of trusted devices. The authors in [60] propose a budgeted number of candidate clients chosen from the best candidate clients in terms of test accuracy to participate in the training process. Optimal selection strategies can verify participants are resourceful and honest before inviting these participants into training.

(iii) Dishonest participation records are not shared between blockchains. Communicating participation records between blockchains can improve scalability. Records of dishonest behavior can warn similar blockchains of a participant's integrity. For example a dishonest mining pool with evidence of performing mining attacks against a blockchain can be blacklisted. Blockchains interoperability can allow participants to be chosen thoughtfully.

BCFL can improve resource-constrained environments, although being considerate of optimal participation further increases overall resource effectiveness. Li et al., [20] suggests that blockchain can further improve FL security and performance, besides increasing its scope of applications.
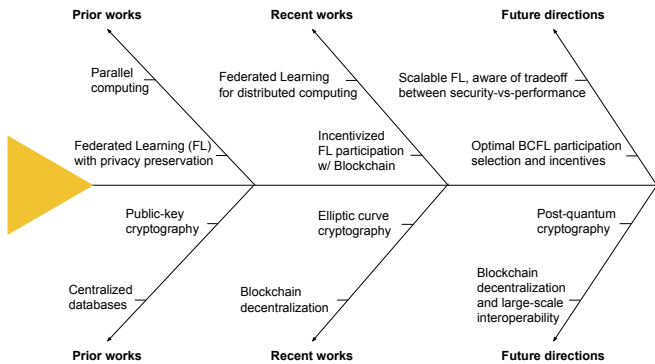
Fig. 8. BCFL flow from prior works towards future directions.

B. Future Research Directions

Future research can be conducted in the areas of BCFL scalability, quantum resilience and AI alignment. The aforementioned directions have challenges that may need to be addressed in the future. Figure 8 displays a timeline of important BCFL research directions. Each direction can disrupt resource management, requiring revolutionary system designs:

- Scalability can cause diminishing returns when resource requirements exceed reasonable blockchain computing limits. The authors in [10] states that as the number of transactions increases, it requires a larger blockchain size to store those transactions. However, mining a large size of blockchain may require more resources, which would be difficult for IoT devices [10]. Dimensional reduction for unreasonably large blockchains can prevent future mining participants from requiring supercomputers, compared to current GPUs.

- Post-quantum cryptography algorithms may be required for resilience from quantum attacks. The authors in [61] mention how the fast progress of quantum computing has opened the possibility of performing attacks based on Grover's and Shor's algorithms. Such algorithms threaten public-key cryptography, and hash functions [61]. Blockchain may require a cryptographic redesign to combat quantum attacks. Quantum computers bring unique threats to classical data management techniques.

- AI alignment is important for building a favorable machine learning model and central server. Corrigibility, the capability of being reparable, can be considered when the machine learning model learns unfavorable behavior. The model may seize to optimize due to model convergence. Further research could evaluate a models curiosity and the relation to model convergence, when significant performance changes are relatively stable. An honest-but-curious central server may or may not be reparable.

V. Conclusion

In this paper, we present legible insights of blockchain-based FL from crucial and time-demanding perspectives. Blockchain provides additional security in a decentralized format that protects FL from various real-world security and privacy issues. We categorize various threat models and point-out the leading vulnerabilities that could be observed during various stages of the FL process and in online environments of blockchain-enabled FL settings. To this end, we show a clear direction on how we can leverage secure and private learning on resource-constrained devices, covering the feasible verification of local model updates, secure global model aggregation, designing fair incentive scheme, and upgrading security and protection. Finally, we present the existing blockchain-based FL applications and highlight the potential solutions of the existing research challenges in the relevant domains. We anticipate that this survey will be helpful for researchers, practitioners, and scientists, developing robust blockchain-enabled FL systems choosing suitable consensus mechanism, identifying security and privacy pitfalls, motivating coherent formations, and following the promising future directions presented in this paper.

VI. Acknowledgment

References

[1] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.

[2] Ahmet Ali Süzen and Mehmet Ali Şimşek. A novel approach to machine learning application to protection privacy data in healthcare: Federated learning. *Namık Kemal Tıp Dergisi*, 8(1):22–30, 2020.

[3] Ahmed Imteaj, Urmish Thakker, Shiqiang Wang, Jian Li, and M Hadi Amini. A survey on federated learning for resource-constrained iot devices. *IEEE Internet of Things Journal*, 9(1):1–24, 2021.

[4] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.

[5] Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4):83–93, 2020.

[6] Ali Raza, Kim Phuc Tran, Ludovic Koehl, and Shujun Li. Designing ecg monitoring healthcare system with federated transfer learning and explainable ai. *Knowledge-Based Systems*, 236:107763, 2022.

[7] Wei Zhang and Xiang Li. Federated transfer learning for intelligent fault diagnostics using deep adversarial networks with data privacy. *IEEE/ASME Transactions on Mechatronics*, 27(1):430–439, 2021.

[8] Hongwei Yang, Hui He, Weizhe Zhang, and Xiaochun Cao. Fedsteg: A federated transfer learning framework for secure image steganalysis. *IEEE Transactions on Network Science and Engineering*, 8(2):1084–1094, 2020.

[9] Sudipan Saha and Tahir Ahmad. Federated transfer learning: Concept and applications. *Intelligenza Artificiale*, 15(1):35–44, 2021.

[10] Ahmed Imteaj, M Hadi Amini, and Panos M Pardalos. *Foundations of Blockchain: Theory and Applications*. Springer, 2021.

[11] Yuanhang Qi, M Shamim Hossain, Jiangtian Nie, and Xuandi Li. Privacy-preserving blockchain-based federated learning for traffic flow prediction. *Future Generation Computer Systems*, 117:328–337, 2021.

[12] Enzo Baccarelli, Michele Scarpiniti, Alireza Momenzadeh, and Sima Sarv Ahrabi. Afafed—asynchronous fair adaptive federated learning for iot stream applications. *Computer Communications*, 2022.

[13] Alberto Blanco-Justicia, Josep Domingo-Ferrer, Sergio Martínez, David Sánchez, Adrian Flanagan, and Kuan Eeik Tan. Achieving security and privacy in federated learning systems: Survey, research challenges and future directions. *Engineering Applications of Artificial Intelligence*, 106:104468, 2021.

[14] Xu Cheng, Weiwei Tian, Fan Shi, Meng Zhao, Shengyong Chen, and Hao Wang. A blockchain-empowered cluster-based federated learning model for blade icing estimation on iot-enabled wind turbine. *IEEE Transactions on Industrial Informatics*, 2022.

[15] Laizhong Cui, Xiaoxin Su, and Yipeng Zhou. A fast blockchain-based federated learning framework with compressed communications. *IEEE Journal on Selected Areas in Communications*, 2022.

[16] Chaosheng Feng, Bin Liu, Keping Yu, Sotirios K Goudos, and Shaohua Wan. Blockchain-empowered decentralized horizontal federated learning for 5g-enabled uavs. *IEEE Transactions on Industrial Informatics*, 18(5):3582–3592, 2021.

[17] Wael Issa, Nour Moustafa, Benjamin Turnbull, Nasrin Sohrabi, and Zahir Tari. Blockchain-based federated learning for securing internet of things: A comprehensive survey. *ACM Computing Surveys (CSUR)*, 2022.

[18] Jing Huey Khor, Michail Sidorov, and Peh Yee Woon. Public blockchains for resource-constrained iot devices—a state-of-the-art survey. *IEEE Internet of Things Journal*, 8(15):11960–11982, 2021.

[19] Hyesung Kim, Jihong Park, Mehdi Bennis, and Seong-Lyun Kim. Blockchained on-device federated learning. *IEEE Communications Letters*, 24(6):1279–1283, 2019.

[20] Dun Li, Dezhi Han, Tien-Hsiung Weng, Zibin Zheng, Hongzhi Li, Han Liu, Arcangelo Castiglione, and Kuan-Ching Li. Blockchain for federated learning toward secure distributed machine learning systems: a systemic survey. *Soft Computing*, 26(9):4423–4440, 2022.

[21] Dinh C Nguyen, Ming Ding, Quoc-Viet Pham, Pubudu N Pathirana, Long Bao Le, Aruna Seneviratne, Jun Li, Dusit Niyato, and H Vincent Poor. Federated learning meets blockchain in edge computing: Opportunities and challenges. *IEEE Internet of Things Journal*, 8(16):12806–12825, 2021.

[22] Dinh C Nguyen, Quoc-Viet Pham, Pubudu N Pathirana, Ming Ding, Aruna Seneviratne, Zihuai Lin, Octavia Dobre, and Won-Joo Hwang. Federated learning for smart healthcare: A survey. *ACM Computing Surveys (CSUR)*, 55(3):1–37, 2022.

[23] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM workshop on artificial intelligence and security*, pages 1–11, 2019.

[24] Shiqiang Wang, Tiffany Tuor, Theodoros Salonidis, Kin K Leung, Christian Makaya, Ting He, and Kevin Chan. Adaptive federated learning in resource constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 37(6):1205–1221, 2019.

[25] Yang Zhao, Jun Zhao, Linshan Jiang, Rui Tan, Dusit Niyato, Zengxiang Li, Lingjuan Lyu, and Yingbo Liu. Privacy-preserving blockchain-based federated learning for iot devices. *IEEE Internet of Things Journal*, 8(3):1817–1829, 2020.

[26] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.

[27] Karan Ganju, Qi Wang, Wei Yang, Carl A Gunter, and Nikita Borisov. Property inference attacks on fully connected neural networks using permutation invariant representations. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 619–633, 2018.

[28] Christopher A Choquette-Choo, Florian Tramer, Nicholas Carlini, and Nicolas Papernot. Label-only membership inference attacks. In *International conference on machine learning*, pages 1964–1974. PMLR, 2021.

[29] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pages 1322–1333, 2015.

[30] Hailong Hu and Jun Pang. Stealing machine learning models: Attacks and countermeasures for generative adversarial networks. In *Annual Computer Security Applications Conference*, pages 1–16, 2021.

[31] Ashwinee Panda, Saeed Mahloujifar, Arjun Nitin Bhagoji, Supriyo Chakraborty, and Prateek Mittal. Sparsefed: Mitigating model poisoning attacks in federated learning with sparsification. In *International Conference on Artificial Intelligence and Statistics*, pages 7587–7624. PMLR, 2022.

[32] Qiang Duan, Shijing Hu, Ruijun Deng, and Zhihui Lu. Combined federated and split learning in edge computing for ubiquitous intelligence in internet of things: State-of-the-art and future directions. *Sensors*, 22(16):5983, 2022.

[33] Yunfei Liu, Xingjun Ma, James Bailey, and Feng Lu. Reflection backdoor: A natural backdoor attack on deep neural networks. In *European Conference on Computer Vision*, pages 182–199. Springer, 2020.

[34] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 387–402. Springer, 2013.

[35] Jie Wen, Zhixia Zhang, Yang Lan, Zhihua Cui, Jianghui Cai, and Wensheng Zhang. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics*, pages 1–23, 2022.

[36] Yann Fraboni, Richard Vidal, and Marco Lorenzi. Free-rider attacks on model aggregation in federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1846–1854. PMLR, 2021.

[37] Qi Xia, Zeyi Tao, and Qun Li. Defenses against byzantine attacks in distributed deep neural networks. *IEEE Transactions on Network Science and Engineering*, 8(3):2025–2035, 2020.

[38] Eric Chan-Tin, Victor Heorhiadi, Nicholas Hopper, and Yongdae Kim. The frog-boiling attack: Limitations of secure network coordinate systems. *ACM Transactions on information and system security (TISSEC)*, 14(3):1–23, 2011.

[39] Lingjuan Lyu, Han Yu, Xingjun Ma, Chen Chen, Lichao Sun, Jun Zhao, Qiang Yang, and S Yu Philip. Privacy and robustness in federated learning: Attacks and defenses. *IEEE transactions on neural networks and learning systems*, 2022.

[40] Shubhani Aggarwal and Neeraj Kumar. Attacks on blockchain. In *Advances in Computers*, volume 121, pages 399–410. Elsevier, 2021.

[41] Jinyuan Jia, Ahmed Salem, Michael Backes, Yang Zhang, and Neil Zhenqiang 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*, pages 259–274, 2019.

[42] Xinle Yang, Yang Chen, and Xiaohu Chen. Effective scheme against 51% attack on proof-of-work blockchain with history weighted information. In *2019 IEEE International Conference on Blockchain (Blockchain)*, pages 261–265. IEEE, 2019.

[43] Guangshun Li, Xinrong Ren, Junhua Wu, Wanting Ji, Haili Yu, Jiabin Cao, and Ruili Wang. Blockchain-based mobile edge computing system. *Information Sciences*, 561:70–80, 2021.

[44] Zhe Yang, Kan Yang, Lei Lei, Kan Zheng, and Victor CM Leung. Blockchain-based decentralized trust management in vehicular networks. *IEEE Internet of Things Journal*, 6(2):1495–1505, 2018.

[45] Marco Barreno, Blaine Nelson, Anthony D Joseph, and J Doug Tygar. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.

[46] Ahmed Imteaj and M Hadi Amini. Distributed sensing using smart end-user devices: Pathway to federated learning for autonomous iot. In *2019 International conference on computational science and computational intelligence (CSCI)*, pages 1156–1161. IEEE, 2019.

[47] Ahmed Imteaj and M Hadi Amini. Fedar: Activity and resource-aware federated learning model for distributed mobile robots. In *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1153–1160. IEEE, 2020.

[48] Ahmed Imteaj and M Hadi Amini. Fedparl: Client activity and resource-oriented lightweight federated learning model for resource-constrained heterogeneous iot environment. *Frontiers in Communications and Networks*, page 10, 2021.

[49] Khandaker Mamun Ahmed, Ahmed Imteaj, and M. Hadi Amini. Federated deep learning for heterogeneous edge computing. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1146–1152, 2021.

[50] Zhanpeng Yang, Yuanming Shi, Yong Zhou, Zixin Wang, and Kai Yang. Trustworthy federated learning via blockchainf. *IEEE Internet of Things Journal*, 2022.

[51] Jungjae Lee and Wooseong Kim. Dag-based blockchain sharding for secure federated learning with non-iid data. *Sensors*, 22(21):8263, 2022.

[52] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (Usenix ATC 14)*, pages 305–319, 2014.

[53] Zhujun Zhang, Dali Zhu, and Wei Fan. Qpbft: practical byzantine fault tolerance consensus algorithm based on quantified-role. In *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 991–997. IEEE, 2020.

[54] Nader Bouacida and Prasant Mohapatra. Vulnerabilities in federated learning. *IEEE Access*, 9:63229–63249, 2021.

[55] Chen Chen, Hong Xu, Wei Wang, Baochun Li, Bo Li, Li Chen, and Gong Zhang. Communication-efficient federated learning with adaptive parameter freezing. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, pages 1–11. IEEE, 2021.

[56] Youyang Qu, Md Palash Uddin, Chenquan Gan, Yong Xiang, Longxiang Gao, and John Yearwood. Blockchain-enabled federated learning: A survey. *ACM Computing Surveys*, 55(4):1–35, 2022.

[57] Saide Zhu, Ruinian Li, Zhipeng Cai, Donghyun Kim, Daehee Seo, and Wei Li. Secure verifiable aggregation for blockchain-based federated averaging. *High-Confidence Computing*, 2(1):100046, 2022.

[58] Chuan Ma, Jun Li, Long Shi, Ming Ding, Taotao Wang, Zhu Han, and H Vincent Poor. When federated learning meets blockchain: A new distributed learning paradigm. *IEEE Computational Intelligence Magazine*, 17(3):26–33, 2022.

[59] Muneeb Ul Hassan, Mubashir Husain Rehmani, and Jinjun Chen. Anomaly detection in blockchain networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 2022.

[60] Ihab Mohammed, Shadha Tabatabai, Ala Al-Fuqaha, Faissal El Bouanani, Junaid Qadir, Basheer Qolomany, and Mohsen Guizani. Budgeted online selection of candidate iot clients to participate in federated learning. *IEEE Internet of Things Journal*, 8(7):5938–5952, 2020.

[61] Tiago M Fernandez-Carames and Paula Fraga-Lamas. Towards post-quantum blockchain: A review on blockchain cryptography resistant to quantum computing attacks. *IEEE access*, 8:21091–21116, 2020.

Shabnam Rezapour is an Assistant Professor of Enterprise and Logistics Engineering at Florida International University. Her research focus is on "Architecting Resilient Societies" that are able to cope with and transform in the face of expected and unexpected chronic stresses (slow-moving disasters that weaken the fabric of societies such as climate change and water shortage) and acute shocks (sudden events that threaten societies such as floods, hurricanes, terroristic attacks, and disease outbreaks). Her research group develops analytical models (e.g., stochastic optimization and simulation models) to understand and analyze the vulnerability of network-centric infrastructures/systems and develop optimal mitigation policies.

Ervin Moore is a Ph.D. student and DHS CAESCIR Fellow at the Knight Foundation School of Computing and Information Sciences at FIU. He is a member of the Sustainability, Optimization, and Learning for InterDependent networks laboratory (solid lab) working with Dr. Amini and Dr. Rezapour. Prior to that, he graduated with a MSc in Artificial Intelligence and Machine Learning from Colorado State University-Global, following his BA in Communication Technology from University of Texas at Arlington. Ervin participated in AI-related hackathons and has won multiple. His applied research involved using machine learning and large language models to improve user experiences. His current research interests are artificial intelligence, machine learning and their application in critical infrastructure security and resilience.

M. Hadi Amini (S'11–M'19-SM'22) is an Assistant Professor at Knight Foundation School of Computing and Information Sciences at Florida International University. He is the director of Sustainability, Optimization, and Learning for InterDependent networks laboratory (www.solidlab.network). He received his Ph.D. in Electrical and Computer Engineering from Carnegie Mellon University in 2019, where he received his M.Sc. degree in 2015. He also holds a doctoral degree in Computer Science and Technology. Prior to that, he received M.Sc. degree from Tarbiat Modares University in 2013, and the B.Sc. degree from Sharif University of Technology in 2011. His research interests include distributed optimization and learning algorithms, distributed computing and intelligence, sensor networks, interdependent networks, and cyber-physical-social resilience. Application domains include smart cities, energy systems, transportation networks, and healthcare. He has received over $3M in funding from various federal and local funding agencies in the United States.

Hadi is a Senior Member of IEEE, and also a life member of IEEE-Eta Kappa Nu (IEEE-HKN), the honor society of IEEE. He served as President of Carnegie Mellon University Energy Science and Innovation Club; as technical program committee of several IEEE and ACM conferences; and as the lead editor for a book series on "Sustainable Interdependent Networks" since 2017. He also serves as Associate Editor of SN Operations Research Forum and International Transactions on Electrical Energy Systems. He has published more than 100 refereed journal and conference papers, and book chapters. He edited/authored six books. He is the recipient of the best paper award from "2019 IEEE Conference on Computational Science & Computational Intelligence", FIU's Knight Foundation School of Computing and Information Sciences' "Excellence in Teaching Award", best reviewer award from four IEEE Transactions, the best journal paper award in "Journal of Modern Power Systems and Clean Energy", and the dean's honorary award from the President of Sharif University of Technology. (Homepage: www.hadiamini.com)

Ahmed Imteaj is an Assistant Professor in the School of Computing at Southern Illinois University, Carbondale. He received his Ph.D. in Computer Science from Florida International University in 2022, where he received his M.Sc. degree with recognition of the Outstanding Master's Degree Graduate Award. Prior to that, he received B.Sc. degree in Computer Science and Engineering from Chittagong University Engineering and Technology. His research interests span Federated learning, cybersecurity, Internet of Things (IoT), and blockchain. Ahmed significantly contributed to the area of privacy-preserving distributed machine learning and internet-of-things and published his research in top-tier conferences and peer-reviewed journals. Ahmed's research has been recognized several times. Ahmed was recognized with the 2022 FIU Real Triumph Graduate Award, "2022 Outstanding Student Life Award: the Graduate Scholar of the Year Award", "2021 Best Graduate Student in Research Award" from the KFSCIS at FIU, Best Paper Award from the 2019 IEEE CSCI'19 conference, and one of the two winners at 2021 FIU GSAW Scholarly Forum. Ahmed is the lead author of the book, "Foundations of Blockchain: Theory and Applications" and published more than 50 peer-reviewed journals and conference papers.