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

Evidence-Centered Clinical Intelligence for Drug Discovery, Trial Recruitment, and Personalized Care

Yinchu Wu

Computer Technology, Shanghai Jiao Tong University, Shanghai, China
Corresponding Email: wuyinchu@qq.com

Abstract

This paper develops an evidence-centered clinical intelligence architecture that links drug-gene reasoning, trial eligibility screening, medical imaging, and personalized care while preserving calibration and privacy. 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*.** clinical intelligence; drug discovery; trial recruitment; pharmacogenomics; medical imaging; personalized care

1 Introduction

Clinical AI is increasingly expected to connect evidence that has traditionally been separated. A genomic profile, an imaging study, an electronic health record, a trial protocol, and a patient's stated preference each describe a different part of the same care pathway. The supplied bibliography includes pharmacogenomic pattern discovery, drug-gene interaction analysis, clinical trial recruitment, biomarker prediction, drug repurposing, medical imaging fairness, disease prediction, and privacy-preserving medical AI. This paper uses that breadth to define a clinical intelligence architecture whose outputs remain traceable to the evidence available at decision time.

The proposed architecture is evidence-centered rather than model-centered. It treats the clinical question, the eligible population, the provenance of each modality, and the uncertainty of the recommendation as first-class objects. A retrieval module can organize evidence from records and

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/)

literature, a multimodal model can align imaging and structured data, and a rule or constraint layer can enforce eligibility and safety conditions. The system does not assume that a high predictive score is sufficient for a treatment or recruitment decision.

The contribution is a literature-grounded framework and comparison protocol. It separates discovery, screening, prediction, and care coordination so that a model can be evaluated for the task it actually supports. Tables map clinical claims to required evidence, and figures show a pathway from molecular and patient data to a human-reviewed decision. No patient records or clinical outcomes are invented; the paper specifies how a real dataset and governance process should be used.

The paper addresses a practical question: how can heterogeneous clinical evidence be combined without hiding uncertainty or weakening patient autonomy? The answer is a staged pipeline with explicit eligibility logic, calibration, privacy controls, and a feedback mechanism for clinicians and patients.

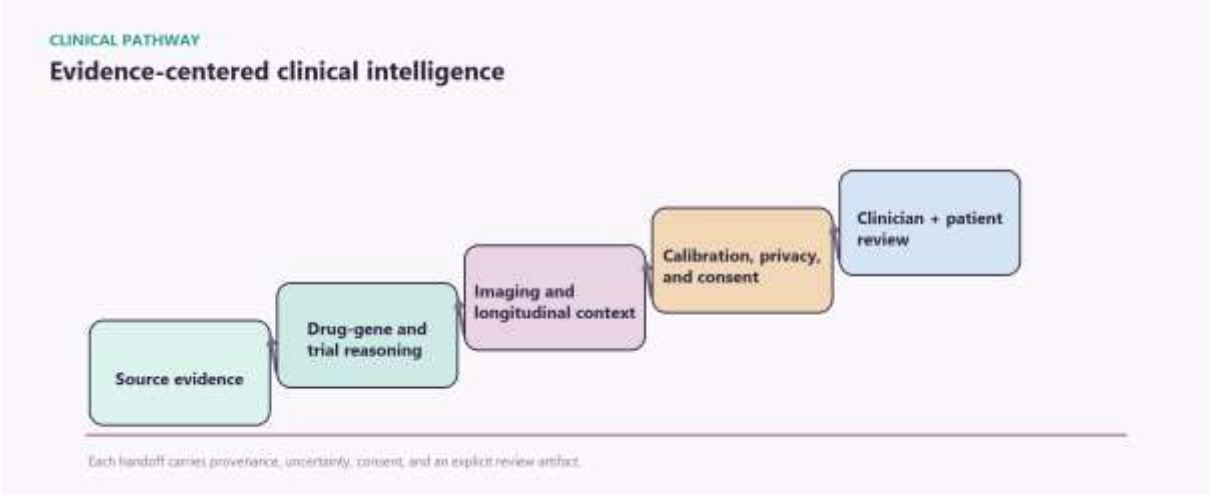


Figure 1. Evidence-centered clinical intelligence pathway.

Table 1. Evidence-centered clinical claim and review contract.

Clinical claim	*Required evidence*	*Review artifact*
Drug-gene association	*Molecular signal and source*	*Evidence grade and version*
Trial eligibility	*Protocol criteria and patient record*	*Eligible, ineligible, unresolved states*
Risk prediction	*Longitudinal outcome and calibration*	*Reliability and subgroup report*
Imaging support	*Image, device, and context*	*Modality contribution analysis*

Clinical claim	*Required evidence*	*Review artifact*
Care recommendation	*Constraints and preference*	*Clinician and patient review*

2 Related Work and Problem Formulation

Let a patient episode consist of multimodal evidence E = {E_m}, a clinical question q, a candidate action a, and a set of constraints C. The system should return an evidence bundle, a calibrated estimate, and an uncertainty statement rather than an unexplained label. The bundle records which observations, time windows, and molecular or imaging features contributed to the recommendation. This formulation supports both clinical prediction and trial recruitment because the same patient evidence can be mapped to different decision contracts.

Clinical eligibility is a reasoning problem with temporal and logical structure. A protocol may require a combination of diagnoses, laboratory ranges, prior therapies, performance status, and exclusion conditions. A screening system must distinguish missing evidence from evidence of exclusion. It should also expose which criteria remain unresolved so that a coordinator can request a new record rather than silently reject a patient.

The data are heterogeneous and often incomplete. Imaging devices, coding systems, molecular assays, and clinical notes use different vocabularies and sampling schedules. The architecture therefore includes a semantic alignment layer and a missingness policy. If a modality is unavailable or outside the validated range, the model should abstain, request review, or use a documented fallback instead of treating absence as a normal value.

The supplied literature cluster [1-4] is useful for the present problem because it connects clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Research on Driving Behavior Risk Identification and Safety Assessment Methods Based on Artificial Intelligence; Comparing Unsupervised Spatial Grouping and Supervised Risk Ranking for County-Level Measles Susceptibility Stratification under Uneven Childhood Immunization C; Accelerating clinical trial recruitment through automated eligibility screening with multi-modal deep learning. Journal of Computing Innovations and Application; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Machine Learning-Based Identification of Anomalous Trading Behavior Patterns Among Asia-Pacific Investors i. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include hoc Explainability Methods for Graph Neural Networks in Node Classification. Journal of Innovative Technologies 2 (2); Enhancing Financial Compliance Transparency through Automated Data Governance and Intelligent Risk Reporting. Journal of Science, Innovation & Social Impact 2 (; Shot Prompting Classify Risky Bitcoin Transactions from Compact One-Hop Narratives? A Small-Scale Exploratory Study. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Aggregated Features Improve Gradient Boosting for Node-Level Fraud Detection? An Empirical Study on Two Multi-Relational Graphs. Journal of Selected Topics in A; Based Credit Risk Assessment for Green Bonds: Climate Factor Integration and Default Prediction Analysis. Journal of Sustainability, Policy, and Practice 1 (2),; Comparing prompting strategies and backbone choices for LLM agents on tabular credit-card fraud detection: An empirical study. Artificial Intelligence and Machi; Comparative Evaluation of Feature Extraction Techniques in Margin Call Cascade Detection: Balancing Accuracy and False Alarm Rates. 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 hoc Explainability Methods for Graph Neural Networks in Node Classification; Quantitative Assessment of Regional Carbon Neutrality Policy Synergies Based on Deep Learning. Journal of Advanced Computing Systems 4 (10), 38-54; Tension Among Fairness Metrics and the Effectiveness of Pre-, In-, and Post-Processing Bias Mitigation: An Empirical Comparison on Public Credit Data; Assisted Human Oversight in Credit-Card Fraud Detection: A Comparative Evaluation of Review-Routing Strategies under Analyst-Capacity Constraints. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Profile Verbalization for Interaction-Free Cold-Start Recommendation: A Cross-Domain Empirical Evaluation; Data Handling Strategies for Customer-Support Ticket Text ClassificationAn Empirical Study of Robustness Across Imbalance Ratios Revised as: Comparing Concept-D; A Comparative Empirical Study of Machine Learning Methods for Credit Card Fraud Detection: Gradient Boosting, Deep Tabular Learning, and Classical Baselines Und. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities; An Empirical Comparative Study of LoRA Fine-Tuning Hyperparameter Configurations on the GLUE Benchmark; Spectral Graph Decomposition for Parameter Coordination in Multi-Task LoRA Adaptation; Based Credit Risk Early Warning System for Small and Medium-Sized Financial Institutions: An Ensemble Learning Approach with Interpretable Risk Indicators. Jour. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Alert Fatigue Mitigation in Anomaly Detection Systems: A Comparative Study of Threshold Optimization and Alert Aggregation Strategies; Driven Identification of Critical Dependencies in US-China Technology Supply Chains: Implications for Economic Security Policy. Journal of Advanced Computing Sy; Empirical Evaluation of Constraint Discovery Techniques for Business Logic Consistency Verification in Payroll Data Migration; An Empirical Comparison of Heuristic and Neural Session-Based Recommendation Methods under Sparse Click-Stream Cold-Start Conditions. 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 Contextual Behavioral Analytics for Proactive Cloud Security Threat Detection. Academia Nexus Journal 3 (2); CarbonShift: Harnessing Grid Carbon Variability for Geo-Distributed Workload Scheduling. Artificial Intelligence and Machine Learning Review 6 (4), 18-31; 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply …. Jou; Driven Drug Repurposing for Neurodegenerative Diseases: A Computational Analysis and Prediction Study; Industrial IoT anomaly detection using improved autoencoder architecture. Artificial Intelligence and Machine Learning Review, 5(1), 67–78; Geometry, Scenario-Diversity, and Failure-History Prioritization for BeamNG Lane-Keeping Simulation Regression Tests under Fixed Simulation-Time Budgets: An Emp. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include aware cooperative scheduling for distributed AI training across geo-distributed data centers. Artificial Intelligence and Machine Learning Review, 5(2), 91–100; Spectrum of Research, 2024 - spectrumofresearch.com ◦ An Empirical Study on the Quality Assessment of Code Comments Generated by Large Language Models for Differ; adaptive feature fusion weight allocation under bandwidth constraints for V2X cooperative 3D object detection. Journal of

Advanced Computing Systems 6 (3), 22-3; Artificial Intelligence and Machine Learning Review, 2024 ◦ Intraday Liquidity Patterns and Their Implications for Market Risk Assessment: Evidence from Global E. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Comparative Study of Machine Learning Approaches for Predicting Cardiotoxicity of Anticancer Drugs; at-Risk and Expected Shortfall Across Calm and Turbulent Markets: An Empirical Comparison of Historical Simulation, GARCH-Family, Filtered Historical …. Annals . 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 Optimal feature selection for market risk assessment: A dimensional reduction approach in quantitative finance. Journal of Computing Innovations and Application; Based Cross-Market Risk Contagion Analysis between US Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress. 2026 2nd International Confere; Drift Detection for Systemic-Risk Early Warning: A Comparative Evaluation of Regime-Conditional Adaptive Thresholding under Distribution Shifts in Financial Str; C LED Dosage Prediction and Optimization Using Neural Networks Under Variable Environmental Conditions in Healthcare Settings. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Time Multi-Risk Early Warning

for Community Banks: An Application of Ensemble Anomaly Detection and Explainable Artificial Intelligence. Journal of Advanced Com; Interpretable Gradient Boosting for Transmission Capital Expenditure Deviation: A Multi-Decade Benchmark on FERC-Regulated Utilities. Academia Nexus Journal 5 (; A Comparative Evaluation of Rejection Inference Corrections for Credit Card Acquisition Screening: Inverse Probability Weighting, Hard-Cutoff Augmentation, and ; Enabled Cultural Feature Recognition and Cross-Cultural Comparison in Historic Architecture. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Objective Deep Reinforcement Learning for Carbon-Aware Spatiotemporal Workload Scheduling in Geo-Distributed Data Centers. Journal of Advanced Computing Systems; Evaluating Tool Selection and Usage Efficiency of LLM-based Agents in Domain-Specific Tasks: A Comparative Analysis; List, Edge-List, and Motif-Summary Prompts for Zero-Shot Illicit-Transaction Screening on Elliptic. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover FinBERT: A Hybrid Deep Learning Framework for Corporate Financial Distress Prediction Integrating Accounting Metrics, Market Signals, and Textual Disclosures. P; ocr: An AI-powered framework for automated recognition and enhancement of conversion elements in digital advertisements. Journal of Sustainability, Policy, and ; The Impact of Financial Data Visualization Techniques on Enhancing Budget Transparency in Local Government Decision-Making. Spectrum of Research 5 (2); Analyzing Foreign Investment Patterns in the US Semiconductor Value Chain Using AI-Enabled Analytics: A Framework for Economic Security. Pinnacle Academic Press. 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 based feature selection methods for high-dimensional data in classification tasks. Journal of Advanced Computing Systems 3 (8), 25-38; based representation learning for financial fraud and anomaly transaction detection. Journal of Computing Innovations and Applications 2 (1), 153-164; time Optimization of Bioprocessing Parameters:Applications and Improvements; Horizon Financial Crisis Detection Through Adaptive Data Fusion. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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.

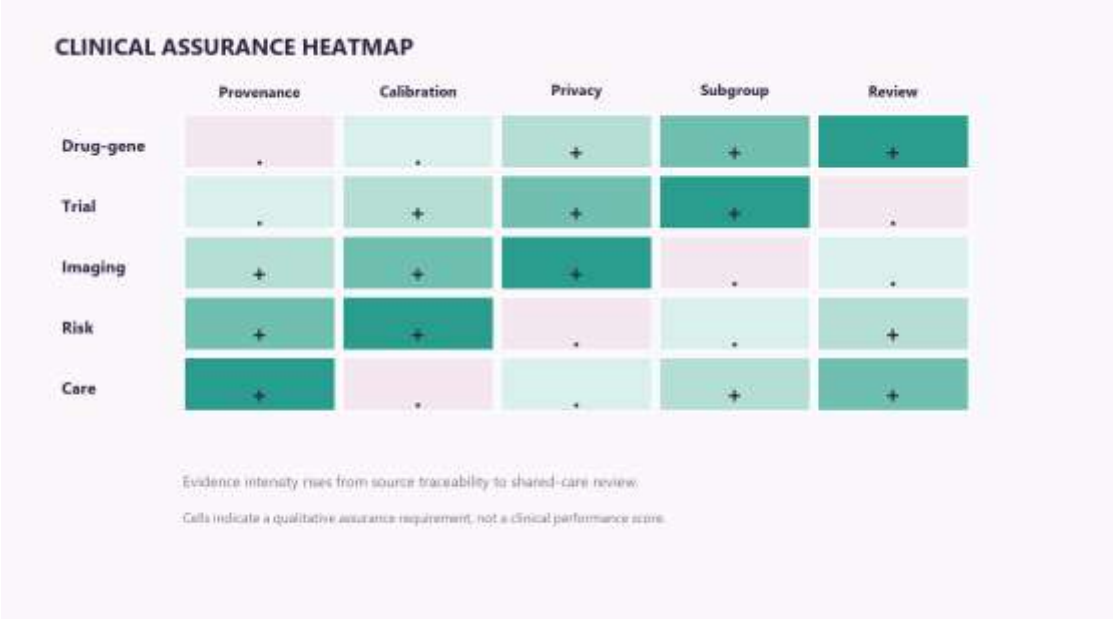


Figure 2. Qualitative assurance matrix for clinical evidence integration.

3 Proposed Framework

The framework has five stages. The evidence layer ingests records, images, molecular measurements, protocol text, and patient preferences with timestamps and provenance. The alignment layer maps modality-specific representations into a patient episode while preserving missingness and confidence. The reasoning layer performs eligibility checks, risk estimation, or drug-gene analysis. The assurance layer calibrates probabilities, audits subgroup behavior, and applies privacy controls. The care layer presents the evidence bundle to a clinician or coordinator and records the final decision.

Drug-gene reasoning is treated as a constrained evidence task. The system separates a molecular association, a mechanistic hypothesis, and a clinically actionable recommendation. Each link is associated with a source, an evidence grade, and a version. This prevents a statistical association from being rendered as a treatment instruction without the required clinical context.

Trial recruitment uses a two-pass design. A fast pass extracts candidate criteria and screens high-confidence exclusions. A second pass checks unresolved criteria against the original protocol text and relevant records. The system reports eligible, ineligible, and unresolved states separately. This design protects recall while giving coordinators a manageable queue for manual confirmation.

Medical imaging and longitudinal prediction are evaluated with calibration and subgroup analysis. The comparison includes simple baselines, multimodal models, and models with missing-modality handling. Tests include site-held-out validation, temporal drift, device changes, and label delays. The goal is to determine whether added modalities improve a clinical decision, not merely whether they lower a training loss.

Privacy and patient agency are embedded in the release process. Data access is minimized, model outputs are logged, and patients or clinicians can challenge a recommendation. A rollback package contains the model version, evidence schema, calibration map, and policy configuration. These controls make a clinical intelligence system auditable after deployment.

The supplied literature cluster [60-62] is useful for the present problem because it connects clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Aware Credit Risk Assessment Using Alternative Data: An Explainable AI Approach for Bias Detection and Mitigation; Effects of Prompt Example Count on Money-Laundering Transaction Classification from Textualized Financial-Graph Neighborhoods: An Empirical Study; Prefixed Text Embeddings Exploit Anonymized Financial-Graph Features? A Low-Label Comparison on T-Finance and Elliptic. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Time Detection Framework for High-Risk Content on Short Video Platforms Based on Heterogeneous Feature Fusion. 2025 International Forum on Computing and Intelli; border capital flow patterns and their implications for national economic security: An empirical analysis. Journal of Advanced Computing Systems 4 (5), 42-54; Augmentation Strategies for Surface-EMG Hand-Gesture Recognition under Small-Sample Conditions; An Empirical Analysis of Click Temporal Features for Automated Ad Fraud Detection in Mobile In-App Browser Environments. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover A Comparative Study of Spatiotemporal Clustering and Classification Approaches for Security Incident Risk Assessment in UN Peacekeeping Operations; Market Risk Contagion Amplifiers via Graph Attention Networks: Empirical Evidence from US Financial Stress Periods. Journal of Computing Innovations and Applica; based framework for privacy-protected cross-border financial risk evaluation: Analyzing US-asia investment flows. Journal of Sustainability, Policy, and Practic; Adaptive bidding strategies for hybrid auction mechanisms in programmatic advertising. Journal of Advanced Computing Systems 4 (4), 13-25. 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 How Much Transaction Structure Should an LLM See? Zero- and Few-Shot Evaluation of Local Transaction-Network Summaries for Money-Laundering Transaction Classifi; Enabled AI Analytics for Urban Green Space Optimization: Enhancing Microclimate Benefits in High-Density Urban Areas. 2025 International Forum on Computing and ; powered customer engagement sequence analysis and conversion funnel optimization in multi-channel e-commerce marketing. Journal of Industrial Engineering and Ap. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Dynamic Reinforcement Learning for Suspicious Fund Flow Detection: A Multi-layer Transaction Network Approach with Adaptive Strategy Optimization. Applied and C; Lightweight Neural Networks with Attention Mechanism for Loop Closure Detection in Visual SLAM; Spectrum of Research, 2025 -

spectrumofresearch.com ◦ AI-Driven Cross-Cultural Consumer Purchase Intention Prediction Methods for E-commerce Platforms: An Applie; Jump Prediction in Systemically Important Financial Institutions' CDS Prices. Spectrum of Research 4 (2). 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Unveiling Pharmacogenomic Patterns in Breast Cancer Through Big Data. Medinfo 2025-Healthcare Smart x Medicine Deep: Proceedings of the 20th World …; Ensemble, and Neural Methods for Electricity CustomerBehavior Analysis: A Comparative Evaluation across Load-ShapeSegmentation, Multi-Horizon Forecasting, and T; Adaptive traffic signal timing optimization using deep reinforcement learning in urban networks. Artificial Intelligence and Machine Learning Review 5 (4), 55-6; modal attention mechanisms for interpretable biomarker discovery and early disease prediction. Journal of Computing Innovations and Applications 2 (2), 111-121. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Term Governance Under Broad-Match Automation: A Comparative Evaluation of AI-Assisted Keyword Decision Approaches; Algorithmic fairness in financial decision-making: Detection and mitigation of bias in credit scoring applications. Journal of Advanced Computing Systems 4 (2),; 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 [85-88], spanning Position Bias and Verbosity Bias in LLM-as-a-Judge: An Empirical Cross-Benchmark Comparison; A Comparative

Evaluation of Class Imbalance Handling Techniques for Credit Card Fraud Detection; Driven Reliability Algorithms for Medical LED Devices: A Research Roadmap; based Data Mining for Cancer Drug Synergy Analysis: Applications in Non-small Cell Lung Cancer Treatment. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Based Cross-Market Risk Contagion Analysis between U.S. Equity Sectors and Treasury Yields during the 2023 Regional Banking Stress; Enhancing Credit Decision Transparency for Small Business Owners: An Explainable AI Approach to Mitigate Algorithmic Bias in Micro-lending; Academia Nexus Journal, 2025 - academianexusjournal.com ◦ AI-Powered Anomaly Detection in Cloud Environments: A Lightweight Security Framework under Zero Trust A; Boosted Trees Versus Classical Time-Series Baselines for 7- to 28-Day-Ahead Emergency-Department Respiratory Visit Pressure During Viral Co-Circulation: A Count. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Preserving Click Pattern Anomaly Detection for Mobile In-App Browser Advertising Fraud. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516 2 (2), 151-161; Comparative Analysis of Foreign Exchange Market Shock Transmission and Recovery Resilience Among Major Economies Under Geopolitical Conflicts: Evidence from the; Based Pattern Recognition and Characteristic Analysis of Cross-Border Money Laundering Behaviors in Digital Currency Transactions. Pinnacle Academic Press Proce. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Order Representation for Drive-by-Download Detection in Linux Audit Graphs: A Reproducible Evaluation; An Empirical Comparison of Missing Value Imputation Methods for Alzheimer's Disease Cognitive Assessments and Their Downstream Impact on Diagnostic Classificati; aware feature selection for user behavior analytics in zero-trust environments. Journal of Advanced Computing Systems 3 (5), 21-33; Timing Flexibility Pay An Empirical Least-Squares Monte Carlo Assessment of Equity Divestment Timing for Utility-Scale Solar in the PJM Market. 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 Deep reinforcement learning for optimizing order book imbalance-based high-frequency trading strategies. Journal of Computing Innovations and Applications 2 (2); Based Cascading Disruption Path Identification in Multi-Tier Rare Earth Processing Networks. Journal of Global Engineering Review 4 (1), 99-112; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Option-Implied Information for Forward-Looking Market Risk Assessment: Evidence from Commodity Derivatives ; Statistical Anomaly Detection Approach for Field Mapping Validation in Enterprise Payroll Data Migration. COMPUTING INNOVATIONS AND APPLICATIONS ISSN: 3068-5516. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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. Clinical intelligence layers and safeguards.

Component	*Clinical role*	*Primary safeguard*
Evidence layer	*Collect and timestamp data*	*Provenance and access control*
Alignment layer	*Connect modalities*	*Missingness policy*
Reasoning layer	*Screen or predict*	*Constraint and uncertainty checks*

Component	*Clinical role*	*Primary safeguard*
Assurance layer	*Calibrate and audit*	*Drift and subgroup monitoring*
Care layer	*Support decisions*	*Override and appeal path*

4 Comparative Analysis and Discussion

The literature suggests that clinical AI benefits from separating discovery from action. A model can help identify a pattern or prioritize a chart while still requiring human confirmation before a consequential decision. The proposed architecture keeps this boundary visible through evidence bundles and unresolved states.

A second implication is that fairness and privacy must be evaluated together. Privacy constraints can reduce the visibility of small subgroups, while biased labels can persist even when protected attributes are removed. The framework therefore uses protected audit data under controlled access and reports calibration and error patterns by clinically meaningful group.

A third implication concerns workflow fit. A technically strong model can fail if it adds documentation burden or returns an answer at the wrong point in the care pathway. Evaluation should include time-to-review, unresolved-case handling, and clinician acceptance, while preserving the ability to ignore or override a recommendation.

The framework cannot resolve causal treatment effects from observational data and does not replace clinical judgment. Its purpose is to organize evidence, surface uncertainty, and make a later validation study reproducible.

The supplied literature cluster [104-106] is useful for the present problem because it connects clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Detecting Fraudulent Click Patterns in Mobile In-App Browsers: A Multi-dimensional Behavioral Analysis Approach. The Artificial Intelligence and Machine Learnin; Gene Interaction in Patients with Prostate Cancer. Envisioning the Future of Health Informatics and Digital Health, 131-135; 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include An Empirical Comparison of XGBoost and LightGBM for Capital Expenditure Deviation Prediction in U.S. FERC-Regulated Rate-Base Transmission Projects; Tail-Latency Trade-off of Continuous Batching and Chunked-Prefill Scheduling under Production LLM Traces: An Empirical Comparison; preserving financial transaction pattern recognition: A differential privacy approach. Preprints; Driven Personalized Web Interface Adaptation Strategies and User Satisfaction Evaluation. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover based framework for real-time anomaly detection in cloud-native container environments. Journal of Advanced Computing Systems 5 (1), 1-17; Dimensional Training Data Bias Detection and Fairness-Aware Augmentation for Equitable AI-Assisted Medical Imaging Diagnosis; powered multilingual audio understanding: A cross-cultural emotion analysis framework for international communication. Journal of Sustainability, Policy, and Pr; powered model for identifying and categorizing high-risk security vulnerabilities in php-based critical infrastructure applications. Journal of Sustainability, . 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 FTAFO: A Federated Transparent Adaptive Financial Optimizer for Reducing Third-Party Dependencies in Workflow Management. Journal of Science, Innovation & Socia; Centrality, Community-Detection, and Cascade-Simulation Analysis of Disruption Propagation and Concentrated Vulnerability in a Real Inter-Sectoral Supply Networ; A Comparative Evaluation of Loss Reweighting and Neighborhood Resampling for Rare-Event Risk Node Detection in Financial Graphs. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Driven Risk Prediction Methods in Retail Supply Chain Disruption Management: A Multi-Enterprise Study; Comparative Evaluation of LSTM, GRU, and Bi-LSTM for 30-Day Heart Failure Readmission and Incident Atrial Fibrillation Prediction from Longitudinal Electronic H; Does a Predicted 0.9 Mean 90 Percent? A Comparative Evaluation of Post-hoc Probability Calibration Methods for Default-Probability Scoring under Class Imbalance; Explainable Attack Path Reasoning for Industrial Control Network Security Based on Knowledge Graphs. Journal of Computing Innovations and Applications 2 (1), 12. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Assisted Disclosure Analysis of EDPR’s Onshore Wind Asset Rotation; Feature Weight Optimization in Machine Learning Classifiers for Conflict Escalation Early Warning: Evidence from Diplomatic Signals and News Text; Cache Management and Chunked-Prefill Scheduling across Mainstream LLM Inference Engines under Production Traces; 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 [126-128] broaden the evidence base for clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Evaluating the Impact of Data Augmentation Techniques on Few-Shot Image Classification Performance; preserving federated learning in medical AI: A systematic review of techniques, challenges, and the clinical deployment gap. Artificial Intelligence and Machine; An Empirical Comparison of Representation-Learning Clustering and Causal Effect Estimation for Electricity Customer Segmentation and Demand Response Targeting. . 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 KNN Framework for Detecting Market Microstructure Anomalies:: Evidence from High-Frequency Jump Behaviors in Credit Default Swap Markets. Journal of Knowledge L; Based Building Energy Consumption Prediction and Carbon Reduction Potential Assessment in US Metropolitan Areas. Journal of Industrial Engineering and Applied S; An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks. Journal of Innovative Technologies ; Incremental Contribution of Linguistic, Behavioral, and Propagation Signals to Early Detection of LLM-Generated Coordinated Inauthentic Campaigns. Spectrum of R. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Source Monitoring Signal Effectiveness and Lead Time for Performance Degradation Prediction in Kubernetes-Based Microservices. Vol. 6(4), pp. 15-26, April 2026 ; Intelligent optimization algorithm for chain restaurant spatial layout based on generative adversarial networks. Journal of Industrial Engineering and Applied S; Financing Date Uncertainty and Liquidity Forecasts for Utility-Scale Solar Projects; Comparative Empirical Evaluation of Monte Carlo Dropout Uncertainty Measures for Failure Capture and Referral Trade-offs in Chest X-ray Disease Detection. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Driven Stress Testing for Small and Medium Financial Institutions: A Variational Autoencoder Approach with Extreme Value Theory for Macroeconomic Scenario Gener; Based Feedback Agent Move the Needle? An Empirical Study of Student Correctness and Hint Dependency on the ASSISTments 2017 Interaction Logs; An Empirical Analysis of the Temperature Hyperparameter in Knowledge Distillation for Image Classification Student Networks. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Comparative Empirical Evaluation of Prompting Strategies for Large-Language-Model Web-Navigation Agents on the WebArena Benchmark; Evaluating Prompt Engineering Strategies for Few-Shot Cyber Threat Intelligence Entity and Relation Extraction from Multi-Source Reports. Journal of Science, In; A Comparative Analysis of Attention Mechanisms in Vision Transformers for Fine-Grained Image Recognition; Enhanced Risk Identification and Intelligence Sharing Framework for Anti-Money Laundering in Cross-Border Income Swap Transactions. Journal of Advanced Computin. 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 Reinforcement Learning for Pattern Recognition in Cross-Border Financial Transaction Anomalies: A Behavioral Economics Approach to AML. Applied and Computationa; Credit risk transmission mechanism and prevention strategies in supply chain finance: A core enterprise perspective. Artificial Intelligence and Machine Learnin; sensitive anomaly detection: Culturally-aware AI models for enhanced AML in US securities trading. 2025 International Conference on Computer, AI, Systems and Au; The Impact of Government Budget Data Visualization on Public Financial Literacy and Civic Engagement. Journal of Economic Theory and Business Management 2 (4), . This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation with a concrete application setting. The titles indicate work on Academia Nexus Journal, 2024 - academianexusjournal.com ◦ Research on Test Data Quality Assessment and Outlier Processing Methods in Semiconductor Chip Manufactu; driven quality assessment and investment

risk identification for carbon credit projects in developing countries. 2025 International Forum on Computing and Intel; Behavioural Feature Analysis for Anomalous Click Detection in Mobile Advertising Environments: Toward In-App Browser-Specific Detection. COMPUTING INNOVATIONS A. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation cannot be optimized in isolation. The associated topics include Factor Attribution in Utility-Scale Power-Project Returns: A SHAP-Based Comparative Study of Ensemble Learners for Unlevered-IRR Distribution Prediction; Comparative Analysis of Deep Learning Algorithms for Disease-Related Protein Function Prediction: Performance Optimization and Computational Efficiency Evaluati; A reinforcement learning approach for adaptive budget allocation in pharmaceutical digital marketing: Maximizing roi across patient journey touchpoints. Journal; Comparative Evaluation of Graph Neural Networks for Cross-Market Risk Contagion Path Identification in Multi-Layer Financial Networks. 2026 2nd International Co. 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 clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation. Their titles cover Comparing Prompting Strategies and Backbone Choices for LLM Agents on Tabular Credit-Card Fraud Detection: An Empirical Study. Journal of Innovative Technologie; 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, spanning Research on Driving Behavior Risk Identification and Safety Assessment Methods Based on Artificial Intelligence; Comparing Unsupervised Spatial Grouping and Supervised Risk Ranking for County-Level Measles Susceptibility Stratification under Uneven Childhood Immunization C; Accelerating

clinical trial recruitment through automated eligibility screening with multi-modal deep learning. Journal of Computing Innovations and Application; Spectrum of Research, 2025 - spectrumofresearch.com ◦ Machine Learning-Based Identification of Anomalous Trading Behavior Patterns Among Asia-Pacific Investors i. This cluster is relevant to clinical evidence integration, eligibility reasoning, uncertainty, privacy, and patient-centered validation 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, 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 Clinical evidence intelligence, where evidence, resources, and user impact must remain visible throughout the lifecycle.

5 Conclusion

This paper presented an evidence-centered clinical intelligence architecture for drug discovery, trial recruitment, and personalized care. The design aligns multimodal records with eligibility logic, calibrated prediction, privacy, and patient-centered review.

A real implementation should begin with one clinical question, a fixed evidence contract, and a prospective evaluation plan. The same controls can then support additional molecular, imaging, and longitudinal models without weakening traceability.

References

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

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

[3] 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

[4] 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

[5] 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)

[6] 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

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

[8] 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)

[9] 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

[10] 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

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

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

[13] 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

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

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

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

[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] 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

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

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

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

[22] 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

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

[24] 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.

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

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

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

[28] 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

[29] 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.

[30] 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)

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

[32] 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

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

[34] 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

[35] 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

[36] 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

[37] 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

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

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

[40] 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)

[41] 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

[42] 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 …

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

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

[45] 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

[46] 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)

[47] 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

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

[49] 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

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

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

[52] Ge, L., & Rao, G. (2025, July). MultiStream-FinBERT: A Hybrid Deep Learning Framework for Corporate Financial Distress Prediction Integrating Accounting Metrics, Market Signals, and Textual

Disclosures. In 2025 International Forum on Computing and Intelligent Systems (pp. 107-122). Pinnacle Academic Press.

[53] 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

[54] 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).

[55] Ju, C., & Rao, G. (2025, May). Analyzing foreign investment patterns in the US semiconductor value chain using AI-enabled analytics: A framework for economic security. In 2025 International Conference on Computer, AI, Systems and Automation (pp. 60-74). Pinnacle Academic Press.

[56] 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

[57] 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

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

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

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

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

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

[63] 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

[64] 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.

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

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

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

[68] 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

[69] 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

[70] 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

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

[72] 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

[73] 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

[74] 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.

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

[76] 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

[77] Rao, G., Trinh, T. K., Chen, Y., Shu, M., & Zheng, S. (2024). Jump prediction in systemically important financial institutions' CDS prices. Spectrum of Research, 4(2).

[78] 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 …

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

[80] 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

[81] 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

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

[83] 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

[84] 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

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

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

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

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

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

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

[91] 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

[92] 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

[93] 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

[94] 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 Crisis. Journal of Computing Innovations and Applications, 2(1), 46-61.

[95] Kang, A., & Ma, X. (2025, July). AI-Based Pattern Recognition and Characteristic Analysis of Cross-Border Money Laundering Behaviors in Digital Currency Transactions. In 2025 International Conference on Digital Intelligence and Computing Technologies (pp. 1-19). Pinnacle Academic Press.

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

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

[98] 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

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

[100] 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

[101] 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

[102] 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

[103] 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

[104] 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

[105] 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

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

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

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

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

[111] 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

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

[113] 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

[114] 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

[115] 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.

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

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

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

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

[120] 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)

[121] 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

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

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

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

[125] Kang, A., Zhang, K., & Chen, Y. (2025, July). AI-Assisted Analysis of Policy Communication during Economic Crises: Correlations with Market Confidence and Recovery Outcomes. In 2025 International Forum on Computing and Intelligent Systems (pp. 159-173). Pinnacle Academic Press.

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

[127] 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

[128] 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)

[129] Rao, G., 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), 3(4), 361-371.

[130] 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

[131] 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)

[132] 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)

[133] 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

[134] 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

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

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

[137] 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

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

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

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

[141] 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

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

[143] 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.

[144] Rao, G., 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(1), 116-127.

[145] 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

[146] 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

[147] 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.

[148] 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

[149] 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

[150] 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

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

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

[153] 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

[154] 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 …

[155] 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)

[156] 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