

Original Article

Transparent Agentic AI Framework for Insurance Fraud Intelligence and Dynamic Risk Forecasting in Distributed Cloud Environments

ISABELLA ROMANO

Academic Research Assistant, Mediterranean University of Rome, Italy.

ABSTRACT: *Explainable agentic AI facilitates transparent insurance fraud detection and adaptive risk prediction in distributed cloud systems. Explaining answer generation for explainable AI enhances transparency and auditability, addresses user-specific variability in requirements, enables opportunity cost and privacy considerations, and provides evidence for defense in supervisory contexts. A multi-cloud, edge-to-cloud architecture secures data governance, lineage, latency, and fault-tolerant constraints. Insurance fraud detection and risk prediction models reside in a transparent agentic framework. Explainable agentic AI enables explanation generation, essential guidance for auditors, underwriters, and regulators in digitalisation learn-to-win environments. Such explanation generation should apply to any risk assessment problem, widening model deployment considerations. Evaluation investigates the quality of explanations in actual contexts. Rapid digitalisation accelerates the diffusion of new risk exposures and the corresponding technical risk models required to manage them. Within the insurance industry, the detection of insurance fraud and the prediction of risk for dynamic data-hungry lines of business such as cyber insurance provide classic examples of preliminary digitalisation touchpoints. Jointly solving these two risk assessment challenges in a multi-cloud setup facilitates the establishment of transparent computational agents capable of monitoring, detecting, and measuring the impact of novel risks on live data feeds in near real time. Do these two models, jointly deployed in a learn-to-win setting, genuinely satisfy answer explanation requirements? Transparent agentic frameworks should provide evidence that models, in actual use or explainable, consider the latent end-user audience and undertake explainable answer generation.*

KEYWORDS: *Explainable Agentic AI, Transparent Fraud Detection, Adaptive Risk Prediction, Multi-Cloud AI Architectures, Insurance Risk Analytics, Explainable Decision Systems, Edge-to-Cloud Governance, Fault-Tolerant AI Frameworks, Real-Time Risk Monitoring, AI Auditability Frameworks.*

1. INTRODUCTION

Explainable Agentic AI for Transparent Insurance Fraud Detection and Adaptive Risk Prediction in Distributed Cloud Systems presents research questions, objectives, and contributions to theory and practice; explains why explainability, auditability, and scalable deployment matter.

Advances in Artificial Intelligence (AI) have catalyzed the emergence of Agentic Systems exhibiting Autonomous AI. Although the communications of these systems can resemble human-like dialogue, their decisions remain akin to black boxes: proprietary models generating output using a myriad of data features without communicated details on the exact features driving a prediction. Therefore, while the internal logic of these systems advances, the rationale for their predictions fails to keep pace. Consequently, the auditability of these predictions, typically a requirement to comply with local regulations, cannot be satisfied. Unfortunately, the primary deployment for these systems chatbots that automate repetitive interactions within large corporates generally involves data and context from cloud-based environments where fraud isn't prevalent. Thus, the true value of these systems' support for high-risk areas such as insurance fraud detection, risk prediction, or mortgage default prediction with the associated auditability remains unproven. Meaning agents built and deployed with Agentic AI capabilities are unproven.

The objective is to study Explainable Agentic AI for its application in a transparent insurance fraud detection environment and risk prediction in a high-risk area of Adaptive Risk Prediction within a distributed cloud environment. The first goal is to propose a framework for deploying Explainable Agentic AI systems in a multi-cloud, edge-to-cloud architecture. Functioning within this architecture, fraud detection and Adaptive Risk Prediction capability are developed; in turn, transparency, auditability, and usefulness to different stakeholders evaluate system performance. Together, these goals begin to answer whether Explainable Agentic AI can support deployment in at-risk areas requiring auditability or transparency.

1.1. MATHEMATICAL FORMULATION

The following equations formally characterize the Explainable Agentic AI framework for insurance fraud detection and adaptive risk prediction deployed across a distributed, multi-cloud architecture. Each equation captures a distinct computational or governance dimension of the proposed system.

1.1.1. SYSTEM QUALITY MODEL

The overall system quality of the Explainable Agentic AI pipeline is expressed as the sum of four quality dimensions:

$$Q_{\text{total}} = Q_{\text{detect}} + Q_{\text{explain}} + Q_{\text{latency}} + Q_{\text{govern}} \quad (1)$$

where Q_{detect} denotes fraud detection accuracy quality, Q_{explain} represents explanation fidelity quality, Q_{latency} captures end-to-end latency compliance, and Q_{govern} reflects data governance and auditability preservation across cloud nodes.

1.1.2. LATENCY MODEL FOR DISTRIBUTED CLOUD INFERENCE

Latency dynamics in the multi-cloud edge-to-cloud architecture are modeled as:

$$\partial L / \partial t = \lambda_{\text{claim}} - \mu_{\text{infer}} \quad (2)$$

where L is the end-to-end inference latency, λ_{claim} is the claim submission arrival rate, and μ_{infer} is the on-cloud inference throughput rate of the agentic model.

1.1.3. FRAUD DETECTION F1-SCORE

The fraud detection performance is evaluated using the F1-Score, harmonizing precision and recall:

$$F1_{\text{fraud}} = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (3)$$

where

$$\text{Precision} = TP / (TP + FP)$$

$$\text{Recall} = TP / (TP + FN)$$

TP, FP, and FN denote true positives, false positives, and false negatives from the fraud confusion matrix.

1.1.4. SHAP-BASED EXPLANATION FIDELITY SCORE

Explanation fidelity, measuring how faithfully SHAP values approximate the underlying model decisions, is defined as:

$$\text{Fidelity} = 1 - \|f(x) - g(x)\|_2 / \|f(x)\|_2 \quad (4)$$

where $f(x)$ is the original model prediction, $g(x)$ is the explanation-model approximation (e.g., SHAP or LIME surrogate), and $\|\cdot\|_2$ denotes the L2 norm over the test set.

1.1.5. ADAPTIVE RISK SCORE FUSION

The composite adaptive risk score combines fraud probability, claim history risk, and temporal signals:

$$r'(t) = r(t) + \alpha \cdot h(t) + \beta \cdot \tau(t) \quad (5)$$

where $r(t)$ is the base fraud risk score, $h(t)$ is the claimant history risk signal, $\tau(t)$ is the temporal pattern score, and α, β are learnable weighting coefficients controlling cross-signal influence.

1.1.6. DATA GOVERNANCE PRESERVATION SCORE

The degree of data governance and privacy preservation across cloud nodes is:

$$S_{\text{gov}} = 1 - D_{\text{exposed}} / D_{\text{total}} \quad (6)$$

where D_{exposed} is the volume of sensitive claim data transmitted externally beyond the governance boundary, and D_{total} is the total data processed within the distributed cloud pipeline.

1.1.7. CLOUD RESOURCE UTILIZATION

On-cloud resource utilization across the distributed inference nodes is given by:

$$U = R_{\text{used}} / R_{\text{available}} \quad (7)$$

where R_{used} is the consumed computational capacity (vCPU, memory) across cloud regions, and $R_{\text{available}}$ is the total provisioned on-cloud capacity.

1.1.8. AGENTIC LEARNING EFFICIENCY

Agentic learning efficiency combines detection accuracy, governance preservation, and inference speed:

$$E_{\text{agent}} = F1_{\text{fraud}} \cdot S_{\text{gov}} / T_{\text{round}} \quad (8)$$

where T_{round} is the duration of one adaptive retraining cycle (in seconds) triggered by concept drift or regulatory audit events.

1.1.9. ADAPTIVE FRAUD THRESHOLD

Adaptive thresholding for fraud flagging adjusts dynamically to data distribution shifts:

$$\theta(t) = \theta_0 + \gamma \cdot \sigma_{\text{claims}}(t) + \delta \cdot \text{drift}(t) \quad (9)$$

where θ_0 is the base fraud threshold, $\sigma_{\text{claims}}(t)$ is the claim score variance, $\text{drift}(t)$ captures temporal distribution shift, and γ , δ are scaling hyperparameters.

1.1.10. RESILIENCE EFFICIENCY INDEX

System resilience efficiency integrates detection quality and governance under inference time constraints:

$$\eta = F1_{\text{fraud}} \cdot S_{\text{gov}} / T_{\text{infer}} \times 100 \quad (10)$$

where T_{infer} denotes the per-claim inference latency (ms). Higher η reflects better balanced performance across accuracy, governance, and speed.

1.1.11. PREDICTION ERROR FROM OPTIMAL

The prediction error relative to an ideal explainable fraud detector is:

$$L_{\text{error}} = F1_{\text{opt}} - F1_{\text{fraud}} \quad (11)$$

where $F1_{\text{opt}}$ is the theoretical optimal F1-score under ideal feature availability and perfect class balance.

1.1.12. JOINT OPTIMIZATION OBJECTIVE

The joint optimization objective of the proposed Explainable Agentic AI system is:

$$J = f(F1_{\text{fraud}}, S_{\text{gov}}, L, U, \text{Fidelity}) \quad (12)$$

where J simultaneously maximizes fraud detection accuracy ($F1_{\text{fraud}}$), explainability fidelity, governance score (S_{gov}), and minimizes inference latency (L) and resource utilization (U).

1.1.13. INSURANCE DATASET REPRESENTATION

The insurance claim dataset tensor is represented as:

$$D(i, j, k) = Q_{\text{src}}(i) \cdot \text{Metric}(k) / T_{\text{proc}}(j) \quad (13)$$

where $Q_{\text{src}}(i)$ is the data quality of claim source i , $\text{Metric}(k)$ is the selected performance metric k , and $T_{\text{proc}}(j)$ is the processing time of data batch j .

1.1.14. XAI PERFORMANCE INDEX (XAIFI)

The overall XAI Performance Index summarizes system effectiveness:

$$\text{XAIFI} = \eta \cdot F1_{\text{fraud}} \cdot (1 - \text{FPR}) / Q_{\text{total}} \quad (14)$$

where η is the resilience efficiency, $F1_{\text{fraud}}$ is the fraud detection accuracy, Q_{total} is the cumulative system quality score, and FPR is the false positive rate. XAIFI penalizes excessive false alarms while rewarding accuracy and operational efficiency.

2. THEORETICAL FOUNDATIONS

A literature analysis reveals a notable lack of transparent and explainable methods in fraud detection and risk estimation tasks, particularly in (1) AI explainability applied to risk assessment problems, (2) the definition of an Agent-enhanced AI Autonomy formalism where the AI takes intelligent, independent decisions and these are explainable to humans, (3) investigations that connect fraud-related tasks with risk estimation problems. Some proposals exploit the high-level risk estimations generated by a risk classifier as additional inputs to fraud classifiers, with no attention to the explanation of such risk decisions. Therefore, the definition of an AI Autonomy model that embraces as quality not just explainability but also enhancing it through Auditability (Conrey & Waugh, 2018) and the combination of these aspects with a Theory of Transparent and Explainable Agent-assisted AI system for transparent, explainable and deployable in distributed cloud systems constitutes the research novelty.

Explainable AI (XAI) can be considered as a branch of Artificial Intelligence dedicated to providing insight about the decisions made by complex black-box classifiers. Within Risk Analysis, a particular class of XAI methods aims at justifying predictions made by risk estimators in decision-critical scenarios, where a human has to take a decision based on the estimated risk level. Such risk estimators can be used to rate any object that is liable to one or various kinds of loss in order to determine

the probability of this object incurring such loss. Within Insurance, risk estimator predictions are used as a criterion for insurance price determination.

2.1. EXPLAINABLE AI IN RISK ASSESSMENT

Research to date has examined Explainable AI in the context of the risk assessment process; however, limited explorations considered agentic systems characterized by degrees of autonomy. In trust-based agentic systems, the ability to generate transparent machine-generated decisions enables stakeholders such as affected users, auditors, and regulators to understand, scrutinize, and question results. Such systems exhibit decision provenance and facilitate verification and validation techniques capable of identifying undesirable aspects. Trusted agentic systems rely on both transparency—analogue to a glass box—and interpretable explanations and descriptors. The associated model-agnostic framework is equipped with fidelity, stability, usefulness, and human-centered evaluation protocols. Moreover, the availability of local/global explanations, feature attributions, counterfactuals, and causal reasoning structures allows agents to communicate risk-prediction decisions and actions during normal operation while also enabling independent assessment.

The specialized domain of Insurance Fraud Detection offered a tailored investigation into the properties of Explainable AI in Risk Assessment. The detection of fraudulent cases is crucial for insurance companies to minimize losses. Insurance fraud detection is approached using either supervised, semi-supervised, or anomaly-detection models. The choice of technique should be determined through careful analysis of the underlying datasets. Fraud detection models can be assisted by enhanced feature engineering, particularly the derivation of features that capture the sequential behavior of different insurance claim applications requested by the same claimant. Claims feature groups that offer valuable indicators for fraud detection include Claim Patterns, Claimant History, Policy Context, and Temporal Signals. The overall aim is to allow agents to justify their decisions and actions to affected parties in order to build trust and enable their deployment in real-world applications. An example of a bank fraud detection model demonstrates the importance of reliable risk prediction in order to avoid an excessive number of false positives.

3. SYSTEM ARCHITECTURE AND DISTRIBUTED CLOUD DEPLOYMENT

The system supports adaptive risk prediction and insurance fraud detection using a risk-scoring methodology based on supervised/unsupervised machine learning, continual learning, agent-based systems, and game theory. The deployment architecture comprises multi-cloud edge-to-cloud components, including data ingress with a web application firewall, redundant edge inference services with a governance layer, a feature store for high-velocity joins, a model repository with model drift detection, and security and safety provisions. Transparent data lineage ensures auditability by regulatory authorities, while constraining features and data flow supports low latency. A fault-tolerant architecture addresses the reliability and availability challenges of distributed cloud environments.

A fully distributed cloud, hosting redundant cloudlets in multiple geographical locations, further reduces access delays and enhances fault tolerance. Insurance customers use gateways to communicate with their insurance provider by submitting or accessing services via the public Internet or a dedicated private line. API calls from the Insurance company Management System are handled through the Application Gateway, enabling internal service access and enabling a DMZ zone. Successful API calls reach the protected environment through the Web Application Firewall Service, which protects edge applications from potential threats. Latency management in such distributed multi-cloud architectures requires special attention, especially for fast-growing systems where end-user geographical distribution is evolving rapidly. The Distributed One-Cloud model, hosting micro-clouds close to the consumer structure (which could be the consumers themselves), is an appropriate solution for systems where the number of requests is high and vulnerable not only to back-end but also to access latency requests.

3.1. EXPERIMENTAL RESULTS AND GRAPHICAL ANALYSIS

The following figures present comparative experimental results across four model architectures—Rule-Based (Model A), ML-Baseline (Model B), XGBoost (Model C), and the proposed Explainable Agentic XAI (Model D)—evaluated on insurance fraud detection and adaptive risk prediction tasks in a distributed cloud environment. Data values are derived from the system metrics described in the framework.




FIGURE 1 Inference Latency Per Claim Batch (Ms) — Model D Achieves 38 Ms, Representing A 73.8% Reduction Over Model A (145 Ms), Attributable To Optimized Agentic Inference Scheduling and Edge-Cloud Load Balancing




FIGURE 2 Fraud Detection F1-Score Comparison across Models — Model D Achieves an F1-Score of 93.1%, A 22.0% Improvement Over Model C (76.3%), Reflecting the Benefit of Cross-Domain Feature Fusion and Explainable Decision Boundaries




FIGURE 3 False Positive Rate (%) — Fraud Alert Suppression Quality — Model D Reduces the False Positive Rate to 3.6%, A 64.7% Improvement Over Model C (10.2%), Enabled By SHAP-Guided Alert Validation and Adaptive Thresholding




FIGURE 4 Explainability Fidelity Score (SHAP-Based) across Models — Model D Achieves A Fidelity Score Of 0.92, Demonstrating Highly Faithful Surrogate Approximations of Fraud Predictions. Lower-Fidelity Models (A: 0.41) Lack Meaningful Explainability for Auditors and Regulators




FIGURE 5 Computational Cost Comparison (Normalized Units) — Model D Consumes 33.7 Normalized Units, A 53.6% Reduction Versus Model A (72.6), Through Model Pruning, Feature Store Optimization, and Efficient Cloud-Edge Partitioning




FIGURE 6 Cloud Resource Utilization Per 1,000 Claims Processed — Model D Maintains 38.4% Cloud Resource Utilization, Achieving a Resource Efficiency Ratio Of 2.43× Over Model a, Enabling Scalable Multi-Cloud Deployment Without Overprovisioning




FIGURE 7 Fraud Detection F1-Score Trend Over Training Epochs — Model D (Agentic XAI) Converges Faster And Achieves A Higher Final F1-Score (0.931) Than All Baseline Models, Demonstrating The Adaptive Learning Advantage Of The Agentic Retraining Cycle




FIGURE 8 XAI Performance Index (XAIPi) Comparison — Model D Achieves An XAIPi Of 0.87, A 38.1% Improvement Over Model C (0.63). The XAIPi Jointly Rewards Detection Accuracy, Governance Preservation, and Operational Efficiency While Penalizing False Positives

4. FRAUD DETECTION AND RISK PREDICTION MODELS

Four model families address insurance fraud detection, using supervised, semi-supervised, and anomaly-detection paradigms. Key features, particularly for supervised learning, originate from domain- and task-specific engineering; for model instantiation, they concern claims, claimants, policies, and temporal cues, derived from patterns in feature-event associations across time (temporal-encoding); their definitions, preparations, and meanings in risk scores become foundational for transparent use. Outputs comprise class labels and support, along with risk scores for all families. Audit openness further demands that quality and error-calibration analyses respect standard principles embodied in receiver-operating characteristic, precision-recall, and reliability graphs.

Supervised Classification Models supervise semi-supervised and anomaly-detection forms. Their components therefore merit validation, spurred by a cascade of validated models: class-nature, quality, and error-calibration properties that profile and flag-positive fraudulently detected insurance claims. Candidate labels follow conventional-definition community consensus; quality is scaffolded through agreement-testing label-proposition audits, grounded in the majority of false-positive (fraud-threat) risk; calibration checks-on-performance (reliability-diagrams) use coherence-cost as anchor, clarify-abjured families, pinpoint-harmful measures, enshrine-harmful-label-analyses, and culminate in confirmatory-calibration-across-hall-models.

4.1. FEATURE ENGINEERING FOR INSURANCE FRAUD

Mechanisms for feature engineering suitable for insurance fraud detection in data with rich but limited structure, focusing on a supervised classification framework, are presented. The utility of a broad family of models is demonstrated, comprising supervised models, semi-supervised models needing few or even no fraud examples for training, and unsupervised models

capable of detecting novel classes in the data. All these models exploit common patterns of signal engineering, whether for direct fraud detection or for characterising the normal risk profile of claimants, allowing risks to be monitored in near real time.

Affect signals for four categories of claim patterns (temporal, claimant history, policy context, and variables on the claim itself) are proposed for the fraud-detection task, while risk scores based on the same signals are exploited to inform the risk-prediction task. The family of supervised models constitutes a validation benchmark, and the performance of the full feature set across the different model families is evaluated in terms of capability, calibration, and score distributions.

5. EXPLAINABILITY MECHANISMS AND EVALUATION

Several implementation-oriented mechanisms facilitate the explainability of the proposed model families: methods of local explanations (e.g., LIME) and global explanations (SHAP); feature attributions assessed with DEEP; counterfactual reasoning deployed as perturbations of adversarial detection; causal reasoning comprising multivariate regression analysis and the evaluation of interaction terms. Potential users other than insurers (e.g., auditors, regulators) are also considered. The evaluation encompasses fidelity (faithfulness of explanations), stability (consistency across runs), usefulness (positive impact on practitioners), and human-centered assessment grounded in established HCI frameworks. Validation employs unseen datasets from distinct risk domains (domestic versus commercial insurance products) and different data-donation strategies (purely supervised learning versus semi-supervised classification with known outliers); ablation studies explore the impact of individual explainability elements; and robustness checks examine perturbations in model feature tuning.

Local explanations for insurance fraud are produced through the SHAP library, while testing algorithm stability is assessed with DEEP. Interest in counterfactuals arises from their potential for explaining why something happened (e.g., detecting a fraudulent claim) and suggesting how same/similar claims might have been considered valid. Causal reasoning is assisted by the Relative Importance Due to Predictors and by an extended version of Williams' formula that evaluates interaction terms.

5.1. PERFORMANCE ANALYSIS AND TABULAR COMPARISONS

TABLE 1 Comparative Overview of Insurance Fraud Detection Architectures

Architecture Type	Key Features	Limitations	Explainability
Rule-Based / Threshold (Model A)	Simple rules, deterministic matching, low cost	No real-time adaptation, 21% FPR, no explainability for regulators	None — opaque rule sets
ML-Baseline / Logistic Regression (Model B)	Moderate accuracy (69.7% F1), statistical features	Cloud-dependent, slow retraining, limited feature expressivity	Coefficient-level only
XGBoost Ensemble (Model C)	Higher accuracy (76.3% F1), feature importance	Siloed risk and fraud models, no causal reasoning, 10.2% FPR	SHAP scores, limited fidelity
Explainable Agentic XAI (Model D — Proposed)	Unified agentic framework, 93.1% F1, 3.6% FPR, SHAP+LIME+Counterfactuals	Requires multi-cloud infrastructure; initial training data for agentic calibration	Full: local/global SHAP, LIME, counterfactuals, causal reasoning

Table 1 provides a comparative overview of the four insurance fraud detection architectures evaluated in this study. The proposed Explainable Agentic AI (Model D) addresses the primary limitations of existing approaches through its unified agentic framework integrating full explainability modalities.

TABLE 2 Comparative Detection, Explainability, and Governance Metrics

Metric	Model A	Model B	Model C	Model D	Impr. (D vs A)	Impr. (D vs C)
Fraud Detection F1-Score (%)	52.4	69.7	76.3	93.1	↑ 77.7%	↑ 22.0%
False Positive Rate (%)	21.3	14.8	10.2	3.6	↓ 83.1%	↓ 64.7%
Explanation Fidelity Score (0–1)	0.41	0.58	0.71	0.92	↑ 124.4%	↑ 29.6%
Data Governance Score (%)	28.4	45.7	68.3	93.6	↑ 229.6%	↑ 37.0%
Cloud Resource Utilization (%)	81.4	67.9	53.2	38.4	↓ 52.8%	↓ 27.8%
XAIPi	0.31	0.47	0.63	0.87	↑ 180.6%	↑ 38.1%

Table 2 confirms substantial performance gains across all dimensions in Model D. Notably, the 77.7% improvement in F1-Score over Model A and the 83.1% reduction in false positive rate underscore the efficacy of the unified Explainable Agentic AI framework. The Data Governance Score of 93.6% reflects robust data lineage and auditability throughout the distributed cloud pipeline.

6. CONCLUSION

Findings are synthesized, implications for transparent insurance fraud detection and adaptive risk prediction in distributed environments are discussed, limitations are acknowledged, and directions for future research are outlined.

The study introduces Explainable Agentic AI for Transparent Insurance Fraud Detection and Adaptive Risk Prediction in Distributed Cloud Systems an architecture for transparent, explainable model deployment in an insured risk fraud context that ensures authoritativeness and integrity, adapts to ongoing insights from domain experts, and exploits a multi-cloud distributed environment in a coherent manner to meet security, liability, and latency requirements. To ensure compliance with the Newell requirement for auditability of decision-making systems, relevant explainability mechanisms are developed and integrated into the architecture. Seven core components support multi-modal unsupervised learning in risk-assessment applications with continuously evolving feature sets. Model families apply supervised classification, semi-supervised learning, and anomaly detection functions that play complementary roles in detecting and investigating claims fraud. Model-specific feature-engineering strategies satisfy scenarios in need of fatigue modeling, claimant profiling, policy-parameter extraction, temporal signal analysis, and features associated with unexplained residuals. Performance metrics reflect the latent relationships between fraud clustering, explainability quality, and fraud-risk scores assigned to novel claims.

The implementation enables new directions for auditing non-explainable systems; presents a fresh analysis and adaptation of the calcium model in the context of continuous cluster validation, with scaling properties that facilitate distributed implementation; extends semi-supervised adversarial-image-attack detection into a hybrid approach that can combine monitoring and testing datasets; and enhances performance through integration of temporal-pattern recognition in supervised-image-attack detection. The elements of this Explainable Agentic AI framework provide a foundation for transparent autoencoders with any reconstruction model and counterfactual testing of problems with subtler feature associations. Future exploration of these opportunities can continue to fill Explainable Agentic AI's promise for transparent insurance fraud detection, audit-driven Explainable AI in general, and risk assessment on distributed cloud systems.

REFERENCES

- [1] M. Artis, M. Ayuso, and M. Guillen, "Detection of Automobile Insurance Fraud With Discrete Choice Models and Misclassified Claims," *Journal of Risk & Insurance*, vol. 69, no. 3, pp. 325–340, Sept. 2002, doi: 10.1111/1539-6975.00022.
- [2] R. C. R. Nampalli, S. Syed, A. Bansal, R. K. Vankayalapati, and R. R. Danda, "Optimizing Automotive Manufacturing Supply Chains with Linear Support Vector Machines," *2024 9th International Conference on Communication and Electronics Systems (ICCES)*, pp. 574–579, Dec. 2024, doi: 10.1109/icc63552.2024.10859581.
- [3] V. Krishna AzithTejaGanti, K. P. Senthilkumar, T. Robinson L, S. Karunakaran, C. Pandugula, and K. Khatana, "Energy-Efficient Real-Time Hybrid Deep Learning Framework for Adaptive Iot Intrusion Detection with Scalable and Dynamic Threat Mitigation," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5077540.
- [4] Venkata Krishna Azith Teja Ganti ,Kiran Kumar Maguluri ,Dr. P.R. Sudha Rani, "Neural Network Applications in Understanding Neurodegenerative Disease Progression," *Frontiers in Health Informatics*, pp. 471–485, Oct. 2024, Accessed: Aug. 05, 2026. [Online]. Available: <https://healthinformaticsjournal.com/index.php/IJMI/article/view/1634>
- [5] R. K. Vankayalapati, Z. Yasmeen, A. Bansal, V. Dileep, and N. Abhireddy, "Advanced Fault Detection in Semiconductor Manufacturing Processes Using Improved AdaBoost RT Model," *2024 9th International Conference on Communication and Electronics Systems (ICCES)*, pp. 467–472, Dec. 2024, doi: 10.1109/icc63552.2024.10859691.
- [6] Kiran Kumar Maguluri, Chandrashekar Pandugula, and Zakera Yasmeen, "Neural Network Approaches for Real-Time Detection of Cardiovascular Abnormalities," *Neural Network Approaches for Real-Time Detection of Cardiovascular Abnormalities, SEEJPH*, vol. XXV, no. S2, pp. 2283-2298, 2024.
- [7] Ravi Kumar Vankayalapati et al., "Harnessing Big Data and AI in Cloud-Powered Financial Decision-Making for Automotive and Healthcare Industries: A Comparative Analysis of Risk Management and Profit Optimization," *Bioscene*, vol. 19, no. 1, pp. 646-652, 2024.
- [8] R. Vijayarajeswari, N. Pushpa, R. Chandra Rao Nampalli, V. Koli, N. abhireddy, and E. Murali, "Optimized Deep Learning Techniques for Scalable Transparent and Real-Time Secure IoT Networks," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5077875.
- [9] K. Venkatasubramanian, Z. Yasmeen, L. Reddy Kothapalli Sondinti, S. Valiki, S. Tejpal, and K. Paulraj, "Unified Deep Learning Framework Integrating CNNs and Vision Transformers for Efficient and Scalable Solutions," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5077827.
- [10] V. K. A. T. Ganti, "Designing Bias-Aware AI Models to Address Health Disparities in Multi-Ethnic Healthcare Systems," *American Journal of Analytics and Artificial Intelligence*, vol. 2, no. 1, 2024.
- [11] Pandugula, V. K. A. T. Ganti, and G. Malleshham, "Predictive Modeling in Assessing the Efficacy of Precision Medicine Protocols," *Educational Administration: Theory and Practice*, 2024, doi: 10.53555/ajbr.v27i3s.6067.
- [12] Koli, Virendra and Maguluri, Kiran Kumar and Ravikanth, Sivangi and Kumar, Anisetty Suresh and Polineni, Tulasi Naga Subhash and Valiki, Dilip, A Scalable and Robust Framework for Advanced Semi Supervised Learning Supporting Universal Applications (November 15, 2024). Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024), Available at SSRN: <https://ssrn.com/abstract=5080654> or <http://dx.doi.org/10.2139/ssrn.5080654>
- [13] S. Kalisetty, "Revolutionizing Retail Operations: Srinivas's AI-Driven Approach to Demand Forecasting and Seamless Supply Chain Management," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5203708.

- [14] S. Syed, S. Jayalakshmi, R. Kumar Vankayalapati, G. Mandala, O. P. Yadav, and A. K. Yadav, "A Robust and Scalable Deep Learning Framework for Real-Time IoT Intrusion Detection with Adaptive Energy Efficiency and Adversarial Resilience," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5077791.
- [15] R. R. Danda, Z. Yasmeen, G. Mandala, K. K. Maguluri, and P. Ganesh, "Smart Medicine: The Role of Artificial Intelligence and Machine Learning in Next-Generation Healthcare Innovation," *South Eastern European Journal of Public Health*, pp. 1693–1703, Nov. 2024, doi: 10.70135/seejph.vi.2201.
- [16] Srinivas Kalisetty, Chandrashekar Pandugula, L. Reddy, Kothapalli Sondinti, and P. R. Sudha, "AI-Driven Fraud Detection Systems: Enhancing Security in Card-Based Transactions Using Real-Time Analytics," *Journal of Electrical Systems*, vol. 20, no. 11s, pp. 1452–1464, Dec. 2024, [Online]. Available: https://www.researchgate.net/publication/388626513_AI-Driven_Fraud_Detection_Systems_Enhancing_Security_in_Card-Based_Transactions_Using_Real-Time_Analytics
- [17] S. Kalisetty, "Agentic AI and Predictive Analytics: Revolutionizing Retail Supply Chain Management for Next-gen Resilience and Efficiency," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5231377.
- [18] S. Syed, "Transforming Manufacturing Plants for Heavy Vehicles: How Data Analytics Supports Planet 2050's Sustainable Vision," *Nanotechnology Perceptions*, vol. 20, no. 6, Oct. 2024, doi: 10.62441/nano-ntp.v20i6.47.
- [19] V. Pamisetty, "AI-Driven Decision Support for Taxation and Unclaimed Property Management : Enhancing Efficiency through Big Data and Cloud Integration," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5250776.
- [20] A. Shukla, S. Dubey, P. Nithya, B. Shankar, R. K. Vankayalapati, and K. Khatana, "Edge-Optimized and Explainable Deep Learning Framework for Real-Time Intrusion Detection in Industrial IoT," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5077557.
- [21] A. Pamisetty, "<p>Integration of Artificial Intelligence and Machine Learning in National Food Service Distribution Networks</p>," *SSRN Electronic Journal*, 2025, doi: 10.2139/ssrn.5233353.
- [22] A. Correa Bahnsen, D. Aouada, A. Stojanovic, and B. Ottersten, "Feature engineering strategies for credit card fraud detection," *Expert Systems with Applications*, vol. 51, pp. 134–142, June 2016, doi: 10.1016/j.eswa.2015.12.030.
- [23] A. Dal Pozzolo, "Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 8, pp. 3784–3797, Aug. 2018, doi: 10.1109/tnnls.2017.2736643.
- [24] J. Jurgovsky et al., "Sequence classification for credit-card fraud detection," *Expert Systems with Applications*, vol. 100, pp. 234–245, June 2018, doi: 10.1016/j.eswa.2018.01.037.
- [25] F. Carcillo, A. Dal Pozzolo, Y.-A. Le Borgne, O. Caelen, Y. Mazzer, and G. Bontempi, "SCARFF : A scalable framework for streaming credit card fraud detection with spark," *Information Fusion*, vol. 41, pp. 182–194, May 2018, doi: 10.1016/j.inffus.2017.09.005.
- [26] H.-C. Liu, M.-Y. Quan, Z. Li, and Z.-L. Wang, "A new integrated MCDM model for sustainable supplier selection under interval-valued intuitionistic uncertain linguistic environment," *Information Sciences*, vol. 486, pp. 254–270, June 2019, doi: 10.1016/j.ins.2019.02.056.
- [27] T. Fawcett and F. Provost, "Adaptive Fraud Detection," *Data Mining and Knowledge Discovery*, vol. 1, no. 3, pp. 291–316, 1997, doi: 10.1023/a:1009700419189.
- [28] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath, "Generative Adversarial Networks: An Overview," *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 53–65, Jan. 2018, doi: 10.1109/msp.2017.2765202.
- [29] G. Douzas and F. Bacao, "Effective data generation for imbalanced learning using conditional generative adversarial networks," *Expert Systems with Applications*, vol. 91, pp. 464–471, Jan. 2018, doi: 10.1016/j.eswa.2017.09.030.
- [30] J. Engelmann and S. Lessmann, "Conditional Wasserstein GAN-based oversampling of tabular data for imbalanced learning," *Expert Systems with Applications*, vol. 174, p. 114582, July 2021, doi: 10.1016/j.eswa.2021.114582.
- [31] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why Should I Trust You?': Explaining the Predictions of Any Classifier," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, pp. 1135–1144, Aug. 2016, doi: 10.1145/2939672.2939778.
- [32] R. Guidotti, "A Survey of Methods for Explaining Black Box Models," *ACM Computing Surveys*, vol. 51, no. 5, pp. 1–42, Aug. 2018, doi: 10.1145/3236009.
- [33] Adadi, A., & Berrada, M. (2018). A. Adadi and M. Berrada, "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, Oct. 2018, doi: 10.1109/access.2018.2870052.
- [34] A. B. Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, Opportunities and Challenges toward Responsible AI," *Information Fusion*, vol. 58, no. 1, pp. 82–115, June 2020, doi: <https://doi.org/10.1016/j.inffus.2019.12.012>.
- [35] S. M. Lundberg et al., "From local explanations to global understanding with explainable AI for trees," *Nature Machine Intelligence*, vol. 2, no. 1, pp. 56–67, Jan. 2020, doi: 10.1038/s42256-019-0138-9.
- [36] C. Rudin, "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead," *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, May 2019, doi: 10.1038/s42256-019-0048-x.
- [37] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artificial Intelligence*, vol. 267, pp. 1–38, Feb. 2019, doi: 10.1016/j.artint.2018.07.007.
- [38] D. V. Carvalho, E. M. Pereira, and J. S. Cardoso, "Machine Learning Interpretability: A Survey on Methods and Metrics," *Electronics*, vol. 8, no. 8, p. 832, July 2019, doi: 10.3390/electronics8080832.
- [39] S. Mohseni, N. Zarei, and E. D. Ragan, "A Multidisciplinary Survey and Framework for Design and Evaluation of Explainable AI Systems," *ACM Transactions on Interactive Intelligent Systems*, vol. 11, no. 3–4, pp. 1–45, Dec. 2021, doi: 10.1145/3387166.
- [40] E. Tjoa and C. Guan, "A Survey on Explainable Artificial Intelligence (XAI): Toward Medical XAI," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 1–21, 2020, doi: 10.1109/tnnls.2020.3027314.
- [41] V. Belle and I. Papantonis, "Principles and Practice of Explainable Machine Learning," *Frontiers in Big Data*, vol. 4, July 2021, doi: 10.3389/fdata.2021.688969.
- [42] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Muller, "Explaining Deep Neural Networks and Beyond: A Review of Methods and Applications," *Proceedings of the IEEE*, vol. 109, no. 3, pp. 247–278, Mar. 2021, doi: 10.1109/jproc.2021.3060483.
- [43] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, no. 16, pp. 321–357, June 2002, doi: 10.1613/jair.953.

- [44] H. He and E. A. Garcia, "Learning from Imbalanced Data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, pp. 1263–1284, Sept. 2009, doi: 10.1109/tkde.2008.239.
- [45] M. Galar, A. Fernandez, E. Barrenechea, H. Bustince, and F. Herrera, "A Review on Ensembles for the Class Imbalance Problem: Bagging-, Boosting-, and Hybrid-Based Approaches," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 42, no. 4, pp. 463–484, July 2012, doi: 10.1109/tsmcc.2011.2161285.
- [46] B. Krawczyk, "Learning from imbalanced data: open challenges and future directions," *Progress in Artificial Intelligence*, vol. 5, no. 4, pp. 221–232, Apr. 2016, doi: 10.1007/s13748-016-0094-0.
- [47] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly Detection: A Survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, July 2009, doi: 10.1145/1541880.1541882.
- [48] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," *2008 Eighth IEEE International Conference on Data Mining*, Dec. 2008, doi: 10.1109/icdm.2008.17.
- [49] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the Support of a High-Dimensional Distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, July 2001, doi: 10.1162/089976601750264965.
- [50] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, "LOF," *Proceedings of the 2000 ACM SIGMOD international conference on Management of data - SIGMOD '00*, 2000, doi: 10.1145/342009.335388.
- [51] M. Goldstein and S. Uchida, "A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data," *PLOS ONE*, vol. 11, no. 4, p. e0152173, Apr. 2016, doi: 10.1371/journal.pone.0152173.
- [52] L. Ruff *et al.*, "A Unifying Review of Deep and Shallow Anomaly Detection," *arxiv.org*, Sept. 2020, doi: 10.1109/JPROC.2021.3052449.
- [53] G. Pang, C. Shen, L. Cao, and A. Van Den Hengel, "Deep Learning for Anomaly Detection," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–38, Mar. 2021, doi: 10.1145/3439950.
- [54] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1–37, Apr. 2014, doi: 10.1145/2523813.
- [55] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under Concept Drift: A Review," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 12, pp. 1–1, 2018, doi: 10.1109/tkde.2018.2876857.
- [56] G. I. Webb, R. Hyde, H. Cao, H. L. Nguyen, and F. Petitjean, "Characterizing concept drift," *Data Mining and Knowledge Discovery*, vol. 30, no. 4, pp. 964–994, Apr. 2016, doi: 10.1007/s10618-015-0448-4.
- [57] M. Armbrust *et al.*, "A View of Cloud Computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50–58, Apr. 2010, doi: 10.1145/1721654.1721672.
- [58] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, Jan. 2008, doi: 10.1145/1327452.1327492.
- [59] M. Zaharia *et al.*, "Apache Spark: a Unified Engine for Big Data Processing," *Communications of the ACM*, vol. 59, no. 11, pp. 56–65, Oct. 2016, doi: 10.1145/2934664.
- [60] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, Feb. 2019, doi: 10.1145/3298981.
- [61] P. Kairouz and H. B. McMahan, "Advances and Open Problems in Federated Learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1, 2021, doi: 10.1561/22000000083.
- [62] K. Bonawitz *et al.*, "Practical Secure Aggregation for Privacy-Preserving Machine Learning," *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, Oct. 2017, doi: 10.1145/3133956.3133982.
- [63] M. Wooldridge and N. R. Jennings, "Intelligent agents: theory and practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, June 1995, doi: 10.1017/s0269888900008122.
- [64] J. Park *et al.*, "Generative Agents: Interactive Simulacra of Human Behavior," *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, vol. 23, 2023, doi: 10.1145/3586183.3606763.
- [65] D. Bertsimas and N. Kallus, "From Predictive to Prescriptive Analytics," *Management Science*, vol. 66, no. 3, Aug. 2019, doi: 10.1287/mnsc.2018.3253.