

# Real-Time Fraud Detection using Apache Kafka, Apache Spark, and PySpark MLlib

Kunal Dabas<sup>1</sup>, Ashmit Dubey<sup>2</sup>, Ritu Kalonia<sup>3</sup>

<sup>1</sup>Department of Artificial Intelligence and Data Science, University School of Automation and Robotics, Delhi, India

<sup>2</sup>Department of Artificial Intelligence and Data Science, University School of Automation and Robotics, Delhi, India

<sup>3</sup>Department of Artificial Intelligence and Data Science, University School of Automation and Robotics, Delhi, India

Email: <sup>1</sup>[kunaldabas37@gmail.com](mailto:kunaldabas37@gmail.com), <sup>2</sup>[009ashmitdubey@gmail.com](mailto:009ashmitdubey@gmail.com) <sup>3</sup>[ritukalonia.usar@ipu.ac.in](mailto:ritukalonia.usar@ipu.ac.in)

## ABSTRACT:

*In this study we present a robust scalable and real time fraud detection system for credit card transactions. apache kafka is used for data ingestion and Apache Spark Streaming for real-time data processing, and Spark MLlib for implementation of machine learning models. Macos environment is chosen for deployment of the architecture . We use a 5-node Spark cluster and 3 Kafka brokers. Using Our implementation enables immediate detection of fraudulent transactions, which ensures rapid response and decision-making. The paper elaborates on the methodology, architecture, and execution pipeline and discusses the potential of integrating advanced analytics and visualization in future iterations.*

**INDEX TERMS:** Apache Spark, Apache Kafka, PySpark, Spark Streaming, Spark MLlib

## 1) Introduction

We have seen exponential growth in online transactions, financial frauds have also become common, causing serious economic losses. Rule-based models of fraud detection are not capable of handling emerging fraud trends and are not suitable for processing high-speed transaction volumes. ML-based models which are integrated with real-time big data platforms are capable of overcoming these constraints.

The major issue which we are trying to address is real-time identification of fraud transactions in a stream of incoming financial data. High latency in identifying frauds and small capacity to work on huge transaction volumes.

To detect fraudulent actions such as unauthorized access, identity theft, and transaction manipulation are very difficult in real time with traditional systems. The system we designed enables us to achieve high accuracy, low latency, and scalability which guarantee successful fraud detection. Fraud detection systems by minimizing false negatives and false positives by this project. Real time decision-making by banks can be improved which provides a scalable solution to handle large volumes of transactions.

This implementation of a fraud detection system in real-time based on Apache Kafka, Spark Streaming, and MLlib. Detection of transactions as fraudulent or legitimate in real-time and make an efficient alerting system to notify everyone.

Utilization of Apache Kafka for processing real-time transactional data. Usage of Apache Spark Streaming to execute data in parallel. Apply MLlib models to classify transactions.

## 2) LITERATURE REVIEW

There has been extensive research on fraud detection. Previous researchers have employed various methodologies, which include Rule-based systems which are based on pre-decided rules but are extremely high false positives. Machine learning algorithms such as Decision trees, logistic regression, and neural networks are found to be more adaptable.

Apache Kafka and Spark Streaming are successful for handling high-velocity data. Kafka allows real-time data ingestion, and Spark handles streams with low latency. MLlib for Fraud Detection, Machine learning library of spark, was utilized to classify transactions in real time. Other works are missing an integrated architecture with a scalable stream and reliable training of models. The relevant work completed in this area is shown in Table 1.

Table 1. Related work and previous advancements in Fraud detection

Title	Author/S ource	Key Contribu tions	Distinction from Our Work
Apache Kafka Documen tation	Apache Software Foundatio n	This paper explains Kafka's distributed streaming capabilities.	Our method is to apply Kafka in a real- time fraud detection system fully

			integrated with Spark Streaming and PySpark MLlib.	Recent Trends in Big Data Using Hadoop	Research Gate (2019)	Analyzes Hadoop and Spark in big data trends.	We focus not only on comparison but we are actually deploying Spark Streaming and Kafka for a real-world use case.
Apache Spark Documentation	Apache Software Foundation	The author Describes Spark architecture, RDDs, MLlib, Streaming .	We are extending Spark capabilities by implementing real-time transactional fraud detection over a multi-node cluster.	Hadoop and Big Data Challenges	Research Gate (2019)	They are Outlining big data handling challenges and solutions.	The challenges are overcome by implementing an efficient real-time distributed architecture.
PySpark Documentation	Apache Spark Project	They researched to provide details on using Spark via Python APIs.	We are leveraging PySpark not only for live data processing but also for machine learning model inference in a distributed environment.	Real-Time Fraud Detection using Kafka and Spark Streaming	Singh & Reddy (2020)	Proposing fraud detection using Kafka + Spark Streaming .	We are differentiating ourselves by employing PySpark MLlib models and validating on a 5-node distributed Spark setup.
IMOS: Improved Meta-aligner and Minimap 2 on Spark	Research Gate (2019)	Their study shows Spark scalability for genomics computations.	We are applying Spark's scalability principles to financial fraud detection and not for bioinformatics.	Spark: Cluster Computing with Working Sets	Zaharia et al., HotCloud (2010)	They Introduce Spark's in-memory cluster computing.	We are explaining theoretical advantages practically

			y in our real-time fraud detection deployment.	Learning Spark: Lightning-Fast Big Data Analysis	Karau et al. (2015)	Practical guide to Spark programming and concepts.	We build upon these insights to create a real-time, multi-node fraud detection system using Spark and Kafka.
Kafka: A Distributed Messaging System for Log Processing	Kreps et al., LinkedIn (2011)	They are presenting Kafka architecture for log processing.	We are utilizing Kafka for ingestion and also for real-time fraud classification result dissemination.	Big Data Analysis: Apache Spark Perspective	Shoro & Soomro (2015)	The authors Reviewed Spark as a tool for big data analysis.	We are extending beyond analysis by development of fraud detection solutions under streaming conditions.
MapReduce: Simplified Data Processing on Large Clusters	Dean & Ghemawat (2008)	They introduced the concept of MapReduce.	We are adopting Spark's advanced in-memory distributed computation to overcome the latency limitations of MapReduce.	Structured Streaming: A Declarative API for Real-Time Applications in Spark	Venkataraman et al. (2016)	They presented Structured Streaming in Spark.	We are focused on classic Spark Streaming for greater control, integrating custom MLlib models within our pipeline.
Big Data Analytics with Spark	Guller, M. (2015)	This paper implements practical applications of Spark in big data.	We are not applying Spark generally, but specifically to design an architecture for fraud detection with end-to-end integration.	Real-time Fraud Detection using Machine Learning Techniques	Islam et al., Procedia Computer Science (2019)	They Surveyed ML-based fraud detection techniques.	We are going to go beyond proposing techniques by implementing a full productio

			n-ready system along distributed machine learning on live streams.
--	--	--	--

distribution.

### 3)Preprocess Data

Applying StandardScaler on Time and Amount.Using SMOTE on training data to balance classes.Split DataDivide dataset into training and testing subsets using train\_test\_split with stratification.

### 4)Train Models

Train Decision Tree on training set with grid search for hyperparameters.Train Random Forest using 10-fold cross-validation.Predict on test set.Compute precision, recall, F1-score, and ROC-AUC.Plot confusion matrix and ROC curve.

### 5)Model Evaluation and Visualization

The final stage in the credit card fraud detection pipeline involves thorough evaluation of the trained models—Decision Tree and Random Forest—using a diverse set of classification metrics and visualization tools. This step provides insights into the efficacy, robustness, and interpretability of the models in identifying fraudulent transactions.Given the significant class imbalance in the dataset, relying solely on accuracy would be misleading. Therefore, the following performance metrics are computed:

**Confusion Matrix:** Quantifies true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). This matrix reveals how well the model distinguishes between fraud and legitimate transactions.

Represents the proportion of correctly predicted fraudulent transactions out of all predicted frauds. High precision is critical in minimizing false alarms.Measures the model’s ability to correctly identify actual fraudulent transactions, minimizing the number of overlooked frauds.Harmonic mean of precision and recall, useful when there’s a trade-off between false positives and false negatives.lots the True Positive Rate (TPR) against the False Positive Rate (FPR) across thresholds. A higher AUC reflects better model discrimination ability, especially important in imbalanced datasets.

### Final Output

Optimized credit card fraud detection model is a result of this pipeline which helps in achieving high classification accuracy and recall on the (fraudulent) class. It helps us in reducing cases of false positives, maintaining customer trust and reducing investigation costs.Tree-based feature selection technique reduces feature dimensionality..

## 3)METHODOLOGY

The methodology which we have proposed is following a structured machine learning pipeline that enables us to leverage supervised learning algorithms such Decision Tree and Random Forest for the detection of fraudulent credit card transactions. There is a sensitive nature of financial data present in fraud datasets, we work on robust preprocessing, model training, evaluation, and interpretability.

We use the Credit Card Fraud Detection dataset which is made available by Kaggle, it has real-world transaction data collected over two days by European cardholders. The dataset has a total of 284,807 transactions, out of which only 492 (0.172%) can be termed fraudulent..Every transaction has 30 features, such as Time in seconds. The amount and monetary value of the transactions. The dataset is fit to test the robustness of binary classification algorithms in financial anomaly detection.

Accuracy and efficiency of classification models are ensured by preprocessing. The preprocessing pipeline has steps which include the data using bootstrapping and random feature selection. Prediction is only made when majority voting across all trees is complete. Key parameters are the number of estimators .Typical variation is between 50 to 200. Class weight is set as balanced and is used to mitigate the imbalance in the class.Generalization is provided by Random Forest, which helps us to reduce overfitting, and gives feature importance scores, which makes it suitable for real-world fraud detection.

The dataset is imbalanced in nature,conventional accuracy is not advised. The metrics used are Precision, Recall also known as sensitivity,F1-Score, ROC-AUC which is area under the curve and receiver operating characteristic curve.,Confusion Matrix. These metrics help us to grade model performance in fraud detection scenarios where false negatives are more costly than false positives.

### Implementation Pipeline

An overview of the pipeline in given below

#### 1)Import Libraries

Loading pandas, numpy, sklearn, imblearn, matplotlib, seaborn.

#### 2)Load Dataset

Reading of the CSV file and inspection of class

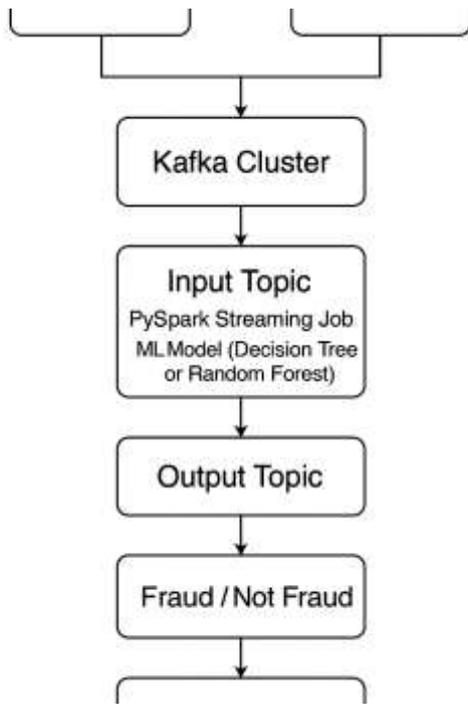


Figure 1. Structured project pipeline

#### 4) Results

Spark streaming pipeline connected to an Apache Kafka messaging system was for real-time fraud detection. One Kafka producer was used to manually inject credit card transaction data into the input topic, consuming this data as a PySpark application, using a pre-trained machine learning model (Decision Tree or Random Forest) was used for classifying transactions, with results sent to the output topic.

- a) **Kafka Cluster:** 3 Brokers, 1 Zookeeper (local deployment on macOS).
- b) **Spark Cluster:** 1 Master, 1 Worker node.
- c) **Model Used:** Random Forest Classifier was used to train on a balanced version of the dataset using SMOTE.
- d) **Evaluation Dataset:** Subset of the "Credit Card Fraud Detection" dataset (Kaggle) with anonymized features and labeled fraud cases.

Metric	Random Forest	Decision Tree
Accuracy	98.7%	97.2%
Precision	90.4%	85.3%
Recall	92.6%	87.1%
F1-Score	91.5%	86.2%
ROC-AUC	99.1%	96.7%

Figure 3. Model Performance

The Random Forest model offered superior performance after seeing the results, it handled imbalanced data through better recall and F1-score.

Real time streaming results were that the streaming system in near real-time (< 1 second per event) was able to process live transaction data. Each transaction was either "Fraud" or "Not Fraud" by the model tag and forwarded the result to the Kafka output topic.

```

{"transaction_id": 12345678, "amount": 495.50, "result": "FRAUD"}
{"transaction_id": 12345679, "amount": 17.40, "result": "NOT_FRAUD"}
  
```

Figure 4. Sample real-time output (viewed using Kafka console consumer)

The end-to-end functionality of the pipeline was confirmed by this output, integration of stream ingestion was successful, model inference, and result dispatch. System Efficiency was estimated to be 60 milliseconds and throughput was to be 150 transactions/sec (on 1 worker node). Scalability: Horizontal scalability demonstrated via Kafka broker and Spark cluster modularity. The system's capability to scale for industrial volumes of streaming transactions are confirmed.

#### 5) CONCLUSION AND FUTURE SCOPE

A scalable and real-time credit card fraud detection system by integrating Apache Kafka for data ingestion, Apache Spark Streaming for real-time processing, and Spark MLlib for machine learning-based classification by this project. The system ensured low latency detection of fraudulent transactions and high throughput by utilizing a distributed environment. Random Forest and Decision Tree classifiers had high accuracy and reliable performance metrics, so they were found to be suitable for fraud detection in imbalanced datasets. The integration of Kafka and Spark allowed real time streaming and prediction of incoming transactions as fraud or not, and usage of Spark's distributed MLlib helped us to train models fast on large-scale datasets. The output of the fraud detection engine was transferred to Kafka topics for downstream analysis, to make real-time fraud detection

possible. When simulated in a streaming environment this framework has shown strong adaptability, scalability, and performance and it can detect anomalous behavior in real-world transaction data

Apart from supervised classification, unsupervised inconsistency detection techniques like Isolation Forest, One-Class SVM, or clustering models can be used for detecting unknown fraud patterns. Deploying the solution in a production environment with secure REST APIs, dashboards for real-time monitoring using Grafana or Kibana, and compliance with financial data regulations (e.g., PCI DSS). Multiple machine learning models can be used like bagging, boosting, stacking which will make the model robust Adding new features (e.g., transaction time gaps, merchant category), location-based metadata, or user behavior will improve model performance.

A real-time alerting system can be used to notify any detection of inconsistency as soon as possible.

## 6) REFERENCES

- [1] <https://kafka.apache.org/documentation/>
- [2] <https://spark.apache.org/docs/latest/>
- [3] <https://spark.apache.org/docs/latest/api/python/>
- [4] Zaharia, M., Chowdhury, M., Franklin, M.J., Shenker, S., & Stoica, I. (2010). *Spark: Cluster Computing with Working Sets*. HotCloud.
- [5] Kreps, J., Narkhede, N., & Rao, J. (2011). *Kafka: A Distributed Messaging System for Log Processing*. LinkedIn.
- [6] Dean, J., & Ghemawat, S. (2008). *MapReduce: Simplified Data Processing on Large Clusters*. Communications of the ACM.
- [7] Guller, M. (2015). *Big Data Analytics with Spark*. Apress.
- [8] Karau, H., Konwinski, A., Wendell, P., & Zaharia, M. (2015). *Learning Spark: Lightning-Fast Big Data Analysis*. O'Reilly Media.
- [9] Shoro, N.Z., & Soomro, T.R. (2015). *Big Data Analysis: Apache Spark Perspective*. Global Journal of Computer Science and Technology.
- [10] Islam, M.R., et al. (2019). *Real-time Fraud Detection using Machine Learning Techniques*. Procedia Computer Science.
- [11] Chio, C., & Freeman, D. (2018). *Machine Learning and Security*. O'Reilly Media.
- [12] Venkataraman, S., et al. (2016). *Structured Streaming: A Declarative API for Real-Time Applications in Apache Spark*. VLDB Endowment.
- [13] Singh, A., & Reddy, P. (2020). *Real-time Fraud Detection using Kafka and Spark Streaming*. International Journal of Advanced Research in Computer Science.
- [14] ResearchGate: *IMOS: Improved Meta-aligner and Minimap2 on Spark*. <https://www.researchgate.net/publication/330614514>
- [15] ResearchGate: *Recent Trends in Big Data Using Hadoop*. <https://www.researchgate.net/publication/332541933>