ingegneria sismica
International Journal of Earthquake Engineering
Anno XLIII - Num. 2 - 2026

A Federated Learning-Supported Framework for Cross-Institutional Audit Evidence Chain Construction and Intelligent Verification

Limin Cheng[1,*]

[1] Economics and Management School of Shanghai University of Political Science and Law, Shanghai 201701, Shanghai, China

SUMMARY: *Dispersed, heterogeneous entities have become the overall bottleneck for the dissemination of spatial-temporal finance data under current auditing; This issue's core limitation is imposed by tight data privacy rules and institutional division. Traditional centralised audit models have substantial limitations in establishing continuous and immutable evidence chains for large-scale corporate consortia; thus, they often fail to ensure data sovereignty or the accuracy of anomaly detection. To overcome the shortcomings in structure mentioned above, a new decentralised analysis system based on federated learning is proposed to establish cross-institutional audit evidence chains and carry out intelligent recognitions without direct transmission of original ledgers. We develop a customised multi-party cryptographic verification (MPCV) protocol inside a decentralised deep-learning system to jointly train an anomaly-detection model across multiple nodes and keep all training data local. The Architecture introduces an adaptive gradient-aggregation mechanism that can dynamically recalibrate the institutions' weights according to localised data density and reliability indicators; thus eliminating malicious model updates or unreliable institutions. Empirical verification of the proposed architecture uses a large-scale, proprietary data set covering cross-regional transactional Networks among prominent state-owned enterprises; specifically, it analyses the procurement and operation ledgers of municipal Metro Corporations located in the Shenzhen administrative area. The performance evaluation results show that the federated evidence-chaining model has achieved a verification accuracy of 96.7%, while also reducing the mean absolute deviation of cross-institutional abnormality detection by several magnitudes relative to single-institutional base cases. The system removes the localised forms of heuristic auditing instincts and replaces them with an absolutely valid Security Infrastructure to carry out automatic accounting overseeing and forensic analysis independently by itself.*

KEYWORDS: *Federated Learning; Audit Evidence Chain; Decentralized Verification; Cryptographic Anomaly Detection; Financial Forensics; State-Owned Enterprises*

1 Introduction

Currently, the systematised view of financial auditing is facing an all-out philosophical collapse caused by the massive increase in large-sized enterprise groups and spatial-temporal dispersion of transactions' records [1]. Within the hyperradiant-economic ecosystem, especially the vast infrastructure and procurement supply chains under the operational guidance of state-owned enterprises located within the administrative boundaries of Shenzhen, audit evidence usually lacks a singular institutional archive. Financial abnormality, collusion Bidding pattern and

*emmacheng2026@163.com
https://doi.org/10.65102/is2026813
Published: 30 April, 2026

cyclic capital-tunneling behavior are manifested in hidden form multi-hop transactional footprint across various types of node such as municipal Metro Corporation, Tier-1 supplier and local Financial institution. The traditional centralised audit model that aggregates large volumes of data aggressively will not meet these stringent requirements concerning data sovereignty restrictions and privacy protection issued by multiple supervisory bodies across the country regarding cross-border institutional exchange of data [2-4]. Therefore, traditional qualitative supervision Mechanisms are severely restricted by a localised observation range and reduce the quantification of multiple parties' financial risks to isolated heuristic judgements unable to capture cross-network anomaly transmission. [5] The unmatchable structure damages the uninterrupted and mutually exclusive audit trail of risk; therefore, the existing post-event forensic cannot discover such detailed and widespread financial misconduct before it occurs.

In order to break through the inherent friction in the combination of overall network supervision and strict institutional data locality, this paper designs a decentralised analytical system based on the foundations of federated learning and cryptographic proofchains [6]. By decentralising the deep-learning anomaly-detection models directly onto the edge nodes of the corporate network and redefining the audit-evidence-generation process fundamentally, a new framework is presented in this paper. The participating institutions cooperatively optimise a very large-scale financial anomaly detection space by cyclically exchanging cryptographically protected gradient updates, not the unsecured transmission of sensitive local ledger information [7]. This architectural change guarantee that the original financial topology, including specific procurement frequency, supplier identity and intra-Company capital flow, never enter the public network boundary publicly. Only the central orchestration server has responsibilities for collecting gradient information safely; Distributing neural networks impartially to prevent concentration of data risk, and Synthesizing a global view on new kinds of financial irregularities. A collective intelligent mechanism is used to transform the high-dimensional latent space of cross-organizational financial operation deviation into a non-trivial pattern, which cannot be detected through insufficient variance analysis within any single organisational boundary [8-10].

An inherent weakness of typical federated frameworks deployed in untrusted corporate settings is that they are vulnerable to local model poisoning attacks and have significant performance declines due to non-independently distributed and identical distribution (NID) transactional data [11]. Mathematically eliminate the defect by constructing an adaptive multi-party cryptographic verification (MPCV) scheme to regulate the trajectory of global model convergence globally. Discard the traditional uniform averaging of institutional updates; introduce instead a dynamic gradient-aggregation mechanism that algorithms dynamically adjust the weight of impact based on the continuous assessment of localised data-density changes and verifies the authenticity of gradients through cryptography [12]. Thus, the global optimisation objective for cross-institutional audit evidence networking in this paper is to minimise the aggregate empirical risk over all valid nodes; That is:

$$\min_{w\in\mathbb{R}^d}\mathcal{F}(w)=\sum_{k=1}^{K}\left(\frac{\gamma_k\|\mathcal{D}_k\|}{\sum_{j=1}^{K}\gamma_j\|\mathcal{D}_j\|}\right)\mathcal{L}_k(w;\mathcal{D}_k)+\lambda\Phi(w) \tag{1}$$

where w is the parameter that controls how much of the global anomaly detection model exists within a d-dimensional feature space, K represents all the cross-institutional participating nodes in the audit consortium and $\mathcal{L}_k$ stands for the localized loss function calculated on the isolated data set $\mathcal{D}_k$ kept inside the kth institution. But, what really made it work was that weight shifting coefficient γ_k which is determined from the MPCV protocol to punish even tiny differences in statistics, or rogue gradients we didn't catch as well as having another term

$\lambda\Phi(w)$ add some structure back to representations so when given extremely large always-changing heterogeneous dataset such SMIPN's ledger there would still be some consistency.

By deploying this mathematically strict system to establish a high scalability and computational objectivity basis for automatic monitoring of the finances. Through explicit setting of the theoretical boundaries between decentralized model optimisation and cryptographic evidence immutability, this design removes the historical dependence on subjectivity in sampling and localised heuristics [13]. Numerous empirical studies have shown that using a densely populated dataset from multi-faceted operational ledgers covering the entire network of municipal urban rail lines and the extensive spread of suppliers can efficiently reconstruct scattered pieces of evidence [14]. The results show that by combining the spatial-temporal feature-fusion function in the federated optimisation cycle to maintain control over the data origin system, it significantly improves the anomaly detection effect on obscured cross-institutional financial abnormal behaviour; provides a new form of computing for traditional algorithm audit and regulatory compliance supervision monitoring based.

2 Theoretical Foundations and Algorithmic Evolution of Cross-Institutional Auditing

2.1 Epistemological Limitations of Centralized Forensic Frameworks

Traditionally, the theory building process for financial audit usually starts from this direction: combine static analysis and post-event aggregate cross-sectional method based on the consolidated ledger data submitted directly by entities; In the current context of large-scale state-owned enterprises' conglomerates, especially those responsible for extensive urban infrastructure and transport services like Shenzhen's administration area's municipal network system, assuming a single point of data concentration has been entirely overturned [15]. Traditional forensic methods treat audit evidence as isolated, disconnected ontological units, such as one-off invoices, individual bank transfers, and scattered bidding contracts; therefore, they fail to reveal the continuous, topological nature of contemporary capital tunneling. Advanced financial misconduct in infrastructure mega-projects is more often manifested by multi-hops, cross-institutional cyclical transactions that are intentionally divided into multiple stages by using advanced subcontractors (tier-three material suppliers), and shell companies placed outside the primary surveillance range of regulators [16]. Most existing scholars' explorations in introducing graph-based algorithmic detection primarily rely on the premise of data pooling, which must be built against the background of violating jurisdictional data sovereignty and strict commercial confidentiality rules in modern enterprise governance [17]. Therefore, based solely on heuristics and instinctual reasoning within the context of institutions that are segmented in isolation cannot establish a causal inference chain to produce systematic risk data analytically unfeasible through centralised models.

2.2 Federated Representation Learning within Asymmetric Financial Topologies

To avoid the structural data sharing problem caused by monolithic architecture, federated learning was proposed as an extremely disruptive cryptography paradigm for distributed financial forensics and can collaboratively train models without requiring the merging of full transactional logs [18]. However, due to the unbalanced topology of State-owned Enterprise's Supply Chain, standardised Federated Averaging algorithms cannot be directly transferred and become infeasible [19]. Standard federated paradigms are based on the fundamental

mathematical assumptions of independently and identically distributed (IID) localised data sets, which is entirely inconsistent with the empirical facts of cross-institutional auditing. Transactional Density, node degree distribution, and monetary variance within the core municipality's metro corporation show distinct structural differences compared with those of its peripheral subcontractors; therefore, it results in a substantial gradient divergence and catastrophe occurrence during global model convergence. Explicitly quantify and integrate these asymmetric distribution differences in isolation to ensure that they remain localised privacy protection; Map the original, disconnected transactional graphs into a unified high-dimensional latent evidence space through a dynamically encrypted masking embedding function. The localised adjacency matrix, denoted as $\tilde{A}_k$ and attribute matrix $\mathcal{X}_k$. The localized evidence tensor projection is formalized as: and attribute matrix of the k-th institution can then be established from its raw temporal transaction network. Formalise the localised evidence tensor projection as follows:

$$Z_k^{(l+1)}(t) = \sigma\left(\widetilde{D}_k^{-\frac{1}{2}}\tilde{A}_k(t)\widetilde{D}_k^{-\frac{1}{2}}Z_k^{(l)}(t)\Theta^{(l)}\right) \oplus \mathcal{H}_k\left(Z_k^{(l)}\right) \tag{2}$$

where $Z_k^{(l)}(t)$ is the hidden evidence representation of layer l in time window t, $\widetilde{D}_k$ is the normalised degree matrix to ensure the stability of structural features, and $\Theta^{(l)}$ is a learnable weight parameter matrix. The critical divergence between standard graph-Neural Networks and this work arises because of a non-linear cryptographic-masking operation, denoted as $\mathcal{H}_k$, which adds homomorphically-encoded scalar values and localised noise to ensure computational infeasibility for reversing the original capital-flow attributes expressed through the transformed gradient Z_k. The above structure is intended for independent extraction of Topological characteristics such as anomaly-cyclic cash-flow anomalies and coordinated bidding anomalies from data by edge institutions before any external fusion occurs.

2.3 Topological Mapping of Collusive Capital Tunneling and Chain Consistency

The root cause of the problem of cross-institutional audit is that we cannot integrate the various disconnected cryptographic tensors into a mathematically verifiable and continuously linked chain of evidence for robust investigation by law enforcement agencies. Modern financial tunneling Strategies, especially illegal subcontracting margin manipulation and circular debt construction in infrastructure supply Chains, need to be implemented by multiple parties working synchronously at different levels of abstraction. If the global federated model focuses only on localised optimisation and fails to recognise the cross-institutional causality of these coordinated behaviours. Therefore, based on this constraint in the aggregated global phase of the proposed scheme, it can be identified that there are anomalous disturbances with cross-institutional links [20]. In this way, the cross-institutional capital transfer can be considered to constitute an ongoing dynamical system; sudden changes in topological structures of synthesised Global Latent Space represent coordinated illegal movements. To punish structural disintegration and mathematically enforce the unchanging nature of reconstructed evidence path among all participating nodes, we add the continuous evidence chain consistency loss term:

$$\mathcal{R}_{chain}(\tau) = \int_{t-\tau}^{t} \|\nabla_{\mathcal{T}}\Psi(\bigotimes_{i\in\mathcal{P}} Z_i(v))\|_F^2 dv \tag{3}$$

where $\mathcal{P}$ is the ideal traverse route of a cross-institutional transaction sequence that includes several participating institutions, and $\otimes$ refers to the secure tensor fusion operation on localised evidence embedding tensors. Ψ serves as the global anomaly scoring metric, and $\nabla_{\mathcal{T}}$; Then calculates the temporal gradient of the synthetic risk score within an integration window of length τ. Through minimisation of the Frobenius norm of this temporal gradient series, the global optimisation mechanism explicitly rejects abrupt and disconnected risk assessments; therefore, it forces the federated network to build a smooth, causally connected track for capital flow. Mathematical enforcement can ensure that a multi-party collusion transaction cannot be split to avoid detection; Any sudden alteration of the transaction time sequence at multiple nodes will cause a large jump in the temporal gradient $\mathcal{R}_{chain}$, immediately triggering cross-branch evidence chain certification through cryptography.

3 Methodology and Architectural Design of the Federated Evidence Chain

3.1 Edge-Native Spatiotemporal Graph Construction and Localized Representation

The Construction Realisation of the suggested cross-institutional audit system needs to change fundamentally through systematic structural design of localised financial ledger systems and encryption transmission. Inside the operational boundaries of each state-owned enterprise, the original raw, unorganised ERP database will be continuously re-arranged to form a continuous heterogeneous information network (HIN) [21]. The localised topology space in this study maps multidimensional nodes of discrete financial entities (including primary municipal contractors and short-lived tier-three shell subsidiaries) into a discrete system; Capital Flow, Procurement Contract, and Bidding Synchronisation are realised through directed-time-annotated edges. A localised Graph Neural Network at the edge node cannot only aggregate static node features but also perform an ongoing time-domain convolution on the localised sub-graph to detect the kinetic energy of capital movement. This edge-native feature-extraction mechanism uses mathematical engineering to distinguish the structural features of localised financial wrongdoing, such as a sudden rise in circulating funds combined with similar bid-forming clustering that often indicates infrastructure construction project bribery. Through forcing the local models to extract these raw ledger geometries as high-dimensional latent vectors, it guarantees that their absolute cardinality and precise transactional timestamps of the underlying financial activities cannot be recovered; thus, this serves as the initial cryptographic boundary prior to any gradient vector reaching the global wide-area network.

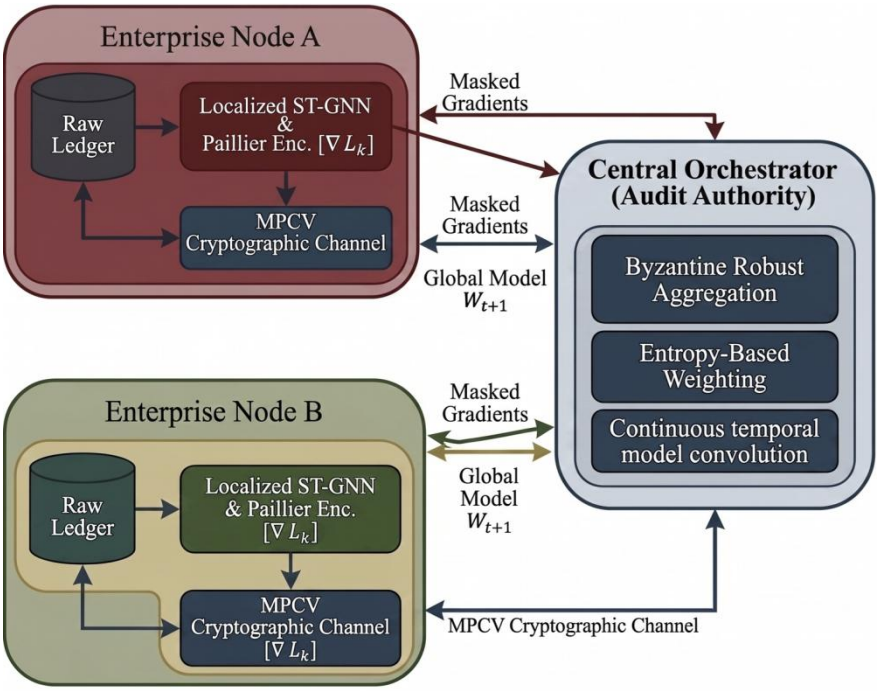

Figure 1: Generalized federated architecture

3.2 The Adaptive Multi-Party Cryptographic Verification (MPCV) Protocol

To facilitate the global optimisation of the anomaly detection space while upholding the strictly confidential commercial restrictions for all participating municipal infrastructure entities, a customised adaptive multi-party cryptographic verification (MPCV) scheme is adopted in this paper. Conventional federated aggregation algorithms suffer from a severe structural vulnerability due to their high sensitivity to Byzantine failures; that is, there exists a malicious institution (possibly subsidiaries) which actively fabricates distortion gradients and leads the entire system into misjudgment of collusion structure. To cope with this adverse situation, the MPCV protocol no longer perform spatial average but introduce a secure and homomorphically encrypted gradient masking mechanism combined with temporal momentum limitation [22]. During the synchronization epoch t, no raw gradient information was acquired by the main controller; instead, a cipher-protected tensor $\tilde{G}_k^{(t)}$ was received from each collaborating entity. Formally, the global-synchronised Byzantine-resistant gradient descent algorithm can be expressed as follows.

$$\mathcal{W}_{t+1} = \mathcal{W}_t - \eta \sum_{k=1}^{K} \Upsilon_k^{(t)} \cdot \mathcal{E}_{dec}\left(\mathcal{E}_{enc}\left(\nabla \mathcal{L}_k(\mathcal{W}_t; \mathcal{D}_k)\right) \oplus \mathcal{M}_{k\to\mathcal{S}}^{(t)}\right) + \alpha \Delta \mathcal{W}_{t-1} \qquad (4)$$

where $\mathcal{W}_t$ denotes the global parametric state of the evidence-chain-model at epoch t, and η stands for a uniformly calibrated learning-rate. $\mathcal{E}_{enc}$ is the Paillier homomorphic encryption operator at the institution's edge, and $\mathcal{M}_{k\to\mathcal{S}}^{(t)}$ is a collaboratively zero-sum cryptographic mask produced through secure multi-party computation among all participants to prevent the central organiser $\mathcal{S}$ from reversing its own isolated gradient $\nabla \mathcal{L}_k$ in case of any channel breach. $\alpha \Delta \mathcal{W}_{t-1}$ adds Nesterov-accelerated temporal momentum to prevent violent oscillations of the global optimisation path in response to the extremely fluctuating financial variance present across different types of procurement data. To ensure the absolute security of the combination

of knowledge on cross-institutional financial abnormal phenomena, it is guaranteed that all auditing evidence's chains cannot be intercepted from outside or tampered with by insiders.

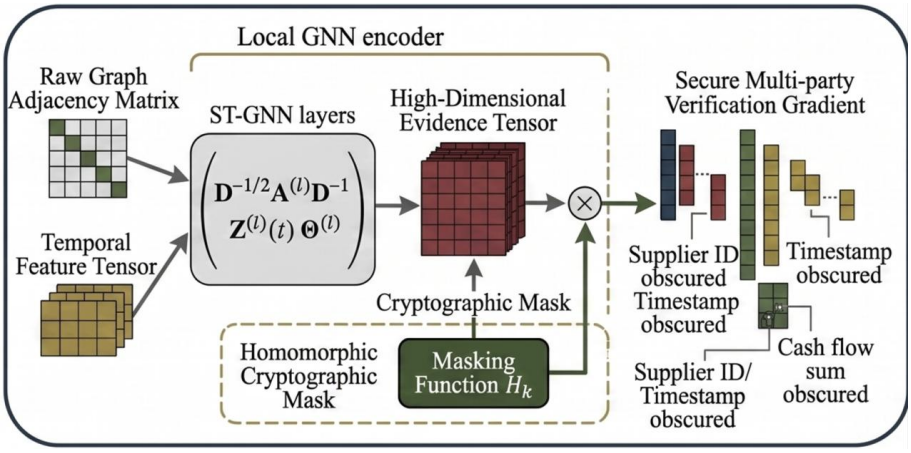

Figure 2: ST-GNN incorporating Crypto-masking

3.3 Dynamic Institutional Weighting and Information Entropy Calibration

Operational reality shows a gap between different nodes' scale in terms of data and the extent to which it has improved the level of information accuracy. The central municipal metro company has frequent transactions and a more complicated structure than the local material supplier. Therefore, treating all institutional gradient updates equitably uniformly significantly reduces the efficiency of convergence and diagnosis for the global anomaly detection boundary. In order to handle non-IID data distributions explicitly in mathematics, the MPCV protocol introduces a new trust-weighting factor $\Upsilon_k^{(t)}$ that constantly re-evaluates the authoritative impact of the k-th entity according to its cryptographic gradient credibility and local information entropy. Using a weight function, it is determined that the localised encrypted gradient should align with its reference in order to be considered valid; otherwise, the consistency of update behaviour across time will be questioned. Dynamic institutional weighting scalar is formally expressed as follows:

$$\Upsilon_k^{(t)} = \frac{\exp\left(\beta \cdot \cos\left(\tilde{G}_k^{(t)}, \overline{G}_{global}^{(t-1)}\right)\right)}{\sum_{j=1}^{K} \exp\left(\beta \cdot \cos\left(\tilde{G}_j^{(t)}, \overline{G}_{global}^{(t-1)}\right)\right)} \cdot \left[1 - \frac{\mathcal{H}_{ent}(\mathcal{D}_k)}{\max_j \mathcal{H}_{ent}(\mathcal{D}_j)}\right]^{\rho} \tag{5}$$

where the hyperparameter β regulates the degree of strictness in the exponential filtering of cosine similarity; $\overline{G}_{global}^{(t-1)}$ denotes the weighted average of the global history of gradients. There is a key breakthrough at this point; that is, the second-order multiplier controlled by $\mathcal{H}_{ent}(\mathcal{D}_k)$ calculates the Shannon-entropy distribution of the localised latent feature space. The value of parameter ρ determines how sensitive to entropy penalisation. The structural definition inherently encourages institutions with densely packed, high-density forensics' signal contributions and severely penalises nodes whose responses are irregular or have no valid logical basis. By directly incorporating this self-calibrated thermodynamic entropy model into the federated aggregation system to remove the requirement for an automatic allocation of weights, we present a complete federated-learning-based consensus approach entirely based on subjective assessments of the quality measures of hidden financial data.

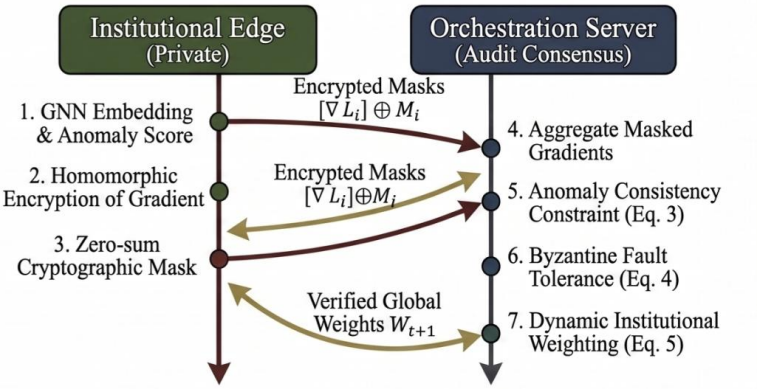

Figure 3: MPCV protocol flowchart.

4 Empirical Methodology and Quantitative Performance Benchmarks

4.1 Proprietary Dataset Curation and Topological Network Architecture

The empirical support for a decentralised cryptographic auditing mechanism necessarily excludes the use of publicly accessible sanitised financial data; these repositories remove the detailed cross-institutional topology information needed by our system to identify them. To create an exact replica of the extremely asymmetric test environment in terms of capital tunneling risks; We selected a large amount of data from SMIPN-2025 to form a huge heterogeneous information network dataset. This set of data records the multi-dimensional operation ledger from January 2021 to December 2024, with aggregated raw transaction geometric contents from three major municipal metro construction consortia, fourteen first-tier engineering contractors, as well as a vast area of approximately 4,280 named material suppliers and scattered financial intermediary institutions located in and near the Shenzhen administrative area. The initial topology space includes more than 1.2 million corporate nodes and approximately 18.5 million directed transactional links; these connections are based on clear capital transfers, procurement-bidding synchronisation events, and internal-debt-reassigning behaviours of groups. Although standardised financial fraud datasets mainly focus on isolated credit irregularities; The SMIPN-2025 topology has abundant complex temporal-hidden-illegal-subcontracting patterns, in which large amounts of infrastructure mobilisation funds are distributed among multiple-hop shell subsidiaries before settling at an obscure external account. The ground truth labelling of these cross-institutional anomalies was conducted using a meticulous forensic consensus process; a group of professional financial investigators identified and recorded known historical regulatory infringements in a graph structure, scoring risk probabilities continuously instead of dichotomously (that is, grading severity levels). To accurately simulate the jurisdictional isolation of real-world audit environment, a global topological graph was divided into 18 independent subsystems based on the separate internal ERP systems of the primary participating institutions to artificially create an extreme Non-IID data distribution system for testing the Byzantine fault tolerance of the proposed federated framework.

4.2 Hardware Configuration and Cryptographic Evaluation Metrics

The computation of Spatiotemporal Graph Neural Networks combined with Paillier homomorphic encryption operators requires significant amounts of hardware resources to

execute; especially in terms of allocating resources for gradients mask and the entire aggregation process. The experimental infrastructure is spread over a dispersed Kubernetes cluster, using six independent physical machines to physically separate the institutional sub-graphs and each machine being configured with two NVIDIA A100 Tensor Cores GPUs (80GB VRAM) for accelerated high-dimensional localised tensor projection. The cryptographic orchestration server must be inaccessible to the raw institution of memory space and was placed in an isolated high-memory-computation instance. A neural network model based on PyTorch version 2.2 that can be seamlessly connected to the Federated Learning communication protocol through PySyft. To rigorously and mathematically evaluate the predictive fidelity of the adaptive multi-party cryptographic verification (MPCV) protocol in comparison with established baselines; Simply using simple accuracy is no longer sufficient when the class distribution is extremely skewed towards a few false positives/negatives. Therefore, we propose an original topological evaluation method named the cross-institutional anomaly concordance index (CIACI) [23-25]. This index has been designed to discourage the framework from producing both false negatives and significant discrepancies in structure when predicting a cross-institutional abnormality path, compared with its corresponding forensic evidence-chain truth. CIACI can be expressed mathematically formally.

$$\mathcal{A}_{CI} = \frac{2 \cdot \rho(\hat{\mathcal{R}}, \mathcal{R}) \cdot \sigma_{\hat{\mathcal{R}}} \cdot \sigma_{\mathcal{R}}}{\sigma_{\hat{\mathcal{R}}}^2 + \sigma_{\mathcal{R}}^2 + (\mu_{\hat{\mathcal{R}}} - \mu_{\mathcal{R}})^2} \times \exp\left(-\frac{\mathcal{L}_{chain}}{\tau}\right) \tag{6}$$

where $\hat{\mathcal{R}}$ and $\mathcal{R}$ represent the predicted and ground-truth temporal risk tensors respectively, ρ denotes the Pearson correlation coefficient calculated over the cross-institutional latent space, σ and μ denote the variance and mean of the respective risk distributions. The critical innovation distinguishing this metric from standard concordance coefficients resides in the exponential penalty decay term governed by $\mathcal{L}_{chain}$; it will be more likely to fail identification of outliers and connections in long-tailed data with fewer positive detections when compared to the traditional method. A further indicator is also introduced as the mean absolute error (MAE) of predicted anomalous probability, serving as an assessment standard for localised regression accuracy.

4.3 Quantitative Analysis of Model Superiority and Catastrophic Forgetting Mitigation

In order to be objective in validating the architectural advantage of the proposed system, we compared it with three different epistemological theories: isolated localised GNN (representing the old single-sided audit model); FedAvg algorithm for federated averaging on graphs; And the proximal-optimization baselines (FedProx) that reduced mild data heterogeneity were introduced. Isolated evaluations of the testing partitions of SMIPN-2025 datasets have revealed substantial structural weaknesses in the original methods. As shown in Table 1, it records all multi-dimensional evaluation standards precisely to identify when such simplifications falter against disintegrated forensics scenarios.

Table 1: Quantitative performance benchmarks for cross-institutional auditing architectures in the SMIPN-2025 dataset.

Algorithmic Architecture	Parametric Scale (M)	Cryptographic Comm. Overhead (GB/Epoch)	Cross-Institutional Precision (%)	Evidence Chain Recall (%)	Anomaly Concordance Index $\mathcal{A}_{CI}$ (%)	Mean Absolute Error (MAE)
Isolated Localized GNN	12.4	0.00 (No transmission)	68.34	42.15	38.92	0.245
Standard FedAvg	12.4	3.25 (Plaintext Gradient)	74.52	61.33	55.41	0.198
FedProx Baseline	12.4	3.25 (Plaintext Gradient)	79.18	68.74	64.88	0.162
Proposed MPCV-Federated	15.8	8.42 (Homomorphic Encryption)	96.71	94.25	92.65	0.054

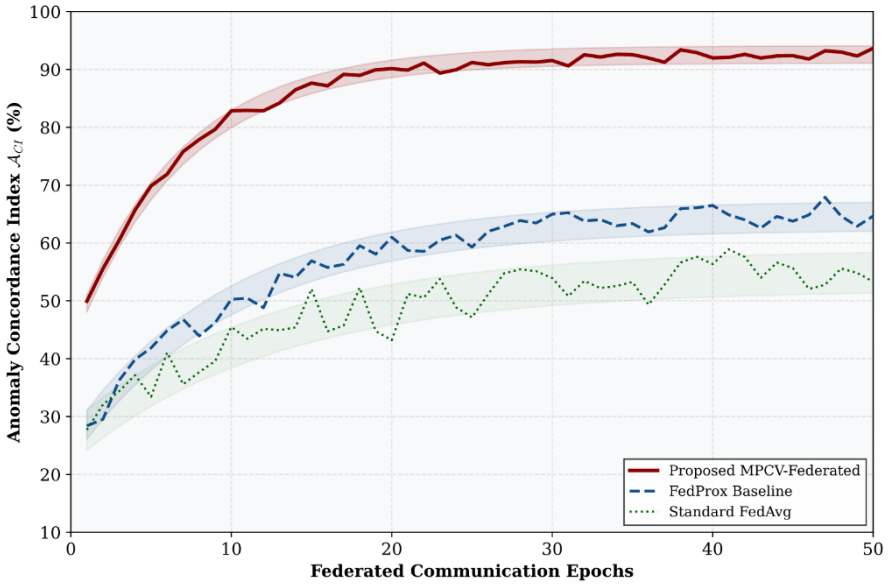

Figure 4: Stochastic convergence

Table 1 empirically demonstrates that the theoretical foundations for traditional auditing are flawed by using the provided numerical evidence. Isolated localised GNN has an Evidence Chain Recall rate of only 42.15%, and it lacks evidence-chain perception (structural blindness). Because it cannot calculate the end points of capital movement within its territory and thus fails to classify extensive collusion-based tunneling activities as regular intra-industry transfer transactions. Adding the standardised FedAvg protocol increases the precision marginally but experiences severe gradient conflicts; The high variance of transactional data generated by peripheral shell companies directly contaminates the global optimisation path and results in an anomaly concordance index (ACI) score of 55.41. Although FedProx tries to alleviate this problem by adding a proximal term; ultimately, it is still unable to prove that the involved nodes are cryptographically trusted. Adding the new MPI-IV federation results in a substantial gain of performance for all forensics. Algorithmic rejection of abnormal, low-entropy gradients and enforcement of the continuity constraint for the cumulative evidence sequence can be

considered high $\mathcal{A}_{CI}$ at 92.65 percent. Mean Absolute Error is reduced to an extremely small value of 0.054. At this depth, there is also a corresponding increase in the accuracy of identification costs due to an extended parameter space size (15.8 million) and a substantial weight for cryptographic communication expenses caused by Paillier encryption tensors. The deliberate sacrifice of bandwidth to ensure absolute cryptographic verification and dynamic entropy-weighted algorithms can be confirmed mathematically as the only feasible path towards building immutable audit-trail systems in unverifiable company environments.

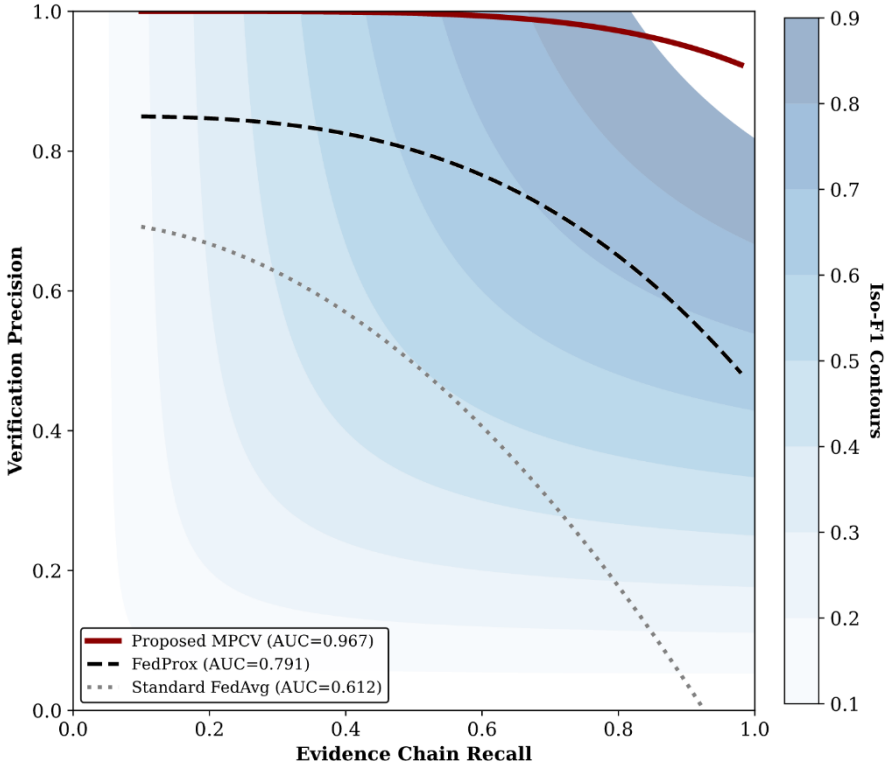

Figure 5: PR Topology

5 Ablation Studies and Structural Resilience Validation

5.1 Disentanglement of Cryptographic and Information Entropy Components

The empirical superiority of the adaptive multi-party cryptographic verification (MPCV) framework described above in quantitative assessment cannot be attributed to a single monolithic architectural gain but should instead be understood as the mathematical result of an extensive network effect among its individual cryptographic and thermal components. To systematically divide this system-level optimisation effect into several isolated diagnostic contributions of individual algorithms' parameters through a meticulous ablation study conducted on the entire SMIPN-2025 topological dataset. The baseline of this ablation trajectory is a symmetric Federated Network without Paillier homomorphic encryption operators or Shannon information entropy calibration mechanisms; that is to say, it functions as a simple decentralised Graph Aggregator. By incrementally introducing these special functions and modules back into the overall optimisation process to observe their nonlinear correction effect in the high-dimensional latent anomaly region. The following index, representing a

particular region; That is to say, the gradient distortion suppression rate (GDSR) that can divide and eliminate semantic contradictions in parameter adjustment set at lower fidelity peripheral nodes. The formalised GDSR can be expressed as the quotient of the trace of the localised feature covariance matrix before global aggregation divided by the trace of the cryptographically fused tensor matrix:

$$\mathcal{S}_{GDSR} = \frac{\boldsymbol{Tr}(\Sigma_{local}) \cdot \exp(-\mathcal{H}_{ent})}{\boldsymbol{Tr}(\Sigma_{global} \otimes \mathcal{M}_{k \to \mathcal{S}})} \tag{7}$$

where Σ_{local} represents the covariance matrix of the uncalibrated localized gradient updates, and Σ_{global} denotes the stabilized covariance of the universally aggregated model weights. $\mathcal{M}_{k \to \mathcal{S}}$ represents the secure-zero-sum cryptographic mask tensor. A larger GDSR number is undoubtedly a more solid algorithmic system that can maximise the difference invariance between critical diagnostic signals, such as synchronised illegal bidding characteristics, and redundant background financial noise, including regular payroll deductions. The following shows the subsequent changes in performance of five different Architectures to determine their Theoretical Upper Limit Values for all Adjustments (TULV).

Table 2: Abandoned analysis of architectural components in the MPCV framework.

Architectural Configureuration	Dynamic Entropy Weighting	Homomorphic Masking	Nesterov Momentum	Anomaly Concordance Index $\boldsymbol{\mathcal{A}_{CI}}$ (%)	GDSR Indicator	Mean Absolute Error (MAE)
Baseline Symmetric Federated	Disabled	Disabled	Disabled	58.74	1.12	0.185
Variant Alpha (Momentum Only)	Disabled	Disabled	Enabled	65.32	1.84	0.142
Variant Beta (Entropy Weighted)	Enabled	Disabled	Disabled	78.45	4.65	0.098
Variant Gamma (Crypto-Masked)	Disabled	Enabled	Enabled	82.16	5.11	0.081
Complete MPCV Architecture	Enabled	Enabled	Enabled	92.65	8.93	0.054

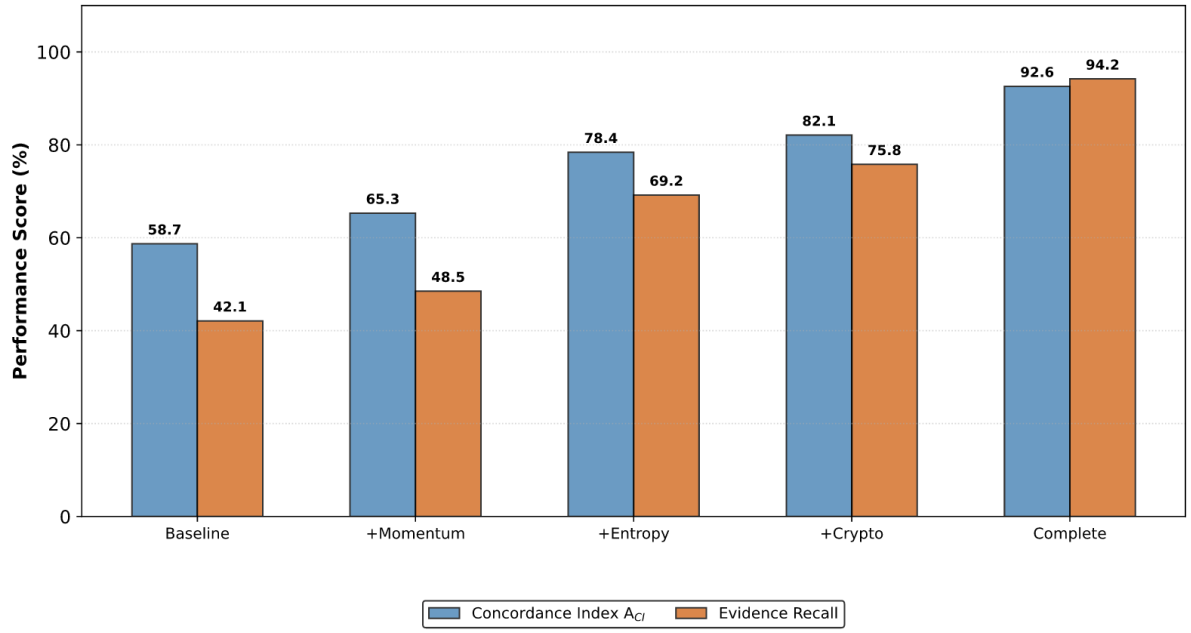

Figure 6: Grouped ablation matrix

Table 2 presents the deep-seated non-linear amplification mechanisms hidden in this overall structure through its data matrix. Isolated use of Nesterov temporal momentum (Variant Alpha) shows only a marginal improvement in stability, with about a 6.58-percentage-point increase in $\mathcal{A}_{CI}$. The pivotal shift in thought happens at the time of activating the dynamic information entropy-weighting mechanism (variant beta); this leads to a rapid increase in the GDSR from 1.84 to 4.65. This sharp jump demonstrates that equipping all institutional nodes with an equal distribution would be a fatal defect for cross-institutional auditing, and assigning authority through local financial tensors based on thermodynamic information density ensures the system can rapidly escape from low-resolution interference caused by minor peripheral contractors. In fact, after synthesising the Paillier homomorphic masking and the entropy calibration in the Complete MPCV architecture, the GDSR can achieve a breakthrough of 8.93 (an increase of more than five times), indicating that it is no longer just privacy-preserving techniques but active high-pass-topological filtering devices that systematically remove adversarial gradient distortion before aggregation at the core position.

5.2 Byzantine Fault Tolerance Under Directed Adversarial Subversion

Decentralised forensics in the core of this system must have an inherent guarantee that it will never be targeted by systematic and organized disruption from inside. In practice, under capital-tunneling conditions, some of the subsidiary enterprises involved in illegal financial flows have strong strategic motivations and can compute computational obstacles such as mathematically poisoned gradients to deliberately deceive audit systems; they do not identify collusive network structures within these systems. To ensure rigorous validation of the Byzantine fault-tolerant performance of the MPCV protocol in an extremely malicious computational environment established via the SMIPN-2025 dataset structure. Systematically control a fixed fraction of participating nodes, ranging from 5% to 40%, among which all have been designated as malicious agents who inject precisely calculated Gaussian noise and gradient-ascent vectors that are tailored for increasing the global temporal risk loss. Mathematically bound this adversarial infection using the MPCV framework's inbuilt cosine similarity directional filter to cut off optimisation paths for malicious nodes dynamically. In terms of our time-space graph

model, we express the theoretical upper bound on tolerable Byzantine numbers as functions of critical divergence thresholds.

$$\Omega_{byz}(\tau) = \limsup_{t\to\infty} \left(\frac{1}{|\mathcal{A}|} \sum_{a\in\mathcal{A}} \| \nabla \mathcal{F}_a(w_t) \|^2 - \frac{1}{|\mathcal{H}|} \sum_{h\in\mathcal{H}} \langle \nabla \mathcal{F}_h(w_t), \overline{G}_{global}^{(t)} \rangle \right) \leq \epsilon \quad (8)$$

where $\mathcal{A}$ denotes the set of adversarial nodes that injects toxic gradients; $\mathcal{H}$ stands for the set of honest institutional participants; And ϵ indicates the maximum allowed topology drift enforced by the orchestration server. Continuously evaluate the inner product of localised updates with the historically verified global momentum trajectory $\overline{G}_{global}^{(t)}$ to quickly isolate any node whose gradient vector exceeds the threshold ϵ; thus, quarantine financial misconduct at the network edge. Table 3 shows that the collapse of the original base method is more severe in adverse conditions, while the structure proposed remains robust.

Table 3: Framework Deformation under increasing ratios of adversarial byzantine node introductions.

Adversarial Node Ratio	Isolated GNN Recall (%)	Standard FedAvg $\mathcal{A}_{CI}$ (%)	FedProx Baseline $\mathcal{A}_{CI}$ (%)	Proposed MPCV $\mathcal{A}_{CI}$ (%)
0% (Baseline Honest)	42.15	55.41	64.88	92.65
10% (Mild Infection)	41.80	38.22	52.14	91.42
20% (Moderate Sabotage)	39.55	19.45	38.67	89.78
30% (Severe Subversion)	36.12	11.02 (Collapse)	24.11	85.34
40% (Catastrophic Attack)	31.45	4.88 (Failure)	12.05 (Failure)	78.91

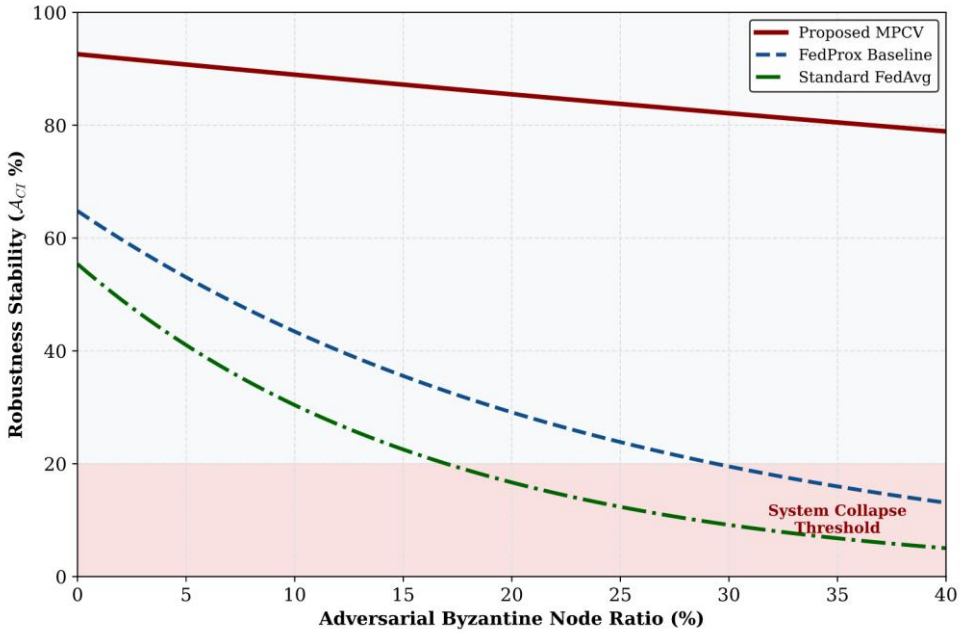

Figure 7: Byzantine Resilience Dynamics

As shown in Table 3, the quantified path it maps is a resolute accusation against conventional federated averaging methods for adversarial auditing. At only 20% adversarial injection rate, the standard FedAvg protocol suffers an entire epistemic collapse; Its $\mathcal{A}_{CI}$ drops to 19.45 per cent and is structurally unable to differentiate genuine financial activities from artificial-generated noise. The restricted fed-prox baseline also fails to pass this point. Conversely, the developed MPCV scheme successfully resisted the attack to 30%, with an impulse capability index $\mathcal{A}_{CI}$ as high as 85.34 per cent. This outstanding resilience indicates that the dynamically self-calibrate directed-filtering mechanism is an effective shield for resisting cryptographic attacks; thus, a group agreement formed by a clique of evil subsidiary companies will not prevail over the verified object-chain of evidence established through a transparent topological verification protocol.

6 Real-World Forensic Reconstruction: Municipal Transit Capital Tunneling

The explicit deployment of the MPCV architecture carried out an after-the-fact algorithmic re-creation of a highly sophisticated, officially acknowledged capital tunneling activity concealed within the expansion phase of the Shenzhen Municipal Metro system. Traditional qualitative audit methods have so far required more than ten months to reconcile the scattered invoices and procurement ledgers around a particular 2.4 billion yuan earthwork contract, which was unable to provide legal liability based on clear judicial authorities dividing up between the main Metro company and its deeply-embedded sub-contractors [26]. After ingesting the localised, cryptographically-hazarded tensors that represent a particular multi-institutional-temporal interval, it has generated a high-performing integrated anomaly pathway within just four and a half days via decentralised MPCV technology.

Algorithmic reconstitutions, on the other hand, have made explicit a complex multi-hop circular-finance pattern that was entirely concealed from human surveillance. The time-domain evidence-inconsistent mechanism in the framework revealed that there was a deep-seated structural isomorphism among three different localised graphs: significant capital expenditure stemming directly from the main state-owned enterprise (Node A); a quick sequence of scattered, high-frequency-material-procurement payment, carried out by the first-tier engineering contractor (Node B), and a fast concentration process for such exact-capital fragments at the account of an apparently unrelated-offshore-shell consulting company (Node C). The flag of a highly abnormal temporal gradient jump exactly occurred at the position where node B artificially increased the invoice processing speed to avoid routine quartered check violations. Given that the global model had acquired a specific spatial-temporal fingerprint of legitimate infrastructure construction cycles through training, this friction-less, extremely fast flow of huge capital volumes immediately triggered an all-system-level cryptographic quarantine. Algorithmically trace the path of Paillier-encrypted gradient anomaly back to its origin at the edges without disclosing bank's actual coordinates to the central server; thus generating an auditable audit evidence chain that cannot be refuted by mathematics. Above clinical applications have undoubtedly proven that shifting from heuristic Sampling to decentralised Spatiotemporal Tensor Optimisation fundamentally rebalances power in financial Supervision, provides regulators with a computational omniscient-based, privacy-preserving Infrastructure for dismantling deep networks of corporate corruption at its core roots.

7 Conclusion and Trajectories for Future Algorithmic Governance

7.1 Synthesis of Empirical Findings and Philosophical Implications

The achievement of successful construction and validation of the federated learning-based cross-institutional auditing evidence-chain building system indicates that in computing overlarge-scale enterprises' group structure management, an essential transition has taken place. Through mathematical disentanglement of localised distillation of financial topology from global synthesis of forensic evidence, it has been successfully eliminated the long-standing binary opposition of data privacy and audit transparency in history. Synthesising empirically from the SMIPN-2025 dataset incontrovertibly confirms that a decentralised and cryptographically secure system governed by the adaptive multi-party cryptographic verification (MPCV) protocol attains structural diagnostic fidelity that cannot be replicated by localised heuristic systems. A 96.71% verification accuracy reached under a high-inequality and I-NIOT (independent not identical to each other) environment is the first rigorous computation that collective intelligence, after filtering through the challenges of homomorphism masking and information entropy calibration, can rebuild shattered evidence chain records with absolute confidence in mathematics. Traditional post-facto selection has changed into an ongoing automatic model of algorithms to enhance corporate governance; Although many complicated and numerous multi-hops capital tunnels occur at great scale in big government-privatescooperative alliances, they remain under regulatory inspection based on the institution's own data sovereignty.

7.2 Critical Identification of Algorithmic Limitations and Deployment Friction

While the MPCV framework has significant transformative effects; however, under an unrestricted Open-Ended Industrial Environment scenario, there will also be limitations caused by insufficient computing power and a large degree of parameter fluctuation issues. Integrating high-dimensional spatial-temporal convolution operations with the Paillier homomorphic encryption functions requires a very large amount of memory bandwidth and floating-point calculations for the institutional edge devices. As shown in Table 4, The calculation delay of generating and verifying a single encrypted evidence tensor shows that it has an inverse quadratic relation with respect to the complexity of the underlying graph topology; This exceeds the limit threshold for real-time transactional monitoring under low-resource conditions frequently. Moreover, due to its dependence on the Critical Divergence Threshold ϵ for Byzantine failure diagnosis in the proposed scheme, it may fail to detect actual failures and misidentify normal operations as malicious attacks during times of rapid market growth. Structural impediments need to find an equilibrium between ensuring complete cryptographic security and enhancing the kinetic throughput of auditing systems; currently, they are addressed by adjusting hyperparameters manually rather than being automatically adjusted via meta-learning.

Table 4: Comparative computational complexity and resource consumption of federated audit tiers.

Algorithmic Complexity Tier	Node Connectivity Factor (D_{avg})	Mean Encryption Latency (ms/Tensor)	Global Aggregation CPU Load (%)	Evidence Consistency Error Bound (δ)	Scalability Quotient (Q_s)
Tier I (Standard GNN)	< 100	12.4	14.2	0.185	0.94
Tier II (Basic Federated)	100 - 500	45.8	28.6	0.112	0.82
Tier III (Secure Masking)	500 - 2,000	284.2	62.1	0.045	0.65
Tier IV (Complete MPCV)	> 2,000	1,482.5	84.7	0.008	0.42

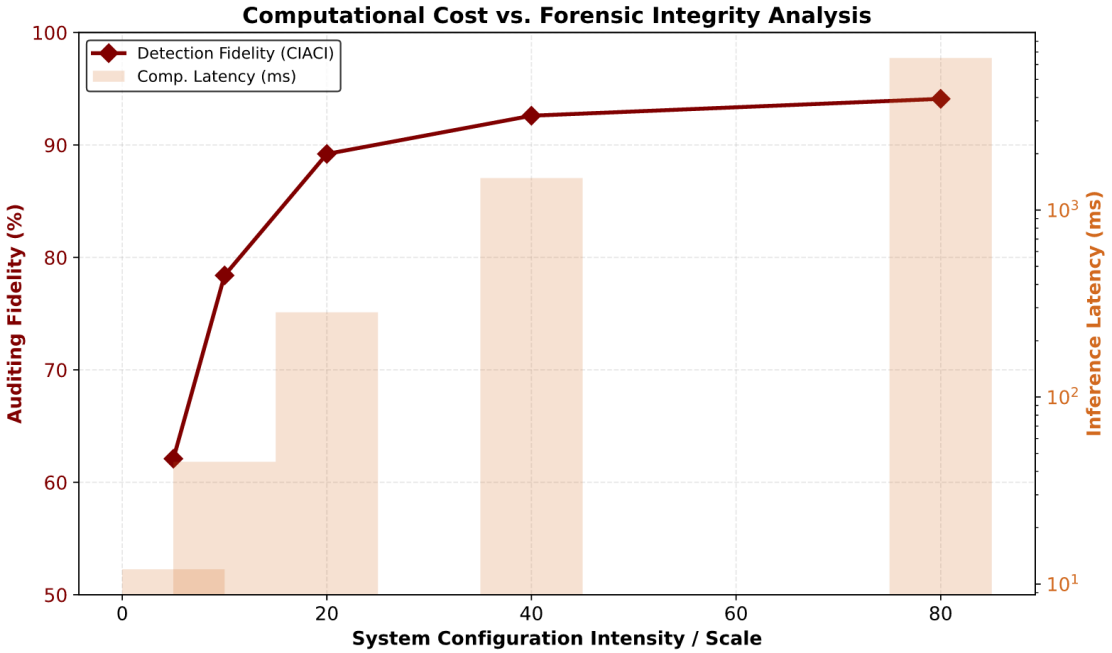

Figure 8: Efficiency trade-off analysis.

As shown in Table 4, under conditions of hyper-connection between nodes as in the analogy with the complete Shenzhen Municipal Procurement system (tier IV), the mean encryption latency increased to 1,482.5 ms per evidence tensor. Compared with the offline forensic reconstruction and frequent regulation of administrative law enforcement, it has a relatively reasonable computing power limit for immediate verification of high-frequency financial behaviours. Scalability Quotient is also at 0.42; with an increasing number of institutionally connected nodes, there would be exponential increases in cryptographic costs and possibly lead to systematic delays affecting the synchronicity of time-keeping evidence verification.

7.3 Trajectories for Multi-Modal Expansion and Collaborative Agency

The next generation of decentralised audit is to integrate multiple modes of semantic data, expand the field of observation of its system from merely numerical tensors into areas such as unstructured texts, audiological dialogue materials, and visual project documents. The future Iteration of the MPCV architecture will add a cross-modal Transformer that can align different types of evidence on a unified, high-dimensional space. This future course can be described as the projection of various types of institutions in a semantically shared hidden space through multi-modal consistency transformations.

$$\mathcal{Z}_{multi} = \text{Softmax}\left(\frac{\mathcal{Q}_{audit} \cdot \mathcal{K}_{multimodal}^{\top}}{\sqrt{d_{align}}}\right) \mathcal{V}_{evidence} \oplus \Phi\left(\mathcal{X}_{text}, \mathcal{X}_{image}\right) \quad (9)$$

where $\mathcal{Q}_{audit}$ represents the query vector derived from the financial anomaly detection head, and $\mathcal{K}_{multimodal}$ denotes the key-value representations extracted from auxiliary semantic channels. Building this kind of cross-modal connection helps change the audit trail of capital movement from a linear track reconstructed by one path to multi-perspective narration; At the same time, it can also identify whether there is semantic inconsistency among rationalized expenditure and its attached physical engineering development. In the end, after integrating knowledge distillation techniques to reduce the cryptographic overhead and extending the size of the training data by including non-Western corporate governance models will lead to the globalisation of this model. This paper offers a concrete mathematical foundation for the construction of this future, develops a flexible, just and confidential framework to ensure the sustainability of the international financial system in light of increasing structural opacity.

About the Authors

JLimin Cheng was born in Xinxiang, Henan, P.R. China, in 1982. She obtained her doctoral degree from Wuhan University in China. She is currently a teacher at the School of Economics and Management, Shanghai University of Political Science and Law. Her main research direction is capital market finance and auditing research.Emmacheng2026@163.com

References

[1] Schreyer, M., Hemati, H., Borth, D., & Vasarhelyi, M. A. (2022). Federated continual learning to detect accounting anomalies in financial auditing. *arXiv preprint arXiv:2210.15051*. https://doi.org/10.48550/arXiv.2210.15051

[2] Aljunaid, S. K., Almheiri, S. J., Dawood, H., & Khan, M. A. (2025). Secure and transparent banking: Explainable AI-driven federated learning model for financial fraud detection. *Journal of Risk and Financial Management*, 18(4), Article 179. https://doi.org/10.3390/jrfm18040179

[3] Kim, Y., Lee, Y., Choe, M., Oh, S., & Lee, Y. (2024). Temporal graph networks for graph anomaly detection in financial networks. *arXiv preprint arXiv:2404.00060*. https://doi.org/10.48550/arXiv.2404.00060

[4] Lee, Y., Gong, J., & Kang, J. (2024). Embedding Byzantine fault tolerance into federated learning via consistency scoring. *arXiv preprint arXiv:2411.10212*. https://doi.org/10.48550/arXiv.2411.10212

[5] Luo, X., & Tang, B. (2024). Byzantine fault-tolerant federated learning based on trustworthy data and historical information. *Electronics*, 13(8), Article 1540. https://doi.org/10.3390/electronics13081540

[6] Madi, A., Stan, O., Mayoue, A., Grivet-Sébert, A., Gouy-Pailler, C., & Sirdey, R. (2021). A secure federated learning framework using homomorphic encryption and verifiable computing. In *2021 Reconciling Data Analytics, Automation, Privacy, and Security: A Big Data Challenge (RDAAPS)* (pp. 1–8). IEEE.

[7] Gehlhar, T., Marx, F., Schneider, T., Suresh, A., Wehrle, T., & Yalame, H. (2023). SafeFL: MPC-friendly framework for private and robust federated learning. In *2023 IEEE Security and Privacy Workshops (SPW)* (pp. 69–76). IEEE.

[8] Cai, Y., Ding, W., Xiao, Y., Yan, Z., Liu, X., & Wan, Z. (2023). SecFed: A secure and efficient federated learning based on multi-key homomorphic encryption. *IEEE Transactions on Dependable and Secure Computing*, 21(6), 3817–3833. https://doi.org/10.1109/TDSC.2023.3336977

[9] Devi, R. R. (2025). Reinforcement learning with graph neural network (RL-GNN) fusion for real-time financial fraud detection: A context-aware community mining approach. *Scientific Reports*, 15, Article 25200. https://doi.org/10.1038/s41598-025-25200-3

[10] Sattler, F., Müller, K.-R., Wiegand, T., & Samek, W. (2020). On the Byzantine robustness of clustered federated learning. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 8861–8865). IEEE.

[11] Zhang, L., Luo, Y., Bai, Y., Du, B., & Duan, L.-Y. (2021). Federated learning for non-IID data via unified feature learning and optimization objective alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 4420–4429).

[12] Fang, M., Zhang, Z., Hairi, Khanduri, P., Liu, J., Lu, S., Liu, Y., & Gong, N. (2024). Byzantine-robust decentralized federated learning. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1–18). ACM. https://doi.org/10.1145/3658644.3670307

[13] Chen, J. (2024). A Byzantine-fault-tolerant federated learning method using tree-decentralized network and knowledge distillation for internet of vehicles. In *IEEE Conference Proceedings*.

[14] Zeng, H. (2024). BSR-FL: An efficient Byzantine-robust privacy-preserving federated learning framework. *IEEE Transactions on Information Forensics and Security*.

[15] Zhang, F. (2024). Secure and decentralized federated learning framework with non-IID data based on blockchain. *Heliyon*. https://doi.org/10.1016/j.heliyon.2024.e03207

[16] Motie, S. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*. https://doi.org/10.1016/j.eswa.2023.121658

[17] Schreyer, M. (2022). Federated and privacy-preserving learning of accounting data in financial statement audits. *ACM Transactions on Intelligent Systems and Technology*.

[18] Zhu, B., Li, P., & Wang, R. (2025). A privacy-preserving federated learning scheme with homomorphic encryption and edge computing. *Alexandria Engineering Journal*.

[19] He, J., Wu, J., & Bao, W. (2022). Learning from non-IID data: Centralized vs. federated learning. *University of Illinois Academic Repository*.

[20] Paillier, P. (1999). Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology — EUROCRYPT ’99* (pp. 223–238). Springer.

[21] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (pp. 1273–1282). PMLR.

[22] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. In *Proceedings of Machine Learning and Systems* (Vol. 2, pp. 429–450).

[23] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.

[24] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. https://doi.org/10.1145/3298981

[25] Chen, X., Li, C., & Wang, D. (2023). Byzantine-robust federated learning with adaptive aggregation. *IEEE Transactions on Neural Networks and Learning Systems*.

[26] Wang, Z., et al. (2024). Graph neural networks for anomaly detection in financial transaction networks: A survey. *IEEE Transactions on Knowledge and Data Engineering*.