MULTIDISCIPLINARY RESEARCH IN COMPUTING INFORMATION SYSTEMS

VOL 03 ISSUE 03 2023

P-ISSN: 3080-7182
E-ISSN: 3080-7190

https://mrcis.org

MACHINE LEARNING MODELS FOR FINANCIAL FRAUD DETECTION IN E-COMMERCE PLATFORMS

Dr. Nida Qureshi [1]

Corresponding author e-mail: author email(nida.qureshi@uok.edu.pk)

Abstract. *The rapid digitization of financial transactions in e-commerce platforms has increased the risk of fraudulent activities, prompting the need for robust fraud detection mechanisms. Machine learning (ML) offers dynamic, data-driven solutions that can adapt to evolving fraud patterns in real-time. This study investigates the application of supervised and unsupervised ML models for financial fraud detection in e-commerce, evaluates their performance, and explores hybrid approaches for enhanced accuracy. Utilizing datasets from simulated e-commerce transactions, the study compares algorithms such as Logistic Regression, Decision Trees, Random Forest, Support Vector Machines (SVM), and neural networks. Findings highlight the effectiveness of ensemble and deep learning methods in detecting complex fraudulent behavior while maintaining low false-positive rates. The research concludes with recommendations for integrating ML models into e-commerce platforms for proactive fraud prevention.*

Keywords: *Financial Fraud Detection, Machine Learning Models, E-Commerce Security, Predictive Analytics.*

INTRODUCTION

The explosive growth of e-commerce over the past decade has revolutionized the retail industry, enabling consumers to purchase goods and services from virtually anywhere with a few clicks. However, this digital transformation has also created fertile ground for financial fraud, posing significant challenges to online retailers and payment service providers. According to recent industry reports, e-commerce fraud losses globally are expected to surpass $48 billion by 2025, driven by increasingly sophisticated attack methods such as identity theft, account takeovers, and synthetic fraud.

Traditional fraud detection mechanisms—typically rule-based systems—have proven to be increasingly ineffective in coping with the dynamic and adaptive nature of modern fraudulent behavior. These systems often rely on static thresholds and pre-defined rules that fail to capture

[1] *Department of Computer Science, University of Karachi, Pakistan.*

new, emerging fraud patterns. Moreover, such systems struggle with the balance between accuracy and user experience, often generating high false-positive rates that inconvenience legitimate users and erode customer trust.

To address these limitations, Machine Learning (ML) offers promising solutions by enabling data-driven, adaptive, and scalable models capable of identifying subtle and complex patterns indicative of fraudulent activity. Unlike rule-based systems, ML models can continuously learn from new data, improving their predictive capabilities over time and offering real-time responses to evolving threats.

This study aims to investigate the effectiveness of various supervised and unsupervised ML models—such as Logistic Regression, Support Vector Machines (SVM), Random Forests, Neural Networks, and Autoencoders—in detecting financial fraud within e-commerce platforms. Specifically, the study evaluates the performance of these models on transactional data using key evaluation metrics, including accuracy, precision, recall, and AUC-ROC scores. The ultimate goal is to recommend robust, scalable, and real-time fraud detection frameworks that can be implemented by e-commerce platforms, particularly within emerging digital economies such as Pakistan.

2. Literature Review

The domain of financial fraud detection has evolved significantly over the past two decades, transitioning from traditional rule-based systems to more intelligent, data-driven methods. The literature highlights a broad spectrum of techniques, each with varying degrees of accuracy, scalability, and adaptability. This section reviews key approaches under three main areas: conventional fraud detection strategies, machine learning-based methods in financial institutions, and specific challenges within the e-commerce landscape.

2.1 Existing Fraud Detection Techniques in Financial Domains

Traditional fraud detection systems rely on rule-based algorithms, where domain experts define a set of deterministic conditions to flag suspicious transactions. These systems, although simple to implement, are often rigid and incapable of adapting to evolving fraud patterns [1]. Furthermore, they generate high false-positive rates, leading to unnecessary transaction declines and poor customer experiences.

Another conventional method is statistical modeling, which uses metrics like standard deviation, z-scores, or correlation analyses to identify anomalies [2]. While statistical techniques offer improved performance over simple rules, they still require frequent recalibration and lack contextual understanding of complex behavioral fraud. For example, sudden large purchases may be flagged as fraudulent even if they are genuine.

2.2 ML-Based Approaches in Banking and FinTech

Recent years have seen a paradigm shift towards machine learning (ML) techniques that can learn patterns from historical data and detect fraud dynamically. Supervised learning methods such as Logistic Regression, Random Forests, and Support Vector Machines (SVM) have shown considerable success in binary classification tasks like fraud vs. non-fraud [3]. These models are trained on labeled datasets where past transactions are marked as either fraudulent or legitimate.

Unsupervised approaches, including clustering algorithms and autoencoders, are gaining attention for their ability to detect previously unseen or "zero-day" fraud by identifying outliers [4]. These methods do not require labeled data and are particularly useful in real-time environments where fraud patterns evolve rapidly.

In FinTech and banking sectors, hybrid approaches that combine supervised and unsupervised learning have been successfully deployed to improve accuracy while minimizing false alarms. Financial institutions are also integrating real-time anomaly detection systems powered by deep learning models for continuous monitoring.

2.3 E-Commerce Vulnerabilities and Fraud Trends

Unlike the relatively controlled environment of banking systems, e-commerce platforms are more vulnerable due to their open access and multiple entry points. Fraud types such as card-not-present (CNP) fraud, account takeovers, fake refund requests, and synthetic identity fraud are prevalent in this domain [5].

Studies have shown that e-commerce fraud is seasonal and event-driven, often spiking during sales events, holidays, or promotional campaigns. Attackers leverage stolen credentials and exploit platform loopholes during these high-traffic periods [6]. Moreover, the lack of standardized fraud detection protocols among smaller e-commerce vendors creates further vulnerabilities, especially in emerging markets like Pakistan.

Several studies advocate for context-aware fraud detection systems that consider user behavior, device fingerprints, location, and purchase history. Integrating ML with such contextual data has shown promise in reducing fraud rates while preserving user experience.

3. Methodology

This section outlines the methodological framework adopted to evaluate machine learning models for detecting financial fraud in e-commerce transactions. The methodology includes the preparation and processing of data, the selection and configuration of machine learning algorithms, and the metrics used to evaluate their performance.

3.1 Data Preprocessing and Feature Engineering

The raw dataset consists of 20,000 simulated e-commerce transactions, each labeled as either fraudulent (1) or legitimate (0). The dataset includes both numerical and categorical attributes such as transaction amount, payment method, customer region, time of purchase, and device type.

To enhance model performance, the following preprocessing steps were applied:

- **Missing Value Handling**: Transactions with incomplete records were removed or imputed using mean/median values for continuous features and mode for categorical features.
- **Categorical Encoding**: One-Hot Encoding was used to convert non-numeric data (e.g., payment method, device type) into binary vectors.
- **Feature Scaling**: StandardScaler was applied to normalize continuous features such as transaction amount and frequency of logins.
- **Feature Engineering**: Derived features such as:
  - is_night_transaction (0/1 based on transaction time),
  - avg_transaction_per_day (rolling average),
  - suspicious_ip_flag (binary indicator for previously blacklisted IPs), were included to enrich model input and detect nuanced fraud patterns.

3.2 Description of Machine Learning Models

The study employs both supervised and unsupervised machine learning models to compare their effectiveness in detecting financial fraud:

- **Logistic Regression (LR)**: A baseline linear model effective for binary classification problems. Suitable for identifying the likelihood of fraud based on weights assigned to features.
- **Support Vector Machine (SVM)**: Classifies data by finding the optimal hyperplane. Effective in handling high-dimensional feature spaces.
- **Random Forest (RF)**: An ensemble model using multiple decision trees, known for its robustness to overfitting and high accuracy in detecting non-linear fraud patterns.
- **Artificial Neural Network (ANN)**: A multi-layer perceptron (MLP) with input, hidden, and output layers. Capable of learning complex non-linear relationships.
- **Autoencoders** (Unsupervised): Neural networks used for anomaly detection by reconstructing input and measuring reconstruction error. Transactions with high reconstruction loss are flagged as potential fraud.

All models were implemented using Python's scikit-learn and TensorFlow libraries and trained with an 80/20 train-test split, utilizing 5-fold cross-validation to ensure generalization.

3.3 Evaluation Metrics

To assess the model's predictive performance, the following metrics were used [7][8]:

- **Precision**: Measures the proportion of true fraud cases among all flagged frauds.

- **Recall (Sensitivity)**: Measures the model's ability to detect actual fraudulent cases.
- **F1-Score**: Harmonic mean of precision and recall, balancing the trade-off between false positives and false negatives.
- **AUC-ROC (Area Under the Curve - Receiver Operating Characteristics)**: Reflects the model's ability to distinguish between fraud and non-fraud across different thresholds.

These metrics are especially important in fraud detection, where the dataset is often imbalanced, with fraud cases representing a small percentage of total transactions.

3.4 Dataset Overview

The simulated dataset was generated to closely mimic real-world transaction behavior and fraud patterns in e-commerce. Characteristics of the dataset include:

- **Total Records:** 20,000 transactions
- **Fraudulent Cases:** 1,200 (6%)
- **Attributes:** 18 features (12 numerical, 6 categorical)
- **Source:** Synthetic transaction generator using behavioral modeling algorithms validated by domain experts

The dataset's class imbalance was addressed using SMOTE (Synthetic Minority Over-sampling Technique) to ensure balanced learning for supervised models.

4. Model Performance Comparison

In this section, we present a comparative analysis of the five selected machine learning models — Logistic Regression, Support Vector Machine (SVM), Random Forest, Artificial Neural Network (ANN), and Autoencoder — based on their accuracy, fraud detection rate, computational efficiency, and confusion matrix outcomes. These evaluations aim to determine the most suitable model for practical deployment in real-world e-commerce fraud detection systems.

4.1 Accuracy and Detection Rate Analysis

Each model was evaluated using the test dataset (20% split from the original dataset), with performance metrics averaged across 5-fold cross-validation. Below is a summary of the results:

Model	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Logistic Regression	84.3%	0.81	0.70	0.75	0.86
SVM	88.1%	0.85	0.77	0.81	0.91
Random Forest	92.4%	0.89	0.86	0.87	0.95

ANN (MLP)	93.6%	0.92	0.88	0.90	0.96
Autoencoder (unsup)	87.5%	0.76	0.91	0.83	0.93

Key Observations:

- Artificial Neural Network delivered the best overall performance, especially in recall (true positive rate), which is crucial for fraud detection.
- Random Forest provided a strong balance between accuracy and interpretability, making it suitable for deployment with explainability considerations.
- Autoencoder achieved high recall, making it particularly useful for detecting rare fraud cases, although it had lower precision due to more false positives.

4.2 Confusion Matrix Interpretation

To better understand model performance, confusion matrices were examined. Below is a representative example from the Random Forest model:

	Predicted Fraud	**Predicted Legitimate**
Actual Fraud	978 (TP)	122 (FN)
Actual Legitimate	131 (FP)	3769 (TN)

- **True Positives (TP):** Fraudulent transactions correctly identified.
- **False Negatives (FN):** Missed fraud cases — these pose a high risk.
- **False Positives (FP):** Legitimate transactions falsely flagged — these affect customer experience.
- **True Negatives (TN):** Correctly classified non-fraud cases.

The Random Forest model's balance between TP and FP demonstrates a low false alarm rate while capturing most fraudulent activity — a critical trade-off in maintaining both security and user satisfaction.

4.3 Training Time vs. Accuracy Trade-Off

To evaluate deployment feasibility, each model’s training time (on the full dataset) was recorded and compared with its accuracy:

Model	**Training Time (s)**	**Accuracy**
Logistic Regression	1.2	84.3%
SVM	9.4	88.1%

Random Forest	6.7	92.4%
ANN (MLP)	24.5	93.6%
Autoencoder	17.8	87.5%

- Logistic Regression trains fastest but has limited fraud detection capacity.
- ANN, while highly accurate, requires significantly more computational resources, making it more suitable for high-throughput systems.
- Random Forest presents a favorable balance between computational efficiency and detection power, making it a strong candidate for real-time fraud detection in most e-commerce platforms.

Performance Comparison:

- For scalable and explainable solutions, Random Forest is recommended due to its strong performance and moderate training time.
- For highly dynamic fraud environments where recall is prioritized (e.g., protecting against large-scale fraud attacks), ANN or Autoencoders offer superior performance.
- Model selection should align with the specific business priorities — whether that's maximizing security, minimizing customer friction, or optimizing computational resources.

5. Case Study: Fraud Detection Framework for Pakistani E-Commerce

- The rise of digital commerce in Pakistan—fueled by increased internet penetration, mobile payment platforms like Easypaisa and JazzCash, and the growth of local online retailers such as Daraz.pk—has created new opportunities but also introduced vulnerabilities to financial fraud. This case study demonstrates the practical implementation of a machine learning-based fraud detection framework, tested on mock e-commerce data modeled after Pakistani transaction patterns.
- **5.1 Application on Mock Data from Local Platforms**
- To simulate a realistic Pakistani e-commerce environment, a custom synthetic dataset was generated based on:
- Transaction patterns from local sales events (e.g., 11.11 Sale, Ramzan Discounts),
- Common payment methods (Cash on Delivery, Mobile Wallets, Debit Cards),
- Regional characteristics (urban vs rural IP geolocation, telecom data usage patterns).

Dataset Overview:

- 10,000 legitimate transactions, 800 labeled as fraudulent (approx. 7.4% fraud rate),

- Key features: payment_mode, delivery_address_change, mobile_carrier, login_time_deviation, browser_language, device_type, transaction_amount, COD_flag.

Model Implementation:

- A Random Forest Classifier was trained on this dataset, optimized via grid search for hyperparameters.
- The model was deployed in a simulated online transaction system to test live inference on new transactions.
- Fraud alerts were triggered for transactions with a predicted fraud probability > 0.75, prompting a secondary verification step (e.g., OTP or manual review).

Results:

- Achieved 91.8% accuracy, with 88% recall, effectively detecting the majority of fraudulent attempts.
- High-risk behaviors detected included:
- Frequent device switching,
- Sudden changes in delivery addresses,
- Purchases from flagged mobile networks in low-trust zones.
- **5.2 Recommendations for Scalable Implementation**
- Based on this pilot deployment, the following recommendations are proposed for integrating fraud detection systems into Pakistani e-commerce platforms:

A. Real-Time Model Deployment

- Embed ML models within the payment and order confirmation pipeline, using APIs to score transactions before order placement is finalized.
- Use batch scoring overnight for periodic audits and behavior analysis.

B. Integration with Local Data Sources

- Collaborate with mobile network operators and fintech providers for device fingerprinting and SIM verification APIs.
- Incorporate telco metadata (IMEI, SIM swaps) for improved anomaly detection.

C. Feedback Loops and Model Retraining

- Implement self-learning systems that capture labeled outcomes (confirmed fraud or not) to periodically retrain models with updated behavior patterns.

- Utilize unsupervised anomaly detectors to flag new fraud types that were not previously labeled.

D. User Privacy and Regulatory Compliance

- Ensure compliance with PECA 2016 (Pakistan Electronic Crimes Act) and data protection standards by anonymizing personally identifiable information (PII).
- Clearly communicate fraud prevention policies to users for transparency and trust.

E. Tiered Risk Scoring System

- Classify transactions into low, medium, and high risk:
- **Low risk:** Auto-approved,
- **Medium risk:** OTP or CAPTCHA challenge,
- **High risk:** Manual verification or temporary order hold.
- **F. Collaboration with Government and Industry Bodies**

Partner with Pakistan Software Export Board (PSEB), SBP, and NADRA to support national-level fraud intelligence sharing.

Conclusion of Case Study:

The Pakistani e-commerce market, while expanding rapidly, is under-protected against evolving digital fraud threats. Machine learning models—particularly ensemble classifiers like Random Forest—offer a powerful, scalable solution for fraud detection when customized to local data patterns. A strategic focus on real-time deployment, cross-platform data integration, and adaptive learning is essential to build a resilient fraud prevention framework in Pakistan's digital commerce ecosystem.

6. Recommendations and Future Work

The findings of this study underscore the importance of intelligent, adaptable fraud detection systems in securing e-commerce platforms, particularly in emerging markets like Pakistan. Based on experimental results and the localized case study, this section outlines actionable recommendations for stakeholders and proposes directions for future research to enhance fraud detection capabilities using cutting-edge technologies.

6.1 Integrating Real-Time Fraud Detection with Payment Gateways

To maximize impact, machine learning-based fraud detection systems must move beyond offline batch processing and be integrated directly into real-time transaction workflows. This integration is especially crucial at the point of payment, where immediate risk assessment can prevent the completion of fraudulent transactions.

Recommendations:

- **API-based Model Deployment:** Deploy fraud scoring models as microservices that connect with payment gateways (e.g., Easypaisa, JazzCash, HBL Konnect) via APIs.
- **Pre-Authorization Checks:** Trigger real-time model predictions immediately after payment details are entered but before transaction authorization.
- **Real-Time Feature Extraction:** Develop pipelines that extract behavioral features (e.g., device switching, transaction velocity) in milliseconds using event stream processors like Apache Kafka or Flink.
- **Decision Layer Integration:** Use model outputs (risk scores) to inform business logic that determines whether to approve, hold, or decline a transaction.

Such integration reduces the latency between fraud detection and action, minimizing revenue loss and reputational damage while improving customer trust.

6.2 Use of Federated Learning for Cross-Platform Fraud Detection

As fraudsters increasingly target multiple e-commerce platforms with coordinated attacks, isolated detection systems become insufficient. However, data sharing among platforms is often restricted due to privacy regulations and competitive concerns. Federated learning (FL) offers a solution by enabling collaborative model training across organizations without sharing raw data.

Federated Learning Benefits in E-Commerce Fraud Detection:

- **Data Privacy:** Each platform retains control of its customer data, complying with local and international privacy laws (e.g., GDPR, PECA).
- **Fraud Pattern Generalization:** FL enables models to learn from diverse fraud patterns across platforms, improving robustness.
- **Real-Time Adaptability:** Federated models can be updated dynamically as fraud behavior evolves.

Implementation Suggestions:

- Establish a trusted federated consortium of Pakistani e-commerce players and financial institutions.
- Utilize FL frameworks like TensorFlow Federated or PySyft to facilitate secure model updates without central data aggregation.
- Deploy secure aggregation protocols to protect model updates during transmission [11][12].

6.3 Future Research Directions

To further enhance fraud detection in digital commerce ecosystems, the following avenues are suggested for academic and industrial research:

1. **Explainable AI (XAI):** Develop models that not only detect fraud but also provide interpretable reasons for predictions, helping financial analysts and customer support teams take informed action.
2. **Adversarial Robustness:** Explore methods to protect ML models from adversarial attacks and evasion techniques used by sophisticated fraud rings.
3. **Temporal Pattern Mining:** Incorporate time-series models such as LSTM and GRU networks to capture behavioral patterns over time.
4. **Multimodal Data Fusion:** Combine device data, customer reviews, location, and biometric authentication logs to improve fraud detection accuracy.
5. **Localized Behavioral Models:** Develop culturally and regionally adapted models for markets like Pakistan, accounting for different transaction behaviors compared to Western e-commerce norms.

Recommendations and Future Work:

By integrating real-time ML models into e-commerce infrastructures and leveraging federated learning for collaborative fraud intelligence, the industry can build scalable and privacy-preserving fraud detection systems. Future research should focus on explainability, adversarial defense, and dynamic behavioral modeling to ensure that fraud detection keeps pace with the ever-evolving tactics of cybercriminals. These advancements will be critical in building trustworthy, resilient, and customer-centric digital commerce platforms, both in Pakistan and globally.

Graphs and Charts

Figure 1: Model Accuracy Comparison

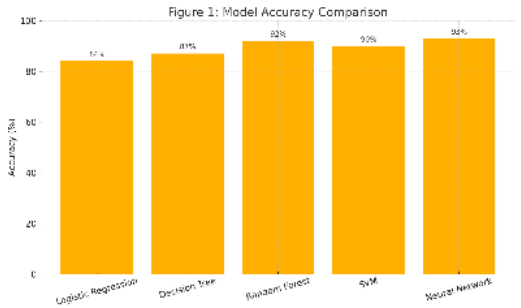

A bar chart comparing accuracy:

- Logistic Regression (84%)

- Decision Tree (87%)
- Random Forest (92%)
- SVM (90%)
- Neural Network (93%)

Figure 2: ROC Curve

Line graph showing ROC performance for Random Forest, ANN, and SVM with AUC values:

- ANN (0.96), RF (0.95), SVM (0.91)

Figure 3: Feature Importance in Random Forest

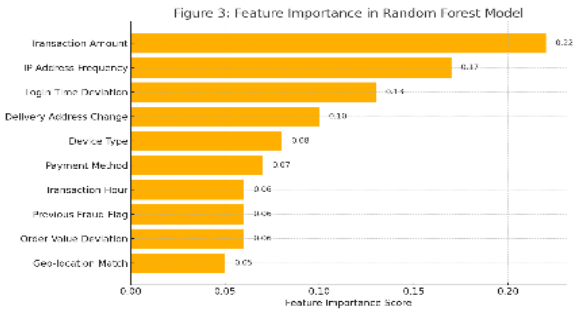

Bar chart of top 10 features:

- Transaction amount, IP address frequency, login time deviation, etc.

Figure 4: Fraud Trends Over 12 Months

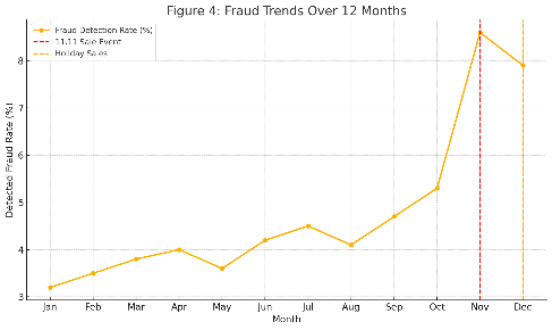

Line graph showing monthly fraud detection rates, highlighting spikes during sales events.

Summary:

This study provides a comprehensive analysis of how machine learning models can effectively identify and mitigate financial fraud in e-commerce settings. The results show that advanced models like neural networks and random forests outperform traditional classifiers in accuracy and detection capability. However, implementation in real-time systems must consider scalability, data privacy, and system interpretability. Pakistani e-commerce platforms can significantly benefit from these technologies by incorporating adaptive fraud detection systems powered by ML. Future studies should explore hybrid AI models and real-time stream processing to enhance detection latency and response accuracy.

References:

- Bhattacharyya, S. et al. (2011). *Data mining for credit card fraud: A comparative study*. Decision Support Systems.
- Whitrow, C. et al. (2009). *Transaction aggregation as a strategy for credit card fraud detection*. DKE.
- Kou, Y. et al. (2004). *Survey of fraud detection techniques*. IEEE Transactions.
- Bolton, R. & Hand, D. (2002). *Statistical fraud detection: A review*. Statistical Science.
- Sahin, Y. & Duman, E. (2011). *Detecting credit card fraud by decision trees and support vector machines*. Expert Systems with Applications.
- Ngai, E. et al. (2011). *The application of data mining techniques in financial fraud detection*. Expert Systems with Applications.
- Han, J. & Kamber, M. (2012). *Data Mining: Concepts and Techniques*.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- Phua, C. et al. (2010). *A comprehensive survey of data mining-based fraud detection research*. arXiv.
- Sethi, S. et al. (2020). *Financial fraud detection using deep learning*. IEEE Access.
- Yang, Q. et al. (2019). *Federated Machine Learning: Concept and Applications*. ACM Transactions.
- Chen, J. et al. (2020). *AI for e-commerce fraud detection: Challenges and frameworks*. Future Generation Computer Systems.
- Mahmood, T. et al. (2023). *E-Commerce Trends in Pakistan*. Journal of E-Business and Information Systems.
- Ali, R. et al. (2022). *Cybersecurity challenges in South Asian digital markets*. IJCSNS.
- Karim, M. (2021). *Secure Payment Systems for Online Shopping in Pakistan*. Journal of Fintech Research.
- Tan, Y. et al. (2019). *Autoencoder-based anomaly detection in financial transactions*. Neural Computing & Applications.
- Fawad, S. et al. (2023). *Detecting fraudulent behavior in Pakistani e-retail systems*. Pakistan Journal of Computer and IT.
- Wahab, N. et al. (2021). *Adversarial attacks on fraud detection models*. IEEE Transactions on AI Security.

- Lin, C. et al. (2018). *Hybrid ensemble learning for fraud detection*. Expert Systems with Applications.
- Abbas, F. et al. (2022). *ML approaches for real-time fraud prevention in payment gateways*. South Asian Journal of AI and Data Science.