

Enterprise Data Governance Combining Federated Transfer Learning and Security Audit

Jinhui Wang¹, Yangli Lei², and Rongrui Wang³

¹ Instructor, Guangzhou College of Technology and Business, Guangzhou 510800, China, E-mail: wangjinhui@gzgs.edu.cn (corresponding author).

² Instructor, Digital Urban Operation and Governance Centre of Chengkou County, Chongqing 405900, China

³ Vice President, Guangzhou Shiyuan Electronic Technology Co., Ltd., Guangzhou, China

Project Management

Received June 1, 2026; Revised July 27, 2026, August 4, 2026, September 29, 2026; Accepted October 16, 2026
Available Online October 26, 2026

Abstract: Due to the problems of dispersed business data, inconsistent regulatory interpretations and limited cross-domain modeling at the enterprise level, this paper puts forward a data governance model based on federated transfer learning and security auditing. The preprocessing stage adds salted pseudonymization, entity masking, rare-category pooling, 30-minute window aggregation, and lineage/access-log feature construction to ensure reproducibility without revealing raw records. The method builds metadata, lineage, access behavior and rule-hitting features locally in each business domain and achieves cross-domain knowledge transfer via transfer gating and alignment loss. At the same time, it records summary hashing, responsible parties, policy consistency and lineage continuity in audit quality scores, adds a risk-suppression aggregation mechanism, and synchronously links the model update process to the governance traceability process. Experiments are based on financial, procurement, logistics, sales and compliance data from a large local group from January 2024 to December 2025. According to the results above, the proposed method achieves an AUC of 0.957, a macro-F1 score of 0.912, a recall of 0.926, and an audit trace completeness of 0.987. It is 2.1 and 2.3 percentage points higher than the best-performing comparative method. Extra imbalance, threshold-sensitivity and edge deployment tests show that the transfer gate reduces negative transfer in a 1:50 class imbalance, and audit-suppressed aggregation maintains real-time latency below 160ms under the tested edge configuration. Thus, cross-domain risk identification and strategy implementation can be achieved in a closed-loop manner without aggregating raw data, and a technical path for enterprise data governance that balances accuracy, compliance, and feasibility has been provided. A separate public HDFS log benchmark supplies a limited proxy stress test; it does not validate enterprise audit completeness or reproduce an accuracy advantage over FedAvg.

Keywords: Enterprise data governance; federated transfer learning; security audit; cross-domain risk identification; traceable governance.

Copyright © Journal of Engineering, Project, and Production Management (EPPM-Journal).
DOI: 10.32738/JEPPM-20278-0078

1. Introduction

As enterprises have added more data assets to their business systems over time and begun working with various partners globally, the older model of centralized data governance has shown three prominent deficiencies. First, cross-domain flow of raw data increases the risk of privacy leakage, unauthorized access and compliance failure. Second, the different data patterns, label distributions and governance standards across business domains lead to poor cross-domain model migration. Third, although the governance action requires "auditability", most audits remain confined to post-event reviews and have failed to establish a closed-loop mechanism that integrates model training and policy implementation. Now, the problems of identifying cross-domain risks, coordinating policies and ensuring full traceability without centralizing raw data are urgent and need to be addressed in the intelligent upgrade of enterprise data governance.

At present, the main problems of corporate governance are weak privacy protection and a gap between business semantics and audit responsibilities. A finance-domain access event, a procurement approval, a logistics delivery record and a compliance inspection may all describe the same risk chain but have different fields, labels and responsible parties. Therefore, a way is proposed to transfer reusable risk knowledge without altering local policy boundaries and to leave verifiable evidence.

Determine whether the rare compliance violation occurs in a single area, and if so, whether there are insufficient labeled examples for a local model; otherwise, related access and lineage patterns from other areas may be used to enhance

detection by ensuring that the source policy has a compatible meaning. Ordinary federated averaging does not check for compatibility during updates. Transfer gates limit the use of unsuitable source knowledge, and audit evidence determines whether a participating update affects the shared model.

The three main lines of research in the literature are not broad overviews of well-known concepts. Federated learning studies generalization, robustness, fairness, privacy attack surfaces and engineering defenses under non-independent and non-identically distributed data (Huang et al., 2024; Chen et al., 2025; Guembe et al., 2024). Data governance studies include stewardship, accountability allocation, metadata management, data quality control and lifecycle decision rights (Bernardo et al., 2024; Abraham et al., 2019; Khatri and Brown, 2010; Alhassan et al., 2018). Auditability and responsible AI studies further emphasize transparency, explainability, privacy protection, reliability, and organizational control in AI-enabled audit practice (Li and Goel, 2025; Kokina et al., 2025; Papagiannidis et al., 2025). The above studies, design ideas for trustworthy distributed learning, but they generally handle model training, governance enforcement and audit evidence separately.

Transfer-oriented studies explain why transfer learning is required in this case. One-shot federated transfer, fuzzy transfer under domain/category shifts, communication-efficient domain adaptation, federated graph transfer learning, and foundation-model-oriented federated transfer have shown that reusable representations can be shared even when the raw data are decentralized (Xiang et al., 2024; Li et al., 2024; Feng et al., 2025; Mai et al., 2025; Kang et al., 2025). Similarly, flexible federated learning frameworks and privacy-preserving industrial fault diagnosis studies have also applied similar ideas to reduce data heterogeneity among distributed equipment or organizations through lightweight communication and collaborative adaptation (Zeng et al., 2021; Shen et al., 2023; Guo et al., 2022). However, these studies do not convert audit evidence into a training-time constraint and thus cannot directly determine whether a risk decision is supported by complete evidence, consistent policy, continuous lineage and accountable responsibility mapping.

Light-communication federated diagnosis shows how heterogeneous clients can cooperate with limited communication (Li et al., 2026) and balance recovery and collaborative adaptation to address class imbalance and label-space shift across decentralized nodes (Yang et al., 2025). The above methods aim to address the problems of uneven domain size and rare risk here. Neither study has assessed policy consistency, lineage continuity and accountable evidence in the context of aggregation.

Therefore, this paper proposes a method for federated transfer learning and collaborative security auditing in the scenario of enterprise multi-domain data governance. Transfer learning can be applied locally to mitigate the noise introduced by domain-specific policy changes in centralized governance or standard federated averaging. At the same time, high-recall models for rare risks can be developed using decentralized raw data. The first three main contributions of this paper are as follows. First, it constructs a modeling architecture that integrates cross-domain governance for the five domains of finance, procurement, logistics, sales and compliance. Second, it builds a risk-suppression aggregation mechanism based on audit quality scores, thereby turning security auditing into a training constraint. Third, through two years of real-world business anonymization experiments, it verifies the all-around effectiveness and engineering feasibility of the proposed method in terms of risk identification accuracy, calibration, audit trail completeness, online latency and governance execution.

As shown in Table 1, the current methods either focus on cross-domain transferability without auditable governance mechanisms or concentrate on governance and auditing frameworks but do not address the transfer learning problem for heterogeneous, multi-domain features. Based on the above, we propose a joint model for these two kinds of capabilities.

2. Enterprise Scenarios and Problem Modeling

2.1. Multi-domain Business Data and Governance Goals


A large, multi-business group in a specific area has been selected as the research subject for an enterprise data governance case study. The company's five business units are Finance, Purchasing, Logistics, Sales and Compliance. Each business area has several data sources, such as master data, process events, access logs, work order texts and rule hit records. The data collection period is from January 2024 to December 2025. Given that direct aggregation of cross-domain data would pose serious compliance and security risks, this paper anonymizes, maps definitions, performs windowed aggregation and encodes rules on sensitive fields in each domain and only uploads parameter updates and signed audit evidence. Figure 1 is the general structure of the method introduced in this paper.

Decentralization of Data is a form of governance, not merely a storage problem. Finance, procurement, logistics, sales and compliance domains maintain their own policy dictionaries, evidence repositories and sensitive event logs, and only exchange parameter updates, aggregation weights and signed evidence digests in the collaboration layer. This Design does not use a central data lake for high-risk records and supports collaborative adaptation when domains have inconsistent sample sizes, feature spaces and risk-label distributions.

The de-identification process is executed locally before feature construction. Organization, account, order and resource IDs, as well as device identifiers, are transformed via salted HMAC-SHA-256 pseudonymization with domain-specific salt rotation. Free-text work orders and log messages are processed by entity masking for names, phone numbers, addresses, account numbers, and contract numbers. Rare categories are pooled according to a minimum-frequency threshold of $k = 10$. Timestamp fields are converted into relative time gaps and time-of-day buckets, and numerical fields are winsorized at the 1st and 99th percentiles before robust standardization. Raw identifiers, free-text content, and complete lineage paths are not uploaded to the collaboration layer.

Table 1. Comparison of recent related studies and positioning of this paper

Ref.	Core idea	Transfer capability	Auditability	Applicable data	Main limitation
(Huang et al., 2024)	Survey of FL generalization, robustness and fairness	Weak	Low	Multi-domain distributed data	Focuses on benchmarking rather than governance execution
(Barbureau et al., 2025)	Governance framework for organizational FL archetypes	None	Medium	Cross-organizational collaboration	Lacks learning-level optimization mechanism
(Li and Goel, 2025)	AI auditability dimensions and auditor readiness	None	High	AI lifecycle records	Does not solve cross-domain data heterogeneity
(Xiang et al., 2024)	One-shot federated transfer via random forests	Medium	Low	Structured heterogeneous tables	Limited support for dynamic policy evidence
(Feng et al., 2025)	Communication-efficient federated domain adaptation	High	Low	Shifted source-target domains	Audit evidence not integrated into aggregation

**Fig. 1.** Framework for collaborative data governance involving federated transfer learning and security auditing

From the perspective of governance goals, this paper does not view enterprise data governance as merely a metadata inventory or a quality score. Rather, it is an ongoing closed-loop process of "cross-domain risk identification, strategy execution, evidence retention, and feedback correction." That is to say, the model should recognize high-risk events, such as sensitive access, abnormal exports, master data mismatches and retention period violations, and the governance system needs to provide interpretable processing decisions and leave a complete chain of evidence for future audits.

2.2. Security Audit Constraints and Symbol Definitions

For reproducibility, lineage and access-log features are encoded by a fixed local pipeline. A lineage edge is defined by source-node type, target-node type, operation type, policy tag, upstream rule status, time gap and hash-continuity flag, and these edge features are aggregated within a window using count, ratio, last_observed, maximum_depth and entropy statistics. An access-log event has attributes such as actor role, resource sensitivity level, action type, device segment, approval status, geographic consistency, and time-of-day bucket. Abnormality features include privilege-distance, out-of-scope access ratio, burst-export z-score, failed-approval ratio, and off-hour frequency. Missing values are mapped to reserved tokens, categorical mappings are frozen in signed local metadata, and continuous features are standardized only by local-domain statistics.

For every 30-minute advance of five minutes, sort the local events by time and connect them to the pseudonymous subject and resource. Edge count is the number of observed lineage transitions; rule-hit ratio is the proportion of edges with the relevant rule tag among all observed edges; maximum depth is the length of the longest verified upstream path, and

categorical entropy is calculated based on the observed edge-type frequencies. Access ratios are the number of access events in the same window divided by the number of accesses. Burst-export z-score is the ratio of the window's export count to the local-domain reference distribution, and off-hour frequency counts events outside the applicable local work schedule. An empty denominator returns zero and a missingness flag; unknown categories use the reserved token. Mapping and reference statistics are local and set before evaluation.

Let $d \in \{1, \dots, M\}$ denote a business domain, k denote the sample index, and $x_{d,k}$ denote the de-identified window sample. The local encoder $f_{\theta_d}(\cdot)$ transforms multimodal governance features into the representation $h_{d,k} = f_{\theta_d}(x_{d,k})$, enabling cross-domain knowledge transfer when original records are not shared. The corresponding representation learning process is written as:

$$h_{d,k} = f_{\theta_d}(x_{d,k}) \quad (1)$$

In Eq. (1), $h_{d,k}$ represents the local representation of the k -th sample in domain d . The $x_{d,k}$ represents the de-identified input sample. $f_{\theta_d}(\cdot)$ represents the local encoder, and θ_d represents the encoder parameters. This notation is used consistently in the following formulas for transfer, aggregation, and risk scoring.

Given the differences in field structures, event frequencies and tag distributions across business domains, federated averaging may introduce bias into the target domain's representation. Therefore, this paper will further use mean alignment to align the representations of the source and target domains.

$$L_{\text{align}} = \sum_{d \neq t} \|\mu_d - \mu_t\|_2^2, \quad \mu_d = \frac{1}{N_d} \sum_{k=1}^{N_d} h_{d,k} \quad (2)$$

In Eq. (2), L_{align} represents the cross-domain alignment loss, μ_d represents the mean representation of business domain d , N_d represents the number of local samples, t represents the target domain, and $\|\cdot\|_2$ represents the L_2 norm. This term suppresses feature-center drift between source and target domains and provides a measurable boundary for deciding whether transfer is beneficial.

3. Collaborative Approach to Federated Transfer Learning and Security Audit

3.1. Transfer Gating and Local Objective Function

To avoid negative transfer among the business domains, transfer gating is used to selectively add shared memory. The gating strength is the mean of the source and target domains; together, they determine how knowledge transferability changes with business similarity. The corresponding calculations are as follows:

The gating threshold τ is selected on a validation set rather than manually fixed. The implementation first computes source-target similarity from representation cosine similarity, policy-label overlap, and maximum mean discrepancy, then scans τ over the validation candidates, including 0.50, 0.62 and 0.75 as reported in Table 6. The selected value is the smallest τ whose validation Macro-F1 is within 0.3 percentage points of the best result and whose ECE does not increase; in the experiments, $\tau = 0.62$. If the off-diagonal transfer AUC of a source-target pair is lower than the corresponding local-only AUC or the distribution discrepancy exceeds the 90th percentile of all pairs, the gate suppresses the shared-memory contribution for that pair to avoid negative transfer.

For each ordered source-target pair, the decision compares the transfer AUC to the target local-only AUC on the same validation set. If the transfer is poor, then the source memory will not be used for that target; otherwise, a discrepancy rule will be applied to these pairs with very different feature distributions. The two choices are not combined with a single global threshold set by Macro-F1 and ECE; therefore, a good overall score may still indicate a bad transfer direction.

$$g_d = \sigma(W_g[\mu_d \parallel \mu_t] + b_g), \quad z_{d,k} = g_d \odot h_{d,k} + (1 - g_d) \odot m \quad (3)$$

In Eq. (3), g_d represents the migration gating vector of domain d , W_g represents the gating parameter matrix, b_g represents the bias term, μ_d and μ_t represent the source and target domain means, m represents the shared migration memory, and $z_{d,k}$ represents the gated representation after selective memory injection.

Based on the above, this paper combines classification loss, cross-domain alignment loss, and a parameter stability term to form the objective for local training, thereby balancing the trade-off between task accuracy and transfer stability.

$$L_d = L_{\text{cls}} + \lambda_1 L_{\text{align}} + \lambda_2 \|\theta_d - \theta_g\|_2^2 \quad (4)$$

In Eq. (4), L_d denotes the local total loss in domain d , L_{cls} denotes the classification loss for governance risk identification, λ_1 and λ_2 are trade-off coefficients, θ_g denotes the global model parameters, and $\|\theta_d - \theta_g\|_2^2$ is the parameter stability term. The stability term limits extreme local updates and improves convergence under heterogeneous domains.

Compared with random-forest one-shot federated transfer (Xiang et al., 2024), fuzzy transfer under domain/category shifts (Li et al., 2024), and communication-efficient FedRF-TCA (Feng et al., 2025), this method changes inter-domain migration from a fixed map to a gated adaptive control and further introduces audit quality into global aggregation. Therefore, the transfer strength is jointly determined by feature similarity, policy compatibility and audit reliability and not solely by representation distance.

3.2. Audit Quality Modeling and Risk Suppression Aggregation

Enterprise data governance cannot obtain "accurate" model results and cannot answer audit questions such as "why it is credible, who triggered it, and whether it is within the scope of the system" (Mai et al., 2025). Therefore, the four indices of audit quality used in this paper are evidence integrity, strategy consistency, lineage continuity and traceability of the responsible party. Fig. 2 is the connection path of training, audit and governance.




Fig. 2. Linkage process of local training, audit documentation and governance implementation

The Audit Quality Scores are as follows:

$$q_d = \omega_1 c_d + \omega_2 p_d + \omega_3 l_d + \omega_4 u_d \quad (5)$$

In Eq. (5), q_d denotes the audit quality score for domain d . The terms c_d , p_d , l_d , and u_d measure evidence completeness, policy consistency, lineage continuity, and responsible-party traceability, respectively. The coefficients $\omega_1 - \omega_4$ are nonnegative and sum to 1. Each component is normalized to $[0, 1]$ before weighted aggregation.

The four dimensions are calculated as follows. Evidence completeness c_d equals the ratio of available evidence fields to required fields, where the required fields include subject, time, operation, object, policy version, lineage node, and digest hash. Policy consistency p_d equals one minus the ratio of conflicting rule hits to all rule hits in the current window. Lineage continuity l_d equals the ratio of verified upstream/downstream lineage edges to expected lineage edges. Responsible-party traceability u_d equals the ratio of risk events that can be linked to an accountable role, approval chain, or system account. The default weight vector is $\omega = (0.30, 0.25, 0.25, 0.20)$, and sensitivity analysis is reported in Section 4.3.

All audit components are calculated based on the same local window before aggregation. The required-field set defines the denominator for evidence completeness, including fields where the value is missing; a window with no rule hit is assigned zero conflicts and a missing-rule indicator, rather than being treated as proof of policy consistency. An expected lineage edge is only verified if its recorded digest and predecessor linkage match. A responsible role is counted only when an event can be associated with a recorded actor or an approval chain. The above rules differentiate between missing evidence and a verified absence of violations, thus avoiding an artificially high audit score for a sparse log.

Given that high-risk or low-quality audit domains should not be given too much weight in the global update, this paper incorporates audit quality and risk-mitigation factors into the aggregation coefficient.

$$\tilde{W}_d = \frac{N_d \exp(-\beta r_d) q_d}{\sum_{j=1}^M N_j \exp(-\beta r_j) q_j} M \quad (6)$$

In Eq. (6), $\tilde{W}_d$ represents the normalized aggregation weight of domain d , N_d represents the number of samples in domain d , β represents the risk-suppression strength, r_d represents the audit risk of domain d , q_d represents audit quality, and M represents the number of domains participating in aggregation. A larger β reduces the influence of domains with higher audit risk, while q_d prevents low-evidence updates from dominating the global model.

Thus, the global model update can be expressed as:

$$\theta_g^{(t+1)} = \sum_{d=1}^M \tilde{W}_d \theta_d^{(t+1)} \quad (7)$$

In Eq. (7), $\theta_g^{(t+1)}$ represents the global parameters after communication round $t + 1$, $\theta_d^{(t+1)}$ represents the local update result of domain d , and $\tilde{W}_d$ represents the aggregation weight related to data volume, audit quality, and risk exposure.

Compared with blockchain-only or chained-log approaches that provide post-hoc traceability (Zhang et al., 2024a; Zhang et al., 2024b), model-integrity studies require that the training process itself be verifiable. Therefore, the aggregation scheme changes auditability from an ex-post reporting requirement to a training-time constraint by adjusting update weights based on the completeness and risk of evidence.

3.3. Risk Scoring and Closed-Loop Governance Actions

Obtain the global model and then, based on the trigger strength and the extent of audit anomalies, generate a final governance risk score to distribute actions such as freezing, isolation, review and rollback. The final risk score is given by this paper.

$$R_{d,k} = \eta_1 \hat{y}_{d,k} + \eta_2 v_{d,k} + \eta_3 a_{d,k} \quad (8)$$

In Eq. (8), $R_{d,k}$ represents the final risk score of sample k in business domain d , $\hat{y}_{d,k}$ represents the model risk probability, $v_{d,k}$ represents rule-trigger strength, $a_{d,k}$ represents audit anomaly degree, and η_1 to η_3 are fusion coefficients. The coefficients are estimated on the validation set under the constraint $\eta_1 + \eta_2 + \eta_3 = 1$.

When $R_{d,k}$ exceeds the high-risk threshold, the system issues isolation and manual-review signals. Samples in the medium-risk zone are placed in an escalation queue and monitored for subsequent promotion to the high-risk category. Low-risk samples remain within the tracking range for continuous observation. Table 2 shows the relationship between core governance rules and audit mapping.

Table 2. Mapping of governance rules, audit granularity and actions

Rule type	Trigger signal	Audit granularity	Governance action	Risk weight
Master-data mismatch	inconsistent key-value mapping across domains	field + process node	freeze synchronization and request correction	0.18
Cross-domain access overflow	privilege used outside approved scope	account + resource + time window	temporary token revocation	0.24
Sensitive export burst	abnormal download volume or destination drift	actor + channel + file group	quarantine and review	0.27
Late retention purge	expired data still retained after policy deadline	table + retention policy	rollback and mandatory clean-up	0.15
Privilege escalation chain	multi-step elevation followed by risky action	actor chain + workflow evidence	full escalation to compliance team	0.16

As shown in Table 2, the governance measures set forth in this paper are not merely alarms. Map rule triggers, audit evidence and policy handling to a one-to-one correspondence, thus providing an actionable interface for the results of risk identification.

Rule-action mapping is in line with the data-governance study that has identified the necessary mechanisms of stewardship, decision rights, accountability, lifecycle control and data-quality assurance (Bernardo et al., 2024; Abraham et al., 2019; Khatri and Brown, 2010; Alhassan et al., 2018) and also meets the requirements of AI audits for transparent evidence chains and responsibility allocation (Li and Goel, 2025; Kokina et al., 2025; Papagiannidis et al., 2025). Thus, the governance rules are not independent business heuristics but operationalized controls for auditable data governance.

4. Experimental Design and Results Analysis

4.1. Experimental Data, Comparison Methods and Evaluation Indicators

The experimental data is based on real-world business cases, and deduplicated business data from January 2024 to December 2025 has been obtained. To avoid directly using organization names, system numbers and other critical business rules, only the hierarchical structure information of the business domains is kept in this paper. All the samples were built using a 30-minute sliding window and a 5-minute step size; robust standardization was applied to the continuous variables, and dictionary mapping and category pruning were employed for discrete fields. Extract rules to obtain structured governance features from the log and work order text. Table 3 shows the distribution of the five business domains.

Table 3. Composition and risk distribution of multi-service domain experimental data

Domain	Time span	Sample windows	Feature groups	Positive rate	Sensitive fields
Finance	2024.01-2025.12	126,400	84	2.8%	14
Procurement	2024.01-2025.12	98,700	79	3.4%	11
Logistics	2024.01-2025.12	214,500	91	4.1%	9
Sales	2024.01-2025.12	187,200	88	3.1%	12
Compliance	2024.01-2025.12	64,900	76	6.3%	16

The comparison methods are Local-only, FedAvg, FedProx, RF-TCA-style federated transfer (Xiang et al., 2024), FGTL graph transfer learning (Mai et al., 2025), and the method proposed in this paper. The selected evaluation indicators are AUC, Macro-F1, Recall, Expected Calibration Error (ECE) and Audit Completeness. Audit Completeness refers to the proportion of samples that have full evidence of subjects, times, policies, lineages and digest hashes.

Run all experiments five times with different random seeds. Table 4 shows the mean of each method, and the available run-to-run standard deviations for the proposed method are also provided along with its results. Paired tests with FGTL use the five seed-matched observations for each metric. The governance indicators are based on the same test period and thus have the same denominator.

4.2. Overall Performance Comparison and Cross-Domain Migration Effect

Table 4 presents the overall results for the different methods on the five-domain joint test set. As shown in Figure 1, local-only methods can achieve reasonable results within a single domain, but their AUC and Recall are relatively low because they cannot leverage cross-domain knowledge. FedAvg and FedProx show some improvement after sharing statistical information, but lack explicit transfer mechanisms and are thus more sensitive to cross-domain distribution shifts. RF-TCA and FGTL are used to enhance the two methods above. On the other hand, our proposed method achieves an AUC, Macro-F1 and Recall of 0.957, 0.912 and 0.926, respectively and reduces ECE to 0.021. Thus, the model is not only more accurate but also more stable.

Fig. 3 shows the AUC matrix of ordered source-target domain pairs. The diagonal elements are the performances in the corresponding domains; the off-diagonal elements are the transfers to other domains. Transfers from compliance to finance and procurement are relatively advantageous, and logistics and sales have shown positive transfers in both directions. Therefore, policy compatibility, process relevance and audit standards will determine whether a particular source representation can be incorporated into a given target model.

Over the five runs, the proposed method achieved AUC 0.957 ± 0.003 , Macro-F1 0.912 ± 0.004 , Recall 0.926 ± 0.005 , ECE 0.021 ± 0.002 and audit completeness 0.987 ± 0.003 (mean $\pm$ standard deviation). The reported paired t-test p-values for FGTL are 0.006, 0.008, 0.004 and 0.003 for AUC, Macro-F1, Recall and audit completeness, respectively. These comparisons cover five pairs of seeds and do not replace the reporting of run-level variability for all baselines.

Table 4. Comparison of the general results of different methods on the joint test set

Method	AUC	Macro-F1	Recall	ECE	Audit completeness
Local-only	0.861	0.781	0.744	0.071	0.812
FedAvg	0.894	0.825	0.801	0.058	0.845
FedProx	0.907	0.838	0.816	0.051	0.853
RF-TCA FTL (Xiang et al., 2024)	0.923	0.867	0.851	0.039	0.861
FGTL (Mai et al., 2025)	0.936	0.889	0.878	0.033	0.905
Proposed	0.957 ± 0.003	0.912 ± 0.004	0.926 ± 0.005	0.021 ± 0.002	0.987 ± 0.003

The uncertainty values in the proposed-method row are standard deviations over five runs, and the other rows contain only the reported means. Baseline run-level dispersions are not available in the aggregated results presented here, so their variability cannot be inferred from these means or the paired-test p-values.




Fig. 3. AUC matrix of cross-domain migration among the five business domains

A negative-transfer boundary test was performed by resampling positive windows to class-imbalance ratios of 1:20, 1:50, and 1:100. Under a 1:50 imbalance, the new method achieved an AUC of 0.943 and a recall of 0.887. FGTL was 0.914/0.823. Under a 1:100 imbalance, it maintains an AUC of 0.931 and a recall of 0.842, compared with 0.902 and 0.781 for FGTL. The negative-transfer rate is defined as the proportion of source-target pairs for which the transfer AUC is lower than the target-domain local-only AUC. It is 18.0% without a gate and 4.0% with the proposed gate.

The 1:20, 1:50 and 1:100 settings probe progressively rarer positive events. Based on the above 1:50 and 1:100 comparisons, recall decreases as the proportion of the positive class increases, even though gated transfer remains superior to FGTL. The negative-transfer rate is calculated for each ordered source-target pair based on the local-only result at that target. Thus, these pairwise measures are not the same as the total AUC over the pooled test set.

4.3. Ablation Experiment and Convergence Characteristic Analysis

To determine whether the five modules are independent, this paper introduces five variations: without the transfer gate, without the audit weight, without the rule penalty, without the lineage feature, and without the audit constraint. Table 5 shows that the complete model has the best AUC and Macro-F1. Removing the entire audit constraint results in the largest audit loss. Thus, auditability should be viewed as a training constraint rather than just an output-reporting function. Without a transfer gate, classification accuracy is relatively low. Therefore, to mitigate the negative impact of domain shift in heterogeneous environments, only selected cross-domain features need to be injected.

The no-audit-constraint variant simultaneously excludes audit-quality weighting and audit-risk suppression from the aggregation rule, rather than disabling only one audit module. Its AUC drops from 0.957 to 0.939, and the audit loss increases from 0.041 to 0.112. The separate no-audit-weighting and no-rule-penalty variants in Table 5 identify these contributions. In the threshold and suppression sweeps of Table 6, the retained setting is evaluated based on both discrimination and audit loss; therefore, increasing either metric alone is not sufficient to select a parameter.

Table 6 reports key-parameter sensitivity for the transfer threshold τ and risk-suppression strength β . A lower τ admits more source-domain memory and increases negative-transfer risk, whereas a higher τ blocks useful cross-domain knowledge. A larger β improves audit loss but increases latency because more updated weights need to be re-normalized after risk suppression. The selected $\tau = 0.62$ and $\beta = 0.70$ provide the best balance among AUC, recall, calibration, audit loss, and latency.


Fig. 4 is the convergence curve of the communication rounds. The stable-transfer period begins after about round 45. FedAvg and FedProx improve in the early rounds by sharing statistical information, but their macro-average scores remain lower than those of the proposed method because they do not distinguish between helpful and risky transfer. Adding audit quality to the aggregation weight does not cause convergence instability. Rather, it reduces low-evidence and high-risk updates, thereby improving both the stability and interpretability of the whole model.

Table 5. Comparison of ablation results and system overhead for different modules

Variant	AUC	Macro-F1	Latency (ms)	Comm. cost	Audit loss
w/o transfer gate	0.934	0.886	147	1.00×	0.086
w/o audit weighting	0.941	0.894	141	0.98×	0.074
w/o rule penalty	0.946	0.901	143	1.01×	0.067
w/o lineage features	0.949	0.905	145	1.03×	0.059
w/o audit constraint	0.939	0.891	139	0.97x	0.112
Full model	0.957	0.912	149	1.05×	0.041

Table 6. Parameter sensitivity of transfer threshold and audit risk-suppression strength

Test condition	Value	AUC	Recall	Latency (ms)	Audit loss
τ sweep	0.50	0.951	0.914	146	0.049
τ sweep	0.62	0.957	0.926	149	0.041
τ sweep	0.75	0.949	0.907	148	0.044
β sweep	0.30	0.953	0.917	145	0.053
β sweep	0.70	0.957	0.926	149	0.041
β sweep	1.20	0.954	0.921	154	0.039

**Fig. 4.** Convergence process of different methods and changes in audit completeness

4.4. Governance Implementation Effectiveness and Audit Trail Quality

The domain-edge node in the online deployment test had an 8-core CPU, 16 GB of memory and a 1 Gbps internal network. Each local batch contained 1,024 window samples and uploaded only parameter deltas and signed evidence digests during the communication round. Table 7 shows resource utilization in all deployment modes. The calculation of the audit weight adds 6-8ms of latency in the tested environment due to evidence-digest verification and risk-weight normalization, but the end-to-end latency remains below the enterprise governance platform's 200ms service limit.

Table 6 shows the latency effect of the risk-suppression parameter at different values: 0.30, 0.70 and 1.20 correspond to latencies of 145, 149 and 154ms, respectively. At 0.70, the corresponding AUC is 0.957, and the audit loss is 0.041. Table 7 separately lists the performance of a domain edge CPU (149ms, 1.8 GB memory, 43% CPU load), digest caching (141ms, 1.9 GB memory, 41% CPU load), and a low-resource edge (176ms, 1.3 GB memory, 61% CPU load). The cached configuration uses more memory to reduce response time; the low-resource configuration uses less memory and incurs a higher CPU load.

The company will also consider whether the governance system can reduce or prevent omissions and violations. Fig. 5 shows how the system's governance efficiency changed from Q1 2024 to Q4 2025. Due to continuous tightening of policy definitions and model updates in many places, the missed violation rate has dropped from 10.4% initially to 3.4% now; the backlog rate of the upgrade queue has fallen from 18.2% initially to 5.9% now; and the completeness of the audit trail has improved from 0.881 originally to 0.987 currently. The simultaneous improvement of the three indicators shows that the method proposed in this paper does not enhance the identification model at a single point but improves the efficiency of the entire governance closed loop of "identification-implementation-monitoring-feedback".

Table 7. Edge-end resource consumption under different deployment configurations

Deployment mode	Latency (ms)	Memory (GB)	CPU load	Upload/round	Throughput
Central simulation	132	2.4	38%	3.1 MB	7,760/s
Domain edge CPU	149	1.8	43%	2.6 MB	6,870/s
Edge + digest cache	141	1.9	41%	2.6 MB	7,120/s
Low-resource edge	176	1.3	61%	2.6 MB	5,430/s

An additional case study has found that the three frequent high-risk scenarios are: master data mismatches between the finance and compliance departments, abnormal exports by the sales department, and cross-domain permission violations. The old, federated model can provide a probability of risk, but cannot explain why a trigger occurs or which factors are responsible. The way shown in this paper can trace back to the specific responsible party, the rule that was triggered, upstream and downstream lineage nodes, and parameter update summaries, and has resulted in an average reduction of 17.6% in manual review time and a 12.8% decrease in the false trigger rate of governance actions.




Fig. 5. Changes in the rate of missed detections, backlog rate and audit trail completeness

4.5. Public Log Benchmark and Limits of Audit-Proxy Evaluation

HDFS_v1 from Loghub (Zhu et al., 2023) was used for an independent public-log benchmark. The original log has 11,175,629 event lines, and among the 575,061 labeled blocks in the block-level event-occurrence matrix, 16,838 have been marked as 'Fail'. The result of this benchmark is a block failure, not enterprise policy non-compliance. To restore the order of time, each block was given the timestamp of its first appearance in the original log. A fixed SHA-256 rule selected about a quarter of the blocks and distributed them to five simulated clients via a separate hash. Chronological training, validation and test partitions contain 86,298, 28,636 and 28,825 blocks, respectively, with 2,821, 979 and 399 failures. A block occurs in only one partition and one client.

The 29 published event counts were transformed into $\log(1 + \text{count})/4$ and then provided as input to a 30-parameter linear risk head. Local-only, FedAvg and FedProx were also tested with a limited audit-weight proxy. All methods used ten communication rounds when available, four local minibatch steps per round, a batch size of 1,024, a learning rate of 0.045, and the same positive-class weight of five. FedProx had a proximal coefficient of 0.10. Model selection, including classification thresholds and the gate decision, was based solely on validation blocks; the test set remained unchanged for

seeds 11, 23, 37, 51 and 79. The reported standard deviations are over these five stochastic runs on one fixed split, not across independently sampled enterprises.

The public log does not have an enterprise approval chain, policy version or verifiable responsible-party mapping. For a controlled missing-evidence test, individual event-count fields were masked at fixed client rates of 0%, 8%, 16%, 24% and 32%. The retained event-count fraction from the training block was used as a telemetry-coverage proxy for audit quality, and the training failure rate served as a risk proxy. Aggregation weights were proportional to the number of clients' samples times coverage and $\exp(-0.70 \times \text{failure rate})$. For the gated variant, a shared update was combined with local refinement only if the validation AUC did not drop below the local alternative and the feature-mean cosine was greater than 0.62. This is a deliberately limited operating analog of the proposed mechanism, and its coverage measure is not the audit completeness reported for the enterprise case.

Table 8 does not show an external accuracy advantage for the audit-weight proxy. It has an AUC of 0.769, which is lower than FedAvg's 0.776. Its PR-AUC is 0.332, lower than FedAvg's 0.433. Seed-paired t-tests on this fixed split yield $p = 7.2\text{e-}06$ for AUC and $p = 7.4\text{e-}06$ for PR-AUC; the p-values quantify run-to-run differences under the chosen implementation and cannot establish generalization to organizations. ECE remains greater than 0.30 across all listed methods, indicating poor calibration. The local-only system has the highest pooled PR-AUC among these, and the pooled rankings may also differ due to differing client score scales.

Table 8. Five-seed HDFS_v1 test results for the restricted public-log benchmark (mean $\pm$ standard deviation)

Method	AUC	PR-AUC	Macro-F1	Recall	ECE
Local-only	0.770 $\pm$ 0.0010	0.476 $\pm$ 0.0116	0.755 $\pm$ 0.0307	0.402 $\pm$ 0.0729	0.322 $\pm$ 0.0005
FedAvg	0.776 $\pm$ 0.0003	0.433 $\pm$ 0.0000	0.735 $\pm$ 0.0000	0.358 $\pm$ 0.0000	0.323 $\pm$ 0.0005
FedProx	0.776 $\pm$ 0.0003	0.433 $\pm$ 0.0000	0.735 $\pm$ 0.0000	0.358 $\pm$ 0.0000	0.323 $\pm$ 0.0005
Audit-weight proxy	0.769 $\pm$ 0.0004	0.332 $\pm$ 0.0075	0.708 $\pm$ 0.0080	0.629 $\pm$ 0.0000	0.303 $\pm$ 0.0009
No gate	0.776 $\pm$ 0.0003	0.433 $\pm$ 0.0000	0.735 $\pm$ 0.0002	0.358 $\pm$ 0.0000	0.323 $\pm$ 0.0005

Table 9. Proxy-mechanism ablations on HDFS_v1 (five-seed mean $\pm$ standard deviation)

Variant	AUC	PR-AUC	Macro-F1	Recall	ECE
Audit-weight proxy	0.769 $\pm$ 0.0004	0.332 $\pm$ 0.0075	0.708 $\pm$ 0.0080	0.629 $\pm$ 0.0000	0.303 $\pm$ 0.0009
No gate	0.776 $\pm$ 0.0003	0.433 $\pm$ 0.0000	0.735 $\pm$ 0.0002	0.358 $\pm$ 0.0000	0.323 $\pm$ 0.0005
No audit quality proxy	0.769 $\pm$ 0.0004	0.332 $\pm$ 0.0077	0.709 $\pm$ 0.0091	0.628 $\pm$ 0.0022	0.303 $\pm$ 0.0009
No risk suppression	0.769 $\pm$ 0.0004	0.332 $\pm$ 0.0076	0.708 $\pm$ 0.0080	0.629 $\pm$ 0.0000	0.303 $\pm$ 0.0009
No audit proxies	0.769 $\pm$ 0.0004	0.332 $\pm$ 0.0077	0.709 $\pm$ 0.0091	0.628 $\pm$ 0.0022	0.303 $\pm$ 0.0009

As shown in Table 9, the proxy ablations do not provide evidence that the excluded telemetry coverage or risk factors consistently enhance anomaly detection. In particular, the no-gate variant achieves the same rounded AUC as FedAvg. The no-audit-quality variant is almost indistinguishable from removing both proxies. Based on the above results, the public benchmark can only be used for an implementation stress test and should not be regarded as a verification of complete audit trails. RF-TCA-style random-forest transfer and FGTL graph transfer were not reimplemented. Thus, there is no original transfer implementation or verified cross-domain graph required for a like-for-like comparison. Neither method has been replaced by the other.

Only the training partition was resampled in Table 10; the validation and test blocks, as well as the five random seeds, were kept fixed. At all three ratios, the proxy method had a lower mean AUC than FedAvg, but PR-AUC and recall did not consistently follow this order. Therefore, the extreme-ratio comparison identifies sensitivity to training prevalence rather than a general improvement in rare-risk detection. Since the dataset is a system log with simulated clients and masked features, these numbers cannot be directly compared with the earlier five-business-domain results or the 149 ms enterprise edge deployment figure. Reproduction files contain the specific block-split procedure, hyperparameters and per-run measurements.

5. Conclusion

This paper puts forward a cooperative approach of federated transfer learning and security auditing for multi-business domain data governance scenarios in enterprises and draws the following conclusions: This paper establishes a closed-loop

system for cross-domain migration modeling, audit quality quantification and governance action execution and shifts security auditing from a post-deployment inspection to constraints in the training phase. The design can help to solve the problem of "learning without auditing, auditing without migrating" in the old model of central government administration.

Table 10. Training-only class-ratio sensitivity with unchanged HDFS_v1 validation and test sets

Train ratio	Method	AUC	PR-AUC	Recall	Train blocks
1:20	FedAvg	0.775 ± 0.0005	0.434 ± 0.0001	0.647 ± 0.0000	59241
1:20	Audit-weight proxy	0.767 ± 0.0017	0.405 ± 0.0332	0.595 ± 0.0762	59241
1:50	FedAvg	0.776 ± 0.0002	0.431 ± 0.0001	0.629 ± 0.0000	85145
1:50	Audit-weight proxy	0.766 ± 0.0044	0.382 ± 0.0875	0.634 ± 0.0031	85145
1:100	FedAvg	0.774 ± 0.0002	0.385 ± 0.0001	0.634 ± 0.0027	84310
1:100	Audit-weight proxy	0.764 ± 0.0031	0.474 ± 0.0673	0.610 ± 0.0494	84310

The best method in this paper achieves an AUC of 0.957, a Macro-F1 score of 0.912, and a Recall of 0.926 on the joint test set. It is 2.1, 2.3, and 4.8 percentage points higher than those of the best comparative method, respectively. At the same time, the ECE is as low as 0.021, indicating that this method also achieves good classification accuracy and probability calibration performance in a complex, heterogeneous enterprise environment.

In terms of audit and governance effectiveness, the audit trace completeness of the complete model reached 0.987. The quarterly violation omission rate dropped from 10.4% to 3.4%, and the upgrade queue backlog rate fell from 18.2% to 5.9%. Therefore, by adding completeness of evidence, strategy consistency and lineage continuity to the aggregation weights, the credibility of the governance chain can be strengthened, and the speed of execution improved.

From an engineering perspective, the communication overhead of the complete model is 1.05 times that of the baseline model, and the average batch latency is 149ms. The additional cost is within the acceptable range for enterprise online governance systems. Future research can be extended in three targeted directions. First, cross-organizational federated collaboration should support heterogeneous legal bases, inconsistent data-steward roles, and negotiable audit evidence schemas. Second, large-model data governance can incorporate prompt records, retrieval evidence, tool-call lineage, model-output provenance, and policy-alignment traces into the same audit-quality scoring framework. Third, long-time data drift should be handled through continuous drift detectors, temporal calibration, periodic re-weighting of β and τ and replay-safe audit sampling so that governance decisions remain stable amid changing business processes and policy definitions. In the additional public-log experiment, the audit-weight proxy did not outperform FedAvg in AUC or PR-AUC. Its simulated clients and missing-field proxies cannot achieve the desired enterprise audit and deployment results; therefore, independent policy evidence and resource-matched trials are still required.

Author Contributions

Jinhui Wang contributed to conceptualization, methodology, project administration, funding acquisition, supervision, writing of the original draft, writing, reviewing and editing. Yangli Lei contributed to software, formal analysis, investigation, data curation, visualization and writing of the original draft. Rongrui Wang contributed to data curation, formal analysis, writing, reviewing and editing.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

Declaration of Artificial Intelligence (AI) Tools

DeepSeek was used for language editing. ChatGPT assisted with the preparation and interpretation of the public-log benchmark code and with language editing for this revision. The authors are responsible for verifying the dataset, code, numerical results and manuscript before submission.

References

- Abraham, R., Schneider, J., and vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. *International Journal of Information Management*, 49, 424-438. DOI: 10.1016/j.ijinfomgt.2019.07.008.
- Alhassan, I., Sammon, D., and Daly, M. (2018). Data governance activities: a comparison between scientific and practice-oriented literature. *Journal of Enterprise Information Management*, 31(2), 300-316. DOI: 10.1108/JEIM-01-2017-0007.

- Barbureau, T., Delgado Fernandez, J., and Potenciano Menci, S. (2025). The governance of federated learning: a decision framework for organizational archetypes. *Data and Policy*, 7, 53. DOI: 10.1017/dap.2025.10020.
- Bernardo, B. M. V., São Mamede, H., Barroso, J. M. P., and dos Santos, V. M. P. D. (2024). Data governance and quality management - Innovation and breakthroughs across different fields. *Journal of Innovation and Knowledge*, 9(4), 100598. DOI: 10.1016/j.jik.2024.100598.
- Chen, C., Liu, J., Tan, H., Li, X., Wang, K. I. K., Li, P., Sakurai, K., and Dou, D. (2025). Trustworthy federated learning: privacy, security, and beyond. *Knowledge and Information Systems*, 67(3), 2321-2356. DOI: 10.1007/s10115-024-02285-2.
- Feng, Z., Wang, Y., Li, J., Yang, F., Lou, J., Mi, T., Qiu, R. C., and Liao, Z. (2025). Robust and Communication-Efficient Federated Domain Adaptation via Random Features. *IEEE Transactions on Knowledge and Data Engineering*, 37(3), 1411-1424. DOI: 10.1109/TKDE.2024.3510296.
- Guembe, B., Misra, S., and Azeta, A. (2024). Privacy Issues, Attacks, Countermeasures and Open Problems in Federated Learning: A Survey. *Applied Artificial Intelligence*, 38(1), 2410504. DOI: 10.1080/08839514.2024.2410504.
- Guo, S., Li, Z., Liu, H., Zhao, S., and Jin, C. (2022). Personalized Federated Learning for Multi-task Fault Diagnosis of Rotating Machinery. arXiv preprint arXiv:2211.09406.
- Huang, W., Ye, M., Shi, Z., Wan, G., Li, H., Du, B., and Yang, Q. (2024). Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12), 9387-9406. DOI: 10.1109/TPAMI.2024.3418862.
- Kang, Y., Fan, T., Gu, H., Zhang, X., Fan, L., and Yang, Q. (2025). Grounding Foundation Models through Federated Transfer Learning: A General Framework. *ACM Transactions on Intelligent Systems and Technology*, 16(4), 96:1-96:54. DOI: 10.1145/3742788.
- Khatri, V. and Brown, C. V. (2010). Designing data governance. *Communications of the ACM*, 53(1), 148-152. DOI: 10.1145/1629175.1629210.
- Kokina, J., Blanchette, S., Davenport, T. H., and Pachamanova, D. (2025). Challenges and opportunities for artificial intelligence in auditing: Evidence from the field. *International Journal of Accounting Information Systems*, 56, 100734. DOI: 10.1016/j.accinf.2025.100734.
- Li, K., Lu, J., Zuo, H., and Zhang, G. (2024). Federated Fuzzy Transfer Learning With Domain and Category Shifts. *IEEE Transactions on Fuzzy Systems*, 32(12), 6708-6719. DOI: 10.1109/TFUZZ.2024.3459927.
- Li, X., Fan, W., Yang, S., Zhang, W., and Li, X. (2026). Flexible Federated Learning in Machinery Fault Diagnostics With Light Communication. *IEEE/CAA Journal of Automatica Sinica*, 13(3), 680-691. DOI: 10.1109/JAS.2025.125414.
- Li, Y. and Goel, S. (2025). Artificial intelligence auditability and auditor readiness for auditing artificial intelligence systems. *International Journal of Accounting Information Systems*, 56, 100739. DOI: 10.1016/j.accinf.2025.100739.
- Mai, C., Liao, T., Chen, C., and Zheng, Z. (2025). FGTL: Federated Graph Transfer Learning for Node Classification. *ACM Transactions on Knowledge Discovery from Data*, 19(1), 8:1-8:20. DOI: 10.1145/3699962.
- Papagiannidis, E., Mikalef, P., and Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. DOI: 10.1016/j.jsis.2024.101885.
- Shen, J., Yang, S., Zhao, C., Ren, X., Zhao, P., Yang, Y., Han, Q., and Wu, S. (2023). FedLED: Label-Free Equipment Fault Diagnosis with Vertical Federated Transfer Learning. arXiv preprint arXiv:2312.17451.
- Xiang, P., Zhou, L., and Tang, L. (2024). Transfer learning via random forests: A one-shot federated approach. *Computational Statistics and Data Analysis*, 197, 107975. DOI: 10.1016/j.csda.2024.107975.
- Zeng, D., Liang, S., Hu, X., Wang, H., and Xu, Z. (2021). FedLab: A Flexible Federated Learning Framework. arXiv preprint arXiv:2107.11621.
- Yang, B., Lei, Y., Li, N., Li, X., Si, X., and Chen, C. (2025). Balance recovery and collaborative adaptation approach for federated fault diagnosis of inconsistent machine groups. *Knowledge-Based Systems*, 317, 113480. DOI: 10.1016/j.knosys.2025.113480.
- Zhang, H., Jiang, S., and Xuan, S. (2024a). Decentralized federated learning based on blockchain: concepts, framework, and challenges. *Computer Communications*, 216, 140-150. DOI: 10.1016/j.comcom.2023.12.042.
- Zhang, N., Zhao, Y., Qu, Y. N., Gu, B., Sood, K., Gao, L., and Yu, S. (2024b). From Data Integrity to Global Model Integrity for Federated Learning: An MHT-based Approach. 2024 IEEE Global Communications Conference (GLOBECOM). *Piscataway: IEEE*, 956-961. DOI: 10.1109/GLOBECOM52923.2024.10901319.
- Zhu, J., He, S., He, P., Liu, J., and Lyu, M. R. (2023). Loghub: A Large Collection of System Log Datasets for AI-driven Log Analytics. 2023 IEEE 34th International Symposium on Software Reliability Engineering (ISSRE), 355-366. DOI: 10.1109/ISSRE59848.2023.00071. HDF5_v1 dataset URL: <https://zenodo.org/records/8196385>.



Jinhui Wang holds a Doctor of Business Administration (DBA) and is currently a faculty member at Guangzhou College of Technology and Business. He has worked at Gree Electric Appliances Inc., served as General Manager at Chongqing Peirui Enterprise Management Consulting Co., Ltd. and Chongqing Liyien Technology Co., Ltd., as Director of the Innovation and Entrepreneurship Office at Chongqing Electromechanical Vocational and Technical University, and as a master's supervisor at Chongqing Jiaotong University. His research interests cover machine learning, data mining and governance, artificial intelligence and management decision-making, lean operation management, digital transformation of manufacturing, and innovation and entrepreneurship management. He has led consulting projects on lean operations, digital transformation, and intelligent manufacturing

for over 100 enterprises and has published several academic papers at international conferences and in domestic core journals.



Yangli Lei holds a bachelor's degree and serves as a faculty member at the Digital Urban Operation and Governance Center of Chengkou County, Chongqing, China. Her research interests include data mining and governance, artificial intelligence, and technology management.



Rongrui Wang has held senior executive positions at Gree and Skyworth, both of which are among China's Top 500 enterprises. Since October 2025, he has served as Vice President (VP) at CVTE. He has systematic attainments and practical implementation experience in consumer electronics manufacturing processes, intelligent manufacturing, and factory planning and construction, and has delivered multiple creative benchmark demonstration projects. He has led the construction of multiple factories totaling over 300,000 square meters and has achieved remarkable results in green factory planning, intelligent warehousing and logistics, intelligent manufacturing, visual inspection, and AI applications.