

Privacy-Aware Financial Risk Control: A Federated Learning Approach with Differential Privacy Optimization

Ziang Yang

Cornell University, Ithaca, USA

ziangyang070@gmail.com

Abstract: This study proposes a financial risk control and privacy protection method based on federated learning (FL) to address the challenges of traditional centralized risk control models in data silos, privacy leakage risks, and cross-institutional collaboration. Financial risk control relies on a large amount of user transaction, credit score, and market behavior data, but due to privacy regulations (such as GDPR, CCPA), it is difficult for financial institutions to directly share data, resulting in limited generalization of risk control models. Federated learning enables multiple financial institutions to collaboratively optimize risk control models without data leaving the local area through distributed training, effectively protecting data privacy. This study constructs different FL architectures, including horizontal FL, vertical FL, and federated transfer learning, and analyzes their impact on risk assessment models. In addition, we introduce a differential privacy (DP) mechanism to evaluate its impact on model performance (AUC, Precision) while protecting user data. The experiment is verified based on the FICO Credit Score Dataset. The results show that FL performs better than traditional centralized learning methods in risk control tasks, and that appropriately adjusting the DP level can strike a balance between privacy protection and model performance. This study provides a secure and efficient data collaboration solution for financial risk control and lays the foundation for the development of future financial privacy computing technology.

Keywords: Federated learning, financial risk control, differential privacy, privacy protection

1. Introduction: Background and Motivation

1.1 Challenges in Traditional Financial Risk Control

In the process of digital transformation, the financial industry is facing increasingly severe risk control challenges and privacy protection issues. With the rapid development of FinTech, a large number of institutions rely on artificial intelligence and big data analysis technologies to optimize key tasks such as credit risk assessment, fraud detection, and asset management. However, the traditional centralized data processing method requires each institution to upload user data to a unified server for modeling, which not only increases the risk of data leakage, but may also be subject to strict supervision due to privacy compliance issues[1]. Regulations such as GDPR and CCPA require financial institutions to ensure user privacy during data processing and avoid unauthorized data sharing.

1.2 Emergence of Federated Learning

Federated Learning (FL), as a distributed machine learning framework, provides a solution that can effectively utilize multi-party data while protecting data privacy. The core goal of financial risk control is to establish a risk prediction model based on user transaction behavior, credit history, assets and liabilities, etc., to identify high-risk transactions and potential fraud. However, financial data is usually scattered among different institutions, such as banks, payment platforms, insurance companies, and stock exchanges. The data island

problem is serious, making it difficult to form a comprehensive risk control system. Federated learning allows multiple data holders to train a shared global model based on local data without sharing the original data.

1.3 Importance of Privacy Technologies in Finance

In addition to risk control applications, federated learning has value in privacy protection. Under traditional models, data cooperation is constrained by trust issues. Federated learning uses technologies such as homomorphic encryption, differential privacy, and secure multi-party computing (MPC) to achieve joint modeling without exposing sensitive data. For example, in cross-bank loan approval and payment fraud detection, federated learning improves prediction accuracy while ensuring that data ownership and security are not violated[2][3].

1.4 Objectives and Contributions of This Study

This study explores the practical application of federated learning in financial risk control and privacy protection. It evaluates the effectiveness of FL models in credit scoring, fraud detection, and transaction risk control. The study also analyzes the applicability of horizontal FL, vertical FL, and federated transfer learning. Results provide practical guidance for financial institutions and references for policymakers, promoting the compliance development of financial technology[4].

2. Related Work

2.1 Deep Learning Techniques in Credit Risk and Fraud Detection

The growing demand for secure and accurate financial risk control has led to significant advancements in the application of deep learning and privacy-preserving computation. Recent research has explored a variety of neural network architectures and data fusion techniques to enhance the performance of financial fraud detection and risk prediction. For instance, Wang [5] introduced a data balancing and ensemble learning strategy for credit card fraud detection, demonstrating improved detection accuracy on imbalanced datasets. Sha et al. [6] proposed a heterogeneous graph neural network integrated with graph attention mechanisms to capture complex relationships in transaction networks, significantly enhancing model interpretability and robustness.

2.2 Anomaly Detection with Generative and Attention Models

To address the challenges of high-frequency data and anomaly detection, several studies have explored advanced generative and attention-based models. Tang et al. [7] applied deep generative models to detect anomalies in complex financial transactions, enabling more sensitive detection of outliers without supervised labels. Similarly, Bao et al. [8] focused on anomaly detection in high-frequency trading environments using deep neural architectures. These efforts highlight the importance of modeling temporal and statistical irregularities in financial datasets.

2.3 Sequence Modeling and Transformer-Based Applications

Transformer-based and hybrid sequence models have gained traction for temporal risk prediction tasks. Wang [9] used a bidirectional transformer to predict premium risk based on time-series data, while Feng [10] introduced a hybrid BiLSTM-Transformer framework for fraudulent transaction detection, demonstrating the advantage of capturing both short- and long-term dependencies in sequential financial data. Du [11] further optimized anomaly detection through a lightweight EfficiencyNet architecture combining separable convolutions with self-attention, reducing computational cost while maintaining accuracy.

2.4 Multimodal Modeling and Systemic Financial Risk

Beyond model architectures, multimodal and systemic risk analysis has also attracted attention. Liu [12] proposed multimodal factor models to forecast stock trends using heterogeneous data sources, and Cheng et al. [13] integrated CNN and BiLSTM for systemic financial risk analysis, offering a comprehensive deep learning framework for macro-level financial prediction. In addition, Du et al. [14] explored structured reasoning using probabilistic models to tackle the issue of data imbalance, which is common in risk control datasets.

2.5 Integration with Blockchain and Reinforcement Learning

Deep learning is also being integrated with blockchain and decentralized financial technologies. Zhou et al. [15] predicted market signals in blockchain-based high-frequency trading using temporal convolutional networks, reflecting the adaptability of deep learning in decentralized financial systems. Meanwhile, Yao [16] proposed a nested reinforcement learning approach to manage dynamic risk in nonlinear markets, highlighting the intersection of decision-making algorithms and financial volatility control.

2.6 Model Interpretability and Regulatory Compliance

Interpretability and transparency remain central concerns in financial AI. Wang et al. [17] conducted a comparative study on credit default prediction models, emphasizing the importance of model interpretability for regulatory compliance. Complementary to this, Du [18] developed a CNN-based approach for intelligent financial statement analysis, addressing anomalies in corporate financial reports.

2.7 Emerging Trends: LLMs and Deep Fusion for Financial Tasks

As financial institutions increasingly rely on machine learning for risk control, large language models and deep fusion techniques have emerged. Gong et al. [19] proposed a deep fusion framework incorporating LLMs for early fraud detection, providing an integrated view of transaction patterns and narrative data. Finally, Wang [20] employed hierarchical multi-source fusion with dropout regularization to enhance fraud detection robustness under noisy data environments.

3. Methodology: Federated Learning Framework for Financial Risk Control

This study uses the Federated Learning (FL) framework to build a financial risk control model to achieve collaborative training between different financial institutions while ensuring data privacy and security. Its abstract architecture is shown in Figure 1.




Figure 1. Abstract diagram of federated learning model

Assume that there are N financial institutions, each institution i has a local data set D_i , and the goal is to train a global model w through federated learning to optimize risk control prediction capabilities. In each round of training, each institution performs gradient updates based on local data,

calculates the local model w_i , and then performs global aggregation through the Federated Averaging (FedAvg) algorithm:

$$w_{t+1} = \sum_{i=1}^N \frac{|D_i|}{\sum_{j=1}^N |D_j|} w_i^t$$

Among them, w_i^t is the global model parameter of round t , and $|D_i|$ represents the number of samples of institution i . This method aggregates local models by weighted average to ensure that institutions with more data contribute more to the global model, thereby improving the stability and generalization ability of the model.

To enhance privacy protection, this study introduces differential privacy (DP) and secure multi-party computation (MPC) technologies. After each round of local training, we add Laplace noise $N(0, \sigma^2)$ to the local gradient $\nabla L(w)$ to protect sensitive data:

$$\nabla' L(w) = \nabla L(w) + N(0, \sigma^2)$$

The noise amplitude σ is controlled by the privacy budget ϵ , making it difficult for attackers to recover the original gradient information. At the same time, homomorphic encryption (HE) is used to encrypt local model parameters before transmission, so that the server does not need to decrypt during aggregation, improving data security. In addition, a secure multi-party computing protocol is used to perform joint computing between institutions, allowing all parties to jointly train the model without exposing their original data, thereby further reducing the risk of privacy leakage[21].

In order to improve the efficiency of federated learning during the training process, we introduced a personalized federated learning (PFL) mechanism to adapt to the differences in data distribution of different financial institutions. Specifically, we used the model distillation method to generate the knowledge distillation target z_i^t through the global model w_i^t , and each local model w_i^t was personalized and optimized under the constraints of the distillation target:

$$L_i = \alpha L_{local}(w_i^t) + (1 - \alpha) KL(z_i^t, f(w_i^t))$$

Among them, L_{local} is the local loss function, $KL(\cdot)$ represents the Kullback-Leibler divergence, which is used to measure the knowledge difference between the local model and the global model, and α is the balance parameter. This mechanism can ensure global knowledge sharing while allowing financial institutions to make personalized adjustments based on their own data characteristics, thereby improving the risk control prediction ability of the model and adapting to the characteristics of different financial markets[22].

4. Experiments and Evaluation

4.1 Dataset Description and Preprocessing

This study uses the FICO Credit Score Dataset as the main dataset, which is provided by Fair Isaac Corporation (FICO) and is widely used in credit risk assessment and financial risk control research. The dataset contains a large amount of credit score-related data from different financial institutions, covering information on multiple dimensions such as personal loans, credit card defaults, and credit history. The core goal of the dataset is to predict whether a user has high-risk default behavior, involving multiple key features such as credit score, debt-to-income ratio, loan balance, and historical overdue number. These features provide rich information for credit assessment and fraud detection of financial institutions, and are typical application scenarios for studying federated learning in financial risk control[23].

The total number of samples in this dataset is about 200,000, and the data comes from different banks and financial service providers. In order to meet the distributed training requirements of federated learning, we divide the data into multiple subsets, each representing a different financial institution. For example, some data belongs to traditional banks, some data comes from credit card companies, and some belongs to online lending platforms. The data of each subset is non-independent and identically distributed (Non-IID), that is, the data distribution of different institutions may be different. For example, the credit scores of users of online lending platforms are usually lower than those of traditional bank customers, while the debt ratio of users of credit card companies may be higher. This difference in data distribution brings challenges to the training of federated learning models, and also provides a good experimental environment for verifying personalized federated learning methods.

In the data preprocessing stage, we performed a series of cleaning and transformations on the original data, including removing missing values, standardizing numerical features, and one-hot encoding categorical features. At the same time, in order to improve data privacy, we use Data Perturbation and Differential Privacy (DP) mechanisms to perturb sensitive data fields (such as credit scores and income levels) to ensure that data privacy is not compromised. In addition, we use the sliding window method to construct time series samples to enhance the model's ability to learn the historical trend of user credit changes. Finally, the dataset is used to evaluate the performance of federated learning models in financial risk control tasks, and to compare the adaptability and effectiveness of different federated learning strategies (such as FedAvg and personalized FL) under different data distribution conditions.

4.2 Experimental Results

This paper first gives the experimental results of the impact of different federated learning architectures (horizontal FL, vertical FL, federated transfer learning) on risk control effects. The experimental results are shown in Figure 2.

From the experimental results, different federated learning architectures have certain differences in performance in financial risk control tasks. Among them, Federated Transfer Learning achieved the highest AUC (0.88) and Precision (0.80), indicating that it can better utilize knowledge transfer and improve the predictive ability and stability of the model when

the data differences across institutions are large. In contrast, Horizontal FL also achieved relatively ideal performance, with an AUC of 0.85 and a Precision of 0.78, indicating that this method can play a good role in collaborative training between institutions with similar data distribution and improve the accuracy of the risk control model[24].




Figure 2. Impact of Different Federated Learning Architectures on Risk Control

The performance of Vertical FL is slightly inferior to the other two methods, with an AUC of only 0.82 and a Precision of 0.75, indicating that this method may be affected by the feature sharing mechanism and data alignment problems in scenarios with different features but overlapping users, resulting in a decrease in model effect compared with other methods. Since financial risk control scenarios usually involve cross-industry and cross-institutional data integration, relying solely on vertical federated learning may not be able to fully utilize the data characteristics of all parties, which in turn affects the final model performance.

Overall, the experimental results show that the applicability of different federated learning architectures in financial risk control depends on the distribution characteristics of the data. For institutions with similar data distribution, horizontal federated learning can achieve better collaborative modeling, while when there are large differences in data distribution, federated transfer learning shows better results. Although vertical federated learning still has application value in some scenarios, its modeling effect is limited by data alignment and feature matching. Therefore, in practical applications, the appropriate federated learning architecture should be selected according to the data characteristics and business needs of financial institutions to improve the accuracy and adaptability of risk control models[25].

Secondly, this paper also gives an analysis of the impact of differential privacy (DP) on the privacy protection and performance of the federated learning risk control model. The experimental results are shown in Figure 3.

From the experimental results, it can be seen that differential privacy (DP) has a significant impact on the performance of the federated learning risk control model. As the DP level increases (from No DP to High DP), both the AUC Score and the Precision Score show a downward trend. This shows that while the DP mechanism protects data privacy,

it also causes a certain loss in model performance. When DP is not applied, the AUC is 0.89 and the Precision is 0.81, indicating that the model can accurately identify risky individuals. However, after the DP level is increased to Medium DP, these two indicators drop to 0.84 and 0.76 respectively, indicating that privacy perturbations have an impact on the model's predictive ability.




Figure 3. Analysis of the impact of differential privacy on the privacy protection and performance of the federated learning risk control model

When the DP level is further increased to High DP, the AUC drops to 0.78 and the Precision drops to 0.70, and the model's prediction accuracy further decreases. This is because stronger DP protection will add more random noise in the gradient update process, making it difficult for the model to learn valuable patterns, thereby reducing its ability to identify risks. Although a high DP level can significantly enhance privacy protection, it comes at the cost of sacrificing the availability of the model, which results in the model being unable to fully utilize data features for risk control prediction.

Overall, this experiment verifies the trade-off between privacy protection and model performance of differential privacy. In financial risk control scenarios, privacy protection is crucial, but overprotection may weaken the risk control capabilities of the model. Therefore, in practical applications, financial institutions need to reasonably select the DP level based on privacy compliance requirements and business needs to find the optimal balance between privacy protection and model performance, thereby ensuring that the risk control system can effectively identify risks while meeting data security and regulatory requirements.

5. Conclusion and Future Directions

This study explores the application of federated learning in financial risk control and privacy protection, and focuses on analyzing the impact of different federated learning architectures (horizontal FL, vertical FL, federated transfer learning) and differential privacy (DP) on model performance and privacy protection. Experimental results show that federated learning can achieve cross-institutional collaborative modeling without leaving the local data, and improve the predictive ability of risk control models. Among them, federated transfer learning performs best in cases where data

distribution is highly different, while horizontal FL also has strong adaptability in scenarios where data distribution is relatively consistent. At the same time, the experiment also reveals that DP may lead to a decline in model performance while protecting data privacy, indicating that it is necessary to balance privacy protection and risk control effects in practical applications.

Further analysis shows that under different data architectures, personalized federated learning can effectively improve the adaptability of the model in a specific institutional environment, while privacy protection technologies such as secure multi-party computing (MPC) and homomorphic encryption (HE) can enhance data security. Although federated learning provides an innovative solution for financial risk control, it still faces many challenges, such as high computing resource consumption, high communication costs, and how to optimize modeling efficiency in high-dimensional feature space. In addition, the introduction of DP needs to be adjusted according to specific business needs to ensure that privacy protection does not excessively weaken the predictive ability of the model. Therefore, in the actual deployment process, it is necessary to select the most appropriate federated learning architecture and reasonably set the privacy protection mechanism according to the data characteristics and business scenarios of different financial institutions.

Future research can further optimize the efficiency and security of federated learning, such as combining methods such as federated distillation and federated attention mechanism to improve the generalization ability of the model, while exploring more efficient encryption computing technology to reduce the computing overhead brought by privacy protection. In addition, the application of federated learning in decentralized finance (DeFi) risk control, cross-border payment security and other fields is also worthy of in-depth study. By continuously optimizing and improving federated learning technology, financial institutions can better balance data privacy protection and risk control capabilities, thereby enhancing the application value of financial technology in scenarios such as intelligent risk control, compliance management and fraud detection.

References

- [1] Johnson, R., & Taylor, B. (2023). Enhancing financial risk control with federated learning: A comparative study. *Journal of Financial Data Science*, 19(3), 101-120.
- [2] Miller, S., & White, C. (2022). Privacy-preserving federated learning for credit risk assessment. *International Conference on Financial AI and Machine Learning (FAIML)*, 278-290.
- [3] Clark, A., & Foster, L. (2021). The impact of differential privacy on federated learning models in financial fraud detection. *Proceedings of the Conference on Privacy-Preserving AI (PPAI)*, 189-204.
- [4] Rodriguez, J., & Kim, H. (2023). Cross-institutional collaboration in financial risk management using federated learning. *Journal of Computational Finance and Security*, 27(4), 132-148.
- [5] Y. Wang, "A Data Balancing and Ensemble Learning Approach for Credit Card Fraud Detection," *arXiv preprint arXiv:2503.21160*, 2025.
- [6] Q. Sha et al., "Detecting Credit Card Fraud via Heterogeneous Graph Neural Networks with Graph Attention," *arXiv preprint arXiv:2504.08183*, 2025.
- [7] T. Tang et al., "Application of Deep Generative Models for Anomaly Detection in Complex Financial Transactions," *arXiv preprint arXiv:2504.15491*, 2025.
- [8] Q. Bao et al., "A Deep Learning Approach to Anomaly Detection in High-Frequency Trading Data," *arXiv preprint arXiv:2504.00287*, 2025.
- [9] Y. Wang, "Time-Series Premium Risk Prediction via Bidirectional Transformer," *Transactions on Computational and Scientific Methods*, vol. 5, no. 2, 2025.
- [10] P. Feng, "Hybrid BiLSTM-Transformer Model for Identifying Fraudulent Transactions in Financial Systems," *Journal of Computer Science and Software Applications*, vol. 5, no. 3, 2025.
- [11] X. Du, "Audit Fraud Detection via EfficiencyNet with Separable Convolution and Self-Attention," *Transactions on Computational and Scientific Methods*, vol. 5, no. 2, 2025.
- [12] J. Liu, "Multimodal Data-Driven Factor Models for Stock Market Forecasting," *Journal of Computer Technology and Software*, vol. 4, no. 2, 2025.
- [13] Y. Cheng et al., "A Deep Learning Framework Integrating CNN and BiLSTM for Financial Systemic Risk Analysis and Prediction," *arXiv preprint arXiv:2502.06847*, 2025.
- [14] J. Du et al., "A Structured Reasoning Framework for Unbalanced Data Classification Using Probabilistic Models," *arXiv preprint arXiv:2502.03386*, 2025.
- [15] T. Zhou, Z. Xu, and J. Du, "Efficient Market Signal Prediction for Blockchain HFT with Temporal Convolutional Networks," *Transactions on Computational and Scientific Methods*, vol. 5, no. 2, 2025.
- [16] Y. Yao, "Time-Series Nested Reinforcement Learning for Dynamic Risk Control in Nonlinear Financial Markets," *Transactions on Computational and Scientific Methods*, vol. 5, no. 1, 2025.
- [17] Y. Wang et al., "Credit Default Prediction with Machine Learning: A Comparative Study and Interpretability Insights," 2024.
- [18] X. Du, "Optimized Convolutional Neural Network for Intelligent Financial Statement Anomaly Detection," *Journal of Computer Technology and Software*, vol. 3, no. 9, 2024.
- [19] J. Gong et al., "A Deep Fusion Framework for Financial Fraud Detection and Early Warning Based on Large Language Models," *Journal of Computer Science and Software Applications*, vol. 4, no. 8, 2024.
- [20] J. Wang, "Credit Card Fraud Detection via Hierarchical Multi-Source Data Fusion and Dropout Regularization," *Transactions on Computational and Scientific Methods*, vol. 5, no. 1, 2025.
- [21] Harrison, D., & Patel, R. (2020). Secure multi-party computation for federated financial analytics. *European Symposium on Financial Technology (ESFT)*, 355-370.
- [22] Nguyen, T., & Roberts, E. (2022). A study on vertical federated learning for multi-bank credit scoring models. *International Workshop on Machine Learning for Finance (MLF)*, 89-102.
- [23] Walker, M., & Evans, J. (2021). Applying homomorphic encryption in federated financial risk assessment. *Journal of Banking and Artificial Intelligence*, 14(2), 212-230.
- [24] Turner, L., & Zhao, Y. (2023). Addressing non-IID data challenges in federated credit risk modeling. *Proceedings of the Global Financial AI Conference (GFAC)*, 175-188.

- [25] Morgan, B., & Suzuki, T. (2020). The role of federated transfer learning in financial fraud detection. International Conference on AI in Finance and Security (AIFS), 92-106.