



Engineering Fraud Detection Pipelines: DevOps and Agentic AI with Cloud-Native Deep Learning

Anumandla Mukesh

Independent Researcher, Malaysia

mukeshanumand@gmail.com

Felix Hartmann

Scientific Research Assistant, Westfalen College of Applied Sciences, Germany

felixhartmaann@outlook.com

Publication History: Received: 03.06.2026; Revised: 27.07.2026; Accepted: 01.08. 2026; Published: 14.08.2026.

ABSTRACT: The study investigates DevOps-driven data engineering for agentic AI in insurance fraud detection with a focus on cloud-native automations. Data are centrally secured, accessible, governed, of reproducible quality, properly profiled, and efficiently prepared. Design, development, and deployment encompass extensive safeguards for risk-sensitive decision-making. A superset of continuous integration/continuous deployment tailored to machine learning as well as data governance underlie the approach, demystifying the technology stack and enabling success at operational scale. The investigation is motivated by regulatory, economic, and ethical pressures on the industry, demand for predictive solutions, and pain points of legacy monitoring systems. Automation of machine-learning model development and deployment mitigates issues of scalability, reproducibility, and security while incorporating risk assessment.

Organizations worldwide are increasingly turning to artificial intelligence in a bid to detect fraudulent activity in all its forms. Such systems, however, generally stem from one-off proof-of-concept initiatives, operate in isolation, and lack sufficient controls for operational deployment in support of everyday decision-making. Nevertheless, agentic artificial intelligence—actant solutions capable of executing decisions without direct human intervention—promises a transformative competitive advantage for business leaders. Integrating governance and control ensures that these automated systems operate within an acceptable risk tolerance, align with organizational values, and minimize unintended consequences. Insurance fraud detection solutions typically fall short of these requirements.

KEYWORDS: DevOps Driven Data Engineering, Agentic Artificial Intelligence, Insurance Fraud Detection, Cloud Native Automation, Machine Learning Operations, Continuous Integration Continuous Deployment, Data Governance Frameworks, Risk Sensitive Decision Making, Automated Model Deployment, Predictive Fraud Analytics, Scalable AI Systems, Reproducible Data Pipelines, Secure Data Architectures, Regulatory Compliance In Insurance, Ethical AI Governance, Legacy System Modernization, Autonomous Decision Systems, Operational AI Control, Model Lifecycle Management, Enterprise AI Adoption.

I. INTRODUCTION

Insurance fraud, encompassing deliberate misreporting of claims, is a costly global problem that erodes corporate earnings and increases annual customer expenses. The prevention of fraudulent activity requires constant vigilance by insurance companies, which dedicate substantial resources to fraud detection. However, the sheer volume of claims makes manual assessment prohibitively expensive. Consequently, organizations deploy analytical models capable of detecting suspicious and potentially fraudulent claims for further investigation—ideally using enterprise data lakes that integrate multiple sources of structured and unstructured information and external databases. Machine Learning (ML) processes are well-suited to this task, as they can analyze the information within, derive predictive patterns, and then automatically apply those patterns to future claims.



Nevertheless, the application of ML models to support fraud detection remains constrained by a range of factors. Data and ML development remain mainly siloed, thus preventing stakeholder-driven, responsive development, while produce-scand-operate processes remain disjointed rather than collaborative and continuous. The resulting tools lack transparent governance and control, leading to security concerns and limiting their deployment in production systems; the empirical evidence for the payoff from such AI still remains patchy. Hence, the models cannot be deployed as services, scaling the analytic capability of the enterprise, or rapidly operationalized via pipelines that enable continuous integration/continuous delivery (CI/CD) through data.

1.1. Objectives and Scope of the Study

The primary goal of this study is to explore, develop, and validate a cloud-native DevOps-automated deep learning system to detect fraudulent insurance claims. While the agentic AI perspective offers the promise of scaling machine learning (ML) model training, monitoring, and evaluation, the principal emphasis here is on how the combination of DevOps and certain elements of agentic AI can bring transparency, security, governance, and auditability to the data preparation and model development life cycles. Additional issues under the DevOps umbrella—reproducibility, continuous integration/continuous deployment (CI/CD), supply-chain security, and quality—are also significant but less critical. Nevertheless, testing and evaluation across nearly all these dimensions will maximize usefulness and reliability.

Research activity in the insurance sector has long focused on fraud, reflecting the extensive financial losses incurred around the world. Despite the gradual implementation of machine learning solutions along the fraud-detection process, the application of state-of-the-art (SOTA) deep learning models remains limited. The main drivers of this gap are the complexity of deep learning networks and the difficulties involved in data preparation. Nevertheless, interest in insurance fraud detection persists, and recent developments in data engineering enabled by a combination of cloud-native DevOps and SOTA models can enhance scalability, reproducibility, and security while providing the infrastructure required for the deployment of agentic AI. Such developments could therefore increase the attractiveness of deep learning, closing the gap between established and emerging methods.





Fig 1: Scaling Deep Learning for Fraud Detection: A Cloud-Native DevOps Framework for Automated Governance and Agentic AI Integration



II. BACKGROUND AND MOTIVATION

Detecting fraud during insurance transactions has grown into a pressing concern for the industry. Reports from several countries confirm the growth of fraudulent transactions, together with a lack of transparency of the underlying data. Countries such as the Netherlands have introduced regulatory measures to tackle these concerns; such measures also increase pressure on the insurance industry to invest in machine learning solutions that support transparency and compliance while improving detection performance. While recent years have seen a substantial rise in interest in the research community, associations and industry, most efforts are still in a proof-of-concept stage – few, if any, results have made it into production.

The growing importance of fraud detection as part of the insurance value chain creates additional pressures: developing models that work at scale using industry-grade machine learning practices. Combining the principles established by DevOps for Data Engineering with industry-graded development and deployment considerations allows the creation of solutions that address these drivers. Scalability of the developed and integrated solutions appears to be the term that has been leant on most heavily; fulfilling the additional requirements of compliance, governance, transparency and reproducibility is a clear by-product of the DevOps-for-Data-Engineering principles. Beyond this, it is expected that the underlying business problems in the case of insurance transaction fraud detection can be sufficiently addressed; using the performance of the models in the support of insurance companies and associations can be a key source of motivation.



Equation 1: Data preparation equations

1.1 Min–Max normalization (MinMaxScaler)

Let:

- original feature value = x
- minimum of the feature column = x_{min}
- maximum of the feature column = x_{max}

Goal: map $x \in [x_{min}, x_{max}]$ to $\hat{x} \in [0,1]$.

Step-by-step derivation

1. Shift so the minimum becomes 0

$$x - x_{min}$$

2. Scale so the maximum becomes 1

The max shifted value is $x_{max} - x_{min}$. Divide by that:

$$\hat{x} = \frac{x - x_{min}}{x_{max} - x_{min}}$$



3. (General range $[a, b]$ if needed)
First compute $\hat{x} \in [0, 1]$ as above, then stretch/shift:

$$x' = a + (b - a)\hat{x} = a + (b - a) \frac{x - x_{min}}{x_{max} - x_{min}}$$

2.1. Key Drivers of Research and Innovation

Scaling, reproducibility, security, and governance are becoming increasingly relevant for deploying ML solutions in production. The insurance sector is facing continuous pressure to deliver models that address key business questions and support operational decisions. However, the industry has yet to embrace innovation at scale, often relying on traditional technologies, bespoke approaches, and partly redundant ML initiatives that are rarely integrated into business processes. Pressure for innovation is mounting in insurance companies and new insurance players are entering the market, forcing the insurance industry to compete through continued innovation.

Agentic AI comprises continuous integration and delivery principles applied to ML workflows, enabling the creation of automated, monitored pipelines to regularly produce and deploy new ML models. Regulatory requirements such as the General Data Protection Regulation are forcing insurance companies to implement stricter data governance processes. Data ownership, access management, data quality, data classification, and data lineage tooling are becoming core data engineering issues. Insurance companies can leverage modernization and cloud services to integrate security into the design and transpose cloud-native principles and patterns into data engineering and ML process automation.

Class	Meaning	Count	Share
0	Legitimate Claims	1129	94.1%
1	Fraudulent Claims	71	5.9%

III. THEORETICAL FRAMEWORK

DevOps principles provide valuable underpinnings for scalable, reproducible data engineering. Agentic AI strives for intelligent self-management of complex systems. Together, these ideas structure the research and guide the design of cloud-native, deep-learning-enabled software that addresses longstanding pain points in insurance fraud detection.

Continuous integration/continuous deployment (CI/CD) for machine learning, data governance, automated monitoring, and ethical considerations shape the investigation. Cloud-native processing meshes streaming and batch data, engenders a data-automation layer that supports ingestion, access, security, quality, lineage, profiling, and transformation, and provides a second automation layer for cloud-native ML pipelines.

Deep-learning-automation-as-a-service solutions ease proof-of-concept work by growing the four Tenet-specific pains and supporting three more—lack of production-ready ML, reproducibility, and security—before being taken up by the business-facing ML crew. Service teams outside the O&M squad own related service performance, stability, and security at service-data pipelines that form the dumb enablers (for safety, et al.).

CI/CD for ML and agentic AI's Monitoring tenets telescope to deliver proof of concept, O&M grows support for six company Operations pain points, and the master plan then steers operationalization of automated FaaS and ML pipelines. The latter deal with workloads Dw and Dd—crafting, hosting, and automating Thessala, probably doing the same for any approved MLOps job-lot as proof of concept, and supporting by touch-related-en)))) and Terraform/Docker-for-Dummies transports.




Fig 2: Synergizing DevOps and Agentic AI: A Cloud-Native Architectural Framework for Scalable Deep Learning and Automated Fraud Mitigation


3.1. Conceptual Overview of Key Principles

Continuous Integration/Continuous Deployment of Machine Learning

Although DevOps was conceived for software development, its principles of continuous integration/continuous deployment (CI/CD) can also be applied to machine learning (ML) models to enable the timely availability of high-quality algorithms. CI/CD focuses on three distinct phases: a) continuous integration, which includes the development, testing, and validation of ML pipelines; b) continuous deployment of a candidate model that is updated regularly; and c) continuous training, which acknowledges that the underlying conditions affect the performance of deployed models. This requires an end-to-end ML pipeline (e.g. data ingestion, cleaning, preparation, modeling) to be in place, which, besides automating the modeling process, allows the management of model accuracy and robustness throughout the entire model development and production lifecycle.

Data Governance

Data are the lifeblood of successful ML solutions, and specific governance strategies, tailored to the scale and complexity of the organization, are needed to ensure that data is always available, secure, and compliant with regulatory requirements. Data governance comprises the design and execution of policies, standards, and processes that enable the acquisition and use of data, guarantee confidentiality, availability, and integrity, and mitigate relevant risks, such as security breaches, data leaks, and regulatory noncompliance. Effective data governance provides an inventory of the data assets of an enterprise, including metadata such as data classification, schema, access rights, lineage, source, quality metrics, documentation, and contact information, to ensure data consistency, security, and quality. It creates a shared understanding of data assets, minimizes risks, and generates a cycle of ongoing improvement.



IV. SYSTEM ARCHITECTURE

The architecture provides a high-level view of the solution, detailing components, data flows, and integration patterns. Each process involved in insurance fraud detection is framed as a pipeline. The choice for cloud-native deployment is justified, and the multi-layer automation keeps complexity in check and security considerations addressed.

A. Data Ingestion Pipelines

Data ingestion covers the collection, access, lineage tracking, profiling, and quality assurance of data in separate processes for batch and streaming feeds. Access is controlled for both data producers and consumers according to governance principles. For batch loads, point-in-time access is provided (with snapshots enabled) and correctness ensured with integrity checks. A quality pipeline checks for records exceeding profile quality thresholds. Data sources include client lists (ground truth), transactional databases, claims history, user sign-up profiles, and social media postings.

B. Cloud-Native ML Pipelines

ML training is also organized as a set of pipelines. It comprises readiness checks, orchestration, containerization, training, monitoring, model registry, and CI/CD elements and supports the online production of models for direct consumption by the business. Readiness checks verify the secure setup of cloud resources, design specifications, and data availability. The orchestration component handles resource provisioning and model training. In terms of model reproducibility, a container image is prepared before training—containing the environment libraries, the model training code with hyperparameter definitions, and the model storage bucket. The scan module periodically verifies the training setup based on CI principles.

Thres hold	T P	F P	F N	T N	Preci sion	Re call (T PR)	F1 Sc ore	F P R
0.30	7 0	3 9 2	1	73 7	0.15	0.9 9	0.2 6	0. 35
0.50	5 9	8 3	1 2	10 46	0.42	0.8 3	0.5 6	0. 07
0.70	3 6	7	3 5	11 22	0.84	0.5 1	0.6 3	0. 01



4.1. Data Ingestion and Governance

Ingestion pipelines ensure quality, governance, and security, delivering trustworthy datasets for consumption. A Batch Loading Zone (BLZ) maintains structured data accessible for business users and analysts. The Cleaning Data Repository (CDR) holds refined datasets ready for machine learning. Stream and batch ingestions differ for incremental fraud monitoring and periodic training data updates. Batch operations handle files staged in the Cloud Storage Landing Zone. A comprehensive audit trail, established via DataFlow, tracks every operational aspect.

Automated data quality checks, covering schema validation, completeness, accuracy, timeliness, and consistency, are vital for archived fraud detection data. For sensitive data, tools like Double-Check create a secondary copy with schema masking, ensuring secure development/testing access. Governance mechanisms enable appropriate authorizations at every stage. Data egress employs Data Transfer Service with fpp-level control, and Data-Lake-Information-Climate exposes additional technical metadata (e.g., data quality scores).

Equation 2: Outlier capping at three standard deviations

Let:

- dataset values: $x_1, \dots, x_n$
- mean:

$$\mu = \frac{1}{n} \sum_{i=1}^n x_i$$

- population standard deviation:


$$\sigma = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2}$$

- bounds:

$$L = \mu - 3\sigma, \quad U = \mu + 3\sigma$$


Capping rule (piecewise)

$$x^{(cap)} = \begin{cases} L, & x < L \\ x, & L \leq x \leq U \\ U, & x > U \end{cases}$$



4.2. Cloud-Native Machine Learning Pipelines

Pipelines ingesting the data into the machine learning (ML) module enable both model training and model orchestration at inference time. ML candidate models are deployed to a Google Cloud Vertex AI model registry where they can be evaluated, versioned, and monitored. Containerization enables new model versions to be generated in CI/CD for ML style,



5.1. Data Cleaning and Transformation Techniques

The data preparation workflow relies on a combination of data-catalog services, ETL orchestrators, and user-developed scripts or notebooks. The first step focuses on the cleansing of the InsuranceCompany Fraud Detection Dataset (ICFDD) (Kaggle 2021) prior to being used in the training, testing, and validation of machine learning models. After almost 1900 records of missing values were removed, the categorical attributes (underwriting, route, make, type, and fraud) had their values transformed into one-hot variables to allow for proper processing by machine learning algorithms. The fraud value was inversely coded via bitwise not and then multiplied by -1 to prepare the dataset for plotting of the Fraud Detection Fractal, thus producing a stronger contrast between fraud and nonfraud patterns. Continuous, discrete, and categorical attributes were then analysed using K-means profiling in order to leverage the data's natural structure for future model training.

A separate data preparation workflow conforms the ICFDD for ingestion into an NLP-based fraud detection pipeline. User-defined Python scripts perform additional cleaning duties, including removal of URL and HTML coding information that could otherwise interfere with the NLP engine's predictive capability. Normalization transforms any anomalous or redundant strings into a common format and part-of-speech tagging supplies structure to every word token in a description to facilitate intelligent selection of target language for the NLP engine. Word-count detection further identifies portions of the comments made by an insurance agent that may be inadequate for high-quality translation into the report's target language.




Fig 3: Bimodal Data Engineering for Insurance Fraud Detection: Integrating Fractal-Encoded Numerical Processing and NLP-Based Linguistic Structuring




VI. MODEL DEVELOPMENT AND EVALUATION

The model development for attack detection is conducted in three high-level stages. The first stage involves the exploration of various deep learning architectures and the selection of promising candidate models suitable for classifying fraud. The second stage covers the definition of training protocols for the identified candidates, an assessment of potential overfitting issues, and the identification of appropriate evaluation metrics as well as a validation plan. The third stage includes a concrete plan for the deployment of the best-performing model into continuous operation.

While deep learning discriminative classifiers are the core candidate architectures, some foundational supervised models have been added and will be evaluated to better understand the data and to provide performance benchmarks. In addition to convolutional networks for images, long short-term memory models for sequences, and spatiotemporal architectures (convolutional long short-term memory networks), it is also planned to explore less common unsupervised / self-supervised / adversarial approaches: generalised autoencoders, GANs and VAEs for both image and time series data anomalies. Predictive deep-fake models and GANs in particular might also indicate impression-manipulating sophistication or socially-aware generation capability. Other less-well-explored strategies like knowledge-distilled models – operationally efficient classifiers derived from knowledge-embodied and heavy transformers – will also be trialled on dataset subsets.

Evaluation metrics will include standard classification measures (precision, recall, F1-score, ROC-AUC, Brier score) but with added sensitivity to the criticality of false negatives. For instance, a relatively high false-positive-to-false-negative cost ratio could justify maximising the false-negative-and-missed-class rate instead of minimising them. Models will also be cross-validated using multiple techniques: standard k-fold, hold-out retroactive sets and time-respecting splits. Attention-based architectures will undergo attention-visualisation techniques to uncover which parts of the data are being used for classification. Sensitivity and explainability of supervised algorithms will be studied to understand what features they are using to identify fraud rapidly and which are “distracting” them during inference.




6.1. Model architectures and algorithms

A variety of model architectures is available for building classification models in insurance fraud detection. Popular supervised learning methods include tree-based algorithms such as random forests, gradient boosting, and XGBoost. Novel unsupervised techniques for fraud detection use autoencoders, deep support vector machines, and graph embedding methods. Unsupervised techniques facilitate novel insights while avoiding known pitfalls of supervised learning, such as the bias introduced by labelling and the dependence on correctly capturing fraud signals in the training data. Moreover, fraud, in particular new types, tends to occur very rarely in insurance data, making supervised learning less suited to circumstances where only historical data can be used.

An alternative approach to avoiding labelled training data is to derive inherent costs from the data and build cost-sensitive classifiers, capable of operating in an environment in which fraud occurs. For controlling misconduct, a new way to



optimize rules by minimizing label-dependent measures with a fast approximation using a reproducible cost-sensitive forest model has emerged. The new rule-induced cost-sensitive model provides an approximate solution to cost-sensitive multi-class problems while combining the advantages of tree-based methods and rule models. Finally, multiple-instance learning approaches tackle the issues stemming from the absence of precise labels and use the complement to classification for building explainable models in complex situations. All methods are equally suited for risk communication and should thus be applied in conjunction with interpretable AI methods.



Equation 3: Bitwise NOT label transform (as described)

“fraud value was inversely coded via bitwise not and then multiplied by -1 ”

For an integer y , bitwise NOT has the identity:

$$\sim y = -y - 1$$

Then multiplying by -1 :

$$-(\sim y) = -(-y - 1) = y + 1$$

So the transformation “bitwise not then multiplied by -1 ” is:

$$y' = -(\sim y) = y + 1$$

6.2. Evaluation metrics and validation

Different metrics communicate different information about the response of models to classified data, and thus, multiple metrics should be defined—relying on just one of them might yield a deceptive response. The metrics used in the development of the models include precision, recall, F1 score, area under receiver operating characteristic and calibration curves, and others. These are sufficient for evaluating whether the models are any good, as well as for suitable comparative analysis.

Precision indicates the percentage of true positives in relation to the number of positives classified by the model. Recall reveals the percentage of true positives in relation to the number of actual positives. The F1 score is the harmonic mean of precision and recall and combines precision and recall into one single score. The area under the receiver operating characteristic curve (ROC-AUC) indicates how well the model differentiates fraudulent transactions from legitimate transactions—an AUC of 0.5 indicates poor performance, while an AUC of 1.0 indicates perfect performance. Calibration indicates how well the model reflects prediction probabilities. Calibration is more important when making risky decisions. The point where calibration curve crosses $y = x$ is known as the ideal state. A model with a good calibration curve is viewed as a trustworthy model.



Raw Value (₹)	Capped ($\mu \pm 3\sigma$)	Min-Max Scaled
60000	42608	1.00
75000	42608	1.00
-5000	-626	0.00
22943	22943	0.54
19327	19327	0.47
21477	21477	0.51
18263	18263	0.44
15730	15730	0.39
25700	25700	0.60
20111	20111	0.49
23589	23589	0.56
16894	16894	0.42

VII. CONCLUSION

The presented work illustrated the theoretical underpinnings and a viable instance of DevOps-supported data engineering for cloud-native agentic AI systems, applied to insurance fraud detection. The architecture, automations, data preparation workflows, and modelling process were documented as a set of design patterns that enable scalable, reproducible, and well-governed data engineering around ML models. Future research will address the development of production-ready models and consider the implications of ML-supported fraud detection in the domain of insurance.

As for direction, automating data engineering for DevOps principles and practices facilitates large-scale adoption of agentic AI, but brings new issues to the fore. Technical automation must be paralleled by normative developments in data governance and monitoring of ML models. Security is a pillar of DevOps, dovetailing with existing regulatory efforts; data privacy emerges as a key consideration; and ethical ML calls for careful attention within organisations before deploying decision-making or risk-assessing models. The ultimate objective is to offer solutions for shoring up insurance against fraud through solutions that align the prevention of malfeasance with the prevention of unfair corruption of the insurance safety net intended for community support.




Fig 4: Distribution of Development Focus



7.1. Final Thoughts and Future Directions

Agentic AI impacts society and economics as intelligent, self-learning systems become ubiquitous. Insurance fraud detection is a prime candidate for expansive deployment because of the amount of data processed, the importance of avoiding financial losses through the detection of fraudulent claims, and the motivation of fraudsters to exploit system weaknesses.

The solution, based on a cloud-native DevOps approach to data engineering, integrates the core components of agentic AI in the continuous integration/continuous deployment of machine learning models. Security, compliance, data quality, and governance are continuously assured throughout the data lifecycle. These principles facilitate the scalable deployment of models tailored for the constant battle against ever-evolving fraudulent techniques, but scalability and operationalization remain to be proven. The commercial readiness of the system with DevOps-driven data engineering for agentic AI is underscored by the data preparation and ML pipelines demonstrating how cloud-native tools can continuously produce clean, well-structured data ready for feeding into ML algorithms and support the continuous training of a variety of models using both supervised and unsupervised techniques.

Continuous deployment of the ML component of agentic AI solutions in a dedicated ML area of the cloud, equipped with a model registry and constant data profiling, would enable the development of a diverse family of models focused on keeping pace with the fast-changing behavior of fraudsters without being caught in the risk-reward trap of real-world operationalization. For the industry, this solution represents not only a way to reduce costs and provide a better service to honest customers but also an opportunity for a real-time ML-driven solution for fraud detection, a follow-on level of thought that could enable customizability for other use cases in the insurance sector itself and the banking industry.

REFERENCES

1. Komaragiri, V. B., Keerthana, G., Fathima, S. K., Vats, S., Priya, R. V., & Sivagami, S. (2025, May). Quantum Computing Integrated Digital Interaction Platforms for Secure Public Engagement Decentralized Data Exchange and Policy Optimization. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 1989-2001). Atlantis Press.
2. Ganti, S., Vankayalapati, R. K., Krishnamoorthy, P., Thakare, P. S., Nayak, U. A., & Vignesh, P. (2025, February). Enhancing IoT-Driven Smart Home Security and Automation with a GCN Model. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
3. Ganti, S., Vankayalapati, R. K., Krishnamoorthy, P., Thakare, P. S., Nayak, U. A., & Vignesh, P. (2025, February). Enhancing IoT-Driven Smart Home Security and Automation with a GCN Model. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
4. Kalisetty, S., Lakkarasu, P., Singireddy, S., Burugulla, J. K. R., Challa, K., & Gadi, A. L. (2025, August). Next-Gen Payment Gateways: Leveraging Federated Learning for Fraud Detection in Cross-Border Transactions. In International Conference on Artificial Intelligence: Theory and Applications (pp. 200-211). Cham: Springer Nature Switzerland.
5. Pandugula, C., Ganti, V. K. A. T., Kalisetty, S., Ande, R., & Mehta, D. (2025). Modeling the Som Kamla Amba Sub-Watershed Using Artificial Neural Networks (ANN) and the SWAT Tool. Available at SSRN 5341614.
6. Challa, S. R., Burugulla, J. K. R., Pamisetty, A., Challa, K., & Paleti, S. (2025, April). AI and ML-Powered Cybersecurity Strategies for Cloud Computing: Ensuring Infrastructure Stability in Financial and Retail Sectors. In International Conference on Smart Computing and Informatics (pp. 315-327). Cham: Springer Nature Switzerland.
7. Meda, R., Motamary, S., Dodda, A., Kummari, D. N., Kaulwar, P. K., & Gadi, A. L. (2025, August). Edge-AI Integration for Accelerated Decision-Making in Industrial Robotics and Manufacturing Plants. In International Conference on Artificial Intelligence: Theory and Applications (pp. 332-344). Cham: Springer Nature Switzerland.
8. Chava, K., Suura, S. R., Burugulla, J. K. R., Malempati, M., Nuka, S. T., & Komaragiri, V. B. (2025, July). Machine Learning Enabled Financial Statements in Assessing a Business's Performance. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS) (pp. 1-6). IEEE.
9. Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying temporal reasoning and agentic machine learning: A framework for proactive fault detection in dynamic, data-intensive environments. *Metallurgical and Materials Engineering*, 31(4), 552-568.



10. Nandal, N., Tyagi, K., Kumar, L., & Gadi, A. L. (2025, November). AI-Driven CV Analysis and Job Matching in HRM Using a Hybrid Fuzzy Logic-Deep Neural Network Approach for Higher Education. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
11. Inala, R. (2025). A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 1614-1628.
12. Pote¹, X. R., Pamisetty, A., Karthikeyan, G., & Gupta¹, D. (2025, May). Artificial Intelligence Enabled Smart Energy Conservation Systems for Intelligent Resource Management and Sustainable Future Power Grids. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 196). Springer Nature.
13. Nandan, B. P., Kumar, M. V. K., Garapati, R. S., Bandi, V. D. V. K., Davuluri, P. N., & Mangalampalli, B. M. (2026, March). AI-Enhanced Semiconductor Yield Optimization Using Hybrid Deep Learning and Edge Data Analytics. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1-6). IEEE.
14. Annapareddy, V. N., Challa, K., Komaragiri, V. B., Sriram, H. K., Kalisetty, S., & Goma, T. (2025, April). Explaining AI Techniques such as SHAP LIME and RISE in Limited Sample Size Neuroimaging Studies. In 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
15. KOTHAPALLI, S. L., SEENU, A., DILEEP, V., & YASMEEN, Z. (2025). THE FUTURE OF CUSTOMER ENGAGEMENT IN RETAIL BANKING: EXPLORING THE POTENTIAL OF AUGMENTED REALITY AND IMMERSIVE TECHNOLOGIES. INTERNATIONAL JOURNAL, 73(1), 72-79.
16. Vankayalapati, R. K. (2025). Zero-trust security models for cloud data analytics: Enhancing privacy in distributed systems. Available at SSRN 5121185.
17. Koti¹, V. M., S, A. J., Vinoth, M., & Bordoloi, D. (2025, May). AI-Driven Gamification Models for Long-Term Educational Engagement. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 73). Springer Nature.
18. Vellela, S. S., Purimetla, N. R., Rao, P. V., Daniel, V. A. A., Koppolu, H. K. R., & Janani, B. (2025). AI-Enabled Wearable Hemodynamic Monitoring System for Early Identification of Thrombotic Events. Vascular and Endovascular Review, 8(16s), 321-336.
19. Mashetty, S. (2025). LEVERAGING DEEP LEARNING, NEURAL NETWORKS, AND DATA ENGINEERING FOR INTELLIGENT MORTGAGE LOAN VALIDATION. INTERNATIONAL JOURNAL OF SOCIAL SCIENCE & INTERDISCIPLINARY RESEARCH ISSN: 2277-3630 Impact factor: 8.036, 14(04), 51-65.
20. Ganti, V. K. A. T. (2025). Federated edge computing for privacy-preserving analytics in healthcare and IoT systems. International Journal of Computer Trends and Technology.
21. Kannan, S. (2025). Transforming community engagement with generative ai: harnessing machine learning and neural networks for hunger alleviation and global food security. Cuestiones de Fisioterapia, 54(4), 953-963.
22. Priyanka, R. P., Annapareddy, V. N., Yenugu, B. C., Nagabhyru, K. C., & Kapila, D. (2025, September). Optimization of Battery Management Systems Using Machine Learning. In 2025 International Conference on Computing and Communications (COMPUTINGCON) (pp. 1-6). IEEE.
23. Recharla, M., Chakilam, C., Kannan, S., Nuka, S. T., & Suura, S. R. (2025). Revolutionizing healthcare with generative AI: enhancing patient care, disease research, and early intervention strategies. American Journal of Psychiatric Rehabilitation, 28(1), 98-111.
24. Sadhukhan, S. K., Jaishankar, R., Adusupalli, B., Revathi, V. M., Kaliappan, S., & Ponnusamy, M. (2025, October). Healthcare Insurance Fraud Detection Using a Blockchain-Deep Neural Network-AdaBoost Hybrid Model. In 2025 2nd Asian Conference on Intelligent Technologies (ACOIT) (pp. 1-6). IEEE.
25. Krishnan, M., Aitha, A. R., Amistapuram, K., Nandan, B. P., Kaulwar, P. K., & Singireddy, J. (2025, November). Human-in-the-Loop Hybrid Neuro-Symbolic AI Model for Reliable Data Engineering in High-Stakes Industrial Systems. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-7). IEEE.
26. Singireddy, S., Pandiri, L., Motamary, S., Kummari, D. N., Somu, B., & Lakkarasu, P. (2025, August). Blockchain-Powered Claims Validation for Enhancing Trust in Health and Auto Insurance Ecosystems. In International Conference on Artificial Intelligence: Theory and Applications (pp. 26-38). Cham: Springer Nature Switzerland.
27. Burugulla, J. K. R., Kannan, S., Malempati, M., Challa, H. A. H. S. R., Pandugula, C., & Goma, T. (2025, April). Federated Learning Based Cloud Computing Solutions with Privacy Preservation for Intelligent Utilities. In 2025 4th



OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.

28. Vadisetty, R., Kannan, S., Vankayalapati, R. K., Sriram, H. K., Chakilam, C., & Malempati, M. (2025, February). AI-Driven Analysis of Urban Heat Islands for Sustainable City Planning. In International Ethical Hacking Conference (pp. 131-147). Singapore: Springer Nature Singapore.
29. Chabokpour, J., Kalisetty, S., Malempati, M., Challa, K., Mandala, V., Kumar, B., & Azamathulla, H. M. (2025). Integrating Hydrological and Hydraulic Approaches for Adaptive Environmental Flow Management: A Multi-Method Approach for Adaptive River Management in Semi-Arid Regions. *Water*, 17(7), 926.
30. Maguluri, K. K., Chandrasekar, M., Gorshetwar, S. D., Prasanna, S., Anusuya, M., & Vallarasu, K. (2025, February). Applications and advancements of nanoparticle-based drug delivery in lung cancer treatment. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
31. Vadisetty, R., Nuka, S. T., Kalisetty, S., Pandugula, C., Burugulla, J. K. R., & Annapareddy, V. N. (2025, February). Generative AI for Advanced Recycling Processes in Polyethylene and Polypropylene Manufacturing. In International Ethical Hacking Conference (pp. 269-284). Singapore: Springer Nature Singapore.
32. Mohan, A. A., Maguluri, K. K., Yasmeen, Z., & Pandugula, C. (2025). Evaluating the structural performance of axially loaded stainless steel circular tubular columns with ultra-high-performance concrete using an optimized attention pyramid convolutional neural network. *Structural Concrete*.
33. Recharla, M., & Nuka, S. T. (2025). Translational Approaches To Commercializing Neurodegenerative Therapies: Bridging Laboratory Research With Clinical Practice. *South Eastern European Journal of Public Health*, 121-144.
34. Consumer, R. E. P. (2025, May). Artificial Intelligence Infused Real Time Interaction Systems for Next Generation Retail Experiences Personalized Consumer Insights and Augmented Shopping Journeys. In Proceedings of the International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 24) (p. 208). Springer Nature.
35. Kannan, S., Raghavendra, C., Kumar, J. S., Narayanasamy, S., Balam, A., & Aravindh, S. (2025, May). Edge Computing and Artificial Intelligence Powered Dynamic Environmental Conservation Strategies for Sustainable Ecosystem Management and Predictive Climate Analysis. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 477-490). Atlantis Press.
36. Manikandan, K., Pamisetty, V., Challa, S. R., Komaragiri, V. B., Challa, K., & Chava, K. (2025). Scalability and efficiency in distributed big data architectures: a comparative study. *Metallurgical and Materials Engineering*, 31(3), 40-49.
37. Challa, K., Challa, S. R., Pamisetty, A., Kaulwar, P. K., & Koppolu, H. K. R. (2025, December). Transforming Payments: The Role of AI and Big Data in Fraud Alerts, Credit Monitoring, and Secure Transactions. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 1406-1414). IEEE.
38. Amistapuram, K., Pandiri, L., Raju, V. R., Paleti, S., Singireddy, S., & Sheelam, G. K. (2025, December). AI-Based Cloud Infrastructure and MLOps Frameworks for Scalable Data Engineering Across Banking and Insurance. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 186-192). IEEE.
39. Narasareddy Annapareddy, V., Pamisetty, A., Malempati, M., Kaulwar, P. K., & Bhardwaj Komaragiri, V. (2025, April). Enhancing Solar Power System Efficiency Through AI-Driven Predictive Maintenance and Cloud-Based Infrastructure Stability Solutions. In International Conference on Smart Computing and Informatics (pp. 328-337). Cham: Springer Nature Switzerland.
40. Koppolu, H. K. R., Munnangi, A. S. M., Reddy, M. B., & Ravulapalli, L. T. (2025). Artificial Intelligence-Based Evaluation of ApoA-I's Anti-Inflammatory and Anti-Fibrotic Roles in Experimental Pulmonary Injury. *Vascular and Endovascular Review*, 8(16s), 299-313.
41. Garapati, R. S., Adusupalli, B., Kaulwar, P. K., Gadi, A. L., Annapareddy, V. N., & Challa, K. (2025, December). The Evolution of Digital Payments: A Study on AI-Powered Transaction Monitoring Systems. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-8). IEEE.
42. Dodda, A., Polineni, T. N. S., Yasmeen, Z., Vankayalapati, R. K., & Ganti, V. K. A. T. (2025, January). Inclusive and Transparent Loan Prediction: A Cost-Sensitive Stacking Model for Financial Analytics. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 749-754). IEEE.
43. Mashetty, S., & Valiki, S. (2025). AI-Powered Data Engineering for Intelligent Retail Stock Management. *International Journal of Research and Applied Innovations*, 8(6), 13096-13109.



44. Kumar, B. H., Nuka, S. T., Recharla, M., Chakilam, C., Suura, S. R., & Pandugula, C. (2025, July). Addressing Ethical Challenges in AI-Driven Health Predictions. In 2025 2nd International Conference on Computing and Data Science (ICCDs) (pp. 1-6). IEEE.
45. Polineni, T. N. S., Yasmeen, Z., Ganti, V. K. A. T., Nirvin, P., & Vignesh, P. (2025, February). Climate Change Perception and Adaptation Strategies for Crop Growth with CNN and BiGRU Models. In 2025 3rd International Conference on Integrated Circuits and Communication Systems (ICICACS) (pp. 1-6). IEEE.
46. Nagamani, A., Nuka, S. T., Iyyanar, P., Suzana, A. A., Nayak, M., & Kumar, S. S. (2025, May). Artificial Intelligence optimized citizen engagement platforms for resilient urban futures inclusive policy making and smart community driven governance. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 221-231). Atlantis Press.
47. Adusupalli, B., Malempati, M., Paleti, S., Mashetty, S., & Singireddy, J. (2025). Integrated financial ecosystems: AI-driven innovations in taxation, insurance, mortgage analytics, and community investment through cloud, big data, and advanced data engineering. *Journal of Information Systems Engineering and Management*, 10, 1103-1117.
48. Endla, P., Kannan, S., Adhiyaman, M., Ramulu, K. P., Kumar, B. B. C., & Hemamalini, M. (2025, May). Sustainable Artificial Intelligence Orchestrated Urban Planning Models with Digital Twin Technology for Smart Cities and Green Development. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 1044-1058). Atlantis Press.
49. Koppolu, H. K. R., Gadi, A. L., Motamary, S., Dodda, A., & Suura, S. R. (2025). Dynamic Orchestration of Data Pipelines via Agentic AI: Adaptive Resource Allocation and Workflow Optimization in Cloud-Native Analytics Platforms. *Metallurgical and Materials Engineering*, 31(4), 625-637.
50. Sirisha, N., Kumar, D. S., Ilanchezhian, P., Komaragiri, V. B., Sri, G. K., & Thamizhkani, B. (2025, May). Energy Aware Artificial Intelligence Models for Scalable Digital Interactions Efficient Resource Management and Carbon Neutral Sustainable Technology Innovations. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 576-586). Atlantis Press.
51. Ande, R., Pandugula, C., Mehta, D., Vankayalapati, R., Birbal, P., Verma, S., ... & Nanavati, N. (2025). Understanding climate change impacts on streamflow by using machine learning: case study of Godavari Basin. *Water*, 17(8), 1171.
52. Sheelam, G. K., Meda, R., Pamisetty, A., Nuka, S. T., & Sriram, H. K. (2025). Semantic Negotiation Among Autonomous AI Agents: Enabling Real-Time Decision Markets for Big Data-Driven Financial Ecosystems. *Metallurgical and Materials Engineering*, 31(4), 587-598.
53. Ceasar, M. J., Karthik, N. L., Paleti, S., & Devaru, S. D. B. Entrepreneurial risk-taking and cultural values: A global behavioral perspective.
54. Annapareddy, V. N. (2025). The Intersection of Big Data, Cybersecurity, and ERP Systems: A Deep Learning Perspective. *Journal of Artificial Intelligence and Big Data Disciplines*, 2(1), 45-53.
55. Sheelam, G. K., & Komaragiri, V. B. (2025). Self-adaptive wireless communication: Leveraging ML and agentic AI in smart telecommunication networks. *Metallurgical and Materials Engineering*, 1381-1401.
56. Kumar, S. S., Garapati, R. S., Segireddy, A. R., Kalisetty, S., Inala, R., & Nagabhyru, K. C. (2026, March). Hybrid Deep Neural Network-DevOps Pipeline Optimization for Risk Prediction in Cloud-Native Workers' Compensation Platforms. In 2026 IEEE International Conference for Convergence in Computing Technology (I3CTCON) (pp. 1-6). IEEE.
57. Suura, S. R., Chava, K., Chakilam, C., Nuka, S. T., Maguluri, K. K., & Goma, T. (2025, April). Blockchain-Based Secure and Scalable Models for Healthcare Network Traffic Monitoring and Optimization. In 2025 4th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1-6). IEEE.
58. Challa, S. R., Kaulwar, P. K., Koppolu, H. K. R., Adusupalli, B., & Suura, S. R. (2025, April). Big Data and AI in Wealth Management: Leveraging Cloud Computing for Secure and Scalable Financial Infrastructure. In International Conference on Smart Computing and Informatics (pp. 307-318). Cham: Springer Nature Switzerland.
59. Sriram, H. K., Gadi, A. L., Challa, K., & Singreddy, S. (2025). Leveraging AI, ML, and Gen AI in Automotive and Financial Services: Data-Driven Approaches to Insurance, Payments, Identity Protection, and Sustainable Innovation. Anil Lokesh and Challa, Kishore and singreddy, Sneha, Leveraging AI, ML, and Gen AI in Automotive and Financial Services: Data-Driven Approaches to Insurance, Payments, Identity Protection, and Sustainable Innovation (March 25, 2025).



60. Challa, K., Singireddy, J., Pamisetty, A., Garapati, R. S., Kannan, S., & Sriram, H. K. (2025, December). Harnessing Agentic AI and IT Infrastructure in Banking to Drive Consumer Insights, Operational Excellence, and Intelligent Financial Innovation. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-7). IEEE.
61. Kumar, B. H., Nuka, S. T., Malempati, M., Sriram, H. K., Mashetty, S., & Kannan, S. (2025). Big data in cybersecurity: Enhancing threat detection with ai and ml. *Metallurgical and Materials Engineering*, 31(3), 12-20.
62. Komaragiri, V. B. (2025). Agentic ai for autonomous network orchestration: A new frontier in telecommunications. *eksplorium-buletin pusat teknologi bahan galian nuklir*, 46(1), 221-241.
63. Paleti, S., Kummari, D. N., Garapati, R. S., Sheelam, G. K., Adusupalli, B., & Singireddy, J. (2025, December). Building a Cyber-Resilient Payment Infrastructure: Transforming Payment Security with Zero Trust Architecture. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-7). IEEE.
64. Maguluri, K. K., Ganti, V. K. A. T., Yasmeen, Z., & Pandugula, C. (2025, January). Progressive GAN framework for realistic chest X-ray synthesis and data augmentation. In 2025 6th International Conference on Mobile Computing and Sustainable Informatics (ICMCSI) (pp. 755-760). IEEE.
65. Chakraborty, S., Pamisetty, A., Chandana, N., & CS, B. (2025, October). Depth-Wise Temporal Convolutional Networks with Layer Normalization for Waste Food Prediction. In 2025 2nd International Conference on Software, Systems and Information Technology (SSITCON) (pp. 1-6). IEEE.
66. Mahadevan, S. (2024). Clouding the Future: Innovating Towards Net-Zero Emissions. *International Journal of Computing and Engineering*, 6(2), 17-23.
67. Pamisetty, A. (2025). The Psychology of Addiction: Understanding Behavioral Patterns. *SCIENTIFIC CULTURE*, 11(1), 1-15.
68. Paleti, S., Kummari, D. N., Garapati, R. S., Sheelam, G. K., Adusupalli, B., & Singireddy, J. (2025, December). Building a Cyber-Resilient Payment Infrastructure: Transforming Payment Security with Zero Trust Architecture. In 2025 3rd International Conference on IoT, Communication and Automation Technology (ICICAT) (pp. 1-7). IEEE.
69. Chakilam, C., Kannan, S., Recharla, M., Suura, S. R., & Nuka, S. T. (2025). The impact of big data and cloud computing on genetic testing and reproductive health management. *American Journal of Psychiatric Rehabilitation*, 28(1), 62-72.
70. Annapareddy, V. N., Kollam, M., Challa, K., Vankayalapati, R. K., Malempati, M., & Yasmeen, Z. (2025). Building resilience in water resource management using artificial intelligence and interdisciplinary strategies for sustainable water security and future governance. *Discover Water*, 6(1), 7.
71. Rani, P. S., Amistapuram, K., Pamisetty, V., Singireddy, S., Kummari, D. N., & Sheelam, G. K. (2025, November). Hybrid Knowledge Graph–Deep Learning Framework for Automated Exception Handling and Investigation in Complex Insurance Claims. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.
72. Suura, S. R., Chava, K., Chakilam, C., Malempati, M., & Nuka, S. T. (2025, April). Revolutionizing Retail Supply Chains with AI and Cloud Computing: A Big Data Approach to Demand Forecasting and Logistics Optimization. In *International Conference on Smart Computing and Informatics* (pp. 271-283). Cham: Springer Nature Switzerland.
73. Kumar, M. V. K., Kannan, S., Annapareddy, V. N., Adusupalli, B., Paleti, S., & Challa, S. R. (2025, July). Transforming Underground Electric Cable Management with AI in Smart Cities. In 2025 2nd International Conference on Computing and Data Science (ICCDs) (pp. 1-6). IEEE. .
74. Mahadevan, S. (2024). Empowering Manufacturing: Generative AI Revolutionizes ERP Application. *Int. J. Innov. Sci. Res*, 9, 593-595.
75. Chava, K., Challa, S. R., Sriram, H. K., Chakilam, C., Kannan, S., & Annapareddy, V. N. (2025, July). Utilizing Query Performance in NoSQL Databases for Applications Based on Big Data. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS) (pp. 1-6). IEEE.