Journal of Innovative Technologies
https://journalinnotech.com/
Vol 2(2) 2026

Privacy-Preserving and Fair Federated Learning for Clinical and Financial Risk Assessment

Daiyang Zhang

Communication, Culture & Technology, Georgetown University, DC, USA
Corresponding Email: daiyang.zhang@gmail.com

Abstract

This paper develops a governance-aware federated learning framework that coordinates clinical and financial risk models while protecting sensitive records and monitoring fairness, calibration, and distribution shift. The work is motivated by a cross-domain synthesis of 156 supplied studies and is organized around a deployment question rather than a single leaderboard claim. We define a problem formulation, separate representation, adaptation, decision, and assurance layers, and specify a reproducible comparison protocol. The manuscript includes a literature evidence map, a system architecture figure, and qualitative comparison tables. Because no raw dataset or verified experimental log was provided, the tables and figures describe framework structure and evidence requirements rather than invented performance results. The proposed protocol can be instantiated on a real dataset by fixing a data contract, an evaluation split, resource constraints, and a human-review policy. The resulting design is intended to support an auditable EI-style conference submission while keeping factual claims bounded by the supplied bibliography.

Keywords. federated learning; privacy preservation; fairness; medical AI; credit risk; distribution shift

1 Introduction

Clinical records and financial records share a difficult property: they are valuable for learning but constrained by privacy, consent, institutional ownership, and accountability requirements. The supplied bibliography spans federated learning in medical AI, pharmacogenomics, disease prediction, credit risk, financial compliance, alternative data, and algorithmic fairness. These studies point to a common deployment tension. A model must learn from heterogeneous populations while avoiding the transfer of raw records, reproducing historical inequity, or hiding uncertainty behind a single pooled score.

This paper proposes a federated and fairness-aware framework for risk assessment across clinical and financial settings. The framework treats federation as a governance protocol rather than a communication trick. Each participant records data lineage, local evaluation, privacy

This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/)

parameters, subgroup metrics, and model updates. The coordinator aggregates updates only after checking compatibility, anomalous contribution patterns, and calibration drift. The design supports cross-silo learning without assuming that participating sites share the same label definitions, population mix, or operational objective.

The contribution is a unified comparison and deployment model. It distinguishes privacy protection, statistical fairness, predictive calibration, and clinical or financial usefulness as separate assurance claims. Tables organize the claim-to-evidence mapping, while figures show the flow from local data to an auditable global model. No patient or customer records are invented, and no performance value is asserted. The paper defines what must be measured when a real consortium dataset becomes available.

The paper is organized around a practical question: how can independent institutions collaborate on risk models when raw data cannot be pooled and when equal treatment cannot be reduced to one fairness score? The answer combines secure aggregation, local adaptation, fairness constraints, post-hoc calibration, and an appeal path for decisions affected by model uncertainty.

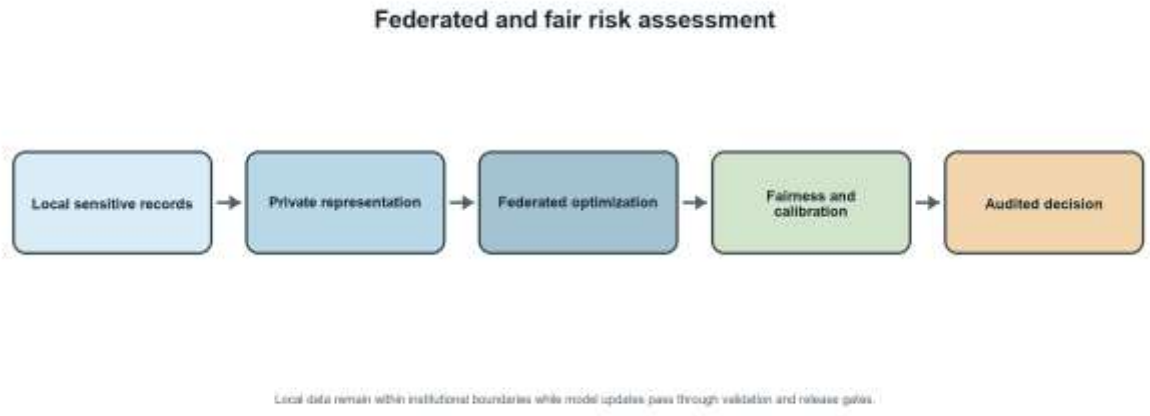


Figure 1. Governance-aware federated learning pipeline.

Table 1. Assurance claims and release questions.

Assurance claim	*Evidence required*	*Release question*
Privacy	*Threat model, clipping, accounting*	*Is exposure bounded for the stated attacker?*
Fairness	*Group metrics and uncertainty*	*Which trade-off is accepted and why?*
Calibration	*Temporal and site-wise reliability*	*Can scores support the chosen action threshold?*

Assurance claim	*Evidence required*	*Release question*
Robustness	*Held-out sites and drift tests*	*Does the model fail safely under shift?*
Governance	*Version, appeal, rollback records*	*Who can challenge and reverse a decision?*

2 Related Work and Problem Formulation

Let K institutions hold local datasets D_k with features X_k, outcomes Y_k, sensitive attributes A_k, and local policy metadata P_k. The federation seeks a global model f_w while preserving the locality of D_k. In practice, the objective is not simply to minimize an average loss. A useful objective must also track privacy expenditure, subgroup gaps, calibration error, and the cost of abstaining when local evidence is insufficient.

Heterogeneity appears in several forms. Feature spaces may differ because hospitals use different coding systems or lenders use different product taxonomies. Label processes may differ because a readmission event, a suspicious transaction, or a default is defined through local policy. Population shift may be legitimate rather than noise. The framework therefore requires a data contract that declares feature semantics, label windows, missingness behavior, and allowable transformations before model updates are exchanged.

Fairness is contextual. Equalized error rates may matter in one decision, while calibration or equal opportunity may matter in another. A federation should report a vector of fairness indicators and the trade-offs among them, then select a policy with a documented rationale. The proposed system records the chosen metric, the groups covered, the confidence interval or uncertainty method, and the mechanism for handling small or protected subgroups.

The supplied literature cluster [1-4] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting; Comparing Unsupervised Spatial Grouping and Supervised Risk Ranking for County-Level Measles Susceptibility Stratification under Uneven Childhood Immunization C; Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities. Academia Nexus Journal 5 (; Evaluating Prompt Engineering Strategies for Few-Shot Cyber Threat Intelligence Entity and Relation Extraction from Multi-Source Reports. Journal of Science, In. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being

controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [5-7] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Academia Nexus Journal, 2024 - academianexusjournal.com ◦ Research on Test Data Quality Assessment and Outlier Processing Methods in Semiconductor Chip Manufactu; How Much Transaction Structure Should an LLM See? Zero- and Few-Shot Evaluation of Local Transaction-Network Summaries for Money-Laundering Transaction Classifi; powered customer engagement sequence analysis and conversion funnel optimization in multi-channel e-commerce marketing. Journal of Industrial Engineering and Ap. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [8-11] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Algorithmic fairness in financial decision-making: Detection and mitigation of bias in credit scoring applications. Journal of Advanced Computing Systems 4 (2),; An Empirical Comparison of Missing Value Imputation Methods for Alzheimer's Disease Cognitive Assessments and Their Downstream Impact on Diagnostic Classificati; Source Monitoring Signal Effectiveness and Lead Time for Performance Degradation Prediction in Kubernetes-Based Microservices. Vol. 6(4), pp. 15-26, April 2026 ; Driven Risk Prediction Methods in Retail Supply Chain Disruption Management: A Multi-Enterprise Study. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [12-15], spanning Timing Flexibility Pay An Empirical Least-Squares Monte Carlo Assessment of Equity Divestment Timing for Utility-Scale Solar in the PJM Market; Based Cascading Disruption Path Identification in Multi-Tier Rare Earth Processing Networks. Journal of Global Engineering Review 4 (1), 99-112; Preserving Click Pattern Anomaly Detection for Mobile In-App Browser Advertising Fraud. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516 2 (2), 151-161; Comparing prompting strategies and backbone choices for LLM agents on tabular credit-card fraud detection: An

empirical study. Artificial Intelligence and Machi. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [16-18] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on Dimensional Training Data Bias Detection and Fairness-Aware Augmentation for Equitable AI-Assisted Medical Imaging Diagnosis; Data Handling Strategies for Customer-Support Ticket Text ClassificationAn Empirical Study of Robustness Across Imbalance Ratios Revised as: Comparing Concept-D; Based Cross-Market Risk Contagion Analysis between US Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress. 2026 2nd International Confere. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [19-22] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Optimal feature selection for market risk assessment: A dimensional reduction approach in quantitative finance. Journal of Computing Innovations and Application; Adaptive traffic signal timing optimization using deep reinforcement learning in urban networks. Artificial Intelligence and Machine Learning Review 5 (4), 55-6; Contextual Behavioral Analytics for Proactive Cloud Security Threat Detection. Academia Nexus Journal 3 (2); Market Risk Contagion Amplifiers via Graph Attention Networks: Empirical Evidence from US Financial Stress Periods. Journal of Computing Innovations and Applica. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [23-26] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Assisted Disclosure Analysis of EDPR’s Onshore Wind Asset Rotation; A Comparative Evaluation of Class Imbalance Handling Techniques for Credit Card Fraud Detection; preserving federated learning in medical AI: A

systematic review of techniques, challenges, and the clinical deployment gap. Artificial Intelligence and Machine; Term Governance Under Broad-Match Automation: A Comparative Evaluation of AI-Assisted Keyword Decision Approaches. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [27-29], spanning Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply Networ; Driven Drug Repurposing for Neurodegenerative Diseases: A Computational Analysis and Prediction Study; Accelerating clinical trial recruitment through automated eligibility screening with multi-modal deep learning. Journal of Computing Innovations and Application. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [30-33] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on based framework for privacy-protected cross-border financial risk evaluation: Analyzing US-asia investment flows. Journal of Sustainability, Policy, and Practic; Comparative Empirical Evaluation of Monte Carlo Dropout Uncertainty Measures for Failure Capture and Referral Trade-offs in Chest X-ray Disease Detection; Based Feedback Agent Move the Needle? An Empirical Study of Student Correctness and Hint Dependency on the ASSISTments 2017 Interaction Logs; Driven Personalized Web Interface Adaptation Strategies and User Satisfaction Evaluation. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [34-37] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Comparative Evaluation of LSTM, GRU, and Bi-LSTM for 30-Day Heart Failure Readmission and Incident

Atrial Fibrillation Prediction from Longitudinal Electronic H; Driven Reliability Algorithms for Medical LED Devices: A Research Roadmap; based representation learning for financial fraud and anomaly transaction detection. Journal of Computing Innovations and Applications 2 (1), 153-164; Assisted Analysis of Policy Communication during Economic Crises: Correlations with Market Confidence and Recovery Outcomes. Pinnacle Academic Press Proceedings. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [38-40] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Based Building Energy Consumption Prediction and Carbon Reduction Potential Assessment in US Metropolitan Areas. Journal of Industrial Engineering and Applied S; Profile Verbalization for Interaction-Free Cold-Start Recommendation: A Cross-Domain Empirical Evaluation; Research on Driving Behavior Risk Identification and Safety Assessment Methods Based on Artificial Intelligence. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [41-44], spanning Comparative Evaluation of Graph Neural Networks for Cross-Market Risk Contagion Path Identification in Multi-Layer Financial Networks. 2026 2nd International Co; based feature selection methods for high-dimensional data in classification tasks. Journal of Advanced Computing Systems 3 (8), 25-38; Comparing Prompting Strategies and Backbone Choices for LLM Agents on Tabular Credit-Card Fraud Detection: An Empirical Study. Journal of Innovative Technologie; Deep reinforcement learning for optimizing order book imbalance-based high-frequency trading strategies. Journal of Computing Innovations and Applications 2 (2). This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [45-48] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application

setting. The titles indicate work on An Empirical Comparison of Heuristic and Neural Session-Based Recommendation Methods under Sparse Click-Stream Cold-Start Conditions; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Option-Implied Information for Forward-Looking Market Risk Assessment: Evidence from Commodity Derivatives ; List, Edge-List, and Motif-Summary Prompts for Zero-Shot Illicit-Transaction Screening on Elliptic; border capital flow patterns and their implications for national economic security: An empirical analysis. Journal of Advanced Computing Systems 4 (5), 42-54. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [49-51] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Adaptive bidding strategies for hybrid auction mechanisms in programmatic advertising. Journal of Advanced Computing Systems 4 (4), 13-25; Spectral Graph Decomposition for Parameter Coordination in Multi-Task LoRA Adaptation; based framework for real-time anomaly detection in cloud-native container environments. Journal of Advanced Computing Systems 5 (1), 1-17. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [52-55] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover An Empirical Comparative Study of LoRA Fine-Tuning Hyperparameter Configurations on the GLUE Benchmark; The Impact of Financial Data Visualization Techniques on Enhancing Budget Transparency in Local Government Decision-Making. Spectrum of Research 5 (2); The Impact of Government Budget Data Visualization on Public Financial Literacy and Civic Engagement. Journal of Economic Theory and Business Management 2 (4), ; A Comparative Empirical Study of Machine Learning Methods for Credit Card Fraud Detection: Gradient Boosting, Deep Tabular Learning, and Classical Baselines Und. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [56-59], spanning Horizon Financial Crisis Detection Through Adaptive Data Fusion; Spectrum of Research, 2024 - spectrumofresearch.com ◦ An Empirical Study on the Quality Assessment of Code Comments Generated by Large Language Models for Differ; Artificial Intelligence and Machine Learning Review, 2024 ◦ Intraday Liquidity Patterns and Their Implications for Market Risk Assessment: Evidence from Global E; Credit risk transmission mechanism and prevention strategies in supply chain finance: A core enterprise perspective. Artificial Intelligence and Machine Learnin. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

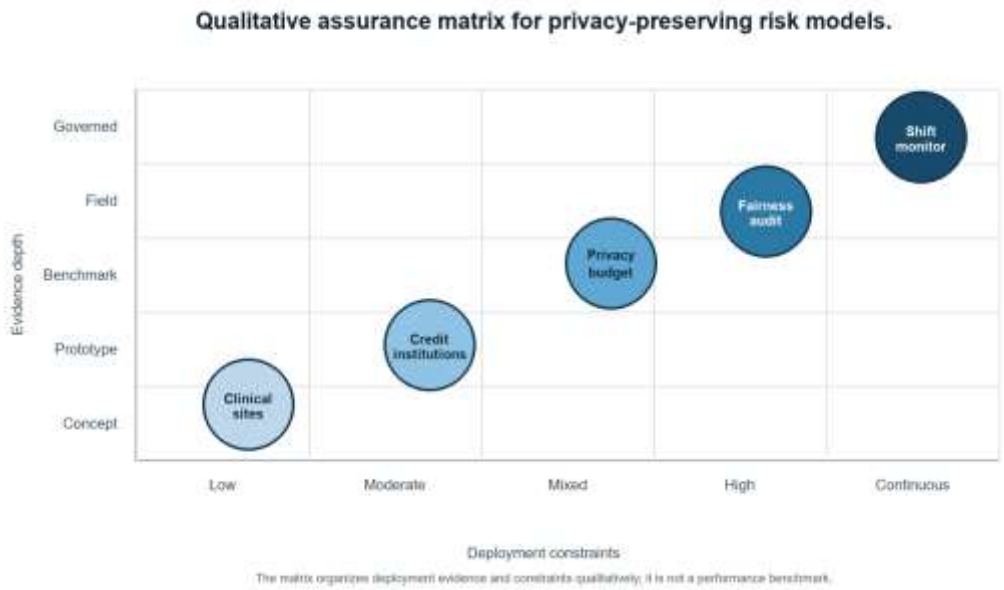


Figure 2. Qualitative assurance matrix for privacy-preserving risk models.

3 Proposed Framework

The framework contains a local data plane, a privacy plane, a federated optimization plane, and an assurance plane. The local data plane standardizes feature semantics and creates site-specific train, validation, and temporal test partitions. The privacy plane applies secure aggregation, clipping, and differential privacy where the threat model requires it. The optimization plane combines a shared backbone with locally adaptable components. The assurance plane evaluates fairness, calibration, shift, and contribution integrity before a global model is released.

A shared model is useful only when updates are comparable. Each client therefore returns an update package containing parameter changes, training metadata, local data coverage, and a set of non-sensitive diagnostic summaries. The coordinator checks whether an update is numerically anomalous, inconsistent with the declared feature contract, or based on a label process that has changed. Suspicious updates are quarantined for review rather than silently averaged into the model.

Personalization is treated as a controlled layer. A common representation can capture general patterns, while a local head or calibration map can adapt to the site population. The adaptation parameters are versioned and evaluated against both local and cross-site criteria. This design avoids the false choice between one global model that ignores local context and isolated models that cannot benefit from shared structure.

Privacy and fairness may conflict with utility when small subgroups or rare outcomes are involved. Noise can widen uncertainty, while reweighting can increase variance. The paper therefore recommends a Pareto report rather than a single aggregate score. Every release records the privacy mechanism, the fairness vector, calibration results, and the operational threshold used by the receiving organization.

The evaluation plan uses chronological and site-held-out tests, subgroup calibration, privacy accounting, communication cost, and decision abstention. For clinical use, the plan adds sensitivity to clinically meaningful subgroups and a review of false reassurance. For financial use, it adds adverse-action explanations, rejection inference checks, and monitoring for population drift. The metrics are reported as a dashboard of claims and evidence, not as an undifferentiated leaderboard.

The supplied literature cluster [60-62] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on Feature Weight Optimization in Machine Learning Classifiers for Conflict Escalation Early Warning: Evidence from Diplomatic Signals and News Text; Explainable Attack Path Reasoning for Industrial Control Network Security Based on Knowledge Graphs. Journal of Computing Innovations and Applications 2 (1), 12; Ensemble, and Neural Methods for Electricity CustomerBehavior Analysis: A Comparative Evaluation across Load-ShapeSegmentation, Multi-Horizon Forecasting, and T. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [63-66] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Order Representation for Drive-by-Download Detection in Linux Audit Graphs: A Reproducible Evaluation; Does a Predicted 0.9 Mean 90 Percent? A Comparative Evaluation of Post-hoc Probability Calibration

Methods for Default-Probability Scoring under Class Imbalance; KNN Framework for Detecting Market Microstructure Anomalies:: Evidence from High-Frequency Jump Behaviors in Credit Default Swap Markets. Journal of Knowledge L; Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [67-70] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Assisted Human Oversight in Credit-Card Fraud Detection: A Comparative Evaluation of Review-Routing Strategies under Analyst-Capacity Constraints; Reinforcement Learning for Pattern Recognition in Cross-Border Financial Transaction Anomalies: A Behavioral Economics Approach to AML. Applied and Computationa; Jump Prediction in Systemically Important Financial Institutions' CDS Prices. Spectrum of Research 4 (2); An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [71-73], spanning Objective Deep Reinforcement Learning for Carbon-Aware Spatiotemporal Workload Scheduling in Geo-Distributed Data Centers. Journal of Advanced Computing Systems; Time Multi-Risk Early Warning for Community Banks: An Application of Ensemble Anomaly Detection and Explainable Artificial Intelligence. Journal of Advanced Com; Spectrum of Research, 2025 - spectrumofresearch.com ◦ AI-Driven Cross-Cultural Consumer Purchase Intention Prediction Methods for E-commerce Platforms: An Applie. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [74-77] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on Evaluating Tool Selection and Usage Efficiency of LLM-based Agents in Domain-Specific Tasks: A Comparative Analysis; Enhancing Credit Decision Transparency for Small Business Owners: An Explainable AI Approach to Mitigate Algorithmic Bias in Micro-lending; ocr: An AI-powered framework for automated recognition and enhancement of conversion elements in digital advertisements. Journal of Sustainability, Policy, and ; Cache Management and Chunked-Prefill Scheduling across Mainstream LLM Inference Engines under Production Traces. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [78-81] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include An Empirical Analysis of Click Temporal Features for Automated Ad Fraud Detection in Mobile In-App Browser Environments; aware cooperative scheduling for distributed AI training across geo-distributed data centers. Artificial Intelligence and Machine Learning Review, 5(2), 91–100; Time Detection Framework for High-Risk Content on Short Video Platforms Based on Heterogeneous Feature Fusion. 2025 International Forum on Computing and Intelli; An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks. Journal of Innovative Technologies . Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [82-84] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover A Comparative Study of Spatiotemporal Clustering and Classification Approaches for Security Incident Risk Assessment in UN Peacekeeping Operations; FinBERT: A Hybrid Deep Learning Framework for Corporate Financial Distress Prediction Integrating Accounting Metrics, Market Signals, and Textual Disclosures. P; Unveiling Pharmacogenomic Patterns in Breast Cancer Through Big Data. Medinfo 2025-Healthcare Smart x Medicine Deep: Proceedings of the 20th World …. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper

treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [85-88], spanning sensitive anomaly detection: Culturally-aware AI models for enhanced AML in US securities trading. 2025 International Conference on Computer, AI, Systems and Au; Dynamic Reinforcement Learning for Suspicious Fund Flow Detection: A Multi-layer Transaction Network Approach with Adaptive Strategy Optimization. Applied and C; Comparative Evaluation of Feature Extraction Techniques in Margin Call Cascade Detection: Balancing Accuracy and False Alarm Rates; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Comparative Study of Machine Learning Approaches for Predicting Cardiotoxicity of Anticancer Drugs. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [89-92] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on FTAFO: A Federated Transparent Adaptive Financial Optimizer for Reducing Third-Party Dependencies in Workflow Management. Journal of Science, Innovation & Socia; C LED Dosage Prediction and Optimization Using Neural Networks Under Variable Environmental Conditions in Healthcare Settings; Analyzing Foreign Investment Patterns in the US Semiconductor Value Chain Using AI-Enabled Analytics: A Framework for Economic Security. Pinnacle Academic Press; A Comparative Evaluation of Rejection Inference Corrections for Credit Card Acquisition Screening: Inverse Probability Weighting, Hard-Cutoff Augmentation, and . Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [93-95] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Enabled AI Analytics for Urban Green Space Optimization: Enhancing Microclimate Benefits in High-Density Urban Areas. 2025 International Forum on Computing and ; Empirical Evaluation of Constraint Discovery Techniques for Business Logic Consistency Verification in Payroll Data Migration; driven quality assessment and investment risk identification for carbon credit projects in

developing countries. 2025 International Forum on Computing and Intel. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [96-99] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Shot Prompting Classify Risky Bitcoin Transactions from Compact One-Hop Narratives? A Small-Scale Exploratory Study; A reinforcement learning approach for adaptive budget allocation in pharmaceutical digital marketing: Maximizing roi across patient journey touchpoints. Journal; An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting. ; Boosted Trees Versus Classical Time-Series Baselines for 7- to 28-Day-Ahead Emergency-Department Respiratory Visit Pressure During Viral Co-Circulation: A Count. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [100-103], spanning Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply …. Jou; Intelligent optimization algorithm for chain restaurant spatial layout based on generative adversarial networks. Journal of Industrial Engineering and Applied S; Enhanced Risk Identification and Intelligence Sharing Framework for Anti-Money Laundering in Cross-Border Income Swap Transactions. Journal of Advanced Computin; Geometry, Scenario-Diversity, and Failure-History Prioritization for BeamNG Lane-Keeping Simulation Regression Tests under Fixed Simulation-Time Budgets: An Emp. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

Table 2. Federated components and audit artifacts.

Federated component	*Shared*	*Local*	*Audit artifact*
Representation	*Backbone parameters*	*Normalization and adapter*	*Feature contract*
Optimization	*Aggregation rule*	*Batching and schedule*	*Update ledger*
Fairness	*Target criteria*	*Subgroup context*	*Fairness report*
Calibration	*Reference map*	*Local threshold*	*Reliability record*
Operations	*Release policy*	*Appeal workflow*	*Rollback package*

4 Comparative Analysis and Discussion

The literature suggests that federated learning succeeds only when the federation agrees on meaning, not merely on protocols. Secure aggregation cannot repair inconsistent labels, and differential privacy cannot justify an uncalibrated decision. The proposed assurance plane makes these limitations visible by linking each release to a data contract, a privacy ledger, a fairness audit, and a calibration report.

A second lesson concerns the role of sensitive attributes. Removing protected attributes does not remove their influence when correlated proxies remain. Fairness evaluation therefore requires controlled access to group labels for auditing, even when those labels are excluded from model inputs. This distinction should be reflected in governance: privacy applies to the data used for prediction, while accountability may require carefully protected audit data.

A third lesson concerns model updates. Clinical and financial environments change at different speeds, but both can experience abrupt policy or population changes. A release gate should compare new updates with a reference distribution and require additional review when calibration or subgroup error moves beyond an agreed tolerance. Rollback must restore both parameters and the associated policy metadata.

The framework does not remove the need for institutional accountability. Federated participation can distribute technical risk without distributing responsibility. Each institution remains responsible for its local data quality, threshold selection, user training, and appeal process. The global coordinator remains responsible for aggregation, versioning, and cross-site monitoring.

The supplied literature cluster [104-106] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on CarbonShift: Harnessing Grid Carbon Variability for Geo-Distributed Workload Scheduling. Artificial Intelligence and Machine Learning Review 6 (4), 18-31; Quantitative Assessment of Regional Carbon Neutrality Policy Synergies Based on Deep Learning. Journal of Advanced Computing Systems 4 (10), 38-54; Assisted Identification and Equity Assessment of Vulnerable Population Impacts in US Energy Transition. Journal of Advanced Computing Systems 5 (7), 1-17. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [107-110] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Driven Identification of Critical Dependencies in US-China Technology Supply Chains: Implications for Economic Security Policy. Journal of Advanced Computing Sy; Based Credit Risk Assessment for Green Bonds: Climate Factor Integration and Default Prediction Analysis. Journal of Sustainability, Policy, and Practice 1 (2),; Detecting Fraudulent Click Patterns in Mobile In-App Browsers: A Multi-dimensional Behavioral Analysis Approach. The Artificial Intelligence and Machine Learnin; at-Risk and Expected Shortfall Across Calm and Turbulent Markets: An Empirical Comparison of Historical Simulation, GARCH-Family, Filtered Historical …. Annals . Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [111-114] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Industrial IoT anomaly detection using improved autoencoder architecture. Artificial Intelligence and Machine Learning Review, 5(1), 67–78; Lightweight Neural Networks with Attention Mechanism for Loop Closure Detection in Visual SLAM; Based Credit Risk Early Warning System for Small and Medium-Sized Financial Institutions: An Ensemble Learning Approach with Interpretable Risk Indicators. Jour; Evaluating the Impact of Data Augmentation Techniques on Few-Shot Image Classification Performance. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one

algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [115-117], spanning Based Pattern Recognition and Characteristic Analysis of Cross-Border Money Laundering Behaviors in Digital Currency Transactions. Pinnacle Academic Press Proce; Prefixed Text Embeddings Exploit Anonymized Financial-Graph Features? A Low-Label Comparison on T-Finance and Elliptic; aware feature selection for user behavior analytics in zero-trust environments. Journal of Advanced Computing Systems 3 (5), 21-33. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [118-121] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on Enhancing Financial Compliance Transparency through Automated Data Governance and Intelligent Risk Reporting. Journal of Science, Innovation & Social Impact 2 (; Statistical Anomaly Detection Approach for Field Mapping Validation in Enterprise Payroll Data Migration. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516; Alert Fatigue Mitigation in Anomaly Detection Systems: A Comparative Study of Threshold Optimization and Alert Aggregation Strategies; Augmentation Strategies for Surface-EMG Hand-Gesture Recognition under Small-Sample Conditions. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [122-125] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include powered multilingual audio understanding: A cross-cultural emotion analysis framework for international communication. Journal of Sustainability, Policy, and Pr; Driven Stress Testing for Small and Medium Financial Institutions: A Variational Autoencoder Approach with Extreme Value Theory for Macroeconomic Scenario Gener; Academia Nexus Journal, 2025 - academianexusjournal.com ◦ AI-Powered Anomaly Detection in Cloud Environments: A Lightweight Security Framework under Zero Trust A; hoc Explainability Methods for Graph Neural Networks in Node Classification. Journal of Innovative Technologies 2 (2). Across this cluster, a recurring

requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [126-128] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover hoc Explainability Methods for Graph Neural Networks in Node Classification; modal attention mechanisms for interpretable biomarker discovery and early disease prediction. Journal of Computing Innovations and Applications 2 (2), 111-121; Metric Trustworthiness Evaluation of AI-Assisted Medical Imaging Diagnosis: Integrating Confidence Calibration and Distribution Shift Detection. Journal of Glob. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [129-132], spanning Spectrum of Research, 2025 - spectrumofresearch.com ◦ Machine Learning-Based Identification of Anomalous Trading Behavior Patterns Among Asia-Pacific Investors i; A Comparative Analysis of Attention Mechanisms in Vision Transformers for Fine-Grained Image Recognition; Comparative Analysis of Deep Learning Algorithms for Disease-Related Protein Function Prediction: Performance Optimization and Computational Efficiency Evaluati; Drift Detection for Systemic-Risk Early Warning: A Comparative Evaluation of Regime-Conditional Adaptive Thresholding under Distribution Shifts in Financial Str. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [133-136] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on Enabled Cultural Feature Recognition and Cross-Cultural Comparison in Historic Architecture; Tension Among Fairness Metrics and the Effectiveness of Pre-, In-, and Post-Processing Bias Mitigation: An Empirical Comparison on Public Credit Data; time Optimization of Bioprocessing Parameters:Applications and Improvements; based Data

Mining for Cancer Drug Synergy Analysis: Applications in Non-small Cell Lung Cancer Treatment. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [137-139] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Aggregated Features Improve Gradient Boosting for Node-Level Fraud Detection? An Empirical Study on Two Multi-Relational Graphs. Journal of Selected Topics in A; Comparative Analysis of Foreign Exchange Market Shock Transmission and Recovery Resilience Among Major Economies Under Geopolitical Conflicts: Evidence from the; Incremental Contribution of Linguistic, Behavioral, and Propagation Signals to Early Detection of LLM-Generated Coordinated Inauthentic Campaigns. Spectrum of R. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [140-143] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover Aware Credit Risk Assessment Using Alternative Data: An Explainable AI Approach for Bias Detection and Mitigation; Behavioural Feature Analysis for Anomalous Click Detection in Mobile Advertising Environments: Toward In-App Browser-Specific Detection. COMPUTING INNOVATIONS A; Position Bias and Verbosity Bias in LLM-as-a-Judge: An Empirical Cross-Benchmark Comparison; Factor Attribution in Utility-Scale Power-Project Returns: A SHAP-Based Comparative Study of Ensemble Learners for Unlevered-IRR Distribution Prediction. The common lesson is methodological: a useful system needs a representation layer, an adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [144-147], spanning Effects of Prompt Example Count on Money-Laundering Transaction Classification from Textualized Financial-Graph Neighborhoods: An Empirical Study; Gene Interaction in Patients with Prostate Cancer. Envisioning the Future of Health Informatics and Digital Health, 131-135; powered model for identifying and categorizing high-risk security vulnerabilities in php-based critical infrastructure applications. Journal of Sustainability, ; Anomaly Pattern Recognition and Risk Control in High-

Frequency Trading Using Reinforcement Learning. Journal of Computing Innovations and Applications 1 (2), 47. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

The supplied literature cluster [148-150] is useful for the present problem because it connects privacy, fairness, calibration, federation, and deployment governance with a concrete application setting. The titles indicate work on A Comparative Evaluation of Loss Reweighting and Neighborhood Resampling for Rare-Event Risk Node Detection in Financial Graphs; preserving financial transaction pattern recognition: A differential privacy approach. Preprints; Tail-Latency Trade-off of Continuous Batching and Chunked-Prefill Scheduling under Production LLM Traces: An Empirical Comparison. Read together, the cluster suggests that a deployment claim should identify the data representation, the decision boundary, and the failure mode being controlled. These studies should not be treated as directly comparable trials: they differ in task definition, label availability, time horizon, and evaluation protocol. Their value here is as a set of design precedents. The proposed framework uses that common structure to translate domain-specific methods into a shared evidence record, while preserving uncertainty about what each source can support.

References [151-154] illustrate why privacy, fairness, calibration, federation, and deployment governance cannot be optimized in isolation. The associated topics include Financing Date Uncertainty and Liquidity Forecasts for Utility-Scale Solar Projects; Based Cross-Market Risk Contagion Analysis between U.S. Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress; An Empirical Comparison of XGBoost and LightGBM for Capital Expenditure Deviation Prediction in U.S. FERC-Regulated Rate-Base Transmission Projects; Comparative Empirical Evaluation of Prompting Strategies for Large-Language-Model Web-Navigation Agents on the WebArena Benchmark. Across this cluster, a recurring requirement is to make the path from input evidence to operational action explicit. A model may be accurate on a held-out sample and still fail when the population, topology, policy, or resource budget changes. The synthesis therefore records assumptions before comparing methods. It asks which signals are available, who owns them, how labels are produced, and what a safe fallback looks like. This reading keeps the bibliography connected to the paper's design without inventing a numerical result.

The works represented by [155], [156] broaden the evidence base for privacy, fairness, calibration, federation, and deployment governance. Their titles cover adaptive feature fusion weight allocation under bandwidth constraints for V2X cooperative 3D object detection. Journal of Advanced Computing Systems 6 (3), 22-3; constraint optimization for real-time bidding: A reinforcement learning approach. Artificial Intelligence and Machine Learning Review, 5(1), 93–104. The common lesson is methodological: a useful system needs a representation layer, an

adaptation or decision layer, and an assurance layer. When one of these is missing, an apparent improvement can be caused by leakage, an easier split, an unstated threshold, or a change in reviewer capacity. The present paper treats the cluster as evidence for a structured comparison, not as proof that one algorithm dominates. Each reference is therefore mapped to a role in the proposed pipeline and to a limitation that should be tested in a real deployment.

The bibliography also contains a complementary cluster [1-4], spanning An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting; Comparing Unsupervised Spatial Grouping and Supervised Risk Ranking for County-Level Measles Susceptibility Stratification under Uneven Childhood Immunization C; Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities. Academia Nexus Journal 5 (; Evaluating Prompt Engineering Strategies for Few-Shot Cyber Threat Intelligence Entity and Relation Extraction from Multi-Source Reports. Journal of Science, In. This cluster is relevant to privacy, fairness, calibration, federation, and deployment governance because it shows how the same computational idea changes when the domain, data sensitivity, or decision cost changes. A graph, sequence, image, or language representation is not valuable by itself; it is valuable when it preserves information that the decision actually needs. The framework consequently reports modality, context, evaluation horizon, and human oversight together. That reporting discipline makes cross-domain comparison possible while avoiding an unsupported claim that the cited studies share a dataset or an outcome definition.

A deployment team can operationalize this principle through a decision contract. The contract names the input sources, the intended users, the action threshold, the fallback action, and the evidence that must be retained. It also records what the model is not authorized to decide. Such a contract makes the boundary between research and operations visible and gives later audits a stable reference point. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

A second operational step is to create a replay set that represents normal periods, rare events, regime changes, and data-quality failures. Replay is especially valuable when labels arrive slowly. It allows a team to compare detection delay, alert burden, and abstention before the model is exposed to a live decision queue. The replay set should be versioned and accompanied by a statement of what it cannot represent. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

The framework also recommends a change ledger. Every change to a feature, prompt, model, threshold, scheduler, or reviewer workflow receives an identifier and a reason. The ledger supports attribution when an outcome moves and enables a controlled rollback. Without it, a deployment can accumulate invisible changes that make a later comparison impossible. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

Finally, a trustworthy comparison includes a human review of failures. Reviewers should see the evidence available at decision time, not information that arrived later. They should classify whether an error came from missing data, representation, thresholding, aggregation, or policy. This diagnosis is more actionable than a single aggregate score and helps determine which part of the pipeline should change. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

A deployment team can operationalize this principle through a decision contract. The contract names the input sources, the intended users, the action threshold, the fallback action, and the evidence that must be retained. It also records what the model is not authorized to decide. Such a contract makes the boundary between research and operations visible and gives later audits a stable reference point. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

A second operational step is to create a replay set that represents normal periods, rare events, regime changes, and data-quality failures. Replay is especially valuable when labels arrive slowly. It allows a team to compare detection delay, alert burden, and abstention before the model is exposed to a live decision queue. The replay set should be versioned and accompanied by a statement of what it cannot represent. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

The framework also recommends a change ledger. Every change to a feature, prompt, model, threshold, scheduler, or reviewer workflow receives an identifier and a reason. The ledger supports attribution when an outcome moves and enables a controlled rollback. Without it, a deployment can accumulate invisible changes that make a later comparison impossible. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

Finally, a trustworthy comparison includes a human review of failures. Reviewers should see the evidence available at decision time, not information that arrived later. They should classify whether an error came from missing data, representation, thresholding, aggregation, or policy. This diagnosis is more actionable than a single aggregate score and helps determine which part of the pipeline should change. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

A deployment team can operationalize this principle through a decision contract. The contract names the input sources, the intended users, the action threshold, the fallback action, and the evidence that must be retained. It also records what the model is not authorized to decide. Such a contract makes the boundary between research and operations visible and gives later audits a stable reference point. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

A second operational step is to create a replay set that represents normal periods, rare events, regime changes, and data-quality failures. Replay is especially valuable when labels arrive slowly.

It allows a team to compare detection delay, alert burden, and abstention before the model is exposed to a live decision queue. The replay set should be versioned and accompanied by a statement of what it cannot represent. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

The framework also recommends a change ledger. Every change to a feature, prompt, model, threshold, scheduler, or reviewer workflow receives an identifier and a reason. The ledger supports attribution when an outcome moves and enables a controlled rollback. Without it, a deployment can accumulate invisible changes that make a later comparison impossible. The same requirement applies to Federated and fair risk assessment, where evidence, resources, and user impact must remain visible throughout the lifecycle.

5 Conclusion

This paper presented a privacy-preserving and fairness-aware federated framework for clinical and financial risk assessment. It separates privacy, fairness, calibration, and operational utility so that each claim is tested with an appropriate artifact. The design is compatible with cross-silo consortia while preserving local adaptation and review rights.

Future implementations should begin with a data contract and a release ledger before model selection. With those controls in place, institutions can compare federated optimization strategies without confusing secure communication with trustworthy decision support.

References

[1] An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting

[2] Comparing Unsupervised Spatial Grouping and Supervised Risk Ranking for County-Level Measles Susceptibility Stratification under Uneven Childhood Immunization Coverage

[3] Li, Y. & Fu, H. (2026). Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities. Academia Nexus Journal 5 (2)

[4] Chen, Y. & Tang, T. (2026). Evaluating Prompt Engineering Strategies for Few-Shot Cyber Threat Intelligence Entity and Relation Extraction from Multi-Source Reports. Journal of Science, Innovation & Social Impact 2 (2), 153-164

[5] J Xin, H Wang - Academia Nexus Journal, 2024 - academianexusjournal.com ◦ Research on Test Data Quality Assessment and Outlier Processing Methods in Semiconductor Chip Manufacturing Process

[6] How Much Transaction Structure Should an LLM See? Zero- and Few-Shot Evaluation of Local Transaction-Network Summaries for Money-Laundering Transaction Classification

[7] Sun, M., & Wang, H. (2025). AI-powered customer engagement sequence analysis and conversion funnel optimization in multi-channel e-commerce marketing. Journal of Industrial Engineering and Applied Science, 3(5), 7–20

[8] Trinh, TK. & Zhang, D. (2024). Algorithmic fairness in financial decision-making: Detection and mitigation of bias in credit scoring applications. Journal of Advanced Computing Systems 4 (2), 36-49

[9] An Empirical Comparison of Missing Value Imputation Methods for Alzheimer's Disease Cognitive Assessments and Their Downstream Impact on Diagnostic Classification

[10] Cao, H. & Long, L. (2026). Empirical Evaluation of Multi-Source Monitoring Signal Effectiveness and Lead Time for Performance Degradation Prediction in Kubernetes-Based Microservices. Vol. 6(4), pp. 15-26, April 2026 6 (4), 15-26

[11] Comparative Analysis of AI-Driven Risk Prediction Methods in Retail Supply Chain Disruption Management: A Multi-Enterprise Study

[12] Does Exit-Timing Flexibility Pay An Empirical Least-Squares Monte Carlo Assessment of Equity Divestment Timing for Utility-Scale Solar in the PJM Market

[13] Chen, Y. & Hu, J. (2026). Graph Neural Network-Based Cascading Disruption Path Identification in Multi-Tier Rare Earth Processing Networks. Journal of Global Engineering Review 4 (1), 99-112

[14] Cao, H. (2024). Privacy-Preserving Click Pattern Anomaly Detection for Mobile In-App Browser Advertising Fraud. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516 2 (2), 151-161

[15] Chen, Y. (2026). Comparing prompting strategies and backbone choices for LLM agents on tabular credit-card fraud detection: An empirical study. Artificial Intelligence and Machine Learning Review, 7(3), 1–11

[16] Multi-Dimensional Training Data Bias Detection and Fairness-Aware Augmentation for Equitable AI-Assisted Medical Imaging Diagnosis

[17] Comparing Imbalanced-Data Handling Strategies for Customer-Support Ticket Text ClassificationAn Empirical Study of Robustness Across Imbalance Ratios Revised as: Comparing Concept-Drift Detection Strategies for Financial Fraud Detection: A Controlled Simulation Study of Detection Latency, False Alarms, and Recovery Across Drift Regimes

[18] Li, Y., Zhang, X., & Hao, C. (2026). Graph Neural Network-Based Cross-Market Risk Contagion Analysis between US Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress. 2026 2nd International Conference on Digital Intelligence and Computing …

[19] Wu, Z., Feng, Z., & Dong, B. (2024). Optimal feature selection for market risk assessment: A dimensional reduction approach in quantitative finance. Journal of Computing Innovations and Applications 2 (1), 20-31

[20] Wu, Z., Wang, S., Ni, C., & Wu, J. (2024). Adaptive traffic signal timing optimization using deep reinforcement learning in urban networks. Artificial Intelligence and Machine Learning Review 5 (4), 55-68

[21] Wu, Z., Feng, E., & Zhang, Z. (2024). Temporal-Contextual Behavioral Analytics for Proactive Cloud Security Threat Detection. Academia Nexus Journal 3 (2)

[22] Li, Y., Zhao, F., & Hu, J. (2026). Identifying Cross-Market Risk Contagion Amplifiers via Graph Attention Networks: Empirical Evidence from US Financial Stress Periods. Journal of Computing Innovations and Applications 4 (1), 164-175

[23] AI-Assisted Disclosure Analysis of EDPR's Onshore Wind Asset Rotation

[24] A Comparative Evaluation of Class Imbalance Handling Techniques for Credit Card Fraud Detection

[25] Wei, C. & Guan, H. (2024). Privacy-preserving federated learning in medical AI: A systematic review of techniques, challenges, and the clinical deployment gap. Artificial Intelligence and Machine Learning Review 5 (3), 124-135

[26] Search-Term Governance Under Broad-Match Automation: A Comparative Evaluation of AI-Assisted Keyword Decision Approaches

[27] A Comparative Network-Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply Network

[28] Artificial Intelligence-Driven Drug Repurposing for Neurodegenerative Diseases: A Computational Analysis and Prediction Study

[29] Wei, C. & Pan, Z. (2026). Accelerating clinical trial recruitment through automated eligibility screening with multi-modal deep learning. Journal of Computing Innovations and Applications 4 (1), 1-11

[30] Yiyi, C. (2025). Federated learning-based framework for privacy-protected cross-border financial risk evaluation: Analyzing US-asia investment flows. Journal of Sustainability, Policy, and Practice, 1(4), 50–65

[31] Comparative Empirical Evaluation of Monte Carlo Dropout Uncertainty Measures for Failure Capture and Referral Trade-offs in Chest X-ray Disease Detection

[32] Does an LLM-Based Feedback Agent Move the Needle? An Empirical Study of Student Correctness and Hint Dependency on the ASSISTments 2017 Interaction Logs

[33] Research on AI-Driven Personalized Web Interface Adaptation Strategies and User Satisfaction Evaluation

[34] Comparative Evaluation of LSTM, GRU, and Bi-LSTM for 30-Day Heart Failure Readmission and Incident Atrial Fibrillation Prediction from Longitudinal Electronic Health Records

[35] AI-Driven Reliability Algorithms for Medical LED Devices: A Research Roadmap

[36] Wei, C., Ge, L. & Brooks, N. (2024). Graph-based representation learning for financial fraud and anomaly transaction detection. Journal of Computing Innovations and Applications 2 (1), 153-164

[37] Kang, A., Zhang, K., & Chen, Y. (2025). AI-Assisted Analysis of Policy Communication during Economic Crises: Correlations with Market Confidence and Recovery Outcomes. Pinnacle Academic Press Proceedings Series 3, 159-173

[38] Zhang, D. & Zheng, Q. (2025). Machine Learning-Based Building Energy Consumption Prediction and Carbon Reduction Potential Assessment in US Metropolitan Areas. Journal of Industrial Engineering and Applied Science 3 (5), 27-40

[39] Profile Verbalization for Interaction-Free Cold-Start Recommendation: A Cross-Domain Empirical Evaluation

[40] Research on Driving Behavior Risk Identification and Safety Assessment Methods Based on Artificial Intelligence

[41] Li, Y. & Fu, X. (2026). Comparative Evaluation of Graph Neural Networks for Cross-Market Risk Contagion Path Identification in Multi-Layer Financial Networks. 2026 2nd International Conference on ArtificialIntelligence and Advanced …

[42] Min, S. & Wei, C. (2023). Comparative analysis of filter-based feature selection methods for high-dimensional data in classification tasks. Journal of Advanced Computing Systems 3 (8), 25-38

[43] Chen, Y. (2026). Comparing Prompting Strategies and Backbone Choices for LLM Agents on Tabular Credit-Card Fraud Detection: An Empirical Study. Journal of Innovative Technologies 2 (1)

[44] Dong, B., Zhang, D., & Xin, J. (2024). Deep reinforcement learning for optimizing order book imbalance-based high-frequency trading strategies. Journal of Computing Innovations and Applications 2 (2), 33-43

[45] An Empirical Comparison of Heuristic and Neural Session-Based Recommendation Methods under Sparse Click-Stream Cold-Start Conditions

[46] Z Feng, C Ni, S Zhou - Spectrum of Research, 2025 - spectrumofresearch.com ◦ Option-Implied Information for Forward-Looking Market Risk Assessment: Evidence from Commodity Derivatives Markets

[47] Comparing Adjacency-List, Edge-List, and Motif-Summary Prompts for Zero-Shot Illicit-Transaction Screening on Elliptic

[48] Kang, A., Xin, J., & Ma, X. (2024). Anomalous cross-border capital flow patterns and their implications for national economic security: An empirical analysis. Journal of Advanced Computing Systems 4 (5), 42-54

[49] Weng, H., Wang, H. & Wei, C. (2024). Adaptive bidding strategies for hybrid auction mechanisms in programmatic advertising. Journal of Advanced Computing Systems 4 (4), 13-25

[50] Spectral Graph Decomposition for Parameter Coordination in Multi-Task LoRA Adaptation

[51] Xiong, K., Wu, Z., & Jia, X. (2025). Deepcontainer: a deep learning-based framework for real-time anomaly detection in cloud-native container environments. Journal of Advanced Computing Systems 5 (1), 1-17

[52] An Empirical Comparative Study of LoRA Fine-Tuning Hyperparameter Configurations on the GLUE Benchmark

[53] Kang, A. & Yu, K. (2025). The Impact of Financial Data Visualization Techniques on Enhancing Budget Transparency in Local Government Decision-Making. Spectrum of Research 5 (2)

[54] Kang, A., Li, C., & Meng, S. (2025). The Impact of Government Budget Data Visualization on Public Financial Literacy and Civic Engagement. Journal of Economic Theory and Business Management 2 (4), 1-16

[55] A Comparative Empirical Study of Machine Learning Methods for Credit Card Fraud Detection: Gradient Boosting, Deep Tabular Learning, and Classical Baselines Under Extreme Class Imbalance

[56] Multi-Horizon Financial Crisis Detection Through Adaptive Data Fusion

[57] P Li, W Liu, Q Zheng - Spectrum of Research, 2024 - spectrumofresearch.com ◦ An Empirical Study on the Quality Assessment of Code Comments Generated by Large Language Models for Different Programming Paradigms

[58] Z Feng, D Zhang, Y Wang - Artificial Intelligence and Machine Learning Review, 2024 ◦ Intraday Liquidity Patterns and Their Implications for Market Risk Assessment: Evidence from Global Equity Markets

[59] Wei, C. & Wu, C. (2024). Credit risk transmission mechanism and prevention strategies in supply chain finance: A core enterprise perspective. Artificial Intelligence and Machine Learning Review 5 (2), 101-115

[60] Feature Weight Optimization in Machine Learning Classifiers for Conflict Escalation Early Warning: Evidence from Diplomatic Signals and News Text

[61] Chen, Y. (2024). Explainable Attack Path Reasoning for Industrial Control Network Security Based on Knowledge Graphs. Journal of Computing Innovations and Applications 2 (1), 128-139

[62] Statistical, Tree-Ensemble, and Neural Methods for Electricity CustomerBehavior Analysis: A Comparative Evaluation across Load-ShapeSegmentation, Multi-Horizon Forecasting, and Theft Detection

[63] Static Structure versus Edge-Order Representation for Drive-by-Download Detection in Linux Audit Graphs: A Reproducible Evaluation

[64] Li, Y. & Mo, Z. (2026). Does a Predicted 0.9 Mean 90 Percent? A Comparative Evaluation of Post-hoc Probability Calibration Methods for Default-Probability Scoring under Class Imbalance. Academia Nexus Journal 5 (2)

[65] Rao, GL., Lu, T., Yan, L., & Liu, Y. (2024). A Hybrid LSTM-KNN Framework for Detecting Market Microstructure Anomalies:: Evidence from High-Frequency Jump Behaviors in Credit Default Swap Markets. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online …

[66] Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities

[67] AI-Assisted Human Oversight in Credit-Card Fraud Detection: A Comparative Evaluation of Review-Routing Strategies under Analyst-Capacity Constraints

[68] Rao, GL., Wang, Z., & Liang, J. (2025). Reinforcement Learning for Pattern Recognition in Cross-Border Financial Transaction Anomalies: A Behavioral Economics Approach to AML. Applied and Computational Engineering 142, 116-127

[69] Rao, GL., Trinh, TK., Chen, Y., Shu, M., & Zheng, S. (2024). Jump Prediction in Systemically Important Financial Institutions' CDS Prices. Spectrum of Research 4 (2)

[70] An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks

[71] Chen, Y. & Chen, Z. (2025). Multi-Objective Deep Reinforcement Learning for Carbon-Aware Spatiotemporal Workload Scheduling in Geo-Distributed Data Centers. Journal of Advanced Computing Systems 5 (10), 18-30

[72] Li, Y. & Ling, Z. (2026). Real-Time Multi-Risk Early Warning for Community Banks: An Application of Ensemble Anomaly Detection and Explainable Artificial Intelligence. Journal of Advanced Computing Systems 6 (2), 15-27

[73] M Sun - Spectrum of Research, 2025 - spectrumofresearch.com ◦ AI-Driven Cross-Cultural Consumer Purchase Intention Prediction Methods for E-commerce Platforms: An Applied Research Approach

[74] Evaluating Tool Selection and Usage Efficiency of LLM-based Agents in Domain-Specific Tasks: A Comparative Analysis

[75] Enhancing Credit Decision Transparency for Small Business Owners: An Explainable AI Approach to Mitigate Algorithmic Bias in Micro-lending

[76] Xin, L. (2025). DeepAd-ocr: An AI-powered framework for automated recognition and enhancement of conversion elements in digital advertisements. Journal of Sustainability, Policy, and Practice, 1(4), 32–49

[77] Benchmarking KV-Cache Management and Chunked-Prefill Scheduling across Mainstream LLM Inference Engines under Production Traces

[78] An Empirical Analysis of Click Temporal Features for Automated Ad Fraud Detection in Mobile In-App Browser Environments

[79] Weng, H., & Li, X. (2024). Renewable-aware cooperative scheduling for distributed AI training across geo-distributed data centers. Artificial Intelligence and Machine Learning Review, 5(2), 91–100

[80] Lei, Y. & Wu, Z. (2025). A Real-Time Detection Framework for High-Risk Content on Short Video Platforms Based on Heterogeneous Feature Fusion. 2025 International Forum on Computing and Intelligent Systems, 93-106

[81] Wang, H. & Sun, Y. (2026). An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks. Journal of Innovative Technologies 2 (1)

[82] A Comparative Study of Spatiotemporal Clustering and Classification Approaches for Security Incident Risk Assessment in UN Peacekeeping Operations

[83] Ge, L. & Rao, G. (2025). MultiStream-FinBERT: A Hybrid Deep Learning Framework for Corporate Financial Distress Prediction Integrating Accounting Metrics, Market Signals, and Textual Disclosures. Pinnacle Academic Press Proceedings Series 3, 107-122

[84] Yang, H. & Finkelstein, J. (2025). Unveiling Pharmacogenomic Patterns in Breast Cancer Through Big Data. Medinfo 2025-Healthcare Smart x Medicine Deep: Proceedings of the 20th World …

[85] Yuan, D. & Zhang, D. (2025). APAC-sensitive anomaly detection: Culturally-aware AI models for enhanced AML in US securities trading. 2025 International Conference on Computer, AI, Systems and Automation, 108-121

[86] Rao, G., Zheng, S., & Guo, L. (2025). Dynamic Reinforcement Learning for Suspicious Fund Flow Detection: A Multi-layer Transaction Network Approach with Adaptive Strategy Optimization. Applied and Computational Engineering 145, 1-11

[87] Comparative Evaluation of Feature Extraction Techniques in Margin Call Cascade Detection: Balancing Accuracy and False Alarm Rates

[88] G Weng, L Yu, L Guo - Spectrum of Research, 2025 - spectrumofresearch.com ◦ Comparative Study of Machine Learning Approaches for Predicting Cardiotoxicity of Anticancer Drugs

[89] Wang, Z. & Kang, A. (2025). FTAFO: A Federated Transparent Adaptive Financial Optimizer for Reducing Third-Party Dependencies in Workflow Management. Journal of Science, Innovation & Social Impact 1 (1), 329-339

[90] Adaptive UV-C LED Dosage Prediction and Optimization Using Neural Networks Under Variable Environmental Conditions in Healthcare Settings

[91] Ju, C. & Rao, G. (2025). Analyzing Foreign Investment Patterns in the US Semiconductor Value Chain Using AI-Enabled Analytics: A Framework for Economic Security. Pinnacle Academic Press Proceedings Series 2, 60-74

[92] A Comparative Evaluation of Rejection Inference Corrections for Credit Card Acquisition Screening: Inverse Probability Weighting, Hard-Cutoff Augmentation, and Semi-Supervised Estimation Revised as: A Comparative Evaluation of Rejection Inference Methods for Credit Application Screening: Inverse Probability Weighting, Hard-Cutoff Augmentation, and Semi-Supervised Self-Training

[93] Wu, Z., Cheng, C., & Zhang, C. (2025). Cloud-Enabled AI Analytics for Urban Green Space Optimization: Enhancing Microclimate Benefits in High-Density Urban Areas. 2025 International Forum on Computing and Intelligent Systems, 123-133

[94] Empirical Evaluation of Constraint Discovery Techniques for Business Logic Consistency Verification in Payroll Data Migration

[95] Zhang, D. & Wang, Y. (2025). AI-driven quality assessment and investment risk identification for carbon credit projects in developing countries. 2025 International Forum on Computing and Intelligent Systems, 76-92

[96] Can Zero-Shot Prompting Classify Risky Bitcoin Transactions from Compact One-Hop Narratives? A Small-Scale Exploratory Study

[97] Zhenghao, P. (2025). A reinforcement learning approach for adaptive budget allocation in pharmaceutical digital marketing: Maximizing roi across patient journey touchpoints. Journal of Sustainability, Policy, and Practice, 1(4), 1–15

[98] Guo, X. & Sun, Y. (2026). An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting. Journal of Selected Topics in Academic Research 2 (1)

[99] Gradient-Boosted Trees Versus Classical Time-Series Baselines for 7- to 28-Day-Ahead Emergency-Department Respiratory Visit Pressure During Viral Co-Circulation: A County-Level Comparative Evaluation

[100] Tian, Y., Pan, H., & Sun, Y. (2026). A Comparative Network-Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply …. Journal of Selected Topics in Academic Research 2 (1)

[101] Xu, S. (2025). Intelligent optimization algorithm for chain restaurant spatial layout based on generative adversarial networks. Journal of Industrial Engineering and Applied Science, 3(3), 32–41

[102] Kang, A., Li, Z., & Meng, S. (2023). AI-Enhanced Risk Identification and Intelligence Sharing Framework for Anti-Money Laundering in Cross-Border Income Swap Transactions. Journal of Advanced Computing Systems 3 (5), 34-47

[103] Road-Geometry, Scenario-Diversity, and Failure-History Prioritization for BeamNG Lane-Keeping Simulation Regression Tests under Fixed Simulation-Time Budgets: An Empirical Comparison

[104] Chen, Y., Chen, Z., & Zou, D. (2025). CarbonShift: Harnessing Grid Carbon Variability for Geo-Distributed Workload Scheduling. Artificial Intelligence and Machine Learning Review 6 (4), 18-31

[105] Zhang, D. & Feng, E. (2024). Quantitative Assessment of Regional Carbon Neutrality Policy Synergies Based on Deep Learning. Journal of Advanced Computing Systems 4 (10), 38-54

[106] Zhang, D. & Zhang, F. (2025). AI-Assisted Identification and Equity Assessment of Vulnerable Population Impacts in US Energy Transition. Journal of Advanced Computing Systems 5 (7), 1-17

[107] Rao, G., Ju, C., & Feng, Z. (2024). AI-Driven Identification of Critical Dependencies in US-China Technology Supply Chains: Implications for Economic Security Policy. Journal of Advanced Computing Systems 4 (12), 43-57

[108] Zhang, D. & Ma, X. (2025). Machine Learning-Based Credit Risk Assessment for Green Bonds: Climate Factor Integration and Default Prediction Analysis. Journal of Sustainability, Policy, and Practice 1 (2), 121-135

[109] Cao, H. (2024). Detecting Fraudulent Click Patterns in Mobile In-App Browsers: A Multi-dimensional Behavioral Analysis Approach. The Artificial Intelligence and Machine Learning Review 5 (2), 130-142

[110] Li, Y. & Mo, Z. (2026). Backtesting Value-at-Risk and Expected Shortfall Across Calm and Turbulent Markets: An Empirical Comparison of Historical Simulation, GARCH-Family, Filtered Historical …. Annals of Applied Sciences 7 (1)

[111] Zhang, S., Wang, Y., & Weng, H. (2024). Industrial IoT anomaly detection using improved autoencoder architecture. Artificial Intelligence and Machine Learning Review, 5(1), 67–78

[112] Lightweight Neural Networks with Attention Mechanism for Loop Closure Detection in Visual SLAM

[113] Li, Y. & Zhang, S. (2025). Machine Learning-Based Credit Risk Early Warning System for Small and Medium-Sized Financial Institutions: An Ensemble Learning Approach with Interpretable Risk Indicators. Journal of Science, Innovation & Social Impact 1 (1), 372-383

[114] Evaluating the Impact of Data Augmentation Techniques on Few-Shot Image Classification Performance

[115] Kang, A. & Ma, X. (2025). AI-Based Pattern Recognition and Characteristic Analysis of Cross-Border Money Laundering Behaviors in Digital Currency Transactions. Pinnacle Academic Press Proceedings Series 5, 1-19

[116] Can Task-Prefixed Text Embeddings Exploit Anonymized Financial-Graph Features? A Low-Label Comparison on T-Finance and Elliptic

[117] Zhang, Z. & Wu, Z. (2023). Context-aware feature selection for user behavior analytics in zero-trust environments. Journal of Advanced Computing Systems 3 (5), 21-33

[118] Li, Y. (2026). Enhancing Financial Compliance Transparency through Automated Data Governance and Intelligent Risk Reporting. Journal of Science, Innovation & Social Impact 2 (1), 299-313

[119] Cao, H. & Shi, W. (2026). Statistical Anomaly Detection Approach for Field Mapping Validation in Enterprise Payroll Data Migration. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516 4 (1), 137-153

[120] Alert Fatigue Mitigation in Anomaly Detection Systems: A Comparative Study of Threshold Optimization and Alert Aggregation Strategies

[121] A Controlled Comparison of Data-Augmentation Strategies for Surface-EMG Hand-Gesture Recognition under Small-Sample Conditions

[122] Ye, L. (2025). Rlhf-powered multilingual audio understanding: A cross-cultural emotion analysis framework for international communication. Journal of Sustainability, Policy, and Practice, 1(4), 66–79

[123] Li, Y. & Long, L. (2026). Lightweight AI-Driven Stress Testing for Small and Medium Financial Institutions: A Variational Autoencoder Approach with Extreme Value Theory for Macroeconomic Scenario Generation. Artificial Intelligence and Machine Learning Review 7 (1), 108-119

[124] H Lian, X Wang, C Zhang - Academia Nexus Journal, 2025 - academianexusjournal.com ◦ AI-Powered Anomaly Detection in Cloud Environments: A Lightweight Security Framework under Zero Trust Architecture

[125] Sun, Y. & Tang, R. (2026). A Comparative Evaluation of Post-hoc Explainability Methods for Graph Neural Networks in Node Classification. Journal of Innovative Technologies 2 (2)

[126] A Comparative Evaluation of Post-hoc Explainability Methods for Graph Neural Networks in Node Classification

[127] Zhang, F., Ye, H. & Wei, C. (2024). Leveraging multi-modal attention mechanisms for interpretable biomarker discovery and early disease prediction. Journal of Computing Innovations and Applications 2 (2), 111-121

[128] Chen, Y. & Lai, J. (2026). Multi-Metric Trustworthiness Evaluation of AI-Assisted Medical Imaging Diagnosis: Integrating Confidence Calibration and Distribution Shift Detection. Journal of Global Engineering Review 4 (1), 113-126

[129] Y Li, S Fan, H Wang - Spectrum of Research, 2025 - spectrumofresearch.com ◦ Machine Learning-Based Identification of Anomalous Trading Behavior Patterns Among Asia-Pacific Investors in U.S. Securities Markets

[130] A Comparative Analysis of Attention Mechanisms in Vision Transformers for Fine-Grained Image Recognition

[131] Comparative Analysis of Deep Learning Algorithms for Disease-Related Protein Function Prediction: Performance Optimization and Computational Efficiency Evaluation

[132] Concept-Drift Detection for Systemic-Risk Early Warning: A Comparative Evaluation of Regime-Conditional Adaptive Thresholding under Distribution Shifts in Financial Stress Indices

[133] AI-Enabled Cultural Feature Recognition and Cross-Cultural Comparison in Historic Architecture

[134] Tension Among Fairness Metrics and the Effectiveness of Pre-, In-, and Post-Processing Bias Mitigation: An Empirical Comparison on Public Credit Data

[135] Machine Learning for Real-time Optimization of Bioprocessing Parameters:Applications and Improvements

[136] Cloud-based Data Mining for Cancer Drug Synergy Analysis: Applications in Non-small Cell Lung Cancer Treatment

[137] Cao, H. & Yu, M. (2026). Do Neighborhood-Aggregated Features Improve Gradient Boosting for Node-Level Fraud Detection? An Empirical Study on Two Multi-Relational Graphs. Journal of Selected Topics in Academic Research 2 (1)

[138] Kang, A., Min, S., & Yuan, D. (2024). Comparative Analysis of Foreign Exchange Market Shock Transmission and Recovery Resilience Among Major Economies Under Geopolitical Conflicts: Evidence from the Russia-Ukraine …. Journal of Computing Innovations and Applications 2 (1), 46-61

[139] Deng, M., Fu, H., & Sun, Y. (2026). Incremental Contribution of Linguistic, Behavioral, and Propagation Signals to Early Detection of LLM-Generated Coordinated Inauthentic Campaigns. Spectrum of Research 6 (1)

[140] Fairness-Aware Credit Risk Assessment Using Alternative Data: An Explainable AI Approach for Bias Detection and Mitigation

[141] Cao, H., Hu, J., & Luo, C. (2025). Behavioural Feature Analysis for Anomalous Click Detection in Mobile Advertising Environments: Toward In-App Browser-Specific Detection. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516 3 (2), 96-105

[142] Position Bias and Verbosity Bias in LLM-as-a-Judge: An Empirical Cross-Benchmark Comparison

[143] Quantifying Risk-Factor Attribution in Utility-Scale Power-Project Returns: A SHAP-Based Comparative Study of Ensemble Learners for Unlevered-IRR Distribution Prediction

[144] Effects of Prompt Example Count on Money-Laundering Transaction Classification from Textualized Financial-Graph Neighborhoods: An Empirical Study

[145] Yang, H. & Finkelstein, J. (2025). Using Big Data to Uncover Drug-Gene Interaction in Patients with Prostate Cancer. Envisioning the Future of Health Informatics and Digital Health, 131-135

[146] Jin, Z. (2025). SecureCodeBERT: An ai-powered model for identifying and categorizing high-risk security vulnerabilities in php-based critical infrastructure applications. Journal of Sustainability, Policy, and Practice, 1(4), 80–94

[147] Liu, W., Rao, G., & Lian, H. (2023). Anomaly Pattern Recognition and Risk Control in High-Frequency Trading Using Reinforcement Learning. Journal of Computing Innovations and Applications 1 (2), 47-58

[148] A Comparative Evaluation of Loss Reweighting and Neighborhood Resampling for Rare-Event Risk Node Detection in Financial Graphs

[149] Wu, Z., Zhang, Z., Zhao, Q., & Yan, L. (2025). Privacy-preserving financial transaction pattern recognition: A differential privacy approach. Preprints

[150] Quantifying the Throughput–Tail-Latency Trade-off of Continuous Batching and Chunked-Prefill Scheduling under Production LLM Traces: An Empirical Comparison

[151] Financing Date Uncertainty and Liquidity Forecasts for Utility-Scale Solar Projects

[152] Graph Neural Network–Based Cross-Market Risk Contagion Analysis between U.S. Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress

[153] An Empirical Comparison of XGBoost and LightGBM for Capital Expenditure Deviation Prediction in U.S. FERC-Regulated Rate-Base Transmission Projects

[154] Comparative Empirical Evaluation of Prompting Strategies for Large-Language-Model Web-Navigation Agents on the WebArena Benchmark

[155] Guo, Y. & Wei, C. (2026). Latency-adaptive feature fusion weight allocation under bandwidth constraints for V2X cooperative 3D object detection. Journal of Advanced Computing Systems 6 (3), 22-31

[156] Weng, H., Zhang, S., & Min, S. (2024). Multi-constraint optimization for real-time bidding: A reinforcement learning approach. Artificial Intelligence and Machine Learning Review, 5(1), 93–104