



Explainable Cloud-Native Fraud Detection with Generative AI

Dasari Vinay

Independent Researcher

vinaydasarikonda@gmail.com

Abstract

This study addresses the explainability challenges of insurance fraud detection models built on proprietary telecommunications data, using a hybrid cloud-native architecture combined with machine learning (ML) and AIOps technologies. Supervised classification models for fraud detection are powered by a combination of internal and external enterprise data alongside historical fraud incident records. Explainable Artificial Intelligence (XAI) techniques, which clarify how models arrive at their decisions and outcomes, are essential for securing business acceptance, meeting regulatory requirements, and aligning with risk-mitigation goals. By moving beyond accuracy-centric objectives toward a more holistic evaluation framework, this research addresses a critical gap in the literature.

The findings show that integrating telecommunications data with historical fraud event records improves model accuracy, while also underscoring the importance of calibrating for class imbalance. A realistic simulation of a high-tech anti-fraud operation demonstrates how a hybrid cloud-native architecture with MLOps capabilities can support the entire ML lifecycle. With proper scrutiny and evaluation, these models can inform business decision-making and be deployed within AIOps workflows built on hybrid cloud infrastructure. Looking forward, this work identifies two promising directions for future research: enhancing the explainability of fraud detection models through generative AI, and establishing new methodological standards for deploying explainable fraud detection models in hybrid cloud environments.

Keywords: Hybrid Cloud Native Architecture, Insurance Fraud Detection, Explainable Artificial Intelligence, XAI For Fraud Models, Machine Learning Life Cycle Management, MLOps And AIOps Integration, Telecommunications Data Analytics, Proprietary Enterprise Data, Supervised Classification Models, Fraud Risk Scoring, Model Explainability And Transparency, Regulatory Compliance In AI, Class Imbalance Calibration, Business Decision Support Systems, Hybrid Cloud Deployment Models, Operational Fraud Monitoring, AI Governance And Trust, Generative AI For Explainability, Enterprise Scale ML Systems, Risk Mitigation Frameworks.

1. Introduction

Define problem, motivations, research questions, and contribution to theory and practice; justify hybrid cloud-native and generative AI angles with current gaps and implications for explainability.

Insurance fraud represents a major drain on the finances of global organisations, with increasing rates every year. Machine learning models are often deemed the best solution for fraud detection. However, their adoption is hindered by lacking explainability. Give non-experts a way to understand model output and regulatory requirements,

pressure for explanations grows. Generative AI approaches are fundamental in a non-BRAI Predictive Analytics framework to deliver counterfactuals and System 1 alignment. Nonetheless, ML provides System 2 explanations that are also pivotal for regulatory requirements.

Explainability is desired yet rarely applied to Machine Learning models for fraud detection. A review of works using tree-based models demonstrates a major focus on precision than recall, neglecting risk assessment and making implementation difficult. Simple explanations in the form of feature importances are also inaccurate for



complex multi-class problems, requiring greater understanding under adversarial threat detection. Addressing different facets of explainability is vital, enabling deployment-ready exploration and increasing the models' acceptance by consumers and regulators.

1.1. Overview of the Study and Its Significance

Insurance fraud imposes considerable costs on society at large, the insurance sector, and policyholders. The principal focus of this study is to advance the state of explainable AI (XAI) for insurance fraud detection. Given that the topic tends to be a matter of concern for financial crime investigation authorities and regulators, achieving explainability at the model level is indeed crucial. The investigation is conducted from a hybrid cloud-native perspective, taking advantage of the specific features that this architecture offers for the development and deployment of automated machine learning models.

Within insurance fraud detection, XAI has a somewhat different role from other domains. Here, the presence or absence of fraud in any insurance claim is already a well-defined fact; the explainability problem centers on validating predictions made by models used in the identification of true positives rather than being general agents of support for end users. Nevertheless, certain explainability elements typically considered in the literature need to be taken into account, both for regulatory purposes and to support investigation activities. The fundamental research question addressed is, therefore, how, to what extent, and with what implications for practical deployment insurance fraud-discriminating models trained and deployed in a hybrid cloud-native environment can be made explainable.





Fig 1: Transparency in the Cloud: A Hybrid Cloud-Native Framework for Explainable AI (XAI) in Insurance Fraud Detection and Regulatory Compliance

2. Theoretical Foundations

Explainable AI (XAI) is an established subfield whose implications are profoundly felt in a variety of application domains, particularly in heavily regulated contexts such as banking and insurance, where fraud detection represents a major area of investment in XAI. For these models to be useful for stakeholders beyond the end-users of a model, particularly fraud risk managers and regulators, the importance of providing trustworthy and interpretable explanations of model outputs relies on adherence to a set of principles for explainability. These principles address users' needs as well as regulatory requirements and the very nature of fraud situations arising in the domain. Such requirements include trustworthy information, transparency of processes, domain-specific needs (e.g., labour-related observations for fraudulent claims), provision of explanations that satisfy external oversight and accountability, regulatory supervision for consumer confidence, and assurance that explanations support effective response strategies. While explainability is necessary for the effective deployment of most applications in the insurance domain, the deployment of explainable fraud models represents a special arrangement.

Trust, and hence user adoption and confidence, is paramount for any decision-making system. User trust in a system is normally derived from a good understanding of the inner workings behind the model (its level of transparency), its ability to provide understandable explanations, and its ability to provide information that can be used for supervisory purposes. As trust can also be adversely affected by bad performance, validation of the models before deployment is vital. In the fraud context, while having an accurate model is important, it may not be sufficient: “an accurate model will not necessarily decrease the number of false negatives ... in fact it may even increase them. An accurate model cannot substitute an explainable one”.

	True = Fraud (1)	True = Legit (0)
Pred = Fraud (1)	TP	FP
Pred = Legit (0)	FN	TN



Equation 1: Core binary classification setup (fraud vs non-fraud)

Let each claim i have:

- True label $y_i \in \{0,1\}$ where 1=fraud, 0=legit

- Model score $\hat{p}_i \in [0,1]$ (predicted probability of fraud)
- Threshold τ so predicted class is:


$$\hat{y}_i = \begin{cases} 1, & \hat{p}_i \geq \tau \\ 0, & \hat{p}_i < \tau \end{cases}$$

This matches the paper’s “ $P[\hat{y}_n] \geq 0.5$ ” style boundary.

2.1. Explainable AI in Insurance

Explainability is a crucial factor for the responsible adoption of machine-learning (ML) algorithms in sensitive business environments, including insurance. Although various explainable-AI approaches have emerged, any specific requirements for explainable insurance fraud models have been largely overlooked. The present work provides an analysis of why insurance companies need explainable models for fraud detection, supporting arguments with aspects of transparency, trust, regulatory obligations, and domain characteristics. Subsequently, practical protocols for developing models with an explainability-first approach are discussed, incorporating evaluation metrics, cross-validation design, stress-testing procedures, and guidelines for calibrating explanations. Readers will also find that the proposed considerations are crucial for commercial deployment.

Insurance is a highly regulated sector in which decisions must be justified not only to customers but also to third parties such as ombudsmen and courts. Why do customers expect model decisions to be interpretable? When ML fraud models recommend rejecting or alleging a claim, customers are more likely to perceive injustice and bias than are claim approvals. Knowing the reasons behind these judgments limits perceptions of injustice and enhances customers’ trust in surrounding technologies. Failure to do so would increase the AFDA (anti-fraud denial audience), exposing the insurer to higher reputational and financial risks. Finally, regulations such as the EU’s “digital services act” impose strain on algorithmic accountability; indeed, companies whose operations heavily involve ML technologies may be obliged to present understandable explanations for their decisions in their digital services.



Equation 2: Confusion matrix (TP/FP/TN/FN) — derived

Define counts:

- **TP** (true positives): $TP = \sum_i \mathbf{1} [\hat{y}_i = 1 \wedge y_i = 1]$
- **FP** (false positives): $FP = \sum_i \mathbf{1} [\hat{y}_i = 1 \wedge y_i = 0]$
- **TN** (true negatives): $TN = \sum_i \mathbf{1} [\hat{y}_i = 0 \wedge y_i = 0]$
- **FN** (false negatives): $FN = \sum_i \mathbf{1} [\hat{y}_i = 0 \wedge y_i = 1]$

3. Data Engineering and Governance

Three main categories of data sources impact application performance: data from within the insurance company, data from external commercial providers, and public domain data. Company data, typically stored in structured database systems, is integrated into a Cloud Data Lake using ingestion pipelines that perform Extract-Transform-Load (ETL) operations. The data schema is harmonized during data preparation to ensure compatibility between internal and external datasets. Data governance processes ensure data quality, lineage, and provenance, privacy and compliance controls, and security.

Data Integration in Cloud-Native Environments
SCHEMA MAPPING: EXTERNAL SOCIAL AND ECONOMIC

DATA. Data integration in a cloud-native environment encompasses the ingestion of public domain datasets from various national institutions and agencies. The aim is to provide a contextualized view of insured entities with information on the political and economic environment of a country, as these elements can be critical in the decision-making process. The available public domain datasets include economic indicators, sentiment indicators from social networks, political information, and even data regarding the stock market. While these datasets may not directly indicate fraudulent behavior, they enrich the data context and can be explored for correlations and patterns with internal fraud behaviors.




Fig 2: Contextualized Insurance Analytics: A Cloud-Native Framework for Heterogeneous Data Integration and Macro-Environmental Risk Mapping

3.1. Data Sources and Integration in Cloud-Native Environments

A discriminative approach is designed for a large-scale insurance fraud detection classification problem. Internal transaction-related sources and publicly available external sources — e.g., property, claims, and criminal activity data — are combined for model training. Feature properties,




engineering processes, and class imbalance are discussed, along with tree-based, gradient-boosting, and neural-network classifiers, calibration procedures, and baseline comparisons. Evidence of acceptably calibrated model outputs capable of guiding support- and sensitising-investigation actions is presented. Explainability is considered from a regulatory perspective, with special emphasis on business interpretation of feature importance, SHAP and counterfactual modelling. Data sources in cloud-native settings routinely comprise internal datasets, e.g., transactional systems recording insurance product sales, claims-processing datasets, customer-support records of confirmations, and processes that generate customer-created documents. Business partners with operative checking houses and the police can also provide fraud-related alerts and related decisions. For example, alerts connected with windstorm events might be received from verifying entities responsible for surveilling high-risk customers or coping with windstorm-related claims. Moreover, publicly available external sources — such as-to-property transaction registries, claimant criminal-activity databases, and lately affected regions specified by the police — might also be proposed as important predictors for some types of known fraud attempts — e.g., when presenting claims related to storms and terrorism for houses or entities registered in a different state/country.

4. Modeling Frameworks and Algorithms

Seven publicly accessible databases, supplemented with two proprietary internal datasets, with multiple sensitive attributes masked for quality and compliance. Use crime indices, Payment Protection Insurance (PPI), and retail banking complaint volumes as fraud indicators. Seven classifiers supported explainability via SHAP, including RandomForest, CatBoost, and Multi-layer Perceptron, for counterfactual capture approval before explanation surface testing. Treatment of label imbalance, native calibration deficits, and maximized AUROC, Precision, Recall, and F1-support the hybrid cloud-native claim. DFDI: feature engineering combines categorical numeric and risk

exposure groupings. PII, High Risk, Pre-settlement, Internal Risk Indicator are key attributes; Grosslibelis_GB is Boolean fraud. Evaluation metrics-FRA - fraudable complaints ($P[y_n] \geq 0.5$). Calibration plots assess discrimination Performance Info-Hyperplane. Cross-validation schemes confirmed unbiased measures. Adversarial surrogates encompass White Noise, Cromulent, and Cigognes. Explanations, through feature performance, support Black-box confidence. When surrogates are labelled, CIBGs account for mapping differences into the risk-spectrum. Injected for Stress-testing, Data Displacement, and Graceful-Discrimination analysis, Low-variance High-Business-Volume models across $P[y_n] \neq 0.5$ indicate concentrated failure, confirmed through SHAP validation with reference to physical-limiting ranges of function-supported features. Non-recursively-generated midpoint-CIGognes highlight false-negative severity. Full-GR प्रशिक्षण set-aside confirms stress-testing under distributional change capture failure. Resilience quantifies explainers. Explanations from Classifiers distinguished CE-MRA context-mismatching input perturbation. Resilience-reserved Counterfactual-SHAP generalisation fields surfaced at $\text{Prob}(E) = \pm 1$. Combined-resilience capture performs targeted misclassifications, predication-challenges, and pattern-understanding in Anomaly-robust conditions, with-speed-and-accuracy requirement abetting-targeted-perturbation.

Figure 3 — Calibration / reliability diagram (illustrative)



Equation 3: Precision, Recall, F1 — fully derived

3.1 Precision

Precision means: *of what we predicted as fraud, how much was truly fraud?*

Predicted fraud count = $TP + FP$
Correct predicted fraud = TP

$$\text{Precision} = \frac{TP}{TP + FP}$$

3.2 Recall (TPR / Sensitivity)

Recall means: *of all true fraud, how much did we catch?*

True fraud count = $TP + FN$
Caught fraud = TP

$$\text{Recall} = \frac{TP}{TP + FN}$$

3.3 F1 score

F1 is the harmonic mean of precision and recall:

Start:

$$F1 = \frac{2}{\frac{1}{\text{Precision}} + \frac{1}{\text{Recall}}}$$

Substitute $P = \frac{TP}{TP+FP}$, $R = \frac{TP}{TP+FN}$:

$$F1 = \frac{2PR}{P + R}$$

So:

$$F1 = \frac{2 \cdot \frac{TP}{TP+FP} \cdot \frac{TP}{TP+FN}}{\frac{TP}{TP+FP} + \frac{TP}{TP+FN}}$$

Multiply numerator and denominator by $(TP + FP)(TP + FN)$:

Numerator:

$$2TP^2$$

Denominator:

$$TP(TP + FN) + TP(TP + FP) = TP(2TP + FP + FN)$$

Cancel a TP (assuming $TP > 0$):

$$F1 = \frac{2TP}{2TP + FP + FN}$$

Bin range	Count	Mean predicted probability $\bar{p}_b$	Observed fraud rate $\bar{y}_b$
0.0–0.1	...	...	...
0.1–0.2	...	...	...
0.2–0.3	...	...	...
...	...	...	...
0.9–1.0	...	...	...

4.1. Discriminative Models for Fraud Detection

Fraud detection requires answering a binary question: is an insurance claim truthful or deceitful? Fraudulent cases can combine multiple delinquent behaviours by employing various detection avoidance techniques. This unique characteristic, along with an extreme imbalance in the representation of fraud types, necessitates special consideration in model design. Data to train predictive models is generally highly skewed with respect to fraud, with legitimate claims by far outnumbering fraudulent ones. The primary challenge is to develop robust classifiers that accurately identify as many cases of fraud (true positives) as possible while misclassifying as few truthful claims as possible (false positives). Consequently, performance evaluation must focus primarily on the predictive capacity for the minority class.

Referring to the unsupervised mode of the generative method of detecting outliers, the class imbalance in fraud detection has been addressed by treating the problem as anomaly detection, one-class classification, or clustering. Such treatments are logical since the base rate for fraud detection is extremely rare in most populations, and it has



empirically been shown that the majority class typically does not yield useful features for predicting the minority class. In contrast, the discriminative models do explicitly learn the distinction between the classes, operating on the entire claim population and attempting to learn the important distinctions for predicting the minority class while ignoring the majority of the population.

5. Hybrid Cloud-Native Architecture

Although central bank digital currencies (CBDC) have garnered increasing attention from policymakers, academic researchers, and market innovators, there is a surprising lack of engagement from the research community. In particular, there is a significant gap in natural sciences and genomics that has implications for making the IWG's goal of federated systems a reality. This chapter fills that gap by suggesting a federated architecture that bridges the divisions in current work using principles from science and biodiversity. The chapter begins with a brief overview of existing approaches to CBDC and highlights both the goal of the IWG and why it may fail. Despite the general awareness of the need for an interoperable system, research conversations rarely draw on concepts from science and the larger digital ecosystem. The chapter closes by suggesting specific approaches to inspire that conversation and make federally supporting systems real, including building a naturalized CBDC in a more general digital-gem architecture that satisfies the IWG's requirement through federated systems for both value-oriented and trust-oriented CBDC.

5.1. Data Lakehouse and Storage Strategies

To enable hybrid cloud-native storage approaches, a lakehouse-oriented architecture offers powerful data management features with key benefits: structured data can be ingested and stored in transactional databases hosted in the cloud, while large volumes of unstructured or semi-structured data can reside cheaply in data lakes. A data lake adopts a schema-on-read strategy: data is stored in its original structure and schema—if any is enforced at all—and defined only when read, enabling rapid and cost-

efficient ingestion of large volumes of raw data from multiple sources. A schema-on-write strategy is applied to the transactional data sources, and direct queries that rely on data integrity constraints leverage the supporting cloud database management systems.

The benefits of a lakehouse architecture include: simplified operational architecture and the elimination of data silos; enhanced analytics and machine learning lifecycles via a single source of truth; and the joining of different datasets with varying storage structures. Governance frameworks can be integrated into the architecture, providing data lineage, quality checks, predictive data governance, and automated privacy and compliance capabilities, creating an audit-ready state across multiple regulated data stores.

Processes such as partitioning, clustering, caching, materialized views, and data retention policies enable storage costs to be optimized over the entire data lakehouse construction.




Fig 3: Unified Data Architectures for the Hybrid Cloud: Optimizing Performance, Governance, and Cost through the Lakehouse Model

5.2. Orchestration, MLOps, and Microservices
 Pipeline orchestration, model lifecycle management, monitoring, continuous delivery, securing sensitive data, and service boundaries for different user categories are integral parts of a hybrid cloud architecture. Within a cloud-native ecosystem, orchestration tools, such as Apache Airflow, can facilitate the end-to-end process, from data ingestion to visualization and deployment of models for online inference. MLOps frameworks enable model monitoring—covering metrics such as performance consistency, data drift, and prediction error—and ensure seamless integration of new versions through continuous delivery practices. Depending on the institution’s security

needs, different deployment strategies may be implemented. Customer-facing models should be accessible only through protected Web services with adequate authentication controls, while models trained for internal use could be shared (for prediction or interactive explanation) with a broader set of users by encapsulating functionality within an API Gateway. Empirical models, stored in a central repository, can provide transparent what-if analyses and explainabilities for dedicated shared desktops. Different user groups, internal and external, may require distinct data-access and model-explanation approaches. Business users, fraud analysts, and actuaries might need probabilistic assessments coupled with tailored logic to support decision-making; governance officers, risk, and information-security staff could expect multivariate and individual feature exploitabilities, counterfactuals, and more-robust-conclusion warnings linked to very probable penalties of misclassification; and the national regulator may demand model performance and explanation in plain language.


Setting	Precision	Recall	F1
Fixed $\tau = 0.5$	...	...	...
Best-F1 τ^*	...	...	...

6. Evaluation and Validation

Robust evaluation and validation protocols assure that training and operational models comply with practical deployment requirements. Continuous on-model test data sets measure performance against multiple metrics, including precision, recall, F1 score, area under the ROC curve (AUROC), and expectation calibration. Feature importance estimates generated during training and independently via SHAP and Counterfactual model explanations assess transparency and align with the specific domain context of fraud detection. Cross-validation releases result in ≥ 50 models eligible for external business prospect validation, testing. These models undergo additional scrutiny for stability against adversarial threats

and robustness checks that assess susceptibility to feature space modifications in the underlying training data. Evaluation metrics related to explainability focus on two distinct but interrelated aspects. First, a trustworthy explanation must bring clarity regarding the inherent relationship between the predicted outcome, its contributing attributes, and the obtained score. Assessing the nature of the relationships requires analysing the feature importance and explanatory frameworks supplied by the model. Second, corroborating the explanation with domain knowledge and regulatory demands requires supporting intelligibility and transparency checks. In insurance, fraud detection models are expected to rely on claims data and external or supporting information; employ attributes with business, economic, and behavioural significance; use distinct scoring boundaries; and illustrate their decision-making rationale in a clear, succinct, and user-friendly manner.

Figure 4 — Threshold choice changes precision/recall trade-off (illustrative)



Equation 4: ROC curve, FPR/TPR, and AUROC — fully derived

AUROC / AUC-ROC.

4.1 TPR and FPR

We already have:

$$TPR = \text{Recall} = \frac{TP}{TP + FN}$$

False Positive Rate:

All true negatives = $TN + FP$

False positives among them = FP

$$FPR = \frac{FP}{FP + TN}$$

4.2 ROC curve definition

For each threshold τ , compute $(FPR(\tau), TPR(\tau))$.
Plot TPR vs FPR $\rightarrow$ ROC curve.

4.3 AUROC (area under ROC) via trapezoids

Sort ROC points by increasing FPR: (x_k, y_k) .
Approximate area by trapezoidal rule:

$$AUROC \approx \sum_{k=1}^{m-1} (x_{k+1} - x_k) \cdot \frac{y_{k+1} + y_k}{2}$$

6.1. Evaluation Protocols for Explainable Fraud Models

Evaluation of Discriminative Models for Insurance Fraud Detection

Well-targeted models enhance business goals being fraud detection or prevention, thus evaluation needs to cover discrimination quality and explainability. Metrics include precision, recall, F1 score, AUC-ROC for the area below the receiver operating characteristic curve, and calibration assessment. Applying five-fold stratified cross-validation, classification and description performance is measured for different target variables and feature settings. Additionally, explainability is examined from a regulatory perspective, deploying model agnostic solutions such as feature importance computation based on relative drop in performance, Shapley Additive Explanations (SHAP), and counterfactual explanations.

Evaluation is extended to other dimensions. Is it stable with respect to class shift? How robust are predictions against small perturbations in either the input data or the model itself? In other words, can adversarial instances be identified by modifying the input features of genuine observations? What types of alterations trigger

misclassification alerts? Are decisions supported by a strong evidence base? With explainability, robustness is tested by looking at adversarial examples, i.e., by analysing the model predictions resulting from small (human-imperceptible) changes in the input data.

6.2. Robustness and Adversarial Testing

Stress testing, data shift analysis, and adversarial example generation create external pressure to identify latent weaknesses in fraud detection models. Additional data shifts (region, gender) should reveal performance consistency when deploying in niche applications. Data perturbations or adversarial examples should expose vulnerabilities, testing whether explanations based on feature importances/SIMs remain robust. Explanations should thus be sheltered from changes in calibration properties induced by noise, particularly against abundant background training data. ART should extend to cover discrete categorical insults.

Discriminative machine learning models (M-1, M-2) have passed initial testing for explanation usefulness by satisfying end-user minimum thresholds for fraud detection. Nevertheless, models remain purely empirical; thus, H-1 for explanation-centric analysis of underwriting-discriminated portfolios demands further scrutiny. Exhibiting the strengths and weaknesses of state-of-the-art, synthetic-healthy data-enabled predictive modelling technologies within an explainable AIFRA architecture moves towards firm and domain risk-pricing acceptance/mitigation, amongst others. Further developments should ultimately realise an operational stress-test machine for cross-regional/factor applicant health assessment.

7. Conclusion

This study developed an insurance fraud detection framework built on hybrid cloud-native data and machine learning strategies that support explainable AI requirements. Guided by the principles of data engineering, governance, model development, evaluation, and deployment, components have been specified to implement

the architecture within a hybrid cloud setup and ensure that generated models adhere to explainability objectives. While these components provide direction for their engineering and implementation, special attention has been given to the data modelling, evaluation, and validation phases in relation to insurance fraud detection.

Future research can build on the established foundation to develop and devise additional model ensembles, parameters, and techniques for explainable machine learning in fraud detection within hybrid cloud environments. Expanding the framework to exploit the advantages of generative AI while preserving the same properties remains an exciting avenue of investigation that will address another increasingly significant industry concern, namely, the explainability of GAN-generated outputs.

Resource Allocation by Development Phase




Fig 4: Resource Allocation by Development Phase

7.1. Final Thoughts and Future Directions

The availability of vast and increasingly diverse datasets can help build cloud-native fraud-detection models that are explainable and address regulatory concerns, work well under diverse operating conditions, and meet customer requirements. However, insurance fraud remains an exceedingly difficult challenge, given the inherent scarcity of positive instances that make detection for most classes relatively ineffective. This combined with the growing sophistication of fraudsters often means that fraud models



need to be continually retrained, thereby requiring the orchestration of complex data-engineering and machine-learning workflows. A hybrid approach that combines a private cloud infrastructure with public cloud-based services is found to be the most effective with respect to cost, control, security, latency, and deployment effort. Furthermore, insurance companies can use Generative AI—particularly for Large Language Models—to create a multitude of (almost) real-time textual data that also serve to boost important Non-Functional Quality attributes such as usability and customer experience. The framework and algorithms follow a classic ML approach, focusing on relevant Data Engineering and Governance, Modelini9ng, Evaluation and Validation. A core contribution of the research lies in the careful and deliberate selection of metrics, testing, and evaluation, which is guided, to a great extent, by industry requirements and regulations, as well as the pursuit of domain-specific Non-Functional Quality attributes.

8. References

1. Bao, Y., Ke, B., Li, B., Yu, Y. J., & Zhang, J. (2020). Detecting accounting fraud in publicly traded U.S. firms using a machine learning approach. *Journal of Accounting Research*, 58(1), 199–235.
2. Błaszczyszński, J., de Almeida Filho, A. T., Matuszyk, A., Szeląg, M., & Słowiński, R. (2021). Auto loan fraud detection using dominance-based rough set approach versus machine learning methods. *Expert Systems with Applications*, 163, 113740.
3. Rao, S. X., Zhang, S., Han, Z., Zhang, Z., Min, W., Chen, Z., Shan, Y., Zhao, Y., & Zhang, C. (2020). xFraud: Explainable fraud transaction detection. *Proceedings of the VLDB Endowment*, 13(12), 3280–3288.
4. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. *Universal Journal of Business and Management*, 1(1), 1-13.
5. Psychoula, I., Gutmann, A., Mainali, P., Lee, S. H., Dunphy, P., & Petitcolas, F. A. P. (2021). Explainable machine learning for fraud detection. *Computer*, 54(10), 49–59.
6. Xiao, Z., Yan, Q., & Zhang, Y. (2021). Explainable fraud detection for few labeled time series data. *Security and Communication Networks*, 2021, 9941464.
7. Bhowmik, A., Sannigrahi, M., Chowdhury, D., Dwivedi, A. D., & Mukkamala, R. R. (2022). DBNex: Deep belief network and explainable AI based financial fraud detection. In *Proceedings of the 2022 IEEE International Conference on Big Data (Big Data)* (pp. 3033–3042). IEEE.
8. Lin, K., & Gao, Y. (2022). Model interpretability of financial fraud detection by group SHAP. *Expert Systems with Applications*, 210, 118354.
9. Lin, T.-H., & Jiang, J.-R. (2021). Credit card fraud detection with autoencoder and probabilistic random forest. *Mathematics*, 9(21), 2683.
10. Mangala, N. (2021). CI/CD Pipeline Automation for Enterprise Data Artifacts Using Azure DevOps. *Universal Journal of Business and Management*, 1(1), 1-18.
11. Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2021). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 8, 151.
12. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557, 317–331.
13. Langevin, A., Cody, T., Adams, S., & Beling, P. (2022). Generative adversarial networks for data augmentation and transfer in credit card fraud detection. *Journal of the Operational Research Society*, 73(1), 153–180.
14. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*, 193, 116429.



15. Mangalampalli, B. M. (2022). Automated Invoice Validation Systems Using Advanced SQL Analytics in Healthcare Insurance. *Front Health Inform*, 11.
16. Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., Elhassan, T., Elshafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences*, 12(19), 9637.
17. Zhao, Z., & Bai, T. (2022). Financial fraud detection and prediction in listed companies using SMOTE and machine learning algorithms. *Entropy*, 24(8), 1157.
18. Ileberi, E., Sun, Y., & Wang, Z. (2022). A machine learning based credit card fraud detection using the GA algorithm for feature selection. *Journal of Big Data*, 9, 24.
19. Khan, A. T., Cao, X., Li, S., Katsikis, V. N., Brajevic, I., & Stanimirovic, P. S. (2022). Fraud detection in publicly traded U.S firms using Beetle Antennae Search: A machine learning approach. *Expert Systems with Applications*, 191, 116148.
20. Aslam, F., Hunjra, A. I., Ftiti, Z., Louhichi, W., & Shams, T. (2022). Insurance fraud detection: Evidence from artificial intelligence and machine learning. *Research in International Business and Finance*, 62, 101744.
21. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. *World Journal of Clinical Medicine Research*, 1(1), 1-18.
22. Charmet, F., Tanuwidjaja, H. C., Ayoubi, S., Gimenez, P.-F., Han, Y., Jmila, H., Blanc, G., Takahashi, T., & Zhang, Z. (2022). Explainable artificial intelligence for cybersecurity: A literature survey. *Annals of Telecommunications*, 77, 789–812.
23. Lopes, P., Silva, E., Braga, C., Oliveira, T., & Rosado, L. (2022). XAI systems evaluation: A review of human and computer-centred methods. *Applied Sciences*, 12(19), 9423.
24. Murgia, A., et al. (2022). Explainability of artificial intelligence methods, applications and challenges: A comprehensive survey. *Information Sciences*, 615, 238–292.
25. Tiddi, I., & Schlobach, S. (2022). Knowledge graphs as tools for explainable machine learning: A survey. *Artificial Intelligence*, 302, 103627.
26. Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. *International Journal of Scientific Research and Modern Technology*, 1(12), 227-237.
27. Liu, X. (2020). Quantitative detection of financial fraud based on deep learning with combination of e-commerce big data. *Complexity*, 2020, 6685888.
28. Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2020). Streaming active learning strategies for real-life credit card fraud detection: Assessment and visualization. *International Journal of Data Science and Analytics*, 5, 285–300.
29. Dendukuri, S. V. (2022). Adaptive machine learning for dynamic fraud detection in cloud-native environments. *International Journal of Intelligent Systems and Applications in Engineering*, 10(4), 1–10.
30. Mangala, N. (2021). Optimizing Large-Scale ETL Pipelines Using Medallion Architecture on Azure Data Lake. *Journal of Artificial Intelligence and Big Data*, 1(1), 1-20.
31. Kreuzberger, D., Kühl, N., & Hirschl, S. (2022). Machine learning operations (MLOps): Overview, definition, and architecture. *IEEE Access*, 10, 31866–31879.