



Federated Learning and Privacy Preserving Ai: A Comprehensive Overview

Prajwa H M¹, Sanjana Kulkarni², Niveditha Y³, Nisha Y M⁴, Sahana M⁵

1, 2, 3, 4, 5 Department of computer science and engineering, Sri Sairam College of Engineering

ABSTRACT

AI has gained currency in recent times, where statistical models seem to be at the forefront. Yet, there are increasing privacy and legal concerns with the conventional centralized training and inference approach to AI technologies. Federated learning (FL) comes onto the set to actually solve such issues revolving around secure distributed AI. It integrates security features at different stages of data including pre-processing, training, evaluation, and deployment using secure multi-party computation, and differential and hardware privacy. FL aims at addressing issues of data privacy as well as enacting measures to mitigate against data isolation. Nonetheless, federated models are yet to evade threats like plagiarism, illegal copying and misuse. In a bid to do this, FedIPR incorporates the use of watermarks onto FL models for proof of ownership and protection of intellectual property rights (IPR). Many papers that emphasize on security still misappropriate FL to refer to unsecured distributed machine learning. For this reason, this paper reemphasizes the concept of FL and suggests SFL to develop honest, privacy-preserving AI systems with secure IP-rights cross-border guarantees. It also gives a succinct picture of threats as well as attacks and defenses applicable in each SFL phase from a life cycle perspective.

Keywords: Federated Learning, Privacy-Preserving AI, Differential Privacy, Secure Aggregation, Decentralized Machine Learning, Data Privacy, Non-IID Data, Privacy-Utility Trade-off, Communication Efficiency, Applications of Federated Learning

1. Introduction:

In the AI area, the centralized approach for training and inference is really gaining momentum, with impressive milestones. However, with privacy concerns growing fast and the legal standards demanding greater and greater scrutiny over this issue, a classical approach is constantly under threat. With privacy coming as the most prominent point in data, tremendous innovation is required to share ideas without revealing sensitive information.

Federated learning has been identified as the new AI paradigm that would overcome the challenges of model training in a distributed fashion without transferring raw data. Federated learning at every step of its lifecycle-from preprocessing, through training, evaluation, to deployment-is characterized by robust data security measures. FL utilizes secure multi-party computation, differential privacy, and hardware-enhanced solutions that guarantee data confidentiality and integrity across all phases. All of this aside, federated models remain vulnerable to possible threats in the form of intellectual property theft, illegal copying, and



misuse. The answer to these and many more lies in emerging innovations that incorporate watermarks in FL models, through FedIPR, hence an option for ownership validation, along with IPR security.

A very practical use case for federated learning is Google's research illustrating how to allow mobile phones to cooperatively learn a common prediction model with all of their training data residing locally, without needing to leave any information on the network. This work also discusses the technical aspects of federated learning implementation, in terms of communication, encrypted, and the aggregation of model updates. One very interesting application is on Gboard for Android; it uses user interactions to improve query suggestions without infringing on user privacy.

The goal of federated learning is to try building secure and private AI systems. However, one should distinguish FL implementations from unsecured distributed machine learning frameworks. It advocates the re-definition of the present state of Federated Learning (FL) from scratch and focuses on kick-starting safe federated learning as a framework in the process of building trusted and reliable AI. Presenting an elaborative review of recent studies, this research delves into the lifecycle of SFL by creating essential threats, vulnerabilities, and defense mechanisms to offer suitable protection against thefts in data privacy and intellectual property.

2 Architectural Frame Work of Federated Learning

Federated Learning (FL) is a paradigm shift for machine learning in the model training across multiple devices or servers without centralized data gathering. This is critical in a privacy-sensitive domain where collective learning is done without transmitting raw data. The architectures of FL can be considered under the following categories,

2.1 Type of Federated Learning Architecture

The architecture for federated learning is meant to be based on the secure, decentralized processing of data. The principal types include:

- **Centralized Federated Learning:** This method involves a central server that coordinates the learning process by aggregating model updates from multiple clients. Devices (clients) locally train the model on their data, then share model parameters—not data itself—with the central server. The server aggregates these parameters to create a global model, ensuring privacy by retaining data locally.
- **Decentralized Federated Learning:** Unlike the central FL, decentralized architectures do not depend on a central server. In this setup, the clients communicate directly with each other either in a fully connected network or in a peer-to-peer fashion. This increases the resilience against single-point failures and decreases the possible risks that could be incurred from the central data aggregation, which is an appealing reason for privacy and robustness.



2.2 Elements of a Federated Learning System

The architecture of FL comprises of a few crucial constituent elements specifically tailored to serve for distributed learning with maintaining private data in mind, for instance,

- **Clients:** Also called devices: it is segmented into distribution equipment such as smartphones, devices associated with Internet of things; even to the institutions such as hospitals and banking institutions; these clients do have some localised storage with the aid of which training of model by every single client goes privately on individual data thus contributing to distributed learning.
- **Model aggregator** in centralized federated learning, a model aggregator or server collects the updates coming from clients. Aggregation in such centralized FL is often performed using techniques including Federated Averaging, where parameters from the local models are averaged to produce the global model. In decentralized setups, peer-to-peer protocols allow clients to share their updates without a central coordinator.
- **Communication Channels:** Communication among clients and the central server (or among clients in decentralized systems) is an essential requirement for model updates. Efficient communication protocols and compression techniques are used to reduce bandwidth requirements and minimize latency.

2.3 Federated Learning Workflow

The basic workflow of a federated learning process includes the following:

1. **Initialization of global model:** This is performed either by the server when in central FL or by the participating client when in the case of decentralized FL.
2. **Client Selection:** A group of participating clients are selected for this training round. The basis for the selection could be a function of the availability of the devices, the time the batteries can last and data quality.
3. **Local Training:** The client selected trains the model with its local data. The model parameters are optimized for privacy-preserving objectives. This is done without the exchange of data and based on local computations.
4. **Model Aggregation:** For a centralized architecture, the server collects updated models from the clients and aggregates the models by using the weighted average of model parameters. For decentralized FL, peers update their models collaboratively.
5. **Global Model Update:** Finally, it uses these aggregate model parameters for updating the global model. This process is implemented through several rounds of training and is done until some specified convergence criteria are satisfied, such as maximum achievable accuracy or loss values.

2.4 Distinctions over Centralized Machine Learning Architecture

Federated learning is quite distinct from centralised ML, which mainly depends on the aggregation of humongous data sets into a central location. Major differences include:


- **Data Distribution:** FL trains on distributed data without relocation, which very



well fits the requirements for privacy and data ownership as opposed to traditional ML, which depends on centralized datasets.

- Privacy and Security: The decentralization nature of FL has reduced privacy risks because the data never leaves the local devices. This is quite a difference from centralized ML, which aggregates raw data.
- System Scalability: Federated architectures need to address the issues of bandwidth and client reliability. For this purpose, client selection protocols and communication-efficient updates can be applied in the management of large-scale, distributed data sources.

All these architectural features make FL a viable and effective solution for privacy-preserving, distributed learning by offering for diverse, large-scale datasets with maintaining contemporary standards of data privacy and regulations.



3 Privacy-Preserving Techniques in Federated Learning

The most prominent privacy-preserving techniques in FL include differential privacy, homomorphic encryption, secure multi-party computation, and federated averaging.



3.1 Differential Privacy (DP)

Differential Privacy refers to a strong mathematical construct that ensures the individual points of data cannot be reinstated even if the adversaries have some auxiliary knowledge on the dataset. This concept is extensively used in FL that protects client updates by adding specially calibrated random noise to parameters or gradients of the models before sharing them with the servers. The key aspects regarding DP in FL are:

- **Privacy Budget:** DP defines the privacy budget, which denotes the amount of privacy used. Low values of privacy budget result in better privacy but affect model accuracy.
- **Noise Injection:** Injecting noise during model updates in DP results in hiding individual contributions in such a manner that these cannot be back-engineered by the attackers.
- **Privacy vs. Utility:** The amount of noise applied needs to be balanced by the model performance so as to maximize both privacy and utility.

3.2 Homomorphic Encryption (HE)

Homomorphic encryption is encryption that allows computations to run directly on ciphertext, that is, the server can update the model by aggregating the encrypted client updates rather than decrypting them first. Thus, in such an FL setting, computation at the central server remains privatised during the training stage. Homomorphic encryption with FL is useful for

- **End-to-End Encryption:** The clients encrypt their updates and send them; results are decrypted only once for the aggregated outcome so that there is no leakage of any information.
- **Safety during Data Transmission and Aggregation:** Since computation occurs over the encrypted data, it saves against eavesdropping and unauthorised access through homomorphic encryption.
- **Computational Efficiency Issues:** Completely homomorphic encryption has higher computational complexity, thus problematic in large-scale applications in FL. Optimisation using partial or approximate homomorphic encryption improves performance.

3.3 Secure Multi-Party Computation (SMPC)

Secure Multi-Party Computation is a method through which multiple parties can jointly compute the output of a function in such a way that no party learns an individual input. In FL, SMPC is applied to update the model securely by splitting private data into encrypted shares sent to clients or servers while preserving privacy. The application of SMPC in FL is as follows:

- **Data Sharing with Zero Knowledge:** SMPC enables aggregation on shared data without revealing actual data contributions from individual agents and thus preserves client's privacy.
- **Secure Model Aggregation:** SMPC distributes and diffuses model updates, ensuring sensitive information is not a single point of vulnerability.
- **Resource Demands:** Just like HE, SMPC also demands resources for its usage, requiring



huge amounts of computational power and bandwidth in the case of real-time FL applications.

3.4 Federated Averaging

Federated Averaging, an algorithm designed to support FL, aggregates the updates on the model, not on the raw data. FedAvg supports preserving privacy in this way:

- **Aggregation Rather Than Sharing of Data Parameters:** Locally, clients get a trained model; however the exposed data by a client occurs much less with regard to data exposure and on average those updates shared only share a parameter with that will become averaged in server-end side.
- **Improved communication cost:** The mechanism supports sending an update frequency on the model in some or as a compression where again it can also lower that cost due to effective transfer.
- **Privacy Enhancement with Locality:** The locality of the process also helps prevent local models from being used in data leak during training: even the feds' learning never gets to peek inside your local model

3.5 Combining Privacy Techniques

Most FL implementations combine the above techniques to maximize both privacy and security. Differential privacy with federated averaging can achieve a balance between security and computational efficiency, while secure multi-party computation and homomorphic encryption may be combined to enhance the robustness of privacy-preserving protocols in federated systems.

3.6 Practical Considerations and Trade-Offs

Each privacy-preserving approach employed in FL introduces some novel trade-offs between privacy, accuracy, computational efficiency, and communication costs as follows:

- **Privacy-Model Trade-Offs:** Some methods, like DP and HE, can be more successful for preserving privacy at the cost of model performance. So careful tuning becomes inevitable.
- **Overheads of Computation and Communication:** Techniques used under HE and SMPC are quite heavy computationally. Thus, these approaches are rather too expensive on resource-scarce devices, where optimization is usually the best choice.
- **Scalability Issues:** FL is implemented on a very large scale so that the security mechanisms would not hinder the system's performance and create communication and computation bottlenecks.

Such privacy-preserving methods ensure that FL can be applied to applications of sensitive data such as health and finance whereby data remains secure while providing collaborative machine learning.



4 Challenges in Privacy-Preserving Federated Learning

Communication and Bandwidth Limitations

In Federated Learning, there must be constant communication between the central server and multiple clients or among clients in a decentralized FL setup. Since every client processes local data and transmits model updates to the server, communication can be a bottleneck, especially in large-scale deployments. Heavy Communication Cost: High sending frequency for model updates tends to result in high bandwidth, especially with a large number of clients. For example, the devices used by these clients are likely to have low connectivity and could be cellular or IoT devices.

- Compression Techniques can be used for mitigation with bandwidth constraints at updates. Update compression, quantization, and sparsification are some of them. They reduce the update size, but may sometimes introduce errors, reduce model precision, etc.

Model Performance Trade-offs

The most critical privacy-preserving techniques for securing FL are very expensive in terms of model performance. A big challenge remains in balancing privacy and accuracy.

- Impact of Differential Privacy (DP): DP techniques add noise to model updates to hide the model's updates with respect to the data fed into the model. Though noise helps protect privacy, too much noise actually lowers the model's accuracy and slows down convergence.
- Convergence Rates of Reduced Models: Privacy mechanisms as well as communication delays, can slow down the rate by which models converge and would thus impact training speed.

The Privacy-Utility Trade-off

The fundamental objective in Federated Learning is to maintain privacy without losing the efficacy of accuracy in the generated models. In contrast to this, with an enhanced privacy constraint, model efficacy deteriorates; however a trade-off occurs between both privacy and utility.

- Privacy budget constraints: DP limits leakage of information. Tight privacy constraints often cost in the accuracy of the model, and the tuning of the budget needs careful consideration to ensure an adequate balance.
- Risk of Overfitting and Underfitting: Because the updates are not transparent, the FL models tend to underfit, especially when updates carry added noise. On the other hand, overfitting may occur if the model captures local data distribution patterns rather than general trends.

Data Heterogeneity and Non-IID Data Issues

In Federated Learning, data is distributed and mostly kept on many devices, often in the form of



a non-IID collection, meaning that data isn't homogeneously split across instances.

- Data imbalance along the clients leads to models suffering from the introduction of bias in their updates since not all have equal parts of data to update model.
- Variability in Local Data Distributions: Non-IID data might cause inconsistent model updates that could make aggregation complicated and affect model accuracy. Some aggregation methods, such as clustering-based aggregation, can mitigate inconsistencies but complicate the FL system further.

Security Threats and Attacks

FL systems, even with privacy-preserving mechanisms, are still susceptible to various security threats: inference attacks and adversarial attacks. Inference Attacks: This may be achieved by inferring private information from model updates, especially in scenarios where protection of the data is weak. These attacks can be partially prevented through differential privacy and secure aggregation techniques.

- Adversarial Attacks: Bad clients can inject wrong information or manipulate updates to distort the performance of the global model. Robust aggregation techniques and anomaly detection techniques are required to mitigate such attacks, but those incur a cost in increased system complexity.

Scalability Challenges

Federated Learning needs to be scaled up to very large numbers of clients, each possibly having heterogeneous devices, distributions, and network conditions, bringing forth logistic challenges.

- Device Reliability and Availability: Most of the FL applications are devised on resource-constrained devices. These include mobile phones or IoT devices, which have inferior processing power, limited space for storage, or unstable connectivity to the network.
- Client Selection: For proper training, a good and representative subset of clients have to be selected. A big challenge is in balancing client availability with fair participation by the clients.

Regulatory and Ethical Considerations

Federated Learning often involves sensitive data and, hence, has to adhere to privacy regulations and ethical considerations.

- Data Privacy Laws: The laws like GDPR or HIPAA may add additional constraints on the handling and privacy techniques of data in FL.
- Ethical Use of Federated Data: FL needs to be designed and implemented so that it respects the rights of users, with transparency regarding data usage and model behavior.

These pose significant complexities that are part of what defines privacy-preserving Federated Learning, making a need to scale up on scalability, security, and efficiency while designing the learning system in practice. This implies that all such limitations will have to be addressed as a starting point toward achieving the true promise of FL. Applications of Federated Learning



with Privacy. Here's an overview of actual-use-case realizations of how FL enables private processing of data across distributed environments.

5 Applications of Federated Learning with Privacy

Applications of FL across various domains with industries, as their wide application will enable cooperative yet privately trainable machine learning applications.

5.1 Healthcare and Medical Research

Healthcare is one of the sectors where data is abundant, scattered across hospitals, clinics, and research institutions. FL enables medical researchers and practitioners to use data collaboratively without violating the privacy of the patients.

- **Predictive Health Models.** FL allows the sharing of models for disease risk prediction in a collaborative fashion without transferring patient records. This can, therefore, create more generalized models by spreading learning across multiple institutions.
- **Drug Discovery and Genomics:** FL can be applied in genomic data analysis or drug discovery by combining information from various genetic sources. Such an approach would enable quick, more accurate drug discovery without involving patient DNA information.
- **Privacy-Sensitive Diagnostics:** Federated Learning can be applied to the development of AI diagnostic tools processing medical images and records on-site. The risks involved with data transfers are diminished, while more accurate diagnosis is achieved using collaborative learning.

5.2 Financial Services and Transaction Security

There is plenty of treasure in terms of information; hence application of privacy preserving techniques holds the future key for these sectors towards attaining and abiding to data protection laws. Especially, those institutions with regard to enhancement of security with improved models without any type of leaking client's details can depend on FL.

- **Fraud Detection:** FL makes it possible for banks and financial organizations to build jointly models detecting fraud and money laundering patterns and helps diminish the potential of fraudulent transactions without any disclosure of the client's information.
- **Credit Scoring and Risk Evaluation:** The banks and credit departments can design FL models on the profile of data obtained from the other clients within the other institutions, ensuring more realistic credit evaluation.
- **Customer Personalization:** FL enables financial institutions to provide personalized offers-loans and investments-by training models locally on individual transaction patterns thus enhancing customer experience without exposing sensitive data.

5.3 Smart Devices and IoT Applications

Federated Learning is most appropriate for Internet of Things environments where data has



been generated across millions of connected devices, such as smartphones and sensors. Here, FL enables devices to train models locally and updates are shared rather than data with the potential of improving the privacy and efficiency.

- **Personalized Device Models:** FL is being used in smart phones and smart devices to personalize experiences such as predictive text, voice recognition, and recommendation systems for the individual user behavior of each user while keeping all data local.
- **Smart Homes and Wearables:** IoT devices in a smart home, such as thermostats and wearables, can collectively learn how to predict user preferences or monitor health indicators without transmitting raw data to a central server.
- **Autonomous Vehicles:** Federated Learning empowers connected cars to learn the traffic pattern and improve its navigation by training on locally collected data. This is a protective measure for passenger privacy while improving autonomous driving safety and efficiency.

5.4 Government and Public Sector Applications

Governments hold extremely sensitive citizen data. FL enables governmental agencies to tap into the knowledge that machine learning will provide without compromising citizens' privacy, especially in public health, law enforcement, and education.

- **Public Health Monitoring:** FL will enable agencies to monitor health trends and outbreaks of epidemics by aggregating anonymous health data collected from hospitals and clinics for tracking the spread of diseases without compromising individual privacy.
- **Smart City Infrastructure:** Smart city initiatives leverage data from distributed sources, including traffic cameras, air quality sensors, and utility data. FL will allow cities to optimize traffic, monitor pollution, and improve security systems in a collaborative and secure manner.
- **Education:** FL in education allows educational institutions to analyze academic performance, behavioral data, and learning preferences without revealing student information. This facilitates personalized learning solutions while keeping the student's information private.

5.5 Retail and E-commerce Personalization

There are retail and e-commerce platforms where voluminous consumer data is stored, making FL an appropriate application for use in instances when a high degree of personalization is required without revealing consumer data.

- **Product Recommendations:** FL can generate personalized product recommendations for consumers based on their browsing habits, purchase history, or regional preferences. This further enhances the user experience but would not require sensitive data to be sent out of the device.
- **Inventory and Demand Forecasting:** Retail chains can pool the sales data of different locations, determine the demand in those locations, and thereby optimize the stock



quantity without any invasion of personal privacy since they don't have access to the data of any one location.

- Customer Segmentation: FL helps customer segmentation models based on local purchase patterns. The business organizations can then make personalized marketing campaigns with the least amount of invasion of personal privacy.

The applications of Federated Learning with privacy-preserving techniques are diverse and cut across sectors, like healthcare, finance, IoT, public governance, and retail. Each application benefits from the core strengths of FL, including distributed data processing, enhanced privacy, and regulatory compliance. These attributes make FL an ideal solution for organizations aiming to innovate while respecting data privacy constraints.

6 Future Research Directions in Federated Learning

Federated Learning is a budding domain with dynamic methodologies and expanding applicability. Future research directions into FL focus on scalability, improve security, data and model quality issues, and architecture with new designs to extend robustness and wide-scaled applicability of the proposed technique.

6.1 Advanced Privacy-Preserving Techniques

Privacy Preserving Mechanisms in FL Future developments include integration between differential privacy, that is, SMPC as well as HE. As discussed above, it provides strong security in data due to the encryption.

Privacy Model Hybrid Privacy Models to get more robust privacy model while minimizing the performance of trade-offs, combining all techniques such as HE can be used for secure computation, and DP for noisy injection.

Adaptive privacy budgets. Current research: adaptive budgets that depend on the sensitivities of data, opening up possibilities to become more agile and more precise in models of FL.

6.2 Communication Cost Efficiency Improvements

This future version of FL further reduces communications costs in the hope to overcome bandwidths, as well as devices with limited band-widths.

- Federated Optimization Algorithms: Optimizations of Federated Averaging and other such algorithms are under development that minimize the need to transfer data, thereby making FL efficient for large-scale deployment.
- Compression and Sparsification Techniques: Techniques for compressing model updates or sparsifying parameters to share selectively are being explored to reduce the data transferred between clients and servers without affecting model quality, thereby conserving bandwidth.

6.3 Non-IID Data and Fairness

Handling non-IID data has been critical in FL research, especially for improvement in terms of model accuracy and fairness.



- **Federated Transfer Learning:** The use of transfer learning in federated settings allows models to learn through diverse data distributions, thus potentially enhancing the model's ability to generalize between heterogeneous data sources.
- **Fairness-Aware FL Models:** Future FL frameworks are already working on finding ways of reducing biases emerging from client-side data imbalance so as to provide equal outcomes on different populations.

7 Ethical and Regulatory Considerations

It brings in new ethical and regulatory issues regarding the adoption of Federated Learning, especially where sensitive data is used, such as in healthcare, finance, and government sectors. Compliance with the relevant privacy regulations and adherence to the highest ethical standards will only be possible if FL applications are to gain trust.

7.1 Compliance with Privacy Laws

FL must be compliant with the privacy regulations like GDPR in the EU, HIPAA in the US, and so on globally.

- **Data Sovereignty:** FL respects data sovereignty since the data remains within local jurisdictions and hence adheres to the privacy regulations.
- **Right to be Forgotten:** Some degree of individual control over personal information can be achieved within FL systems by providing mechanisms to delete or cease sharing data at the client's request. This fulfills some of the requirements for laws such as GDPR.

7.2 Transparency and Accountability

Transparency in the use of data, model behavior, and client involvement is what ethical use of FL calls for.

- **Explainability in FL Models:** Models need to be developed so that end-users understand precisely how predictions are made so that high-stakes users like those in healthcare can trust the approach.
- **Accountability and Governance:** FL systems need to have some governance frameworks guiding client participation, data sharing permissions, and ethical use of model insights.

8 Conclusion

Federated Learning offers a new way of collaborating on machine learning by ensuring that insights are gained on decentralized data while maintaining individual privacy. FL has applications in healthcare, finance, IoT, and public sectors, giving a scalable, privacy-preserving alternative to centralized models. Yet there are challenges ahead, communication costs, heterogeneity of data, and model performance trade-offs being some of them. Future advancements in such aspects as privacy techniques, communication efficiencies, and alignment with regulations are key drivers for FL's future developments. Continued innovation



and a focus on responsible business practices will determine how FL impacts data-driven industries globally.

References

1. Kairouz; Mahan; Year: 2021, “Advances and Open Problems in Federated Learning”, Foundations and Trends in Machine Learning. This paper provides a comprehensive overview of FL, discussing challenges, applications, and future directions.
2. Geyer; Klein; Nabi, Year: 2017, “Differentially Private Federated Learning: A Client- Level Perspective”, Proceedings of the International Conference on Machine Learning. Discusses differential privacy for FL and its impact on model performance.
3. Yang; Liu; Chen; Tong; Year: 2019, “Federated Machine Learning: Concept and Applications”, ACM Transactions on Intelligent Systems and Technology. Covers key applications of FL and the benefits of privacy-preserving machine learning.
4. Hard; Rao; Mathews; Year: 2018, “Federated Learning for Mobile Keyboard Prediction”, This research demonstrates FL’s application in smart devices and privacy-centric model personalization.
5. Bonawitz; Eichner; Year: 2019, “Towards Federated Learning at Scale: System Design”, Proceedings of the 2nd SysML Conference. Focuses on scalability, security, and communication efficiency in large-scale FL systems.
6. Mahan; Moore, Year: 2017, “Communication-Efficient Learning of Deep Networks from Decentralized Data”, Proceedings of the 20th International Conference on Artificial Intelligence and Statistics. Introduces the Federated Averaging algorithm, a foundational method for efficient FL.
7. Li; Sahu; Talwalkar; Smith, year: 2020, “Federated Learning: Challenges, Methods, and Future Directions”, IEEE Signal Processing Magazine. A detailed review of challenges in FL, including privacy, data heterogeneity, and regulatory compliance.
8. Shokri; Shmatikov, year: 2015, “Privacy-Preserving Deep Learning”, Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. Discusses foundational privacy-preserving techniques that are applicable in FL.