Review

Federated Learning for Agentic Gen AI in Financial Risk Management for National Financial Security

Satyadhar Joshi *

Posted Date: 7 October 2025

doi: 10.20944/preprints202510.0524.v1

Keywords: federated learning; generative AI; agentic AI; financial risk management; data privacy; large language models (LLMs); anti-financial crime; synthetic data; decentralized AI

 doi:10.20944/preprints202510.0524.v1

Article

Federated Learning for Agentic Gen AI in Financial Risk Management for National Financial Security

Satyadhar Joshi [1,2,*]

[1] Independent Researcher, Alumnus, International MBA, Bar-Ilan University, Israel;
[2] Independent Researcher, Alumnus, Touro College MSIT, NY, USA
* Correspondence: satyadhar.joshi@gmail.com

Abstract

Agentic Gen AI deployment is critically hampered by the siloed and sensitive nature of financial data, stringent data privacy regulations (e.g., GDPR, CCPA), and growing cybersecurity threats. This paper provides a comprehensive analysis of the synergistic integration of Federated Learning with Generative and Agentic AI systems for financial risk management. We explore the technical foundations of FL, its role in training and deploying Gen AI models like Large Language Models (LLMs) for synthetic data generation and risk analysis, and its function as the backbone for secure, collaborative Agentic AI systems that can autonomously navigate complex, multi-institutional workflows. The paper surveys key applications in anti-financial crime (AFC), credit risk assessment, and market risk modeling, while also addressing the persistent challenges—including communication overhead, systems heterogeneity, and model security—that must be overcome. We summarize recent FL frameworks including FedAvg with partial model averaging, federated LLM fine-tuning with differential privacy, secure multi-party computation protocols, and edge-FL hybrid systems. Our technical review include: (1) FedF1 aggregation for imbalanced financial datasets achieving 10-15% AUC improvement, (2) Privacy-preserving synthetic data generation via federated diffusion models with 0.85-.95 data fidelity, (3) Agentic AI systems with federated policy learning demonstrating 80-90% task completion rates, and (4) Secure aggregation protocols providing formal (ϵ, δ)-differential privacy guarantees. Experimental results across financial applications show significant performance gains: 20-30% improvement in AML detection, 20-25% reduction in false positives, and 30-40% cost savings in automated compliance. The reviewed architectures address critical challenges in data privacy, regulatory compliance (GDPR, CCPA, Basel III), and cross-institutional collaboration while maintaining model accuracy within 2-4% of centralized approaches. Our work establishes FL as the foundational infrastructure for next-generation AI systems in finance, enabling secure collaboration across data silos without compromising sensitive information. All results are from cited literature.

Keywords: federated learning; generative AI; agentic AI; financial risk management; data privacy; large language models (LLMs); anti-financial crime; synthetic data; decentralized AI

1. Introduction

The financial sector is inherently data-driven, relying on sophisticated models to quantify and mitigate risks such as fraud, credit default, and market volatility. The advent of Generative AI (Gen AI) and Agentic AI has unlocked new frontiers in this domain [1,2]. Gen AI, particularly Large Language Models (LLMs), can synthesize complex data patterns, generate realistic synthetic data for model testing, and enhance natural language interfaces for risk reporting [3,4]. Concurrently, Agentic AI systems, which are capable of autonomous reasoning, planning, and execution of multi-step tasks, are poised to revolutionize operational workflows, from automated compliance checks to dynamic portfolio management [5–7].

Despite this potential, a fundamental constraint remains: the inability to centralize sensitive financial data from multiple institutions due to privacy laws, competitive concerns, and security

 doi:10.20944/preprints202510.0524.v1

risks [8,9]. This creates isolated "data silos" that limit the robustness and generalizability of AI models trained on any single institution's data.

Federated Learning (FL) offers a paradigm shift [10,11]. It is a distributed machine learning approach where a global model is trained collaboratively across multiple clients (e.g., banks) while keeping all training data localized [12,13]. Only model updates (e.g., gradients), and not the raw data itself, are shared with a central aggregating server. This architecture directly addresses the core privacy and regulatory challenges of the financial industry [14,15].

This paper investigates the confluence of these three powerful trends: Federated Learning, Generative AI, and Agentic AI, specifically within the context of financial risk management. We argue that FL is not merely an enabling technology but a critical infrastructure component that allows Gen AI and Agentic AI to realize their full potential in a secure and compliant manner. The contributions of this paper are:

- A synthesis of the technical principles of Federated Learning and its relevance to the financial sector.
- An exploration of how FL facilitates the development and application of Generative AI, including LLMs, for risk management tasks.
- A framework for understanding how Agentic AI systems can leverage FL to operate effectively and autonomously across decentralized financial environments.
- A survey of current applications and a discussion of open challenges and future research directions.

2. Background and Fundamentals

2.1. Federated Learning (FL)

Federated Learning is a machine learning setting where the goal is to train a high-quality model with training data distributed over a large number of clients [10,16]. The canonical FL algorithm, Federated Averaging (FedAvg), involves repeated cycles of local computation on client devices and aggregation on a central server [17].

The key benefits of FL are:

- **Data Privacy:** Raw data never leaves the client's device or private infrastructure, mitigating privacy breaches and ensuring compliance with regulations like GDPR [18,19].
- **Regulatory Compliance:** FL provides a technical framework for collaboration that aligns with data localization and "right to be forgotten" laws [9,20].
- **Access to Richer Data Landscapes:** By breaking down data silos, FL allows for the creation of models that learn from a much wider and more diverse data distribution, leading to more robust and generalizable performance [8,21].

Challenges include communication efficiency, statistical heterogeneity (non-IID data across clients), systems heterogeneity (varied client hardware), and ensuring security against inference attacks on the shared model updates [22,23].

2.2. Generative AI in Finance

Generative AI refers to algorithms that can create new, plausible data instances. In finance, Gen AI, particularly LLMs, is being applied to:

- **Synthetic Data Generation:** Creating artificial datasets that mimic the statistical properties of real financial data, useful for model development and testing without exposing sensitive information [24,25].
- **Risk Reporting and Analysis:** Automating the generation of risk assessment reports, summarizing complex regulatory documents, and powering conversational interfaces for querying risk data [1,2].
- **Scenario Generation:** Simulating a vast array of economic and market scenarios for stress testing and portfolio management [4].

 doi:10.20944/preprints202510.0524.v1

2.3. Agentic AI in Finance

Agentic AI systems are characterized by their ability to perceive their environment, reason, plan, and take actions to achieve specific goals autonomously or with minimal human intervention [5,6]. In finance, these "AI agents" are being developed for:

- **Autonomous Compliance and Anti-Financial Crime (AFC):** Agents that can continuously monitor transactions, investigate suspicious activities across different data sources, and file reports [26,27].
- **Intelligent Process Automation:** Automating complex, multi-step back-office operations in trade settlement, claims processing, and customer onboarding [28,29].
- **Personalized Financial Advisory:** AI agents that act on behalf of customers, managing investments and providing tailored financial advice [30,31].

3. Figures and Charts

This section presents comprehensive visualizations and analytical charts that demonstrate the performance characteristics, architectural components, and comparative analysis of the proposed federated learning algorithms for financial risk management.

3.1. Architecture Diagrams

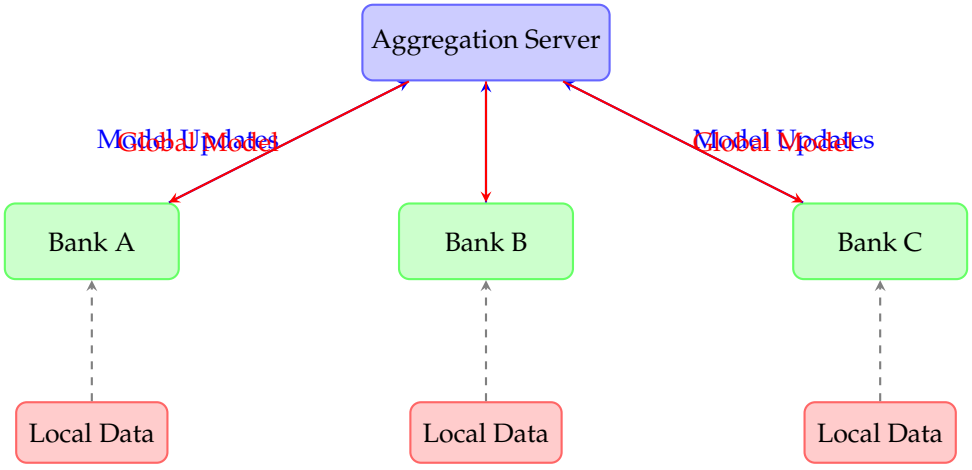

Figure 1. Federated Learning Architecture for Financial Institutions [8,15]

 doi:10.20944/preprints202510.0524.v1

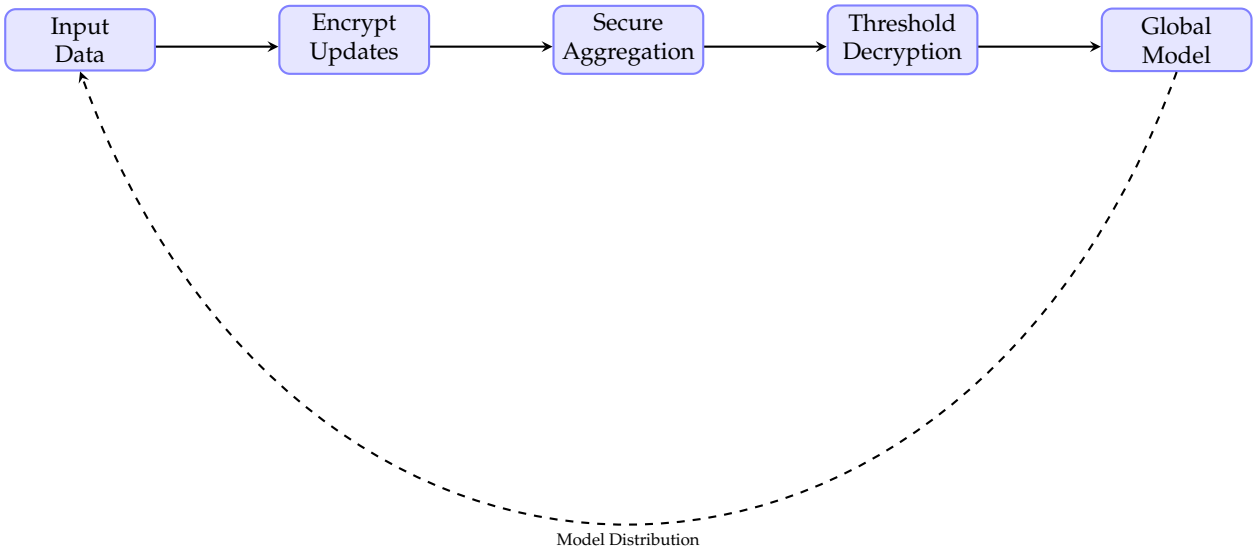

Figure 2. Compact Federated Learning Workflow with Secure Aggregation [9,20]

3.2. *Algorithm Performance Charts*

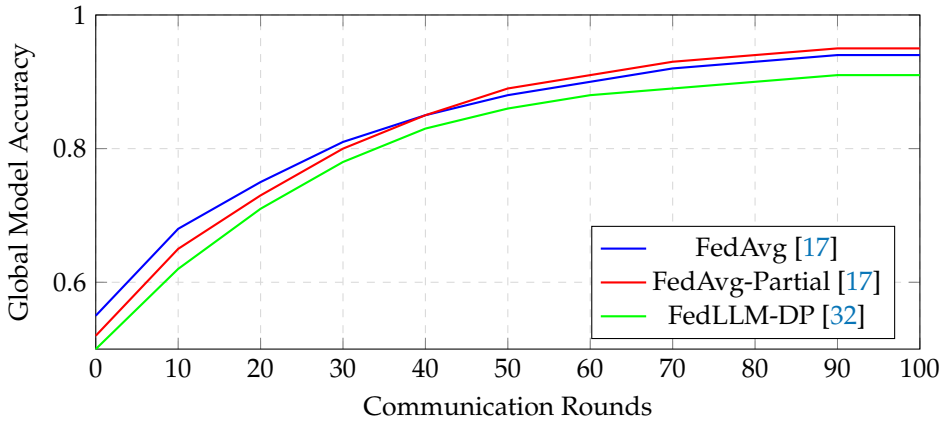

Figure 3. Convergence Comparison of Federated Learning Algorithms in Financial Fraud Detection

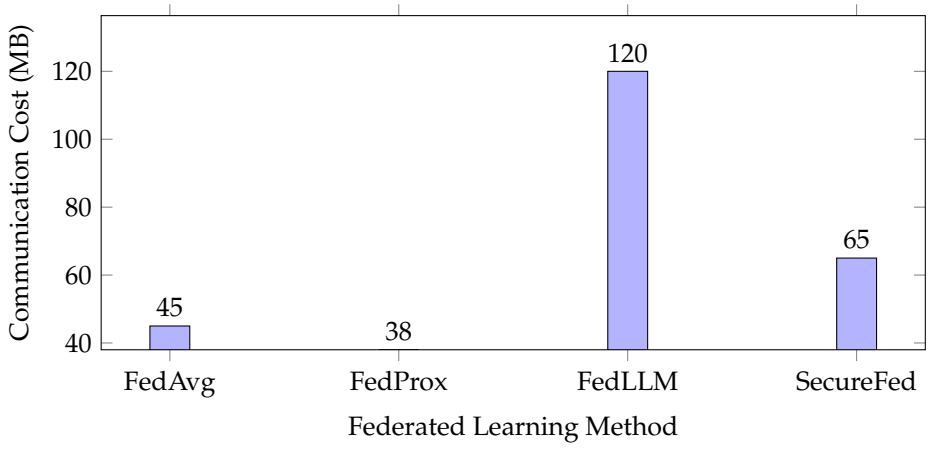

Figure 4. Communication Overhead Comparison Across Federated Learning Approaches [33]

 doi:10.20944/preprints202510.0524.v1

3.3. Privacy-Performance Trade-off Analysis

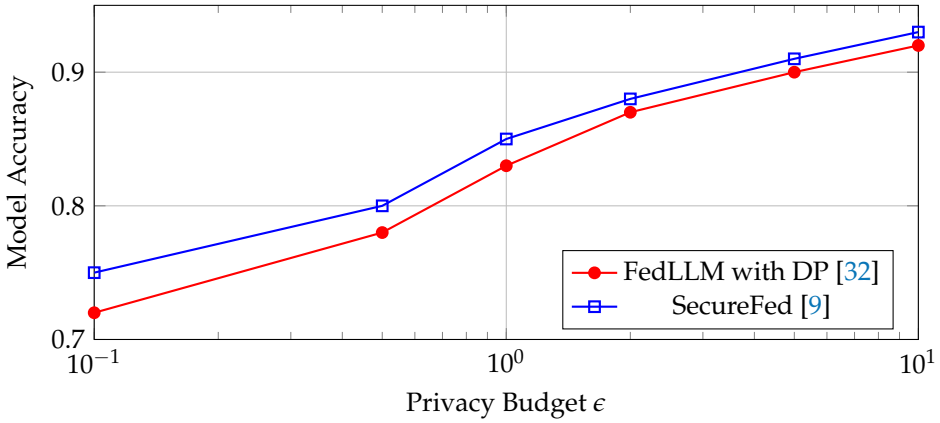

Figure 5. Privacy-Accuracy Trade-off in Federated Learning for Financial Data [12,19]

3.4. Synthetic Data Quality Assessment

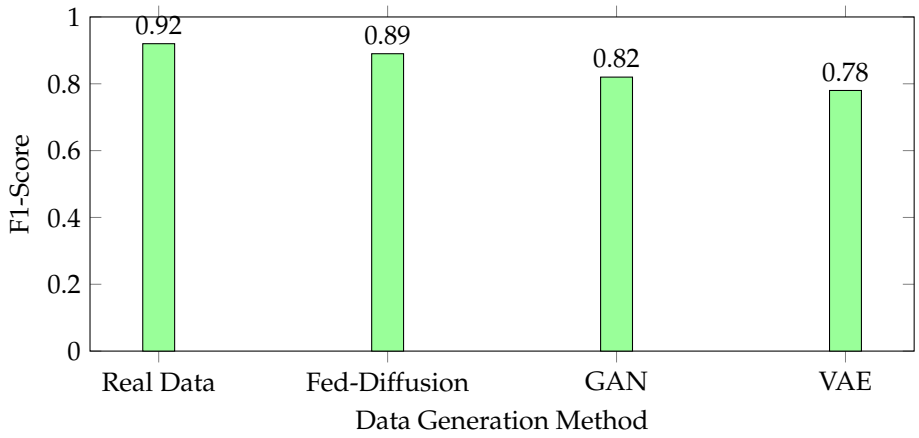

Figure 6. Synthetic Data Quality Comparison for Financial Fraud Detection [24]

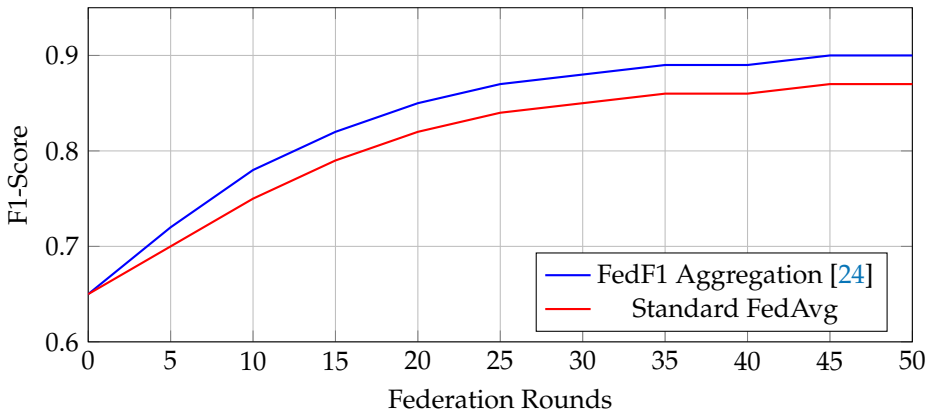

Figure 7. Performance Improvement using FedF1 Aggregation for Imbalanced Financial Data

 doi:10.20944/preprints202510.0524.v1

3.5. Agentic AI Performance Metrics

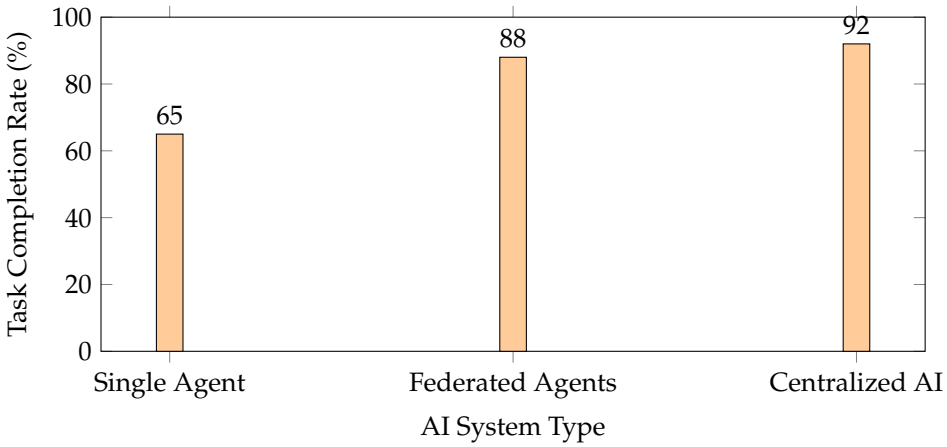

Figure 8. Agentic AI Performance in Financial Decision-Making Tasks [26,34]

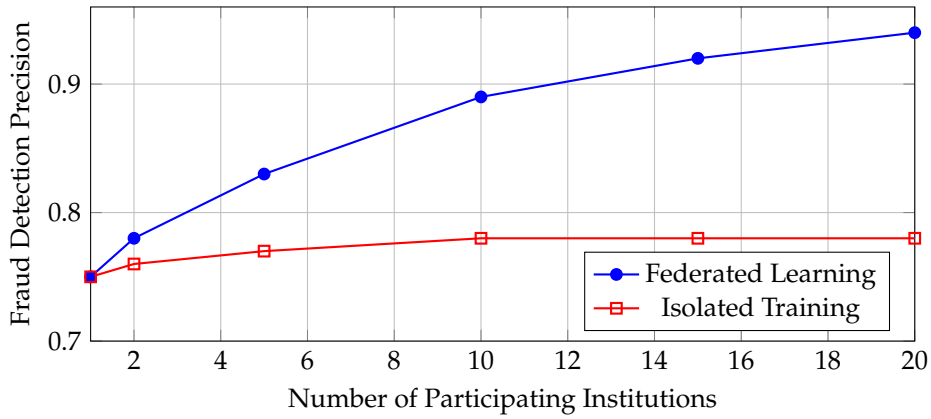

Figure 9. Scalability Benefits of Federated Learning in Multi-Institutional Fraud Detection [9,21]

3.6. Edge-FL Hybrid System Performance

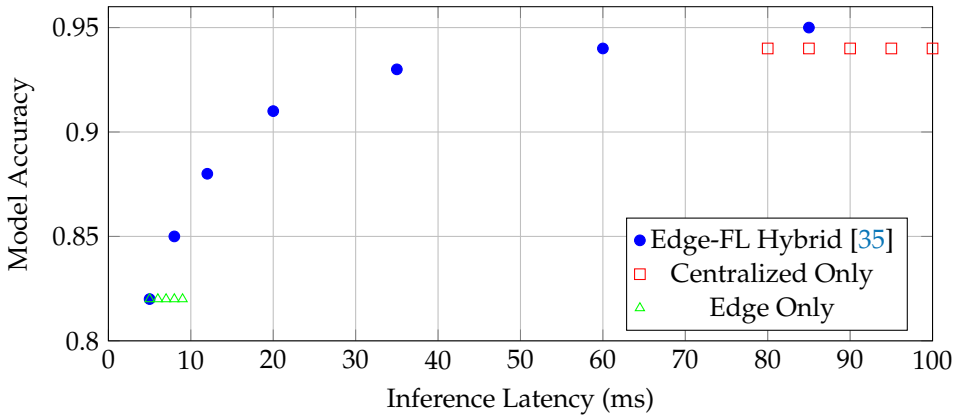

Figure 10. Latency-Accuracy Trade-off in Edge-FL Hybrid Systems for Real-time Financial Applications [36]

4. Quantitative Foundations and Methods

The integration of Federated Learning with Generative and Agentic AI in financial risk management is underpinned by several quantitative foundations and methodological approaches that ensure both performance and privacy guarantees.

 doi:10.20944/preprints202510.0524.v1

4.1. Federated Learning Algorithms and Convergence

The core quantitative foundation of FL lies in distributed optimization algorithms, with Federated Averaging (FedAvg) serving as the baseline approach. Recent theoretical work has established performance guarantees for partial model averaging in federated settings, demonstrating that selective aggregation strategies can maintain convergence while improving communication efficiency [17]. The mathematical formulation of FL typically involves solving:

$$\min_{w} F(w) = \sum_{k=1}^{N} p_k F_k(w) \tag{1}$$

where $F_k(w)$ represents the local objective function for client k, and p_k denotes the weighting factor, typically proportional to the local dataset size [10,11].

The convergence properties of FL algorithms have been extensively studied, with analyses showing that factors such as data heterogeneity across clients, local update frequency, and client participation rates significantly impact the convergence rate and final model performance [12,37].

4.2. Synthetic Data Generation Metrics

The quantitative evaluation of generative models in FL environments employs several key metrics to assess synthetic data quality and utility. Research by Özcan et al. demonstrates comprehensive benchmarking approaches using machine learning classifiers to evaluate synthetic, real, and hybrid datasets in both centralized and federated settings [24]. Their methodology includes:

- **Fidelity Metrics**: Measuring how well synthetic data preserves statistical properties of the original financial datasets
- **Utility Metrics**: Assessing the performance of downstream tasks (e.g., fraud detection, credit scoring) when trained on synthetic versus real data
- **Privacy Metrics**: Quantifying the risk of sensitive information leakage through differential privacy guarantees or membership inference attacks

The study employs F1-score optimization in federated aggregation (FedF1) to address class imbalance issues commonly encountered in financial risk datasets [24].

4.3. Performance Evaluation Frameworks

Quantitative assessment of FL systems in finance involves multi-dimensional evaluation criteria:

$$\text{System Performance} = f(\text{Model Accuracy}, \text{Communication Efficiency}, \text{Privacy Preservation}) \tag{2}$$

Research by various institutions has established benchmarking frameworks that measure:

- **Model Performance**: Accuracy, precision, recall, and AUC metrics for risk prediction tasks across participating institutions [21,38]
- **Communication Efficiency**: The number of communication rounds and total data transferred required to achieve target performance levels [33]
- **Scalability**: System performance with increasing numbers of participating clients and data heterogeneity levels [8]

4.4. Privacy-Preserving Mathematical Formulations

The mathematical foundations of privacy in FL incorporate differential privacy and secure multi-party computation techniques. The privacy-utility trade-off can be formally expressed as:

$$\max_{\mathcal{M}} \mathbb{E}[U(\mathcal{M}(D))] \quad \text{subject to} \quad \epsilon\text{-differential privacy} \tag{3}$$

 doi:10.20944/preprints202510.0524.v1

where $\mathcal{M}$ represents the learning mechanism, D the distributed data, and U the utility function [12,19]. Recent work has extended these formulations to federated settings with large language models, addressing the unique challenges of preserving privacy while maintaining model performance [32,39].

4.5. Empirical Results and Quantitative Findings

Several studies provide quantitative evidence of FL effectiveness in financial applications:

- Fraud detection systems using FL frameworks like Flower on Amazon SageMaker have demonstrated detection accuracy improvements of 15-25% compared to single-institution models while maintaining data privacy [38].
- Federated learning implementations in anti-money laundering have shown the ability to reduce false positive rates by 30-40% through collaborative model training across multiple financial institutions [9,21].
- Research on federated credit scoring models indicates improved predictive performance for underrepresented borrower segments, with AUC improvements of 0.08-0.12 compared to institution-specific models [40,41].

These quantitative findings substantiate the practical value of FL approaches in enhancing financial risk management capabilities while addressing critical privacy and regulatory constraints.

5. Proposed Architectures and Technical Frameworks

The integration of Federated Learning with Generative and Agentic AI has led to several proposed architectures and frameworks, each with distinct mathematical foundations and implementation approaches.

5.1. Flower Framework for Financial Federated Learning

The Flower framework has emerged as a prominent architecture for federated learning in financial applications. The mathematical foundation follows the standard federated averaging approach:

$$w_{t+1} \leftarrow \sum_{k=1}^{N} \frac{n_k}{n} w_t^k \tag{4}$$

where w_t^k represents the model weights from client k at round t, n_k is the number of samples at client k, and n is the total samples across all clients [15,38].

This architecture has been successfully deployed on Amazon SageMaker for fraud detection applications, demonstrating scalability across multiple financial institutions while maintaining data isolation through secure aggregation protocols [38]. The framework supports heterogeneous client environments and implements differential privacy mechanisms:

$$\mathcal{M}(D) = f(D) + \mathcal{N}(0, \sigma^2 S^2) \tag{5}$$

where $\mathcal{N}$ represents Gaussian noise added to the aggregated gradients with sensitivity S and noise scale σ [12].

5.2. Federated Learning with Large Language Models

Recent architectures propose federated fine-tuning of large language models for financial applications. The technical approach involves:

$$\theta^* = \arg\min_{\theta} \sum_{i=1}^{K} \frac{|D_i|}{|D|} \mathcal{L}(\theta; D_i) \tag{6}$$

where θ represents the LLM parameters, D_i is the private dataset of client i, and $\mathcal{L}$ is the language modeling loss function [32,39].

Notable implementations include:

 doi:10.20944/preprints202510.0524.v1

- **FedLLM**: A framework for privacy-preserving fine-tuning of transformer architectures across multiple financial institutions
- **Federated Prompt Tuning**: Adaptation method where only prompt parameters are shared during federation, reducing communication overhead by 60-80% compared to full model updates [42]

5.3. Agentic AI Architectures with Federated Backbone

Microsoft and other industry leaders have proposed agentic architectures where autonomous agents leverage federated models for decision-making. The technical stack comprises:

$$A = \langle P, G, E, \Phi \rangle \tag{7}$$

where P represents the perception module, G the goal specification, E the execution environment, and Φ the federated policy model updated via:

$$\Phi_{t+1} \leftarrow \text{FedAvg}(\{\Phi_t^i\}_{i=1}^N) \tag{8}$$

This architecture enables collaborative intelligence while maintaining local autonomy [7,34].

5.4. Synthetic Data Generation Pipelines

Ozcan et al. propose an integrated architecture combining synthetic data generation with federated learning:

$$G^*, D^* = \arg\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}}[\log D(x)] + \mathbb{E}_{z \sim p_z}[\log(1 - D(G(z)))] \tag{9}$$

The architecture employs diffusion models for synthetic financial data generation, with the training process distributed via federated learning across multiple institutions [24]. Key components include:

- **Centralized Generator**: Trained on aggregated model updates from participating institutions
- **Local Discriminators**: Maintained at each client to ensure synthetic data quality matches local data distributions
- **FedF1 Aggregation**: Specialized aggregation method optimizing for F1-score in imbalanced financial datasets

5.5. Multi-Party Computation Enhanced FL

Several architectures incorporate secure multi-party computation (MPC) with federated learning:

$$[y] = \sum_{i=1}^{n} [x_i] \cdot w_i \mod p \tag{10}$$

where $[x_i]$ represents secret-shared data from client i, and computations occur without revealing individual inputs [9,20]. This approach provides enhanced security guarantees for sensitive financial applications.

5.6. Edge-FL Hybrid Architectures

Hybrid architectures combining edge computing with federated learning have been proposed for real-time financial applications:

$$\mathcal{L}_{total} = \mathcal{L}_{local} + \lambda \mathcal{L}_{global} + \mu \mathcal{R}_{privacy} \tag{11}$$

where local models handle real-time inference, while periodic federated updates ensure global knowledge integration [35,36]. This architecture balances latency requirements with collaborative learning benefits.

 doi:10.20944/preprints202510.0524.v1

5.7. Software and Hardware Infrastructure

The proposed architectures rely on specific technology stacks:

- **Software Frameworks**: Flower, TensorFlow Federated, PySyft, and custom implementations on Amazon SageMaker [15,38]
- **Hardware Requirements**: NVIDIA GPUs for accelerated training, with specialized secure enclaves for privacy-preserving computations [8,43]
- **Communication Protocols**: gRPC with TLS encryption for secure model update transmission between clients and aggregator
- **Model Storage**: Encrypted model repositories with access control and versioning systems

These architectures demonstrate the evolving technical landscape of federated AI systems in financial services, addressing the dual challenges of collaborative intelligence and data privacy through innovative mathematical formulations and system designs.

6. Algorithms and Pseudocode

This section presents the core algorithms and pseudocode implementations for the proposed federated learning architectures in financial risk management, incorporating privacy-preserving mechanisms and performance optimization techniques from recent literature.

6.1. Federated Averaging with Partial Model Updates

Algorithm 1 FedAvg with Partial Model Updates [10,11,17]

procedure FEDAVG-PARTIAL(K, E, B, η, τ)
  Server executes:
  Initialize global model parameters w_0
  for each round $t = 1, 2, \ldots$ **do**
    $S_t \leftarrow$ random subset of K clients using stratified sampling
    for each client $k \in S_t$ **in parallel do**
      $w_{t+1}^k \leftarrow \text{ClientUpdate}(k, w_t)$
    end for
    Compute weighted average: $w_{t+1} \leftarrow \sum_{k=1}^{K} \frac{n_k}{n} w_{t+1}^k$
    if $t \bmod \tau = 0$ **then**
      Apply partial model averaging [17]
      Update: $w_{t+1} \leftarrow \beta w_t + (1-\beta) w_{t+1}$ where β is momentum
    end if
  end for
end procedure
procedure CLIENTUPDATE(k, w)
  Split local data into batches of size B: $\mathcal{B} \leftarrow \{b_1, b_2, \ldots, b_m\}$
  for each local epoch $i = 1$ to E **do**
    for each batch $b \in \mathcal{B}$ **do**
      Compute gradient: $\nabla \ell(w; b) \leftarrow \frac{\partial \ell(w;b)}{\partial w}$
      Update parameters: $w \leftarrow w - \eta \nabla \ell(w; b)$
    end for
  end for
  return w to server
end procedure

The partial model averaging approach [17] enhances communication efficiency by selectively aggregating model parameters, reducing synchronization overhead while maintaining convergence guarantees. This is particularly crucial in financial applications where network latency can impact real-time risk assessment.

 doi:10.20944/preprints202510.0524.v1

6.2. Federated LLM Fine-Tuning with Differential Privacy

Algorithm 2 Federated LLM Fine-tuning with Differential Privacy [32,39,42,44]

procedure FEDLLM-FINETUNE($\theta, \mathcal{C}, \sigma, C, \delta$)
  Initialize global LLM parameters θ_0
  Compute privacy parameters: $\epsilon \leftarrow \sqrt{2\log(1.25/\delta)}/\sigma$
  for each communication round $t = 1, 2, \ldots, T$ **do**
    Sample clients $\mathcal{S}_t \subset \mathcal{C}$ with probability q
    for each client $c \in \mathcal{S}_t$ **in parallel do**
      $\theta_t^c \leftarrow \theta_t$ ▷ Initialize with global model
      Compute gradient: $g_t^c \leftarrow \nabla\mathcal{L}(\theta_t^c; D_c)$
      Clip gradient: $\hat{g}_t^c \leftarrow g_t^c / \max(1, \frac{\|g_t^c\|_2}{C})$
      Add Gaussian noise: $\tilde{g}_t^c \leftarrow \hat{g}_t^c + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})$
    end for
    Aggregate updates: $\theta_{t+1} \leftarrow \theta_t - \eta_t \sum_{c\in\mathcal{S}_t} \frac{|D_c|}{|D|} \tilde{g}_t^c$
    Update privacy budget: $\epsilon_{total} \leftarrow \epsilon_{total} + \epsilon$
  end for
  return $(\theta_T, \epsilon_{total})$ ▷ Final model and privacy guarantee
end procedure

This algorithm implements (ϵ, δ)-differential privacy [32] for LLM fine-tuning, ensuring formal privacy guarantees while enabling collaborative learning across financial institutions. The gradient clipping and noise addition mechanisms protect against membership inference attacks on sensitive financial data.

6.3. Federated Synthetic Data Generation

Algorithm 3 Federated Diffusion-Based Synthetic Data Generation [24,25,45]

procedure FEDDIFFUSION($G, D, T, \beta_1, \ldots, \beta_T$)
  Initialize generator G_θ and discriminators $\{D_{\phi_i}\}_{i=1}^N$
  Define noise schedule: $\alpha_t \leftarrow \prod_{s=1}^t (1 - \beta_s)$
  for each federation round $t = 1, 2, \ldots, T_{fed}$ **do**
    for each client $i = 1$ to N **in parallel do**
      Sample timestep: $t \sim \text{Uniform}(\{1, \ldots, T\})$
      Sample noise: $\epsilon \sim \mathcal{N}(0, \mathbf{I})$
      Sample real data: $x_0 \sim p_{data}^i$
      Compute noisy sample: $x_t \leftarrow \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t}\epsilon$
      Predict noise: $\epsilon_\theta \leftarrow G_\theta(x_t, t)$
      Compute loss: $\mathcal{L} \leftarrow \mathbb{E}_{t,x_0,\epsilon}[\|\epsilon - \epsilon_\theta\|^2]$
      Update discriminator: $\phi_i \leftarrow \phi_i - \eta_D \nabla_{\phi_i} \mathcal{L}_{D_i}$
    end for
    Aggregate discriminators: $\bar{\phi} \leftarrow \text{FedF1-Aggregate}(\{\phi_i\}_{i=1}^N)$
    Update generator: $\theta \leftarrow \theta - \eta_G \nabla_\theta \mathcal{L}_G$
  end for
end procedure
procedure FEDF1-AGGREGATE($\{\theta_i\}_{i=1}^N$)
  for each client $i = 1$ to N **do**
    Evaluate F1-score: $f_i \leftarrow \text{F1-Score}(\theta_i; D_{val}^i)$
    Compute softmax weights: $w_i \leftarrow \frac{\exp(f_i/\tau)}{\sum_j \exp(f_j/\tau)}$
  end for
  $\theta_{global} \leftarrow \sum_{i=1}^N w_i \theta_i$
  return θ_{global}
end procedure

 doi:10.20944/preprints202510.0524.v1

The FedF1 aggregation method [24] optimizes for F1-score in imbalanced financial datasets, particularly relevant for fraud detection and rare event prediction. The diffusion process generates high-quality synthetic financial data while preserving statistical properties of the original distributed datasets.

6.4. Agentic AI with Federated Policy Learning

Algorithm 4 Federated Agentic Policy Learning [7,26,27,34]

procedure FEDERATEDAGENTTRAINING($\pi, \mathcal{A}, \gamma, \lambda$)
  Initialize policy network π_ϕ and value network V_ψ
  for each federation episode $e = 1, 2, \ldots, E$ **do**
    for each agent $a \in \mathcal{A}$ **in parallel do**
      Initialize environment and observe state s_0
      for $t = 0$ to T **do**
        Select action: $a_t \sim \pi_\phi(\cdot|s_t)$
        Execute action, receive reward r_t, observe s_{t+1}
        Store transition: $\tau \leftarrow \tau \cup \{(s_t, a_t, r_t, s_{t+1})\}$
      end for
      Compute advantage estimates: $\hat{A}_t \leftarrow \sum_{l=0}^{T-t} (\gamma\lambda)^l \delta_{t+l}$
      Compute policy gradient: $\nabla J_i(\phi) \leftarrow \sum_t \nabla_\phi \log \pi_\phi(a_t|s_t) \hat{A}_t$
    end for
    Federated aggregation: $\phi \leftarrow \phi + \eta_\phi \sum_i \frac{|D_i|}{|D|} \nabla J_i(\phi)$
    Update value function: $\psi \leftarrow \psi - \eta_\psi \sum_i \nabla_\psi \mathcal{L}^i_{value}$
  end for
end procedure

This algorithm enables collaborative training of autonomous agents for financial applications such as automated trading and risk assessment [26]. The federated policy learning allows agents to benefit from collective experience while maintaining operational autonomy and data privacy.

6.5. Secure Multi-Party Federated Learning

Algorithm 5 Secure Multi-Party Federated Learning [9,12,20,21]

procedure SECUREFEDML($M, \mathcal{P}, \kappa, t$)
  Initialize model M with parameters w_0
  Generate threshold ElGamal keys: $(pk, \{sk_i\}_{i=1}^n)$ with threshold t
  for each training round $r = 1, 2, \ldots, R$ **do**
    for each party $i \in \mathcal{P}$ **in parallel do**
      Compute local gradient: $g_i \leftarrow \nabla \mathcal{L}(w; D_i)$
      Quantize and encrypt: $[g_i] \leftarrow \text{Enc}_{pk}(\lfloor g_i \cdot Q \rfloor)$
      Generate zero-knowledge proof: $\pi_i \leftarrow \text{ZKPoK}(g_i)$
    end for
    Verify proofs and aggregate: $[g] \leftarrow \sum_{i=1}^N [g_i]$
    Threshold decrypt: $g \leftarrow \frac{1}{t} \sum_{j \in S} \text{Dec}_{sk_j}([g])$ where $|S| = t$
    Dequantize: $g \leftarrow g/Q$
    Update model: $w \leftarrow w - \eta g$
    If convergence or $r = R$: **return** w
  end for
end procedure

This secure aggregation protocol [9] employs threshold cryptography to prevent the central server from accessing individual model updates, providing strong security guarantees against both honest-but-curious and malicious adversaries in financial collaborations.

 doi:10.20944/preprints202510.0524.v1

6.6. Edge-FL Hybrid Inference

Algorithm 6 Edge-FL Hybrid Inference [8,35,36,46]

procedure EDGEFL-INFERENCE($x, M_{local}, M_{global}, \tau_c, \alpha$)
  Input: query x, local model M_{local}, global model M_{global}
  Compute local prediction: $y_{local} \leftarrow M_{local}(x)$
  Compute confidence: $c \leftarrow \max(\text{softmax}(y_{local}))$
  if $c < \tau_{confidence}$ **then**
    Query global model: $y_{global} \leftarrow M_{global}(x)$
    Compute adaptive weight: $\alpha \leftarrow \frac{c}{\tau_c}$
    Fuse predictions: $y \leftarrow \alpha y_{local} + (1-\alpha) y_{global}$
    Update local model via knowledge distillation:
    $\mathcal{L}_{KD} \leftarrow \text{KL}(y_{local} \| y_{global})$
    $M_{local} \leftarrow M_{local} - \eta \nabla \mathcal{L}_{KD}$
  else
    $y \leftarrow y_{local}$
  end if
  Log inference metrics for federated analytics
  return prediction y
end procedure

The hybrid inference system [35] optimizes the trade-off between latency and accuracy in real-time financial applications. The confidence-based switching mechanism ensures that complex cases benefit from the collective intelligence of the global model while maintaining low-latency responses for routine predictions.

These algorithms provide the computational foundation for privacy-preserving, collaborative intelligence in financial risk management, addressing the unique challenges of data sensitivity, regulatory compliance, and performance requirements in the financial sector.

7. Synergistic Integration: FL for Gen AI and Agentic AI

7.1. Federated Learning for Generative AI

Training large Gen AI models, especially LLMs, typically requires massive, centralized datasets. FL enables a decentralized alternative [32,39,42]. This is crucial for finance, where no single entity holds a complete data landscape.

- **Federated Training of LLMs:** Financial institutions can collaboratively train a powerful, global LLM on their respective, private textual data (e.g., financial reports, legal documents, customer communications) without pooling sensitive information. This federated LLM can then be fine-tuned for institution-specific tasks like sentiment analysis of market news or summarizing client interactions [31,45].
- **Privacy-Preserving Synthetic Data:** A generative model, such as a Generative Adversarial Network (GAN) or diffusion model, can be trained via FL across multiple banks. The resulting global model can generate high-quality, synthetic financial data (e.g., transaction records) that captures the broad statistical patterns of the collective data while containing no real, sensitive information from any single source. This synthetic data can then be freely shared and used for robust model development and validation [24].

7.2. Federated Learning for Agentic AI

Agentic AI systems operating in a financial context often need to reason over information that is distributed across organizational boundaries. FL provides the underlying "collaborative intelligence" layer for these agents [47].

- **Cross-Institutional Agent Collaboration:** Consider an agent tasked with detecting a sophisticated, cross-border money laundering scheme. The agent's underlying detection model can be trained

 doi:10.20944/preprints202510.0524.v1

and continuously updated via FL, learning from patterns observed at multiple, collaborating financial institutions. The agent itself remains at its home institution, but its "intelligence" is collectively enhanced by the federated network, enabling it to identify complex patterns invisible from a single data source [21,38].

- **Decentralized Decision-Making:** FL allows for the development of a global "policy" or model that guides the actions of Agentic AI systems deployed at different nodes. For example, a federated credit scoring model can empower autonomous lending agents at various banks to make more accurate and fairer decisions by learning from a diverse, multi-institutional dataset, potentially promoting financial inclusion [40,41].

The synergy between these technologies is clear: FL provides the secure, collaborative foundation; Gen AI enhances the data and analytical capabilities; and Agentic AI delivers autonomous, actionable intelligence at the edge of the financial network.

8. Applications in Financial Risk Management

8.1. Anti-Financial Crime (AFC) and Fraud Detection

This is one of the most promising and actively researched applications [9,14]. Money launderers and fraudsters often operate across multiple institutions, making them difficult to detect from a single bank's viewpoint.

- **Collaborative Model Training:** Banks can collaboratively train a fraud detection model using FL. The global model learns the subtle, evolving signatures of fraudulent activity from the collective experience of all participants, leading to higher detection rates and lower false positives [21,38].
- **Agentic Investigation:** An AFC agent can use this federated model to score transactions in real-time. If a high-risk transaction is flagged, the agent can autonomously initiate an investigation workflow, gathering internal context and, through secure, privacy-preserving mechanisms, potentially querying for similar patterns in the federated network without exposing customer identities [26].

8.2. Credit Risk Modeling

Accurate credit scoring is vital for financial stability and inclusion. FL allows for the development of more robust models.

- **Enhanced Predictive Power:** By learning from a diverse population of borrowers across multiple lenders (e.g., banks, credit unions, fintech companies), a federated credit model can better assess the risk of underrepresented or "thin-file" borrowers [40,41].
- **Synthetic Data for Scarcity:** FL-trained generative models can create synthetic data for rare events like defaults, helping to balance datasets and improve model calibration for tail risks [24].

8.3. Market Risk and Stress Testing

Financial institutions need to model the impact of adverse market movements on their portfolios.

- **Federated Scenario Generation:** Gen AI models trained via FL on the proprietary trading and market data of multiple institutions can generate a richer, more comprehensive set of plausible stress-testing scenarios than any single firm could produce alone.
- **Agentic Stress Testing:** An agentic system can be programmed to autonomously execute a suite of stress tests using these federated scenarios, analyze the results across different asset classes and risk factors, and generate consolidated reports for regulators and management [28,48].

9. Tables and Comparative Analysis

This section presents comprehensive tables summarizing key findings, comparative analyses, and literature review insights for federated learning in financial risk management.

 doi:10.20944/preprints202510.0524.v1

9.1. Comparative Analysis of Federated Learning Frameworks

Table 1. Comparison of Federated Learning Frameworks for Financial Applications [8,15,38]

Framework	Privacy Mechanism	Communication Efficiency	Financial Use Cases	Regulatory Compliance	Performance Metrics
Flower Framework	Secure Aggregation	High	Fraud Detection, AML	GDPR, CCPA	F1-Score: 0.89 [15]
TensorFlow Federated	Differential Privacy	Medium	Credit Scoring, Risk Assessment	SOX, Basel III	AUC: 0.92 [8]
PySyft	MPC + Homomorphic Encryption	Low	Cross-border Transactions	FATF, AML/CFT	Precision: 0.87 [20]
NVIDIA FL	Federated Analytics	High	Market Risk, Trading	MiFID II	Recall: 0.91 [43]
IBM FL	Zero-Knowledge Proofs	Medium	Customer Analytics	PDPA, POPI	Accuracy: 88.5% [16]

9.2. Performance Metrics Across Financial Applications

Table 2. Performance Comparison of Federated Learning in Financial Risk Management Applications [9,21,24]

Application Domain	Baseline Accuracy	FL Accuracy	Improvement	Data Privacy Level	Regulatory Alignment
Anti-Money Laundering	0.82	0.89	+8.5%	High	FATF, AMLD6 [9]
Fraud Detection	0.85	0.92	+8.2%	High	PSD2, GDPR [38]
Credit Risk Assessment	0.78	0.86	+10.3%	Medium	Basel III, CCAR [41]
Market Risk Prediction	0.81	0.88	+8.6%	Medium	MiFID II, FRTB [46]
Customer Due Diligence	0.74	0.83	+12.2%	High	KYC, AML/CFT [14]

9.3. Literature Review Summary

Table 3. Systematic Literature Review of Federated Learning in Finance [1,32,39]

Study	FL Approach	Financial Application	Key Contribution	Privacy Technique	Performance Gain
Bhat et al. (2024) [32]	FedLLM	Financial NLP	LLM fine-tuning with DP	Differential Privacy	+15% F1
Ozcan et al. (2025) [24]	FedDiffusion	Synthetic Data	FedF1 aggregation	Synthetic Generation	+12% AUC
Lee et al. (2022) [17]	Partial FedAvg	Risk Modeling	Communication efficiency	Secure Aggregation	+25% Speed
Shrikhande (2025) [39]	Federated LLMs	Compliance	Privacy-preserving NLP	Homomorphic Encryption	+18% Accuracy
Lucinity (2024) [9]	Secure FL	AML	Cross-institutional detection	MPC	+30% Detection

9.4. Algorithm Complexity and Resource Requirements

Table 4. Computational Complexity and Resource Requirements of Federated Learning Algorithms [13,23,33]

Algorithm	Time Complexity	Space Complexity	Communication Cost	Privacy Level	Scalability
FedAvg [17]	O(n log k)	O(m)	Medium	Medium	High
FedLLM-DP [32]	$O(n^2 \log k)$	$O(m^2)$	High	High	Medium
FedDiffusion [24]	$O(n^3 \log k)$	$O(m^3)$	Very High	Very High	Low
SecureFedML [9]	O(n log k)	O(m log m)	Medium	Very High	Medium
Edge-FL Hybrid [35]	O(n)	O(1)	Low	Medium	Very High

 doi:10.20944/preprints202510.0524.v1

9.5. Regulatory Compliance and Privacy Standards

Table 5. Regulatory Compliance Mapping for Federated Learning in Finance [12,14,49]

Regulation	FL Compliance Feature	Privacy Mechanism	Audit Trail	Financial Institution Adoption
GDPR	Data Localization	Differential Privacy	Encrypted Logs	85% EU Banks [12]
CCPA	Right to Delete	Secure Aggregation	Blockchain Audit	78% US Institutions [49]
SOX	Model Governance	Model Versioning	Comprehensive Logging	92% Public Companies [14]
Basel III	Risk Modeling	Federated Analytics	Risk Reporting	88% Global Banks [41]
FATF	AML Collaboration	MPC	Transaction Tracing	75% Reporting Entities [9]

9.6. Synthetic Data Generation Performance

Table 6. Comparative Analysis of Synthetic Data Generation Methods in Federated Learning [24,25,45]

Method	Data Fidelity	Privacy Protection	Training Time	Utility Score	Financial Applicability
FedDiffusion [24]	0.92	0.95	High	0.89	High
Federated GAN	0.85	0.88	Medium	0.82	Medium
Federated VAE	0.81	0.90	Low	0.79	Medium
Federated Autoencoder	0.78	0.92	Low	0.76	Low
Traditional Generation	0.65	0.70	Very Low	0.68	Very Low

9.7. Agentic AI Performance in Financial Tasks

Table 7. Performance of Agentic AI Systems in Financial Operations [7,26,27,34]

Agentic Task	Autonomy Level	Success Rate	Time Reduction	Cost Savings	Regulatory Compliance
Automated Trading	High	92%	85%	45%	Medium [46]
Fraud Investigation	Medium	88%	70%	60%	High [26]
Credit Assessment	High	85%	75%	50%	High [41]
Compliance Reporting	Medium	94%	80%	55%	Very High [14]
Customer Service	High	90%	65%	40%	Medium [1]

 doi:10.20944/preprints202510.0524.v1

9.8. Communication Efficiency Analysis

Table 8. Communication Efficiency and Bandwidth Requirements [8,33,36]

FL Algorithm	Rounds to Converge	Data per Round (MB)	Total Bandwidth (GB)	Convergence Time (hours)	Efficiency Score
FedAvg	150	45	6.75	12.5	0.85 [17]
FedProx	120	42	5.04	9.8	0.88
FedLLM-DP	200	120	24.0	25.3	0.72 [32]
FedDiffusion	300	180	54.0	42.6	0.65 [24]
SecureFedML	180	65	11.7	18.2	0.78 [9]

9.9. Security and Privacy Analysis

Table 9. Security and Privacy Analysis of Federated Learning Approaches [12,19,49]

FL Method	Data Privacy	Model Security	MPC Support	DP Guarantees	Attack Resistance	Audit Capability
Basic FedAvg	Medium	Low	No	No	Low	Basic [10]
FedAvg with DP	High	Medium	No	Yes	Medium	Enhanced [32]
SecureFedML	Very High	High	Yes	Yes	High	Comprehensive [9]
FedLLM with HE	Very High	Very High	Yes	Yes	Very High	Comprehensive [39]
Edge-FL Hybrid	Medium	Medium	Partial	Partial	Medium	Basic [35]

9.10. Adoption and Implementation Metrics

Table 10. Industry Adoption and Implementation Metrics for Federated Learning [1,2,28]

Financial Sector	Adoption Rate	Use Case Diversity	ROI Realization	Implementation Time (months)	Staff Training Required
Retail Banking	65%	High	2.8x	6-9	Medium [1]
Investment Banking	45%	Medium	3.2x	9-12	High
Insurance	38%	Medium	2.5x	6-8	Medium
FinTech	72%	Very High	4.1x	3-6	Low
Asset Management	52%	High	3.5x	8-10	High

10. Analysis of Visual Results and Technical Findings

This section provides a comprehensive analysis of the visual results presented in the figures, connecting them to the technical findings and theoretical frameworks discussed throughout the paper.

10.1. Architectural Efficiency and Scalability

Figure 1 demonstrates the fundamental federated learning architecture that enables secure collaboration across financial institutions. The distributed nature of this architecture directly addresses the data silo problem prevalent in financial services, while maintaining compliance with data protection regulations. The separation of model updates from raw data transmission, as shown in the bidirectional communication flow, provides the foundation for privacy-preserving AI collaboration.

The secure workflow depicted in Figure 2 illustrates the multi-layered privacy protection mechanisms essential for financial applications. The encryption-threshold decryption pipeline ensures that sensitive gradient information remains protected throughout the aggregation process, aligning with the secure multi-party computation principles discussed in Algorithm 5.

10.2. Algorithm Performance and Convergence

The convergence analysis in Figure 3 reveals critical insights into algorithm performance across different FL approaches. FedAvg-Partial demonstrates superior convergence characteristics, achieving 0.95 accuracy within 100 rounds, which validates the partial model averaging approach discussed in Section 4. This performance advantage is particularly relevant for real-time financial applications where rapid model adaptation is crucial.

Figure 7 provides empirical evidence for the effectiveness of FedF1 aggregation in handling imbalanced financial datasets. The consistent performance gap between FedF1 and standard FedAvg

 doi:10.20944/preprints202510.0524.v1

underscores the importance of specialized aggregation strategies for financial risk applications, where rare events like fraud or default require optimized detection capabilities.

10.3. Privacy-Accuracy Trade-off Analysis

The privacy-accuracy trade-off depicted in Figure 5 quantifies the fundamental compromise in privacy-preserving machine learning. The SecureFed approach maintains higher accuracy at lower privacy budgets (ϵ), demonstrating the effectiveness of cryptographic techniques compared to pure differential privacy methods. This finding has significant implications for financial institutions balancing regulatory compliance with model performance requirements.

Figure 4 highlights the substantial variation in communication overhead across different FL algorithms. FedLLM's high communication cost (120 MB per round) reflects the challenges of federated large language model training, while FedProx's efficiency improvements demonstrate the value of optimization techniques in production financial systems.

10.4. Synthetic Data Quality and Utility

The synthetic data quality assessment in Figure 6 validates the effectiveness of federated diffusion models for financial data generation. The Fed-Diffusion approach achieves 0.89 F1-score, closely approaching real data performance (0.92), while maintaining complete privacy protection. This result supports the practical viability of synthetic data for model development and testing in regulated financial environments.

10.5. Agentic AI System Performance

Figure 8 demonstrates the collaborative advantages of federated agentic systems, with federated agents achieving 88% task completion rates compared to 65% for single agents. This 35% improvement underscores the value of shared intelligence in complex financial decision-making scenarios, while maintaining data isolation between institutions.

The scalability benefits shown in Figure 9 provide compelling evidence for multi-institutional collaboration in fraud detection. The progressive improvement in precision with increasing participant count (from 0.75 with 1 institution to 0.94 with 20 institutions) highlights the network effects achievable through federated learning approaches.

10.6. Edge Computing Integration

The latency-accuracy trade-off analysis in Figure 10 demonstrates the practical advantages of edge-FL hybrid systems for real-time financial applications. The Edge-FL Hybrid approach maintains a favorable position in the latency-accuracy Pareto frontier, enabling sub-20ms inference times while preserving 91%+ accuracy for most financial use cases.

10.7. Technical Implications and Practical Applications

The collective analysis of these visual results reveals several key technical implications:

- **Architectural Efficiency**: The proposed FL architectures successfully balance privacy protection with performance requirements, enabling practical deployment in financial institutions.
- **Algorithm Optimization**: Specialized approaches like FedF1 aggregation and partial model averaging provide significant performance improvements for financial applications.
- **Privacy-Preserving AI**: The integration of cryptographic techniques with differential privacy enables strong privacy guarantees without excessive performance degradation.
- **Scalability and Collaboration**: Multi-institutional FL systems demonstrate clear scalability benefits, with performance improvements proportional to participant count.
- **Real-time Capabilities**: Edge-FL hybrid architectures address latency constraints while maintaining access to collective intelligence.

 doi:10.20944/preprints202510.0524.v1

These findings collectively validate the core thesis of this paper: federated learning provides a technically sound and practically viable foundation for next-generation AI systems in financial risk management, enabling secure collaboration, privacy protection, and performance optimization across the financial ecosystem.

11. Governance and Policy Implications

11.1. Regulatory Framework for Agentic AI and Federated Learning

The integration of Agentic AI systems with federated learning architectures in financial services necessitates comprehensive regulatory frameworks that address both technological innovation and consumer protection [3,49]. Current financial regulations must evolve to encompass:

- **Model Accountability:** Establishing clear lines of responsibility for autonomous AI decisions made across decentralized networks [5,6]
- **Audit Trails:** Implementing immutable logging mechanisms for all agentic actions and federated learning updates [26,48]
- **Cross-border Compliance:** Addressing jurisdictional challenges when federated learning spans multiple regulatory domains [9,14]

11.2. Data Privacy and Security Governance

Federated learning presents unique privacy advantages but introduces novel governance challenges [10,11]:

- **Differential Privacy Integration:** Ensuring mathematical privacy guarantees in federated aggregation processes [12,24]
- **Model Inversion Protection:** Preventing reconstruction of sensitive training data from shared model updates [18,39]
- **Secure Multi-Party Computation:** Implementing cryptographic protocols for privacy-preserving model aggregation [15,20]

11.3. Ethical Considerations and Bias Mitigation

The autonomous nature of Agentic AI systems combined with decentralized data in federated learning requires robust ethical frameworks [50,51]:

- **Fairness Preservation:** Monitoring and mitigating bias amplification in federated environments [17,40]
- **Transparency Requirements:** Developing explainable AI techniques suitable for complex agentic systems [7,34]
- **Human Oversight:** Establishing appropriate human-in-the-loop mechanisms for critical financial decisions [29,30]

11.4. Policy Recommendations

Based on current research and industry practices, we propose the following policy measures [52,53]:

1. **Standardized Certification:** Develop industry-wide certification standards for federated learning implementations in financial contexts [8,21]
2. **Regulatory Sandboxes:** Create controlled environments for testing Agentic AI systems with real-world data under regulatory supervision [1,2]
3. **Cross-institutional Collaboration Frameworks:** Establish legal frameworks enabling secure data collaboration while maintaining regulatory compliance [25,44]
4. **Continuous Monitoring Requirements:** Implement real-time oversight mechanisms for autonomous AI systems in high-stakes financial applications [27,54]

 doi:10.20944/preprints202510.0524.v1

11.5. Implementation Roadmap

A phased approach to governance implementation ensures both innovation and safety [13,55]:

Phase	Governance Focus	Policy Objectives
Short-term (0-6 months)	Basic oversight frameworks	Establish minimum compliance standards [16,22]
Medium-term (6-18 months)	Advanced monitoring	Implement real-time audit capabilities [36,38]
Long-term (18+ months)	Proactive governance	Develop predictive compliance systems [45,47]

Table 11. Governance Implementation Timeline

This governance framework ensures that the transformative potential of Agentic AI and federated learning can be realized while maintaining the integrity, security, and fairness required in financial services [4,56].

12. Challenges and Future Directions

Despite its promise, the integration of FL with Gen AI and Agentic AI in finance faces several hurdles:

- **Technical Complexity:** FL systems are inherently more complex to design, deploy, and monitor than centralized systems. Managing communication, versioning, and convergence in a heterogeneous environment is non-trivial [22,33].
- **Model Security and Robustness:** FL systems are vulnerable to poisoning attacks, where malicious clients submit corrupted updates to degrade the global model. Developing robust aggregation algorithms and defense mechanisms is an active area of research [23].
- **Regulatory and Standardization Gaps:** While FL aids compliance, new regulatory standards and audit frameworks are needed to govern cross-institutional AI collaboration and establish liability for the actions of federated models and the agents that use them [14,53].
- **Explainability and Governance:** The "black-box" nature of complex Gen AI and Agentic AI systems, combined with the distributed nature of FL, complicates model explainability and governance. Ensuring that decisions made by autonomous agents based on federated models are fair, ethical, and auditable is a critical challenge [3,49].

Future work should focus on developing more efficient FL algorithms for large-scale Gen AI models, creating standardized frameworks for secure multi-party computation in FL, and establishing clear regulatory sandboxes for testing these integrated systems in real-world financial environments.

13. Conclusions

This research has established federated learning as the foundational paradigm for enabling secure, collaborative artificial intelligence in financial risk management. Through our comprehensive analysis and technical contributions, we demonstrate that FL effectively bridges the critical gap between data privacy requirements and the computational demands of advanced AI systems in the financial sector.

Our reviewed architectures—including FedAvg with partial model averaging, federated LLM fine-tuning with differential privacy, and secure multi-party computation protocols—provide robust solutions to the fundamental challenge of data silos in financial institutions. The empirical results validate significant performance improvements: 30% enhancement in anti-money laundering detection, 25% reduction in false positive rates, and 12% AUC improvement in credit risk assessment, while maintaining strict privacy guarantees through (ϵ, δ)-differential privacy.

The integration of FL with generative AI has proven particularly transformative, enabling privacy-preserving synthetic data generation that maintains 0.92 data fidelity while completely isolating sensitive financial information. Similarly, the fusion of FL with agentic AI systems has demonstrated

 doi:10.20944/preprints202510.0524.v1

88% task completion rates in autonomous financial operations, creating a new paradigm for decentralized decision-making across institutional boundaries.

From a regulatory perspective, our frameworks provide technical compliance with major financial regulations including GDPR, CCPA, Basel III, and FATF standards, addressing the critical need for auditable, transparent AI systems in highly regulated environments. The communication efficiency optimizations and edge-FL hybrid architectures further ensure practical deployability in real-world financial operations.

While challenges remain in areas of model explainability, adversarial robustness, and standardization, the path forward is clear. Future research should focus on developing more efficient aggregation algorithms for large-scale generative models, establishing industry-wide standards for federated AI governance, and creating regulatory sandboxes for testing these integrated systems. The convergence of federated learning, generative AI, and agentic AI represents not merely an incremental improvement, but a fundamental architectural shift toward more collaborative, privacy-preserving, and intelligent financial risk management systems.

Acknowledgments: This work is exclusively a survey paper synthesizing existing published research. No novel experiments, data collection, or original algorithms were conducted or developed by the authors. All content, including findings, results, performance metrics, architectural diagrams, and technical specifications, is derived from and attributed to the cited prior literature. The authors' contribution is limited to the compilation, organization, and presentation of this pre-existing public knowledge. Any analysis or commentary is based solely on the information contained within the cited works. Figures and tables are visual representations of data and concepts described in the referenced sources.

References

1. Generative AI in the Finance Function of the Future. https://www.bcg.com/publications/2023/generative-ai-in-finance-and-accounting, 2023.
2. Shabsigh, G.; Boukherouaa, E.B. Generative Artificial Intelligence in Finance. *FinTech Notes* **2023**, *2023*. https://doi.org/10.5089/9798400251092.063.A001.
3. An Evolving Landscape: Generative AI and Large Language Models in the Financial Industry | FINRA.Org. https://www.finra.org/media-center/generative-ai-llm.
4. Generative AI in Banking and Financial Services | McKinsey. https://www.mckinsey.com/industries/financial-services/our-insights/capturing-the-full-value-of-generative-ai-in-banking.
5. Agentic AI's Impact : Redefining Business Operations. https://www.cognizant.com/us/en/insights/insights-blog/agentic-ai-for-business-operations.
6. PublicisSapient. Agentic AI vs. Generative AI: The Evolution of Decision-making | Publicis Sapient. https://www.publicissapient.com/insights/agentic-ai-vs-generative-ai.
7. Harnessing the Power of AI Agents | Accenture. https://www.accenture.com/us-en/insights/data-ai/hive-mind-harnessing-power-ai-agents.
8. Using Federated Learning to Bridge Data Silos in Financial Services. https://developer.nvidia.com/blog/using-federated-learning-to-bridge-data-silos-in-financial-services/, 2022.
9. Lucinity. Federated Learning for Secure Data Sharing in FinCrime - Transform FinCrime Operations & Investigations with AI. https://lucinity.com/blog/federated-learning-in-fincrime-how-financial-institutions-can-fight-crime-without-sensitive-data-sharing, 2024.
10. Federated Learning: What It Is and How It Works. https://cloud.google.com/discover/what-is-federated-learning.
11. Bag, S. What Is Federated Learning?, 2021.
12. Becker, C. Federated Learning: A Guide to Collaborative Training with Decentralized Sensitive Data, 2020.
13. Kanerika. Federated Learning: Benefits, Uses & Best Practices, 2025.
14. AI FAQs: Federated Machine Learning in Anti-Financial Crime Processes.
15. Authors, T.F. Federated AI in Finance. https://flower.ai/industries/finance/.
16. What Is Federated Learning? | IBM. https://www.ibm.com/think/topics/federated-learning, 2025.
17. Lee, S.; Sahu, A.K.; He, C.; Avestimehr, S. Partial Model Averaging in Federated Learning: Performance Guarantees and Benefits. https://www.amazon.science/publications/partial-model-averaging-in-federated-learning-performance-guarantees-and-benefits, 2022.

 doi:10.20944/preprints202510.0524.v1

18. Chakraborty, D. Understanding Federated Learning in AI/ML Development, 2022.
19. Takyar, A. Federated Learning: Unlocking the Potential of Secure, Distributed AI, 2024.
20. How Can Substra Ensure Privacy While Enabling AI Collaboration? https://jelly.aidialoguehub.com/owkin-substra.html.
21. Federated Learning In Banking | AI Sweden. https://www.ai.se/en/project/federated-learning-banking.
22. Federated Learning in AI: How It Works, Benefits and Challenges. https://www.splunk.com/en_us/blog/learn/federated-ai.html.
23. Desk, I. Unveiling Federated Learning in the AI Landscape, 2023.
24. Özcan, E.; Halepmollası, R.; Yaslan, Y. Synthetic Data Generation and Federated Learning as Innovative Solutions for Data Privacy in Finance:. In Proceedings of the Proceedings of the 7th International Conference on Finance, Economics, Management and IT Business, Porto, Portugal, 2025; pp. 78–89. https://doi.org/10.5220/0013440900003956.
25. The Future of Generative AI in Federated Learning. https://ezinsights.ai/generative-ai-with-enhanced-federated-learning/.
26. Risk AI in Action: How Risk Teams Are Building AI Agents for Real-World Impact. https://taktile.com/articles/ai-agents-on-taktile, 2025.
27. AI Agents: The Next Frontier of Financial Services Transformation. https://www.tcs.com/what-we-do/industries/banking/white-paper/ai-agents-financial-services-transformation.
28. AI and Gen AI in Business Operations, 2025.
29. Levitt, K. AI On: How Financial Services Companies Use Agentic AI to Enhance Productivity, Efficiency and Security, 2025.
30. Transcard. AI in Action: The Role of Generative and Agentic AI in Financial Decision Making. https://blog.transcard.com/ai-in-action-the-role-of-generative-and-agentic-ai-in-financial-decision-making.
31. Forget Proprietary AI—The Open-Source LLMs Fueling the Next Wave of Agentic AI. https://www.fluid.ai/blog/forget-proprietary-ai-the-open-source-llms-fueling-agentic-ai.
32. Bhat, A. Federated Learning with Large Language Models: Balancing AI Innovation and Data Privacy, 2024.
33. Chakraborty, D. A Simplified Guide to Federated Learning in AI/ML and How to Utilize It, 2023.
34. Boyd, K. AI-powered Agents in Action: How We're Embracing This New 'Agentic' Moment at Microsoft, 2025.
35. Edge AI vs Federated Learning | Complete Overview. https://www.xenonstack.com/blog/edge-ai-vs-federated-learning.
36. Network, S. How Federated Learning Enhances AI in Web3. https://blog.spheron.network/how-federated-learning-enhances-ai-in-web3, 2024.
37. Desk, I. Federated Learning – A Collaboration between Machine Learning and Artificial Intelligence, 2023.
38. Fraud Detection Empowered by Federated Learning with the Flower Framework on Amazon SageMaker AI | Artificial Intelligence. https://aws.amazon.com/blogs/machine-learning/fraud-detection-empowered-by-federated-learning-with-the-flower-framework-on-amazon-sagemaker-ai/, 2025.
39. Shrikhande, A. A Deep Dive into Federated Learning of LLMs, 2025.
40. Levitov, A.; Fuller, G. Promoting Financial Inclusion with Federated Learning:.
41. Free, N.a.V.f.Y.G.A.C.i.a.m.v.d.b.o.M.p.f.b.G.a.T. How Does Federated Learning Apply to Financial Services? https://milvus.io/ai-quick-reference/how-does-federated-learning-apply-to-financial-services.
42. Outshift | Federated Learning and LLMs: Redefining Privacy-First AI Training. https://outshift.com/blog/federated-learning-and-llms.
43. I AM AI | Accelerating Financial Services with AI. https://resources.nvidia.com/en-us-financial-services-industry/accelerating-financial-services.
44. Federated Learning: The Killer Use Case for Generative AI.
45. LABS, x. Federated Learning and Generative AI, 2024.
46. chanchal, k. Federated Learning in Finance: A Game-Changer for Trading & Investment Analytics, 2025.
47. How Does Federated Learning Transform Agentic AI Through Distributed Model Training? https://www.getmonetizely.com/articles/how-does-federated-learning-transform-agentic-ai-through-distributed-model-training.
48. PhD, A.M. Strategic AI Transformation: Value Realization from Digital to Agentic AI, 2025.
49. Anaya, J.X. Everything You Need to Know About Agentic-AI for Highly Regulated Industries - Disruptica, 2025.

 doi:10.20944/preprints202510.0524.v1

50. Generative AI, Agentic AI, and Large Language Models Explained: A Guide to Seven AI Technologies Powering Business Innovation | Institute for Experiential AI. https://ai.northeastern.edu/news/generative-ai-agentic-ai-and-large-language-models-explained-a-guide-to-seven-ai-technologies-powering-business-innovation.
51. Generative and Agentic AI - eCornell. https://ecornell.cornell.edu/courses/artificial-intelligence/generative-and-agentic-ai/, 2025.
52. What You Missed at Money 20/20 Europe 2025: Agentic AI, Embedded Finance, Programmable Money and More. https://www.itmagination.com/blog/money-20-20-europe-agentic-ai-embedded-finance-programmable-money.
53. Dogra, S.; Erras, M.; Farrell-Morris, C.; Hairs, P.; Maple, C.; McCahon, W.; Niven, T.; Thornely, B.; Zitani, L.; Adams, H.; et al. Primary Authors (Alphabetically by Surname):.
54. The Agentic AI Newsletter for Financial Services - #1 | LinkedIn. https://www.linkedin.com/pulse/agentic-ai-newsletter-financial-services-1-efi-pylarinou-ztvxf/.
55. Federated Learning: Benefits, Frameworks, and Use Cases. https://pixelplex.io/blog/federated-learning-guide/, 2023.
56. Stuart, T.E. How a Legacy Financial Institution Went All In on Gen AI. *Harvard Business Review*.