Enhancing Financial Security: Real-Time Anomaly Detection in Financial Transactions Using Machine Learning

Authors : Ali Kazemi

Abstract : The digital evolution of financial services, while offering unprecedented convenience and accessibility, has also escalated the vulnerabilities to fraudulent activities. In this study, we introduce a distinct approach to real-time anomaly detection in financial transactions, aiming to fortify the defenses of banking and financial institutions against such threats. Utilizing unsupervised machine learning algorithms, specifically autoencoders and isolation forests, our research focuses on identifying irregular patterns indicative of fraud within transactional data, thus enabling immediate action to prevent financial loss. The data we used in this study included the monetary value of each transaction. This is a crucial feature as fraudulent transactions may have distributions of different amounts than legitimate ones, such as timestamps indicating when transactions occurred. Analyzing transactions' temporal patterns can reveal anomalies (e.g., unusual activity in the middle of the night). Also, the sector or category of the merchant where the transaction occurred, such as retail, groceries, online services, etc. Specific categories may be more prone to fraud. Moreover, the type of payment used (e.g., credit, debit, online payment systems). Different payment methods have varying risk levels associated with fraud. This dataset, anonymized to ensure privacy, reflects a wide array of transactions typical of a global banking institution, ranging from small-scale retail purchases to large wire transfers, embodying the diverse nature of potentially fraudulent activities. By engineering features that capture the essence of transactions, including normalized amounts and encoded categorical variables, we tailor our data to enhance model sensitivity to anomalies. The autoencoder model leverages its reconstruction error mechanism to flag transactions that deviate significantly from the learned normal pattern, while the isolation forest identifies anomalies based on their susceptibility to isolation from the dataset's majority. Our experimental results, validated through techniques such as k-fold cross-validation, are evaluated using precision, recall, and the F1 score alongside the area under the receiver operating characteristic (ROC) curve. Our models achieved an F1 score of 0.85 and a ROC AUC of 0.93, indicating high accuracy in detecting fraudulent transactions without excessive false positives. This study contributes to the academic discourse on financial fraud detection and provides a practical framework for banking institutions seeking to implement real-time anomaly detection systems. By demonstrating the effectiveness of unsupervised learning techniques in a real-world context, our research offers a pathway to significantly reduce the incidence of financial fraud, thereby enhancing the security and trustworthiness of digital financial services.