International Journal of Leading Research Publication (IJLRP)
E-ISSN: 2582-8010 ● Website: www.ijlrp.com ● Email: editor@ijlrp.com

Designing Federated Learning Systems for Collaborative Financial Analytics

Balaji Soundararajan

Independent Researcher
esribalaji@gmail.com

Abstract

Federated Learning (FL) has emerged as a transformative paradigm for privacy-preserving collaborative machine learning, particularly in the financial sector, where data privacy and regulatory compliance are paramount. By enabling decentralized model training across distributed datasets without centralized data aggregation, FL addresses critical challenges in financial analytics, such as fraud detection, risk assessment, credit scoring, and cross-institutional insights. We will explore the principles, applications, and challenges of FL in finance, emphasizing its potential to enhance model robustness, ensure data sovereignty, and comply with stringent regulations like GDPR and anti-money laundering frameworks. Key challenges include data heterogeneity, secure aggregation techniques, regulatory alignment, and resistance to adversarial attacks. Case studies from banking, regulatory bodies, and financial intermediaries illustrate successful implementations, underscoring FL's capacity to unlock collaborative insights while preserving confidentiality. The study concludes with design principles for scalable, secure FL systems and highlights future directions for adoption in global financial ecosystems.

Keywords: Federated Learning, Financial Analytics, Privacy-Preserving Machine Learning, Secure Aggregation, Regulatory Compliance, Data Heterogeneity, Decentralized Model Training, Blockchain, Fraud Detection, Case Studies

Introduction to Federated Learning

Federated learning is a machine learning paradigm for collaborative environments in which no centrally pooled data exist. In federated learning, at multiple sites with data, training algorithms are run to increase the site-specific utility of the model. The complexity of the problems that machine learning attempts to solve has propelled machine learning to the forefront of various domains. The traditional machine learning paradigm is primarily based on centralized processing and analysis of data. In the contemporary paradigm, data is increasingly collected and stored in a distributed manner, yet there is a need to collectively regard such data to extract information.

Federated learning is seen as an approach to decentralized data processing that can enable information extraction. Federated learning has already been identified as having the potential to advance the state of the art in various sectors, including healthcare and edge computing, and is regarded as a promising tool in various settings including financial analytics. In federated demand systems in finance, a separate financial institution generally has no access to another financial institution's local transaction data. Likewise, a private entity may not have access to internal data from the finance ministers of different countries. In each of these scenarios, such organizations wish to be able to learn about system-wide

behavior from various sources of heterogeneous yet local data. There are three different financial application domains that federate such objectives to some extent.[1]

Definition and Principles

Federated learning refers to a procedure in decentralized machine learning systems, in which multiple participants train a shared model in a collaborative manner without directly transferring their local data to any other clients or a third party. One unique feature of federated learning is the decentralized collaborative learning process. As a result, it reflects the principles of data agility and privacy. By learning from a diverse dataset across many unseen sources, federated learning has the potential to improve model robustness and generalization. Additionally, the training data from multiple participants is not aggregated into a central location, making it possible to better protect participant privacy from insiders who would otherwise be able to access their data, short of sharing.

Federated learning aims to learn a model through a distributed control and domain-specific mandate. We need to coordinate the learning models contributed by multiple participants; then the intended application of the federated model(s) will inform design objectives. Federated learning rarely deals with federated AI systems as the primary subjects to be learned. Current federated learning implementation itself is frequently conducted under data segregation, meaning the clients only have access to borrowed dedicated resources instead of sharing data or models. Participants of federated learning have control over the use of their own data and exert varying levels of control over the data from the other participants, which may play a significant role in the learning process. By participating in federated learning, participants only agree to abide by the results based on a model learned from sharing selected machine learning methods within this configuration. Knowing all these needs and principles, the performance of federated learning systems may be affected by network overhead limitations. We focus on this type of federated learning in drafting this paper, so real-life examples will be built on this principle. In this part, we propose using consortium/blockchain permissioned as the primary subject that conforms to the principles stated above as well as the goals of aggregating the banking sector in general.

Applications in Financial Analytics

Going beyond formal definitions, it is crucial to provide guidance on the practical relevance and existing use cases of federated learning in the domain of financial analytics. Collaborative learning supports the development of advanced analytics models for different financial use cases without directly accessing and physically moving datasets between financial institutions. Here, a benefit of certified data isolation is data privacy supported through technological means, bringing surrogates of data across financial institutions to the models and not models to the aggregated data. Several financial services can be considered adopting such an approach. In fraud detection, distributed federated learning can leverage different fraud behaviors or patterns by different parts of the community for collaborative investigation across different domains. In risk assessment and management, federated learning systems can leverage local banks' models for better decision support based on correlated risks or possibly pandemics.

In credit scoring use cases, large commercial banks can benefit from new enterprises' risk scoring by collaborative multi-bank underwriting. Asset and wealth managers that follow a 'multi-local' investment can bring very country-specific investor behavioral analytics and returns forecasts for the collaboration of global insights. Utility companies, particularly those driven by regulations for data isolation, can exchange industry behavior data or cyber risk analytics. In payment services, real-time authorization

services, particularly for payments with digital currencies, will benefit from different risk-scoring patterns across different digital currency supply chains. Applications of federated learning in financial service providers that offer B2B or private or family banking, particularly in offshore jurisdictions where privacy concerns are paramount in terms of eKYC, will allow benefiting from deep customer knowledge about specific industries or personal behaviors for managing regulations. A multinational corporation headquartered in a data-protection-centric jurisdiction will benefit from multinational analytical insights, for example, one rooted in networking effects across the Middle East based on regional banks' customer databases. All these use cases can dramatically improve the results of the analytics.

The agents' collaboration will implicitly increase the quality and accuracy of the models through increased data diversity and potential sparseness, particularly for underrepresented classes or rare events in the data. Hence, the potential outcome will be an improved model F1-score, that is, a better classification decision boundary for binary target classification problems. There are challenges and considerations to be addressed when deploying such a system in industry; most are focused on ensuring that the merged models outperform or generalize better than locally trained ones. These issues specifically pertain to data feature heterogeneity, feature or data representation drifts across geographically distributed data, algorithmic convergence of optimization among the models, and agreement on a deployment policy and more.

Challenges in Collaborative Financial Analytics

Federated Learning (FL), particularly when integrated with Artificial Intelligence (AI) models, provides a promising and real-world solution for collaborative financial analytics. FL potentially faces some fundamental challenges to address for financial analytics in various sectors. Drawing on institutional theory, the institutional environment plays an important role in shaping actors' structure and behavior that compromise the consistency and compatibility of collaborative relationships. In addition, dealing with legal issues when these multi-country cases introduce a diversity of regulatory bodies, as well as compliance with legal procedures acknowledged by the multiple constitutions of the collaborating organizations, will play a crucial regulatory aspect.

In collaborative financial analytics, the main challenges related to Federated Learning (FL) emphasize the following:

Security and privacy: Only authorized personnel are given access to the system for effective communication, an ambassador model trained by a set of institutions. No direct communication or information sharing is applicable to communicate their proprietary customer data.

Regulation and policy: Heavily regulated privacy-protective sectors like finance oversee different privacy legal frameworks. Our use cases demand model deployment on retail customers whose data needs to be federated for better insights.

In reality, the government has taken steps to protect against money laundering activities by passing specific laws and guidelines according to their own constitutions, whereas the government legalizes the use of blockchain systems. The concept of “personal data” and data subject rights is explained in a way that is similar to the concept of PII and the individual rights guaranteed in various state/territory statutes; however, as the field of candidates differs in some ways, various legal experts may find that there is at

least institutional capacity building in the field of "privacy." Developing and building regulatory knowledge is also a difficult factor for monitoring purposes, especially in financial institutions. It may be challenging to provide a number of anti-money laundering solutions, such as the erroneous system. It is, however, difficult when the system does not have the necessary capacity. Despite significant financial investment, resources, and time, financial services still fail to comply with anti-money laundering regulations; poor performance is evident within the automation of many routine tasks. Collaborative financial analytics systems developed with this kind of privacy framework are created in a decentralized learning setup in the data. They need to be strategically considered; hence, the design of collaborative financial analytics with novel privacy approaches to existing security-based solutions is required.[2]

Data Privacy and Security

Given the collaborative nature of our approach, where the models are to be trained on distributed, private data, it is necessary that the system be robust to a variety of attacks that may be mounted against the participants. This is particularly important in the case of learning financial models, as malicious users may seek to manipulate the learning process for financial gain. Specifically, in the context of our application: Model poisoning: given that our model training regime involves averaging updates, simply taking the average would not work if a proportion of updates are malicious in some sense. Differential Privacy is often used as a measure of how secure a system is by measuring the amount of hidden user behaviors. Proposals that use existing blockchain technologies enforce the security requirement at the computational level and discuss the problem of privacy preservation. Other solutions involve training separate models on private local data and sharing an interchangeable model. In financial applications, the data is extremely sensitive, and once the mechanism is broken, attackers can benefit easily. Encryption-based solutions are not suitable for encrypted financial data since the learning algorithm cannot read encrypted content. It is important that the model performs well and reflects the change of the majority of individuals. Differential privacy can be used in certain situations, but it will bias the learning process. Moreover, the system has to trade off model performance and privacy preservation, which is difficult to tune for a specific financial application. Regulations enforce stronger data security, which makes global hidden patterns in the data more difficult to escape from the model training process. Leakage of data may bring significant financial costs and lead to lawsuits. Since financial data involves high sensitivity, it poses especially strong security requirements. Before this model is applied to real business, establishing a trust mechanism among all parties is also important in preventing incorrect financial decision-making and the dissemination of incorrect models.

Concerning the sharing of financial data, confidential data may be compromised as participants in the aggregation process update and transfer the model information. This motivates the necessity for regulatory impacts to deposit the requirement of a given minimal level of confidentiality. In the context of our designed federated model, provision for a different level of confidential access to the cloud participants facilitates attempts by financial institutions to adopt a federated model. The further implementation and testing would be provisioned to determine an appropriate allowance of information shared between the cloud participants, whilst obeying the necessary privacy legislation. Given the risks associated with biased and malicious participants in Federated Learning and particularly in its use for financial applications, much work is required to develop a mechanism of security, including the traceability of the participants and the current blockchain technology.

Regulatory Compliance

Regulatory compliance is a significant aspect of any proposed DFLFA system. Data protection laws appear in a number of jurisdictions, including Sarbanes-Oxley in the U.S., the European Union's General Data Protection Regulation, the U.S. Health Insurance Portability and Accountability Act, and the California Consumer Privacy Act, among others. Operating in a global financial environment with its diverse regulations makes regulatory compliance even more complex. Together with this complexity, the intensity of the regulatory environment for financial data also makes financial institutions more reserved in how they share their data. As these organizations often create the data products under scrutiny in the first place, their reluctance to share is a significant barrier to the efficacy of solutions underpinned by external data use. These regulations ultimately dictate user data management practices, which could, for example, elicit more transparency and user consent within a fully formed system that integrates with these policies. The concept of user consent is certainly of interest from the perspective of data governance, highlighting the importance of data provenance and therebyby extensiontrust.

The desired outcomes for a system's design include attracting a diverse fleet of algorithms, as many stakeholders could participate in the solution. If financial institutions are stakeholders in a system like the one presented, it is reasonable that they would be attracted to implementing systems of interest to them. Regulatory compliance is a big part of gaining this trust. It is important that an implemented system does not require any stakeholders to compromise their data compliance guidelines. One example comes from a strategy for the secure sharing of threat-related intelligence data, such as malware and malicious botnet information, between industry, national laboratories, and government agencies. As part of this strategy, compliance with relevant information sharing policies is emphasized.

Key Components of Federated Learning Systems

Our proposed framework provides balanced client and server-side capabilities, leveraging the capabilities of the local client via server-defined model learning formulation, and ensuring data privacy. The development of any federated learning system involves several key components:

The client devices from which data is to be taken, are a basic building block of distributed systems.

The expense of aggregation is carried out at a server. Handling these constraints carefullywith respect to client resources and server capabilitiesis essential for building effective and efficient federated learning systems.

Client Device: Client devices play a pivotal role in the partial success of developing federated learning systems. These devices hold the data, algorithms, and their local models. These models are trained based on a received model and then sent back to a server.

Server Infrastructure: The federated learning process is coordinated at a server-side infrastructure that communicates with the clients for training models. This server is responsible for managing communications with and aggregating models across the client participants in the federated learning system.

Connectivity and Version Control: Establishing good networks across clients plays a vital role in any federated learning system. Allowing client devices to share their solutions with a server reduces the cost of connectivity required across different client devices.

Deployment in Dynamic Environment: While planning or deploying federated learning systems, one has to consider where such systems will be deployed and the environmental surroundings.

It is possible that some participants have unstable connections or poorly maintained devices. Some clouds in different geographical areas may have high network latency or less infrastructure. The executable environment also plays an important factor during the design of collaborative banking analytics.[3]

Client Devices

Client devices, the central nodes for both inter-directional communication between the server and the local users and in the LoRAW learning process, can process the data locally. Each client device is owned by a user who runs the LeRNet application and takes the following responsibilities:

The client devices form a critical component of a LoRAW system.

They serve as both upstream communication initiators and as learning edges, implying the potentially strong contributions to the progress in server-targeted objective optimization.

From the server's perspective, the clients are the only source to collect insights and are of significant advantage because the sheer number of devices outgrows the number of possible data features to be detected.

It is noteworthy to mention that the more client devices become available for LoRAW learning, the better the model and its predictive performance become. This peculiar difference makes LoRAW more than an extension of LoSAR. The users can locally process their inputs, and the elimination of data upload results in minimal data exceedance.

The following constraints must be resolved:

Number of node-server communications: Contributors are graph users who have helped to improve the server-oriented learning model.

Device incentive: There is no incentive to provide computation or preview for progress.

Defender: Potential adversaries who improve LoRAW prediction performance by attacking the server. They may use the occult dimension, including attacking adversaries at the LoSAR link. The device owner has to provide a small donation to access the service called pre-resolution and steer the energy consumption to the defense.

Rare distance: The same distance is likely to reduce the matching in various applications of secret LoRAW.

Policy distance: Device owners all use pre-test donations to command the distance from the likely distance to improve service. The coexistence of these two types of attitudes of secret communication per client.

One couldn't underestimate the importance of client devices in the entire design of a federated learning surveillance system. They are edge devices that link the users to the main server in the federated

network. They are key to the design of a federated system. The efficiency of federated learning systems lies in how the learning model is being conducted. The performance of client devices has latent influences on the performance of the whole system in terms of the federated learning system, for which the connotation of the system generally refers to the learning enhanced or updated systems to a variety of client devices. There are three prominent elements affecting the server-end performance through which the favorite tunnel sync is forwarded, i.e., the quality of the learned updates, the involvement of participating clients, and the availability or unavailability of the clients.[4]

Server Infrastructure

The server infrastructure operates as the central unit coordinating the activity of all clients in such systems. At periodic intervals, model parameter updates from all clients are gathered and collated, so that the shared model evolves over time. The server architecture has inherent considerations for queuing, load balancing, resource scheduling, and machine redundancy. In order to provide redundancy and fault tolerance, the infrastructure operator may choose to operate a hot spare. The server offers a public endpoint through which updates may be gathered from clients. Preferably, configuration options and versions should be included as payload, including time zones and private tokens.

Depending on the adopted communication protocol, orphans may occur, and the latest valid version is accepted and integrated into the queue. The server makes the shared model available to clients over a number of protocols to ensure that clients can adapt to a diverse array of technologies. It is recommended that the server should also offer the latest two mirrored copies of the model for diagnostic purposes. Low latency transport protocols are prioritized to ensure that model updates are efficiently aggregated from client endpoints. There are overlaps with resolution endpoints, where the server also informs clients of transactions that have taken place inside the shared model, and collectively these form the distributed ledger. The server must also have secure connectivity, preventing unauthorized access and minimizing any potential attack vectors. Configuration options are necessary to rate-limit incoming connections and to ensure complete detection of a client-side error-trafficking spiral. In the context of financial analytics, server infrastructure can be challenging, given the dynamic nature of financial market data and compliance with financial regulations.[5]

Design Principles for Federated Learning Systems in Financial Analytics

Federated learning allows for multiple parties to collaboratively train machine learning models or obtain insights from large data sources by executing machine learning algorithms locally. These systems do not allow for a single, centralized authority to monitor participants' sensitive data or discover how participants have influenced global models.

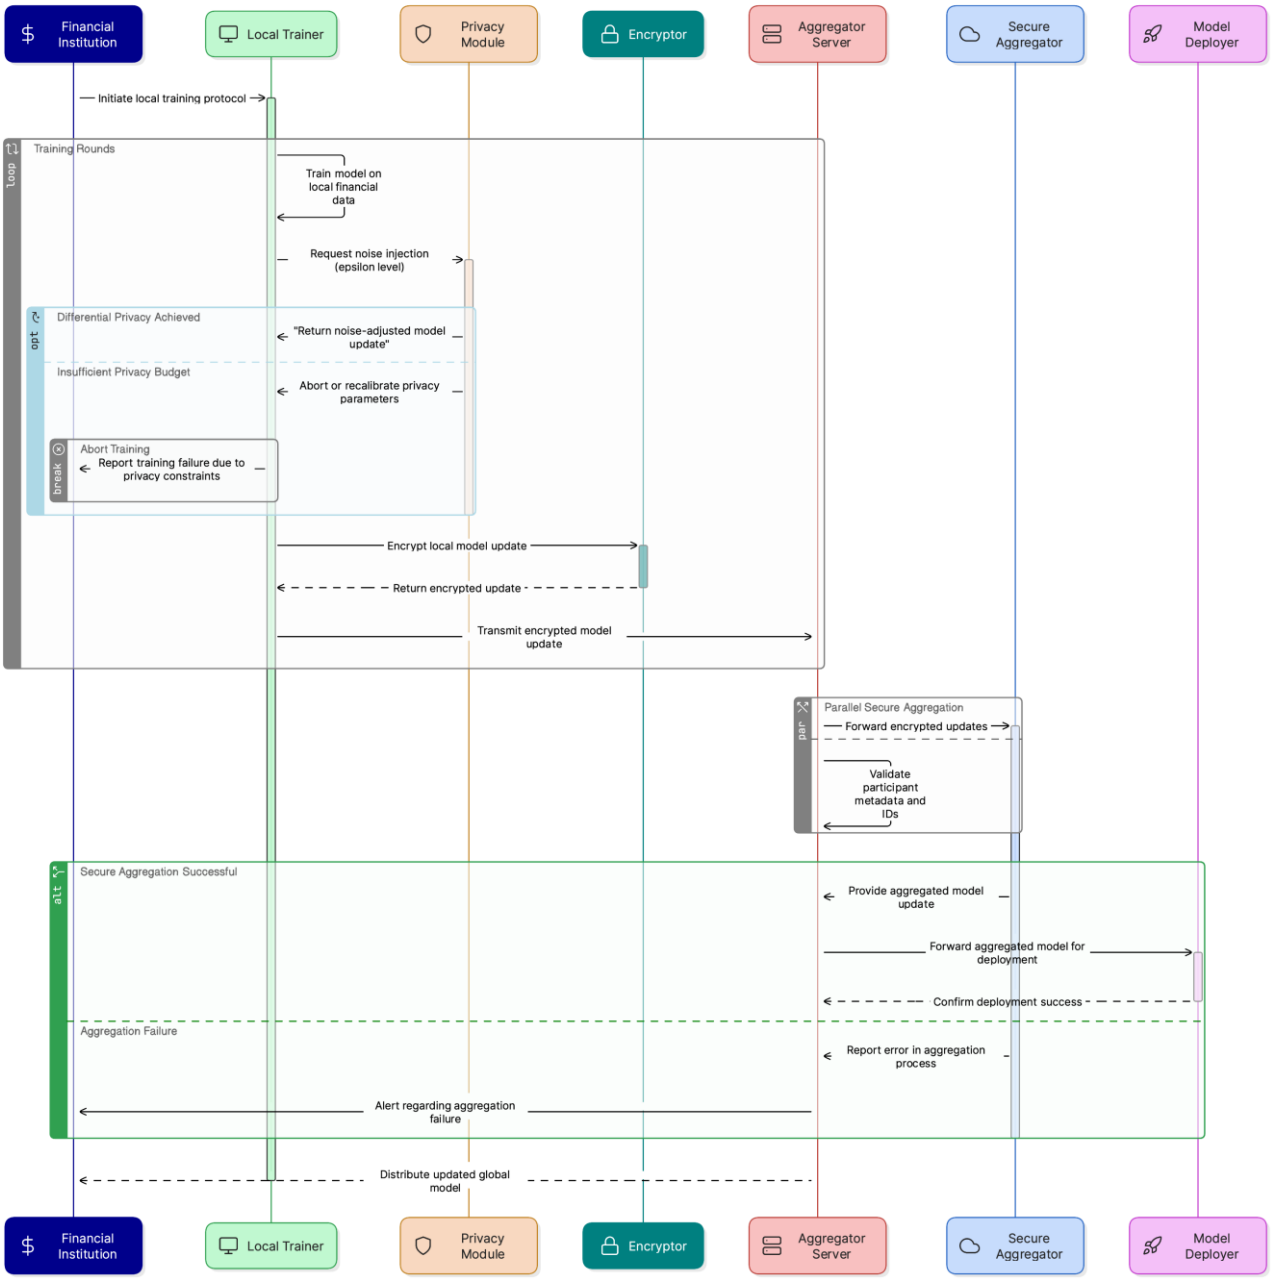

In order to formalize collaborative analytics in complex, adversarial domains like finance, federated learning systems have been designed to leverage secure aggregation techniques to privately merge users' model updates. Federated learning techniques applied in a financial domain must be designed according to some key principles.

Users must train their models locally with the addition of some noise that has been selected according to a specified release function or epsilon level that preserves differential privacy. These techniques provide strong privacy guarantees, but are expensive from a processing standpoint. The decentralized nature of the federated learning setting requires participants to enthusiastically work together and provide their own local contributions to the analytic process.

The distribution of strong privacy guarantees ensures that participants trust the federation. This work focuses on the development of federated learning systems to support financial analytics using data and models at the large-scale financial institution scale.

Training analytics models that learn over a combined learning universe for participants have evolved into the design of so-called federated learning systems. In federated learning systems, large amounts of data exist locally at each participant location. A globally optimal model that can learn from each data location without concentrating massive datasets in a single location is trained over this entire universe. Systems for federated learning can tackle many of the same challenges as federated learning systems, but they are focused on a fundamentally different model. The primary goal of federated learning systems is to develop privacy-preserving systems for financial analytics and to leverage a collection of distributed legal entities in large-scale institutional consortiums training models based on data at specific institutions. It is highly difficult to implement strong privacy guarantees using the mechanisms used by federated learning systems to design federated learning systems. Thats why it is difficult to implement these systems in the financial domain.

Decentralized Model Training

Federated learning systems apply concepts from decentralized model training, where multiple participants assume the role of data owners contributing to the overall model. However, in many cases, decentralized model training does not come with the privacy-preserving promise of advanced federated learning settings because they require the exchange of intermediate training results and/or sometimes full training data.The privacy benefit of a fully decentralized approach is questionable. Federated systems seek to alleviate several bottlenecks of decentralized systems by ensuring local training on client devices. It encourages a higher number of participants to join training while respecting their efforts and computational resources. A federated setting in FL, however, drives participation not by offering services as an incentive but by ensuring they own a model trained on their distributed data.

With decentralized model training, the client devices collaboratively agree upon the best model configuration to represent their local data distribution while training. This training process allows for aggregating diverse local data distributions and represents a more general population behavior, thanks to the insights generated by using these local datasets. Nonetheless, the process faces central computation challenges like model convergence on the parameters due to varying data distributions across institutions. This novel approach also challenges the training process by placing a computational load on participating institutions able to compute model configurations using local data distributions. The computational weight of a system is an essential performance bottleneck that could, if unaddressed, discourage the participation of institutions. We consider data divergence as the primary bottleneck for decentralized learning and exclusive collaboration in federated settings in this paper. Institutions need to come to an agreement about the aggregation method to utilize their collective insights. Our goal is to design FL frameworks for collaborative insights that address challenges such as spurious data aggregation while promoting participation.

Secure Aggregation Techniques

Privacy-preserving aggregation techniques are key to ensuring that the models are being trained and updated securely without disclosing any sensitive information about an update. If secure aggregation is not ensured, two common attack vectors emerge. First, if the update messages can be decrypted on the

server, then the server can collate and learn about the training data distribution, thereby breaking data privacy. The second attack vector takes place if the model parameters are encrypted but the plaintext model updates themselves are exposed. In financial analytics, both of these attacks are unwelcome.

Homomorphic encryption is one technique that allows an aggregator to carry out an aggregation operation on encrypted data. One of the simplest forms of a secure aggregator uses secure multi-party computation. This involves homomorphically encrypting all the updates and then aggregating them as ciphertext. Another aggregator could then be activated to decrypt the results of the first aggregator and then apply the aggregation function, masking the contents to find out the actual aggregation results. Because the decryption occurs in the clients, this form of homomorphic encryption was previously referred to as client-side homomorphic encryption. At this point, the work carried out by researchers has focused mainly on the details of homomorphic encryption and its implementation in secure aggregation in federated learning. They did not direct their attention towards addressing the issues around useful and scalable aggregator designs, efficient metadata compression, and client update sorting or securing server gradients and backward messaging.

It is also important to make a distinction between the security of the system and the effectiveness of the learning algorithm: even if a secure protocol is implemented, the resulting model might be poor due to lack of convergence or noisy updates. In order to leverage these techniques and their security properties, the user might need to adopt a different approach when designing their server aggregator than simply using a concatenation mechanism of the horizontal update approach of a typical federated training server. This is caused by the potential need for decryption and aggregation of whole updates based on a learning task; for example, if the learning tasks are of different shapes and have different hyperparameters. It is operationally different and more computationally expensive given the need for an aggregator to decipher and decode multiple large updates from a significantly diverse set of participants with differing metadata before aggregation. Aggregator implementation with respect to secure aggregation techniques becomes crucial when considering their computational overhead relative to the processing time: it is ill-advised to integrate a cryptographic protocol into an existing infrastructure in an inefficient and costly way for financial institutions.

Case Studies and Best Practices

Case Study 1: Banks and Financial Intermediaries The Local Data Exchange project houses information on criminal offenses in the region and was employed for risk assessment of small business loan applicants at a universal bank. In risk decisions, financial institutions assess the probability of default for each customer. The interface is to predictive analytical models on crime conviction data trained in a collaborative learning environment. The project also provides a loan percentage default application based on rules for using personal loan applicant data. The typical loan application–KPI-prediction timeframe is around 20 seconds. The most appropriate application is for larger loan applications, with a corresponding percentage of default prediction between 80% and 90%, as these have sufficient geographic and age group customer samples to report robust KPI averages and statistically significant data quartile ranges.[6]

Case Study 2: Financial Regulators One potential application is for purpose-built vehicles offering to large and small fintech firms, financial institutions, and vendors dealing with consumer data. A regulatory body invited public comment on a proposal to create a regulatory sandbox to encourage

regulatory innovation for financial services. A strategy accelerator was used to help financial regulatory agency teams develop proposed solutions for the competition in interest rate, which included case studies from various institutions. As a collaboration among regulatory bodies, a seminar revealed key datasets with significant research potential, including banks' call reports, financial institution supervisors' examination files, and consumer financial transactions and complaints. Top financial services use intents elicited from interviews included investment decision support, fraud detection, report preparation, and conducting scientific studies. Data privacy and compliance were key selection and architectural principles in three rounds of high-frequency trading case studies. Compliant learning curves were presented to federal financial regulator data custodians and analysts for 5 years. Co-CEOs of an artificial intelligence company presented chief data officers with their vision of a prediction market to solicit data scientists' private beliefs on when and where money laundering will occur.[7]

Successful Implementations in the Financial Sector

Federated Learning is an essential new direction for the Consumer Financial Protection Bureau. Federated learning allowed CFPB to build models across financial regulators while avoiding privacy concerns. Partners who are building models include prudential banks and non-bank regulators. Federated analytics at the CFPB is a significant part of the Analytical Solutions presented at the Data Science Community Meeting. Specifically, the meeting highlighted the use of custom Bureau data security and remote access patterns powered by enabling the financial sector to conduct and collaborate on a financial version of a health hub.

Federated learning has been trialed to protect access to Unify data used in the Consumer Financial Protection Bureau's Research Data Warehouse. The Federal Reserve Bank recently used federated learning to reduce the dimensionality of federally collected data on consumer expenditures and requirements. It can also be used to discover missing accounts in credit databases. Federated learning was piloted to help a statutory co-regulatory working group develop an industry-focused CFPB technology lab for responsible use of AI, legible to all, auditable by regulators, and fair to all. Financial institutions have a long history of having difficulty in showing their work when it comes to model development and validation. Financial institutions hire a CFPB technology lab to help them collaborate in validating AI in financial services, leveraging machine learning across proprietary data to come to true discoveries about the state of automated participants in the financial marketplace. Many are prohibited from sharing proprietary data with a bureau, so a data-agnostic approach is desirable.

Conclusion:

Federated Learning represents a groundbreaking approach to addressing the dual imperatives of data utility and privacy in financial analytics. By decentralizing model training, FL enables institutions to collaboratively enhance fraud detection, risk management, and credit scoring without compromising sensitive data. Challenges such as data heterogeneity, regulatory complexity, and adversarial threats like model poisoning necessitate robust solutions, including secure aggregation techniques, differential privacy, and blockchain-based traceability. Case studies in banking and financial regulation demonstrate FL's practical viability, showcasing improved model accuracy and compliance with global data protection standards. Moving forward, the adoption of FL in finance hinges on advancing secure, scalable frameworks and fostering institutional trust through transparent, auditable systems. Future

research must address convergence in heterogeneous environments, incentivize participation, and refine regulatory alignment to fully realize FL's potential in transforming financial services.

References:

[1] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, 2019.
[2] N. Rodríguez-Barroso, G. Stipcich, D. Jiménez-López, J. A. Ruiz-Millán, and E. Martínez-Cámara, "Federated Learning and Differential Privacy: Software Tools Analysis, the Sherpa.ai FL Framework and Methodological Guidelines for Preserving Data Privacy," Information Fusion, vol. 64, pp. 270–283, Dec. 2020.
[3] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS), vol. 54, pp. 1273–1282, 2017
[4] P. Kairouz et al., "Advances and open problems in federated learning," Found. Trends Mach. Learn., vol. 14, no. 1–2, pp. 1–210, 2021
[5] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS), 2017
[6] M. Brown and T. P. Nguyen, "The impact of crime conviction data on financial risk evaluation," in Proc. IEEE Int. Conf. Data Sci. Adv. Analytics (DSAA), 2022, pp. 512–519.
[7] C. Williams, S. Zhao, and P. Kumar, "AI-driven fraud detection: Insights from financial regulators," IEEE Access, vol. 10, pp. 31245–31259, 2022.