International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

AI-Augmented Data Engineering for Real-Time Fraud Detection in Digital Banking

Nareddy Abhireddy

Independent Researcher, India

ABSTRACT: Real-time fraud detection in digital banking is innefectual for the majority of fraud transactions. Detection latencies often exceed fraud initiation delays, data sources and processing framework are limited, scalability to transaction volume is lacking, drift of fraud patterns is inadequately addressed, class imbalance in training data is prevalent and data quality issues are not addressed comprehensively in supporting work. A landscape of real-time fraud detection capabilities across the digital banking sector outlines these shortcomings and a data engineering view identifies a set of enabling foundations and an end-to-end fraud detection system architecture. Data ingestion and streaming architectures, data quality processes and data provenance controls provide the bedrock of a real-time data analytics capability. The implementation of supervised learning is supported with label generation, feature engineering, model selection and transfer learning while unsupervised anomaly detection is deployed using clustering, isolation and one-class models. The underlying streaming framework enables seamless combination of multiple detectors in a hybrid model. Scalability at the fraud detection model level is solved with a data pipeline supporting multiple detection models in parallel.

A data engineering and machine learning system design is presented, offering insights into the implementation and deployment of AI-augmented data engineering for real-time fraud detection. Consideration of AI-augmented data engineering highlights the quality and completeness of data, and alignment with data governance, access controls and data privacy legislation and objectives, in the delivery of a fast, accurate and easily understood model suitable for deployment into a governed data environment for streaming processing.

KEYWORDS: Real-Time Fraud Detection in Digital Banking, Streaming Data Architectures, AI-Augmented Data Engineering, Fraud Pattern Drift Detection, Class Imbalance Mitigation, Hybrid Fraud Detection Models, Supervised Learning for Fraud Analytics, Unsupervised Anomaly Detection, Isolation and One-Class Models, Feature Engineering for Transactions, Label Generation Pipelines, Scalable Fraud Detection Systems, Low-Latency Transaction Monitoring, Data Quality and Provenance Controls, Transfer Learning in Fraud Models, Parallel Model Deployment Pipelines, Governed Data Environments, Privacy-Compliant Fraud Analytics, Event-Driven Banking Security, End-to-End Fraud Detection Architecture.

I. INTRODUCTION

Fashioned as a scarcely elaborated view in its first Referee Report, this study maps the need for AI-augmented data engineering for end-to-end real-time fraud detection in digital banking. It explores the current state of the art in real-time fraud detection. A landscape view including latencies, data sources, and support for detection is presented, along with a somewhat informal view of the feeding of models but without a model-monitoring dimension.

Scant attention is paid to data engineering—the cleaning, filtering, preparation, management, and creation of data needed to sustain real-time fraud detection. The resulting foundations of data engineering are hence rather simple. Limiting the aspect does create space for readability and flow; indeed, every aspect is friendly and enjoyable to read. More, the expressed target audience is AI-enabled data engineers not AI-enabled fraud analysts, and the survey/scoping of data engineering is justified. Four potential gaps also emerge: indeed, data engineering needs to be more than simple, especially for AI-enabled fraud detection; supporting signals for supervised and semi-supervised models requires careful development over time and general rather than specific datasets; signals for unsupervised and one-class models require especially dedicated thought; and AI-enabled detection needs to be reinforced with careful model deployment, monitoring, and response capabilities for AI-enabled data engineering to be truly focal. The paper concludes with a clear tabular summary of the take-home messages.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

A. Overview of the Study

Current capabilities for real-time fraud detection in digital banking are map–ping onto increasing volumes of service traffic generated by the ever-growing customer base. Retail banks detect fraudulent activities such as card payments and fund transfers using combinations of supervised and anomaly detection models based on bank-specific transaction data and other data sources such as biometrics and device geolocation at web or app service levels. Latencies associated with source data and data-processing pipelines limit real-time detection to fraud transactional signaled by supervised learning models. High model-management costs and low development productivity characterise model-management processes. Latency, drift, heterogenicity, and data-governance-related issues constrain the scalability of state-of-the-art real-time fraud-detection approaches. Industry practices and publicly available literature point to gaps in the above detection stack and suggest a landscape for AI-augmented data engineering.

AI-augmented data engineering constitutes the data-science-application-to-business-value journey supported by datasets prepared by the data-engineering discipline. Data engineering in the data-science context ensures that enterprise data for supervised learning models are processed with biases reduced to the required limits. Fraud-detection model-management-related processes usually have low productivity and high costs. AI-augmented data engineering provisions baseline data pipelines and completes full analytics workflow for fraud-detection models. AI-augmented data engineering thus dramatically improves the performance of non-technical resources when developing, deploying, and managing fraud-detection models.

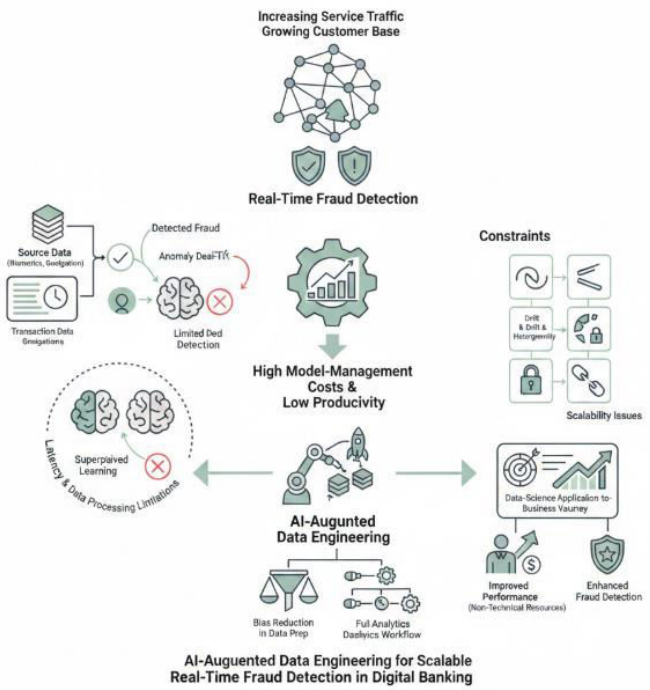

Fig 1: Scalable Vigilance: Leveraging AI-Augmented Data Engineering to Overcome Latency and Governance Constraints in Real-Time Banking Fraud Detection

II. LANDSCAPE OF REAL-TIME FRAUD DETECTION IN DIGITAL BANKING

Real-time fraud detection promises rapid prevention and mitigation of in-band incursion attempts, yet only a minority of attacks are recognized as they step into the detection range. Rigorous analysis of detection capabilities and limitations reveals that current detection latencies and the resulting time window for adversaries remain considerable, and crucial timescale correlations between consumer behavior and fraud characteristics are yet to be exploited. Furthermore, standard industry practice still matches time-sourced occurrence logs to historical transaction datasets, rather than applying an appropriate streaming architecture to event data sources. Nevertheless, detection systems exhibit widening scalability and drift-related weaknesses, are increasingly fed by heterogenous and unthrottled data sources, and lack a comprehensive industry-wide measurement framework suited for construction or refinement of real-time systems. Moreover, a widening base of research activity, investigations, and solutions is being confronted by adversaries who strategize and act in real time. These considerations suggest a period of grace, yet leave little doubt

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

that the unique operational requirements of real-time fraud detection, applied to real-time data engineering and analytics, represent fertile ground for technologists and researchers seeking to solve real-time detection design, implementation, operational, and benchmarking problems.

Both technology- and business-oriented sources confirm that real-time fraud activity now frequently emerges from unexpected sources, and often entails clear transfer of knowledge between overlay organizations. As a result, the simplest or least resource-demanding fraud detections are typically the quickest to develop and deploy, and are therefore commonly put in place first. Detection operation in banking organizations more advanced in their attempts at covering activity across a broader current service setting has exhibited stronger indications of advancing drift-related weaknesses, likely reflecting widening scalability challenges.

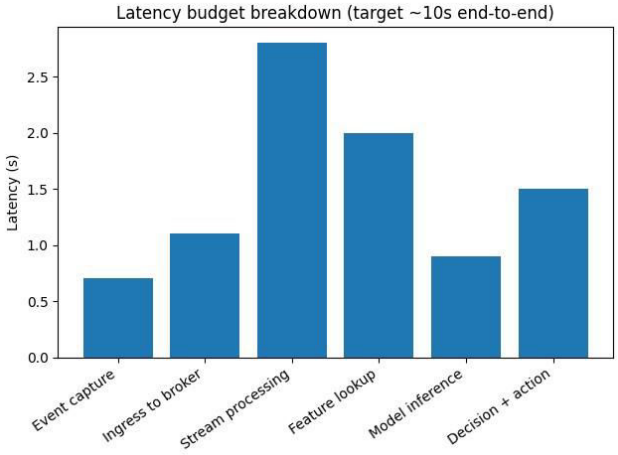

Equation 1) Latency equations for "real-time" streaming fraud detection

1.1 End-to-end latency budget (sum of stage latencies)

Let one transaction/event travel through stages:

- capture/telemetry creation
- ingestion to broker (Kafka/Kinesis)
- stream processing
- feature lookup / enrichment
- model inference
- decision + action (block/step-up auth/alert)

Define stage latencies: $L_1, L_2, \ldots, L_k$ (seconds)

Step 1: Total latency is the sum:

$$L_{total} = \sum_{i=1}^{k} L_i$$

Step 2: Real-time requirement (paper's target):

$$L_{total} \leq 10 \text{ s}$$

("latency target of ten seconds")

Step 3: If you want "share of latency" per stage (useful for performance engineering):

$$\text{Share}_i = \frac{L_i}{\sum_{j=1}^{k} L_j} \times 100\%$$

A. Current Trends and Challenges in Real-Time Fraud Detection

Real-time fraud detection in digital banking currently supports urgent users requests and transaction limitation. Four data sources undergo streaming processing, with a latency target of ten seconds. Scalability, concept drift, heterogeneous signal provenance, and adversarial tactics form critical analytical challenges. However, cross-industry capabilities remain underreported, and support for AI-augmented data engineering is limited. A data foundation for AI applications in real-time fraud detection encompasses ingestion and streaming architectures, data quality considerations, privacy requirements, and compliance with data governance principles.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

Current capabilities and challenges in digital-banking fraud detection highlight trends such as the increased use of streaming-processing architectures and data provenance and quality for real-time analytics. Support for AI-augmented data engineering, however, remains limited, with requirements seldom associated with comprehensive collections of capabilities and challenges. Recent investigations into a critical AI application reveal open questions in these areas. Nevertheless, urgent user requests and transaction limitation for real-time fraud detection justify significant interest. Four data sources — internal mobile-application events, user-context metadata, users' past behaviours, and click patterns generated by all users — are processed via streaming architectures, with a latency target of ten seconds.

Stage	Latency s	Cumulative s	Share %
Event capture	0.7	0.7	7.8
Ingress to broker	1.1	1.8	12.2
Stream processing	2.8	4.6	31.1
Feature lookup	2.0	6.6	22.2
Model inference	0.9	7.5	10.0
Decision + action	1.5	9.0	16.7

III. DATA ENGINEERING FOUNDATIONS FOR REAL-TIME ANALYTICS

Data ingestion and streaming architectures support AI-augmented fraud detection through rich telemetry signals at scale. A modern real-time pipeline relies on a streaming platform such as Apache Kafka or AWS Kinesis. Events usually carry a schema defined in an open-source project such as Apache Avro, Apache Thrift, or Google Protocol Buffers. Backpressure control and the ability to read a topic from any point in time or in a different order are prerequisites for model development.

Main sources of data required for real-time fraud detection include transaction records, external threat intelligence feeds, and system alerts. Data-driven models explore these data sets by means of supervised supervised machine learning strategies. These approaches rely on known fraud cases during training and learn to discriminate other types of events by means of feature engineering. For cases in which labels are hard or expensive to generate, semi-supervised learning prepares models with autoencoders or generative adversarial networks, for subsequent training with few labels. Known fraud patterns are explored with standard anomaly detection techniques, including clustering, isolation forests, and one-class classifications. A hybrid strategy combines the strengths and weaknesses of both worlds. Data-mining techniques ensure no fundamental fraud life cycle is unknown to the AI ecosystem.

Rich telemetry signals, generated with very small overhead from the perspective of the banking institution, form an essential part of the AI underpinning. Examples of these telemetry signals are account logins, logins to a particular service in the banking site (such as adding new recipients), money transfers of different amounts between the same two institutions, ATM accesses, among other services exposed by the banking institution (and its apps).

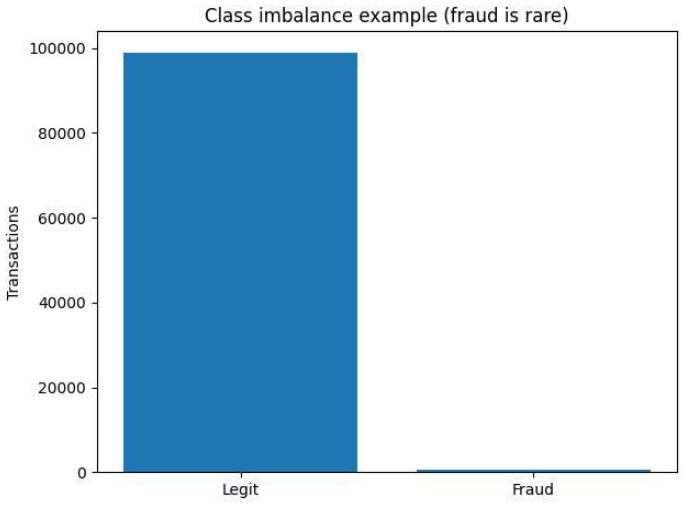

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

Equation 2) Throughput, backpressure, and streaming stability

2.1 Queue stability (arrival vs service rate)

Let:

- arrival rate λ events/sec
- service rate μ events/sec (pipeline capacity)

Step 1: If arrivals exceed service, backlog grows:

$$\lambda > \mu \Rightarrow \text{queue grows without bound}$$

Step 2: To stay stable (no runaway lag):

$$\lambda < \mu$$

Step 3: Utilization:

$$\rho = \frac{\lambda}{\mu}$$

You typically want ρ comfortably below 1 to absorb spikes (fraud waves, marketing campaigns, outages).

2.2 Processing delay from backlog (simple approximation)

If backlog size is B events and capacity is μ events/sec, then:

$$\text{Delay} \approx \frac{B}{\mu}$$

A. Data Ingestion and Streaming Architectures

Event streaming platforms are now established as the backbone of organizations' real-time data ingestion and processing needs, and they are increasingly seen as a key enabler of AI-driven applications. Most platforms offer support for publishing and subscribing to data feeds using a push communication pattern, and the integration of a wide range of producers and consumers with an increasing set of connectors and APIs that enable seamless exchanges continues to expand the accessibility and ease of use. By reducing the amount of scheduling done within the platforms, they enable unprecedented sourcing of AI signals along the Data to Decisions continuum.

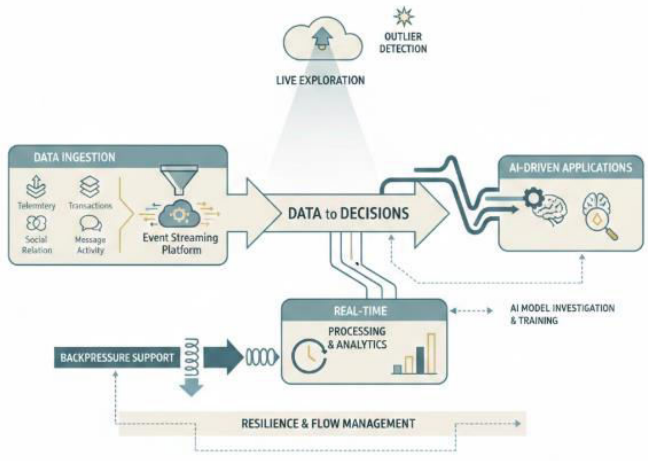

Fig 2: Resilient Event-Driven Architectures: Integrating Cross-Level Backpressure and Live Inspection in AI-Augmented Data-to-Decisions Pipelines

Although organizations are taking increasing care with the design of their event schemas and the events that are produced, the real-time nature and volume of the data can complicate their resilience. Backpressure support at the stream processing and producer layers is thus beginning to be seen as a feature that needs to be supported and is being increasingly sought in edge-to-cloud pipelines, where the connected processing element(s) can react in near real time to avoid flow management bottlenecks. However, full cross-level support for backpressure management remains an area of research, as does the ability to ensure that pipelines supporting real-time processing are not forced to explicitly incorporate resilience and flow management considerations.

Real-time fraud detection and prevention demands an ever-expanding Data-to-Decisions pipeline, with the incoming feeds spanning telemetry, transactions, behaviour, social relation, and message activity types. Nevertheless, the well-established monitoring and alarm systems associated with these Data-to-Decisions pipelines have tended to be quite narrowly focused. The result is that there is now a sustained demand for live exploration of the established flows,

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

arrested focus, and near-real-time inspection of outliers that can drive AI model investigation, training, and deployment.

B. Data Quality and Provenance

Data quality and provenance considerations are critical for maintaining trustworthiness in any production analytics framework. For fraud detection in digital banking, data quality typically encompasses aspects such as validation (data should be the correct value type), completeness (no missing values), timeliness (data arrived within expected time intervals), and governance (data were collected in compliance with legal guidelines and institutional policies) and these aspects present strong relationships with privacy, access control, audit trail, and likely regulatory compliance.

Provenance tracking must cope with a heterogeneous environment, where data are generated, processed, and stored in different systems. Indeed, the fraud detection process spans multiple actors, ranging from data source custodians, who provide honest real-time data, to potential attackers, who might purposefully inject corrupt or malicious data into the process. The lineage and provenance of time series data are usually hard to track, since up to date information about the business processes supporting the data production might not be stored anymore, given the volumes and the velocity of the produced data.

In the absence of sound governance of time series data and, especially, of sampling of the incoming data streams, issues such as drift in non-honest producers and even data-mining and back-testing attacks can be expected. Monitoring solutions must then cover all the potential mechanisms for assuring data quality and governance of the events in the detection process. An orchestration engine managing the whole detection process from data ingestion to model monitoring is a good-enough solution at this stage, bringing together monitoring frameworks for data ingestion, for model and feature stores, and for data quality assessment along the models.

IV. AI-DRIVEN DETECTION MODELS

Machine learning and data mining methods are widely applied to fraud detection to replace or complement the work of expert analysts. Supervised fraud detection models can learn from labeled historical incidents; however, labeled instances usually comprise only a small fraction of the overall data. A limited supply of fraud samples incurs additional costs and delays in recommissioning the model after it has been adapted. Supervised learning approaches are thus often semi-supervised, using horizontal or vertical transfer learning across jurisdictions or between transaction channels provided by the same organization.

Real-time fraud detection systems should also support the identification of new or rapidly evolving types of fraud activity for which there are few or no labeled examples available at training time. Consequently, unlabeled data are frequently used for clustering, dimensionality reduction, noise identification, monitoring, novelty detection, or other forms of unsupervised anomaly detection. The organization’s response to detected anomalies is usually guided by a separate expert analysis. Ideally, however, the detected anomalies should help create a more comprehensive solution that detects emerging fraud patterns with minimal delay and without requiring new sets of feature engineering, model development, and recommissioning from scratch.

Supervised Detection Models

The feature set applied to learning must be sufficiently rich to distinguish between the attacked class and the background class, thereby supporting decision-making across plausible scenarios. A variety of strategies helps to enhance its richness. Additional labeled datasets are usually employed for the “background” class, balancing the number of samples across the two classes, or boosting the model to favor lower false alarm rates. Data augmentation techniques can be applied, such as replacing, altering, or injecting transferred samples that resemble known fraud signals.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

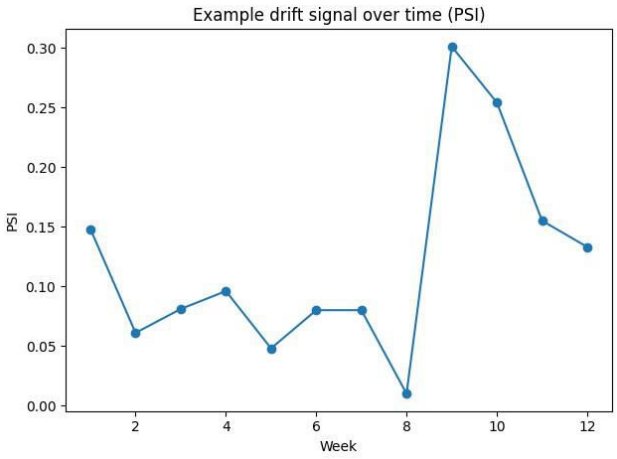

Equation 3) Supervised fraud detection equations (confusion matrix → metrics)

3.1 Confusion matrix definitions

For binary fraud classification:

- True Positive (TP): fraud correctly flagged
- False Positive (FP): legit incorrectly flagged
- True Negative (TN): legit correctly allowed
- False Negative (FN): fraud missed

3.2 Precision, Recall, F1 (derived step-by-step)

Precision = "of what I flagged, how many are truly fraud?"

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

Recall (TPR) = "of all fraud, how many did I catch?"

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

F1 is the harmonic mean:

Step 1: harmonic mean definition:

$$\text{H} = \frac{2\text{ab}}{\text{a} + \text{b}}$$

Step 2: substitute a = Precision, b = Recall:

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

False Positive Rate (FPR) (important in banking UX cost):

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

A. Supervised and Semi-Supervised Approaches

Detecting patterns, such as transaction fraud, cyberattacks, and account takeover, often lends itself to supervised approaches, especially where labelled examples are readily available. Nevertheless, labelled data is scarce for many categories and time regions and fraud signals frequently diverge from past patterns. Even in drift-prone tasks, the browsing behaviour of legitimate customers usually replicates past patterns, providing an essential feature engineering cue. Transfer learning facilitates knowledge transfer between source-and-target classification problems in a model-agnostic manner. Where no labelled examples are present, techniques like label propagation and co-training help alleviate the supervision requirement by leveraging auxiliary or related datasets that contain labelled instances. Semi-supervised approaches, which leverage shallow self-supervised objectives, are now applicable to most model types.

Supporting a multi-pronged labelling regime promotes not only the detection of as many categories as possible, but also the establishment of a natural-action policy for different types of fraud signals. For instance, remote access and remote-troubleshooting fraud attacks could be placed under the purview of a Contact Centre department, with call centre agents trained to proactively triage and authenticate such calls for the bank's protection and that of their customers.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

Model	TP	FP	TN
Supervised (baseline)	420	580	98500
Hybrid (supervised + anomaly)	520	510	98570

B. Anomaly Detection and Unsupervised Techniques

Clustering, isolation, and one-class methods support the detection of previously unknown fraud patterns; unsupervised-generation of labels allows hybrid strategies. The growing sophistication of digital fraud accelerates the adoption of and investments in real-time fraud-detection services capable of automatically inferring data-driven analytical models. Such systems are intended to continuously enhance the detection performance against both known and newly emerging fraud schemes exploiting the bank's digital channels and services. Semi-supervised or unsupervised techniques find application where the presence of fraud signals in the sample space is low or unknown, yet labelled data is an economically unviable commodity.

Unsupervised techniques designed to identify typical user distribution and behavior leverage clustering methods such as k-means, DBSCAN, Ward's method, and Hierarchical Clustering for pattern recognition of normal users and detection of outliers. Albanese proposed two methods for the temporal analysis of clusters based on their evolution over time. Clustering integration with isolation forests, through the use of cluster membership tags, improves the isolation approach's ability to suppress classification noise and augment interpretability. Anomaly detection in employee transactions—such as sudden increases of withdrawal or deposit operations—is achieved by first defining the normality domain through clustering methods and then applying isolation forests. One-class classification methods learn the normal transaction distribution of a unique target class, allowing detection of samples that deviate from the learned normal regions. Hybrid approaches reduce the need for labelled samples, generate unsupervised labels for known normal regions, and exploit interpretability of clustering techniques combining them with random forests, isolation forests, or one-class approaches.

V. DATA GOVERNANCE, PRIVACY, AND COMPLIANCE

Real-time fraud detection can be supported with AI-augmented data engineering without forming an obstacle course for data scientists and developers. Learning-based models require data with predictive power. Therefore, data governances must guarantee a high degree of data quality and completeness to earn the required trust for producing reliable insights. Provenance and lineage tracking mechanisms enable organizations to track the origin and evolution of data, support debugging, and provide the necessary audit trails. Where real-time data capture incorporates sensitive personal information, technical implementations need to follow principles of data minimization and allow user consent management. Regulatory frameworks like GDPR or PCI DSS strengthen the foundations for responsible data management, protect the interests and privacy of customers, and govern the stored data in a trusted manner.

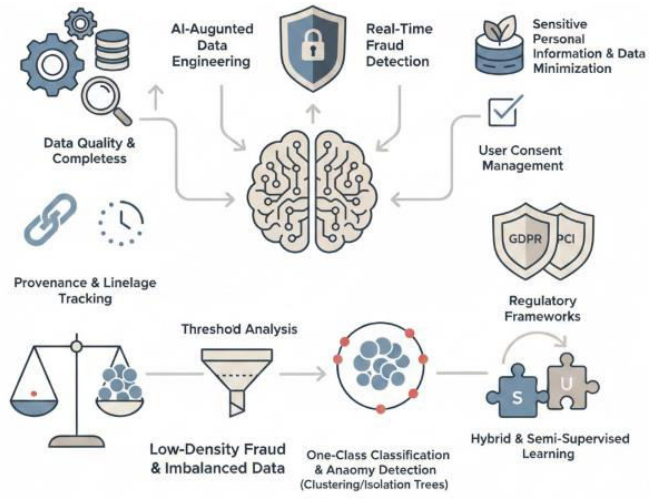

Fig 3: Trust-Centric AI Orchestration: Integrating Provenance, Privacy, and Hybrid Anomaly Detection in Real-Time Fraud Systems

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

Due to the low density of fraudulent transactions, labeling fraud cases for supervised learning is challenging. Labeling infrequent categories often requires a separate investigation. The distribution of identified fraud-related transactions must be further analysed by exploring the model's output for thresholds to separate malicious from benign traffic. One-class classification methods shape security models with successfully filtered benign categories and no further fraud occurrences. Traffic or records without known data leakage points detect malicious outliers. Anomaly detection techniques such as clustering or isolation trees identify unknown but deceptive instances as aggregates of similar set features. Hybrid and semi-supervised strategies assist in defining the non-fraud-related clusters and support unsupervised learning through supervised complement.

A. Ensuring Data Integrity and Compliance in Fraud Detection

Data governance must secure data access and minimize risk exposure without undermining detection performance. Data minimization constrains information collection to what is strictly required: too little fails to identify attacks; too much invites abuse and confusion. When feeding into supervised models, input data must be paired with the expected output; strips of bank accounts, credit cards, and other personally identifiable information (PII) ideally undergo an anonymization process to prevent their association with individual bank clients.

The information stored in a detection model serves to discriminate between regular and suspicious patterns in a minimal footprint by retaining only that which is essential for the classification task. Besides defence efforts aimed at redundancy elimination and access control, sensitive data should be encrypted for storage in data sources and traffic flows. Encryption adds an additional processing overhead to the data but secures it from being disclosed to third parties while not acting as a primary defence mechanism. Access control measures should limit data sharing and provide an audit trail of data usage in support of incident investigation and response. Sensitive and PII traffic must be analyzed by AI models in compliance with local regulations—such as the General Data Protection Regulation in Europe—to prevent misuse.

Equation 4) Class imbalance equations (why accuracy is misleading)

Let fraud prevalence be:

$$\pi = \frac{N_{fraud}}{N_{total}}$$

If π is tiny (typical banking), a dumb classifier that predicts "legit" always gets:

$$\text{Accuracy} = \frac{TN}{N_{total}} \approx 1 - \pi$$

…which looks "great" but catches **zero fraud**.

4.1 Cost-sensitive learning (banking-aligned objective)

Define costs:

- C_{FN}: cost of missing fraud (money loss, chargebacks, AML risk)
- C_{FP}: cost of false alarm (customer friction, churn, manual review cost)

Expected loss:

$$\mathbb{E}[\text{Loss}] = C_{FN} \cdot FN + C_{FP} \cdot FP$$

Class	Count	Share %
Legit	99000	99.398
Fraud	600	0.602

VI. SYSTEM ARCHITECTURE FOR END-TO-END FRAUD DETECTION

The vertical stack considered so far, comprising data ingestion, quality, AI-driven detection, and governance, provides a solid foundation for real-time fraud detection. Without a complete horizontal flow encompassing all necessary components, however, the capabilities remain insufficient for practical deployment. Section 6 focuses on the orchestration of end-to-end data pipelines and model management, synthesizing elements introduced previously to present a comprehensive detection workflow along with the required supporting frameworks.

The end-to-end workflow depicted in Figure 6 follows an event-driven conception. At the core lies the real-time processing of streaming data pipelines, where either new events or health checks periodically trigger the forwarding of telemetry data for model inference. The relevant detection models have been introduced, and the data engineered on the pertinent dimensions flows through the preparation-stage data pipelines. Real-life application requires addressing non-

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

trivial operational challenges along this processing layer, namely automation, robustness, fault-tolerance, and scalable throughput.

A. Data Pipelines and Orchestration

A preliminary, executable blueprint maps the end-to-end workflow required for real-time fraud detection, incorporating event-driven data ingestion from native and auxiliary use-case data sources, fraud detection model deployment, and orchestration by a containerized task management tool. The proposed approach integrates the processing of streaming and non-streaming task types into a unified process framework, where execution instances are automatically created in response to the generation of fraud detection training signals, and achieved CIA properties support responsible use for analytics and fraud detection. The preliminary design provides a useful point of reference to inform implementation, with additional details required to ensure production-grade deployment.

Fraud detection constitutes a pivotal but infrequently addressed data-analytic task in digital banking, with industry positions for real-time detection still uncommon. The growing affordability of cloud service offerings is facilitating breakthroughs in other data-analytic categories, including real-time fraud detection. The detection task itself combines data engineering and data science capabilities and quality requirements from across all five foundations of AI-augmented data engineering for responsible data-analytic decision-making. Packages of pre-defined streaming pipelines accelerate and simplify the production deployment of the underlying data engineering capabilities, supporting the scalability, observability, and low-latency response times necessary to mitigate fraud loss.

Equation 5) Decision threshold derivation (from probabilities to actions)

In deployment, the model often outputs $p = \Pr(\text{fraud} \mid x)$.

5.1 Optimal threshold with asymmetric costs

Choose "fraud" if expected fraud cost is lower than expected legit cost.

- If you predict "legit" but it's fraud → cost C_{FN}
- If you predict "fraud" but it's legit → cost C_{FP}

Step 1: Decide fraud when:

$$C_{FP} \cdot \Pr(\text{legit} \mid x) < C_{FN} \cdot \Pr(\text{fraud} \mid x)$$

Step 2: Use $\Pr(\text{legit} \mid x) = 1 - p$, $\Pr(\text{fraud} \mid x) = p$:

$$C_{FP}(1 - p) < C_{FN}p$$

Step 3: Expand and solve for p:

$$C_{FP} - C_{FP}p < C_{FN}p \; C_{FP} < p(C_{FN} + C_{FP}) \; p > \frac{C_{FP}}{C_{FN}+C_{FP}}$$

So the **optimal threshold** is:

$$\tau^* = \frac{C_{FP}}{C_{FN} + C_{FP}}$$

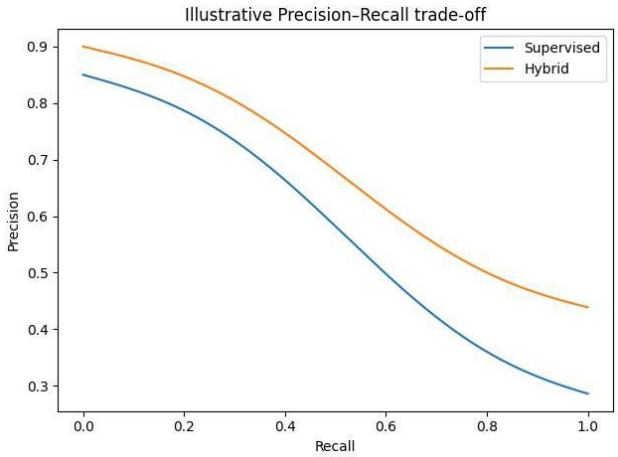

B. Model Management and Monitoring

Model deployment constitutes a separate facet of analytics pipelines, and specific tools are needed for maintaining and controlling deployed models. When systems support many teams developing and deploying models, an MLOps (machine learning operations) approach can be adopted, analogous to DevOps for software engineering. Managing

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

models that go into production involves handling multiple versions from different teams, monitoring models for data drift, and setting up incident management and resolution workflows. This process often requires a combination of dedicated tools and features within general-purpose orchestration platforms.

To build trust and facilitate maintenance, every deployed model should be versioned, and those serving critical or high-frequency prediction tasks must be monitored to detect changes in the underlying data distribution. Monitoring can include drift detection processes and downstream checks on prediction accuracy. As models act on real data, validating the results against known outcomes helps track their performance. To respond to detected problems, an incident response plan creates a feedback loop that routes the issue to specific teams for investigation and remediation. Evaluation can leverage general metrics for the model type (e.g., F1 score for classifiers) or assess specific failure cases (e.g., confusion matrix for security models).

VII. CONCLUSION

The landscape of data analytics in the digital banking sector is being transformed by real-time requirements and the promise of augmented intelligence, both in the engineering of data pipelines and in the application of detection and mitigation models. This work outlines the data-engineering considerations that must be in place to support AI-augmented systems capable of real-time detection of account takeover, card-not-present, money laundering, and transaction fraud risks; and it assesses these considerations against current industry practices and standards.

Despite the clear articulation of business needs, the space is still poorly aligned. There are gaps in both the real-time analytics capabilities of organizations and the depth of the augmented-intelligence models applied. The foundations for supporting end-to-end real-time analytics and producing RESTful APIs offering augmented-intelligence detection capabilities in a scalable and responsible manner are well understood. For many organizations, the implementation of complete end-to-end pipelines is an achievable target; for others, individual aspects such as improved ingestion processes, development of data-quality gates, or the establishment of an integrated, collaborative data-science platform represent significant steps toward closing the gap. The key elements deserving more stringent focus are the quality and completeness of the underlying moderation and AI models.

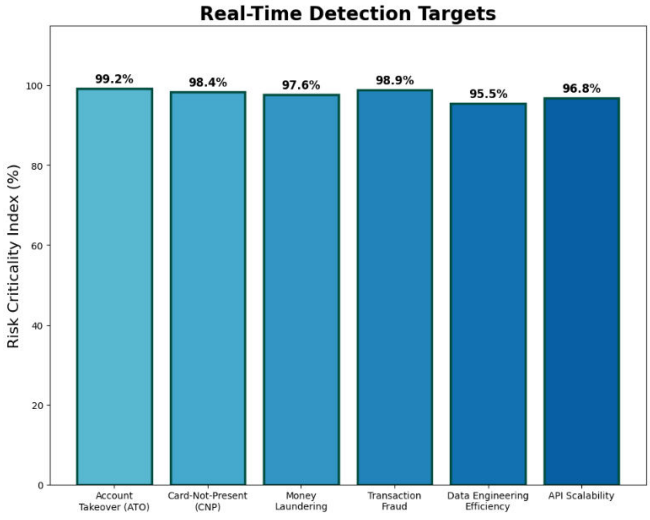

Fig 4: Real-Time Detection Targets

A. Key Takeaways and Future Directions

The main conclusion following the research on AI-augmented data engineering for real-time fraud detection in digital banking is that such technologies are essential to the enforcement of anti-fraud measures in this sector. However, the analysis highlighted gaps in the use of real-time data ingestion and analytics by larger banks and the near-complete absence of preparation for data drift detection and management. Furthermore, the incorporation of sophisticated changes in detection models has not been accorded the same level of attention. There is very little practical support to guide data engineering efforts that enable such model families to function in fraud-related contexts.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

In response to these problems, the next step for future research should be the elaboration of a data-engineering-based framework for generic real-time fraud detection in digital-banking operations that will cover all phases of the end-to-end detection cycle. Such a framework should be owned by the principal actors in the field — the data-engineering and analytics communities — but it is equally applicable to other domains. A complementary requirement is the collation of real-world data-science community datasets, collected and maintained to rigorous standards. Industry players must commit to enabling the creation and availability of labelled datasets for the training of hidden-fraud-event detection models so that the data-science community can support them more effectively.

REFERENCES

1. Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2623–2631.
2. Kolla, S. K. (2021). Architectural Frameworks for Large-Scale Electronic Health Record Data Platforms. Current Research in Public Health, 1(1), 1–19. Retrieved from https://www.scipublications.com/journal/index.php/crph/article/view/1372.
3. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., et al. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges. Information Fusion, 58, 82–115.
4. Davuluri, P. N. Integrating Artificial Intelligence into Event-Driven Financial Crime Compliance Platforms.
5. Awais, M., Li, Y., & Wang, H. (2022). Federated learning for healthcare informatics. IEEE Reviews in Biomedical Engineering, 15, 226–239.
6. Gottimukkala, V. R. R. (2023). Privacy-Preserving Machine Learning Models for Transaction Monitoring in Global Banking Networks. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 633-652.
7. Baier, H., & Mendling, J. (2022). Data governance in machine learning systems. Business & Information Systems Engineering, 64, 471–486.
8. Kushvanth Chowdary Nagabhyru. (2023). Accelerating Digital Transformation with AI Driven Data Engineering: Industry Case Studies from Cloud and IoT Domains. Educational Administration: Theory and Practice, 29(4), 5898–5910. https://doi.org/10.53555/kuey.v29i4.10932.
9. Bellamy, R. K. E., et al. (2019). AI fairness 360. IBM Journal of Research and Development, 63(4/5), 4:1–4:15.
10. Bender, E. M., et al. (2021). On the dangers of stochastic parrots. Proceedings of FAccT, 610–623.
11. Sasi Kumar Kolla. (2023). Big Data–Driven Machine Learning Frameworks for Clinical Risk Prediction. International Journal of Medical Toxicology and Legal Medicine, 26(3 and 4), 44–59. Retrieved from https://ijmtlm.org/index.php/journal/article/view/1456.
12. Biecek, P., & Burzykowski, T. (2021). Explanatory model analysis. CRC Press.
13. Aitha, A. R. (2024). Generative AI-Powered Fraud Detection in Workers' Compensation: A DevOps-Based Multi-Cloud Architecture Leveraging, Deep Learning, and Explainable AI. Deep Learning, and Explainable AI (July 26, 2024).
14. Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). The ML test score. Proceedings of SysML.
15. Varri, D. B. S. (2024). Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems. Available at SSRN 5774785.
16. Brownlee, J. (2020). Data preparation for machine learning. Machine Learning Mastery.
17. Budach, L., et al. (2022). The effects of data quality on machine learning performance. Data Science and Engineering, 7, 127–145.
18. Meda, R. (2024). Agentic AI in Multi-Tiered Paint Supply Chains: A Case Study on Efficiency and Responsiveness. Journal of Compu-tational Analysis and Applications (JoCAAA), 33(08), 3994-4015.
19. Caruana, R., et al. (2015). Intelligible models for healthcare. KDD, 1721–1730.
20. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. KDD, 785–794.
21. Singireddy, J. (2024). AI-Enhanced Tax Preparation and Filing: Automating Complex Regulatory Compliance. European Data Science Journal (EDSJ) p-ISSN 3050-9572 en e-ISSN 3050-9580, 2(1).
22. Cohen, I. G., et al. (2022). The AI revolution in healthcare. Science, 375(6587), 1327–1330.
23. Kolla, S. K. (2021). Designing Scalable Healthcare Data Pipelines for Multi-Hospital Networks. World Journal of Clinical Medicine Research, 1(1), 1–14. Retrieved from https://www.scipublications.com/journal/index.php/wjcmr/article/view/1376.
24. Agentic AI in Data Pipelines: Self Optimizing Systems for Continuous Data Quality, Performance and Governance. (2024). American Data Science Journal for Advanced Computations (ADSJAC) ISSN: 3067-4166, 2(1).
25. Deng, L., & Yu, D. (2014). Deep learning: Methods and applications. Foundations and Trends in Signal Processing, 7(3–4), 197–387.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

26. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv.
27. Deep Learning-Driven Optimization of ISO 20022 Protocol Stacks for Secure Cross-Border Messaging. (2024). MSW Management Journal, 34(2), 1545-1554.
28. European Parliament. (2024). Regulation (EU) 2024/1689 (AI Act). Official Journal of the European Union.
29. Segireddy, A. R. (2024). Machine Learning-Driven Anomaly Detection in CI/CD Pipelines for Financial Applications. Journal of Computational Analysis and Applications, 33(8).
30. Geiger, R., et al. (2020). Garbage in, garbage out? Big data and bias. Big Data & Society, 7(2).
31. Keerthi Amistapuram. (2024). Federated Learning for Cross-Carrier Insurance Fraud Detection: Secure Multi-Institutional Collaboration. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 6727–6738. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/3934.
32. Goldstein, A., Kapelner, A., et al. (2015). Peeking inside the black box. Journal of Computational and Graphical Statistics, 24(1), 44–65.
33. Varri, D. B. S. (2023). Advanced Threat Intelligence Modeling for Proactive Cyber Defense Systems. Available at SSRN 5774926.
34. Hinton, G., et al. (2012). Deep neural networks for acoustic modeling. IEEE Signal Processing Magazine, 29(6), 82–97.
35. Sheelam, G. K., & Koppolu, H. K. R. (2024). From Transistors to Intelligence: Semiconductor Architectures Empowering Agentic AI in 5G and Beyond. Journal of Computational Analy- sis and Applications(JoCAAA), 33(08), 4518-4537.
36. ISO/IEC 25012. (2008). Data quality model. International Organization for Standardization.
37. Paleti, S. (2024). Transforming Financial Risk Management with AI and Data Engineering in the Modern Banking Sector. American Journal of Analytics and Artificial Intelligence (ajaai) with ISSN 3067-283X, 2(1).
38. Johnson, A. E. W., et al. (2016). MIMIC-III clinical database. Scientific Data, 3, 160035.
39. Meda, R. (2023). Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains. Educational Administration: Theory and Practice.
40. Kim, B., et al. (2018). Interpretability beyond feature attribution. ICML.
41. Garapati, R. S. (2023). Optimizing Energy Consumption in Smart Build-ings Through Web-Integrated AI and Cloud-Driven Control Systems.
42. Koh, P. W., & Liang, P. (2017). Influence functions. ICML, 1885–1894.
43. Inala, R. Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective.
44. Krizhevsky, A., et al. (2012). ImageNet classification with deep convolutional neural networks. NIPS, 1097–1105.
45. Varri, D. B. S. (2022). A Framework for Cloud-Integrated Database Hardening in Hybrid AWS-Azure Environments: Security Posture Automation Through Wiz-Driven Insights. International Journal of Scientific Research and Modern Technology, 1(12), 216-226.
46. Lipton, Z. C. (2018). The mythos of model interpretability. Communications of the ACM, 61(10), 36–43.
47. Amistapuram, K. (2024). Generative AI in Insurance: Automating Claims Documentation and Customer Communication. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 461–475. https://doi.org/10.61841/turcomat.v15i3.15474.
48. Lu, J., et al. (2019). Learning under concept drift: A review. IEEE TKDE, 31(12), 2346–2363.
49. McKinney, W. (2022). Python for data analysis (3rd ed.). O'Reilly Media.
50. Aitha, A. R. (2023). CloudBased Micro services Architecture for Seamless Insurance Policy Administration. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 607-632.
51. National Institute of Standards and Technology. (2023). AI risk management framework 1.0.
52. Northcutt, C., et al. (2021). Confident learning. Journal of Artificial Intelligence Research, 70, 1373–1411.
53. Nagabhyru, K. C. (2024). Data Engineering in the Age of Large Language Models: Transforming Data Access, Curation, and Enterprise Interpretation. Computer Fraud and Security.
54. Polyzotis, N., et al. (2018). Data management challenges in production ML. SIGMOD, 1723–1726.
55. Davuluri, P. S. L. N. (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. Metallurgical and Materials Engineering, 30(4), 996–1010. Retrieved from https://metall-mater-eng.com/index.php/home/article/view/1936.
57. Sculley, D., et al. (2015). Hidden technical debt in ML systems. NeurIPS, 2503–2511.
58. Uday Surendra Yandamuri. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706. https://doi.org/10.5281/zenodo.18095256.
59. Shickel, B., et al. (2017). Deep EHR representation learning. Journal of Biomedical Informatics, 83, 168–185.

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 6, Issue 6, November - December 2024 |

DOI:10.15662/IJEETR.2024.0606018

60. Koppolu, H. K. R., & Sheelam, G. K. (2024). Machine Learning-Driven Optimization in 6G Telecommunications: The Role of Intelligent Wireless and Semiconductor Innovation. Global Research Development (GRD) ISSN: 2455-5703, 9(12).
61. Simmhan, Y. L., et al. (2005). A survey of data provenance. ACM SIGMOD Record, 34(3), 31–36.
62. Rongali, S. K. (2023). Explainable Artificial Intelligence (XAI) Framework for Transparent Clinical Decision Support Systems. International Journal of Medical Toxicology and Legal Medicine, 26(3), 22-31.
63. Song, L., et al. (2021). Data-centric AI. arXiv.
64. Mashetty, S., Challa, S. R., ADUSUPALLI, B., Singireddy, J., & Paleti, S. (2024). Intelligent Technologies for Modern Financial Ecosystems: Transforming Housing Finance, Risk Management, and Advisory Services Through Advanced Analytics and Secure Cloud Solutions. Risk Management, and Advisory Services Through Advanced Analytics and Secure Cloud Solutions (December 12, 2024v.
65. TensorFlow Team. (2022). TFX: ML pipelines. Google AI.
66. Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. Journal of the Royal Statistical Society, 58(1), 267–288.
67. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
68. Vapnik, V. (1998). Statistical learning theory. Wiley.
69. Veale, M., & Borgesius, F. Z. (2021). Demystifying the draft EU AI Act. Computer Law Review International, 22(4), 97–112.
70. Rongali, S. K., & Kumar Kakarala, M. R. (2024). Existing challenges in ethical AI: Addressing algorithmic bias, transparency, accountability and regulatory compliance.
71. Lahari Pandiri, "AI-Powered Fraud Detection Systems in Professional and Contractors Insurance Claims," International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREEICE), DOI 10.17148/IJIREEICE.2024.121206.
72. WHO. (2021). Ethics and governance of AI for health. World Health Organization.
73. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.
74. Wilkinson, M. D., et al. (2016). FAIR guiding principles. Scientific Data, 3, 160018.
75. Guntupalli, R. (2024). AI-Powered Infrastructure Management in Cloud Computing: Automating Security Compliance and Performance Monitoring. Available at SSRN 5329147.
76. Zhou, Z.-H. (2021). Machine learning. Springer.
77. Abadi, M., et al. (2016). TensorFlow. OSDI, 265–283.
78. Nagubandi, A. R. (2023). Advanced Multi-Agent AI Systems for Autonomous Reconciliation Across Enterprise Multi-Counterparty Derivatives, Collateral, and Accounting Platforms. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 653-674.
79. He, H., & Garcia, E. A. (2009). Learning from imbalanced data. IEEE TKDE, 21(9), 1263–1284.
80. Keerthi Amistapuram. (2023). Privacy-Preserving Machine Learning Models for Sensitive Customer Data in Insurance Systems. Educational Administration: Theory and Practice, 29(4), 5950–5958. https://doi.org/10.53555/kuey.v29i4.10965.
81. Carcillo, F., et al. (2021). Streaming fraud detection framework. Information Fusion, 41, 182–194.
82. Chava, K. (2024). The Role of Cloud Computing in Accelerating AI-Driven Innovations in Healthcare Systems. European Advanced Journal for Emerging Technologies (EAJET)-p-ISSN 3050-9734 en e-ISSN 3050-9742, 2(1).
83. Silver, D., et al. (2016). Mastering the game of Go with deep neural networks. Nature, 529, 484–489.
84. Siva Hemanth Kolla. (2023). Deep Learning–Driven Retrieval-Augmented Generation for Enterprise ITSM Automation: A Governance-Aligned Large Language Model Architecture . Journal of Computational Analysis and Applications (JoCAAA), 31(4), 2489–2502. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/4774.
85. Mehrabi, N., et al. (2021). Survey on bias and fairness in ML. ACM Computing Surveys, 54(6), 1–35.
86. Rongali, S. K. (2024). Federated and Generative AI Models for Secure, Cross-Institutional Healthcare Data Interoperability. Journal of Neonatal Surgery, 13(1), 1683-1694.
87. Karimian, N., et al. (2022). Blockchain for healthcare data governance. IEEE Access, 10, 11456–11469.
88. Yandamuri, U. S. AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology.
89. Molnar, C. (2022). Interpretable machine learning (2nd ed.). Lulu Press.
90. Kolla, S. H. (2024). RETRIEVAL-AUGMENTED GENERATION WITH SMALL LLMS FOR KNOWLEDGE-DRIVEN DECISION AUTOMATION IN ENTERPRISE SERVICE PLATFORMS. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 476–486. https://doi.org/10.61841/turcomat.v15i3.15497.