AI-Driven Risk Management and Fraud Detection in High-Frequency Trading Environments

Yash Jani

Sr. Software Engineer, Fremont, California, USA
Email: *yjani204[at]gmail.com*

Abstract: ***High-frequency trading (HFT) environments, characterized by the rapid execution of trades and large volumes of data, demand sophisticated and real-time risk management and fraud detection solutions. Traditional systems need help to keep up with the velocity and complexity of data, leaving gaps that can be exploited. This paper proposes an AI-driven architecture leveraging machine learning models to enhance risk management and fraud detection in HFT environments. Implemented using Amazon SageMaker for AI/ML processing, the architecture is designed to be scalable, efficient, and capable of real-time decision-making. The study presents the detailed architecture, discusses each component's role, and evaluates the system's performance in mitigating risks and detecting fraud in live trading scenarios.***

Keywords: AI-driven risk management, fraud detection, high-frequency trading, machine learning, Amazon SageMaker, real-time data processing, predictive modeling, anomaly detection, algorithmic trading, decision support systems

1. Introduction

1.1 Overview of High-Frequency Trading (HFT)

1) **Definition of High-Frequency Trading (HFT)**
High-frequency trading (HFT) uses sophisticated algorithms and robust computing systems to execute many trades at extremely high speeds. These trades often occur in milliseconds or microseconds, allowing traders to capitalize on small price movements in the market.[1]
2) **Importance of HFT in financial markets**
HFT plays a critical role in modern financial markets by providing liquidity, reducing bid-ask spreads, and contributing to market efficiency. However, the speed and volume of transactions can also lead to significant risks, including market manipulation and systemic disruptions.[2]
3) **Challenges associated with HFT**
The rapid pace of HFT presents several challenges, particularly in risk management and fraud detection. Traditional risk management systems, designed for slower, less complex trading environments, often fail to provide the necessary real-time oversight required in HFT. This gap increases the potential for market manipulation, such as spoofing or layering, which can lead to significant financial losses and undermine market integrity.[3]

1.2 The Role of AI in Financial Services

1) **Introduction to AI and machine learning in finance**
Artificial Intelligence (AI) and Machine Learning (ML) are revolutionizing the financial sector by enabling the analysis of large datasets, identifying patterns, and making predictions with greater accuracy than ever before. These technologies are precious in HFT, where processing data quickly and accurately is paramount.[4]
2) **Benefits of AI in risk management and fraud detection**
AI offers several advantages over traditional methods in risk management and fraud detection, including:
   a) **Speed**: AI can analyze data and make decisions in real time, which is essential in the fast-paced world of HFT.[5]
   b) **Accuracy**: Machine learning models can identify complex patterns and correlations in data that may not be apparent to human analysts.[6]
   c) **Adaptability**: AI systems can learn from new data and adjust their models accordingly, making them more effective at detecting emerging threats and changes in market conditions.[4]
3) **Current applications of AI in HFT**
AI is already being used in HFT for various purposes, including:
   a) **Algorithmic Trading**: Developing trading strategies that can adapt to real-time market conditions.[7]
   b) **Sentiment Analysis**: Natural language processing (NLP) analyzes news and social media for market sentiment.[8]
   c) **Market Prediction**: Predicting price movements based on historical and real-time market information.[9]

1.3 Objective of the Study

1) **Purpose of proposing an AI-driven architecture**
The primary objective of this study is to propose an AI-driven architecture specifically designed for risk management and fraud detection in HFT environments. The proposed architecture leverages Amazon SageMaker for AI/ML processing, offering a scalable and efficient solution that can operate in real time.[10]
2) **Overview of Amazon SageMaker's role in the system**
Amazon SageMaker is a cloud-based platform that enables the development, training, and deployment of machine learning models at scale. In the proposed architecture, SageMaker plays a crucial role in processing the vast amounts of data generated by HFT, training models to predict risks and detect fraud, and deploying these models to allow for real-time decision-making.[11]

3) Scope and goals of the research

This paper aims to:

a) Present a detailed architecture for AI-driven risk management and fraud detection in HFT.

b) Discuss the implementation of the architecture using Amazon SageMaker and other AWS services.

c) Evaluate the architecture's performance in real-world HFT scenarios, focusing on its ability to mitigate risks and detect fraud.

2. Architecture Overview

2.1 Data Sources

1) Types of data sources in HFT

a) **Market Data Feeds**

Market data feeds provide real-time information on the prices of stocks, bonds, commodities, and other financial instruments. These feeds are essential for HFT algorithms, which rely on up-to-the-second data to make trading decisions.[12]

b) **Trade Logs**

Trade logs contain detailed records of every transaction executed in the market, including the time, price, volume, and order type. These logs are invaluable for analyzing market behavior and detecting anomalies that may indicate fraud.[13]

c) **Financial News and Social Media**

Unstructured data from news outlets, financial blogs, and social media platforms can provide insights into market sentiment and potential factors that could influence prices. By incorporating this data, AI models can make more informed predictions and detect possible risks associated with sudden market changes.[14][15]

2) Importance of data diversity in AI models

The effectiveness of AI models in HFT depends on the diversity and richness of the data on which they are trained. By incorporating multiple data sources—from structured data like market feeds to unstructured data like news articles—AI models can develop a more comprehensive understanding of the market, leading to more accurate predictions and better risk management.[16]

3) Challenges in collecting and managing HFT data

Collecting and managing data in HFT is challenging due to the sheer volume and speed at which data is generated. Ensuring data quality, dealing with missing or noisy data, and integrating data from diverse sources are significant hurdles that must be overcome to create effective AI models.[2]

2.2 Data Ingestion

1) Process of data ingestion in HFT environments

a) AWS Glue for ETL processes AWS Glue is a managed ETL (Extract, Transform, Load) service that simplifies the process of preparing data for analysis. In the proposed architecture, AWS Glue is used to extract data from various sources, transform it into a suitable format, and load it into a data lake for further processing.[17]

b) Kinesis Data Streams for real-time data ingestion Amazon Kinesis Data Streams is a real-time data streaming service that enables the ingestion and processing of large data streams. This service benefits HFT, where real-time data processing is critical for timely trading decisions and risk identification.[18]

2) Pre-processing of data for AI/ML models

Before data can be used to train AI/ML models, it must be pre-processed to ensure it is clean, consistent, and in the correct format. This includes handling missing values, normalizing data, and converting unstructured data into a structured format that machine learning algorithms can analyze.[19]

3) Ensuring data quality and consistency

Maintaining high data quality is essential for the accuracy of AI models. This involves regularly validating data sources, monitoring for anomalies, and implementing processes to correct errors or inconsistencies in the data.[20]

2.3 Data Storage

1) Storage solutions for high-volume HFT data

a) Amazon S3 for raw and processed data Amazon S3 is a scalable object storage service ideal for storing large volumes of raw and processed data. In the proposed architecture, S3 stores data collected from various sources and data processing and analysis results.[21]

b) Amazon Redshift for data warehousing Amazon Redshift is a fully managed data warehouse that allows for fast querying and analysis of large datasets. It is used in the architecture to store and analyze historical trading data, which can be used to train AI models and generate insights.[22]

c) Amazon DynamoDB for real-time data access Amazon DynamoDB is a NoSQL database service providing low-latency real-time data access. It is used in the architecture to store and retrieve data that needs to be accessed quickly, such as real-time trading information and model predictions.[23]

2) Integration of storage solutions for efficient data retrieval

The proposed architecture integrates these storage solutions to ensure data can be retrieved and processed efficiently. Data flows seamlessly from ingestion to storage and is readily available for analysis and model training.[24]

3) Scalability and performance considerations in data storage

As the volume of data in HFT grows, scalability becomes a critical consideration. The architecture is designed to scale quickly by leveraging AWS's cloud-based storage solutions, ensuring performance remains high even as data volumes increase.[25]

2.4 AI/ML Processing with Amazon SageMaker

1) Overview of Amazon SageMaker's capabilities

Amazon SageMaker is a fully managed service that allows every developer and data scientist to quickly build, train, and deploy machine learning models. It supports various machine learning frameworks and tools, making it versatile for AI applications.

2) Model training and deployment processes

a) Supported frameworks (TensorFlow, PyTorch, etc.) SageMaker supports popular machine learning frameworks such as TensorFlow, PyTorch, and Scikit-learn, enabling developers to build their models using familiar tools and libraries.[26]

b) Real-time inference and model updates SageMaker facilitates real-time inference, allowing models to be deployed and used to make predictions on live data streams. It also supports continuous model updates, ensuring that models remain accurate and relevant as new data becomes available.[27]

3) Benefits of AutoML in continuous model improvement

AutoML capabilities in SageMaker allow for the automated tuning and retraining of models, improving their performance over time. This feature is precious in HFT, where market conditions change rapidly, and models must adapt to new patterns.[28]

4) Integration with other AWS services for seamless operation

SageMaker integrates seamlessly with other AWS services, such as S3 for data storage and Kinesis for data ingestion, creating a cohesive and efficient pipeline for AI/ML processing in HFT environments.[29]

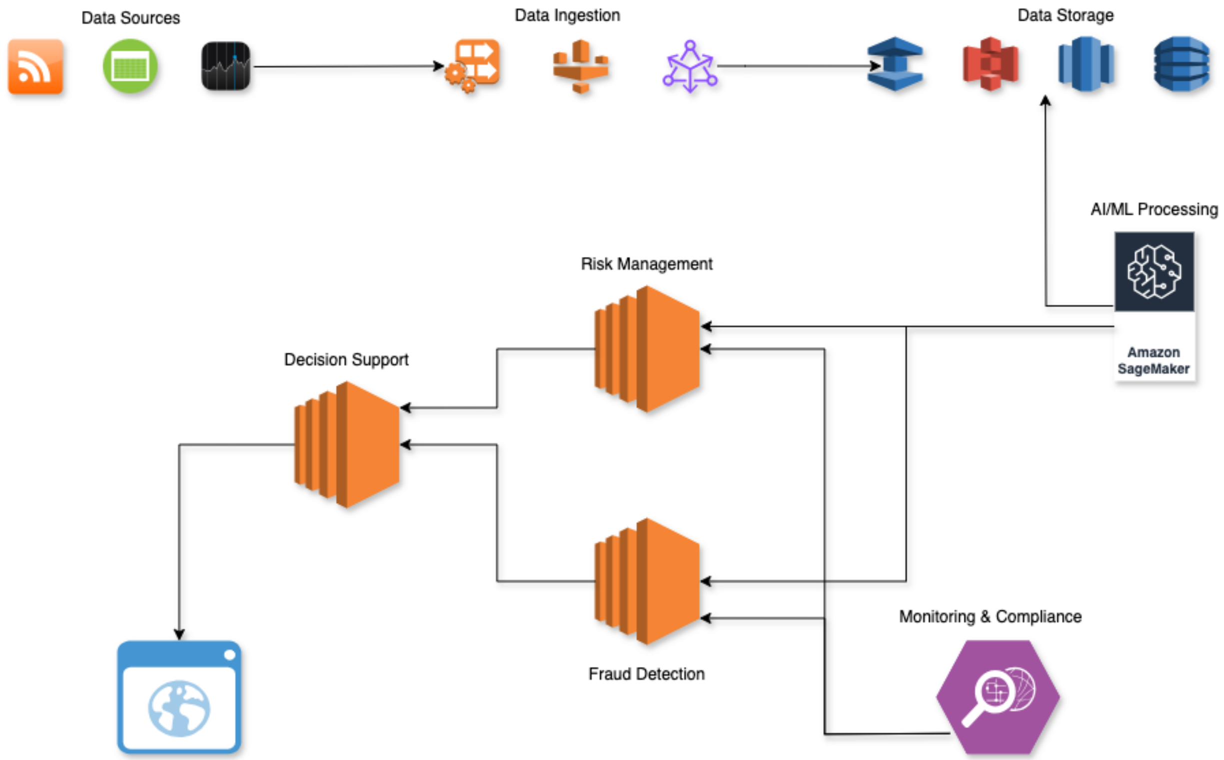

Figure 1: Architecture Daigram

3. Implementation Details

3.1 Risk Management Module

1) Design of the risk management system

The risk management module in the proposed architecture is designed to monitor and assess risk in real time. It uses machine learning models to predict potential market movements and assess the risk associated with various trading strategies. The module is built to handle the high-speed data processing requirements of HFT, ensuring that risks are identified and managed promptly.[30]

2) Predictive modeling for risk assessment

a) Factors considered in risk prediction The AI models used in the risk management module consider various factors, including market volatility, trading volume, historical price movements, and external factors such as news events. These factors are used to predict potential risks and guide trading decisions.[31]

b) Techniques used for predictive modeling The models leverage advanced machine learning techniques, such as time series analysis and deep learning, to make accurate predictions. These techniques allow the system to identify patterns in the data that may not be immediately apparent to human analysts.[32]

3) Risk scoring and its impact on decision-making

Each trade is assigned a risk score based on the model's predictions. This score reflects the likelihood of the trade resulting in a loss or exposing the firm to significant risk. Traders and risk managers can use these scores to decide whether to proceed with a trade or adjust their strategies.[33]

4) Scenario analysis for proactive risk management

The risk management module also includes scenario analysis capabilities, allowing traders to simulate various market conditions and assess their potential impact on the portfolio. This helps traders prepare for different market scenarios and mitigate risks before they materialize.[34]

3.2 Fraud Detection Module

1) Challenges in fraud detection within HFT

Detecting fraud in HFT is particularly challenging due to the speed and volume of trades. Fraudulent activities, such as spoofing, layering, and wash trading, can occur within milliseconds, making them difficult to detect using traditional methods.[35]

2) Anomaly detection using AI models

a) Identification of abnormal trading patterns
The fraud detection module uses AI models to identify abnormal trading patterns that may indicate fraudulent activity. These models are trained on historical trading data and can detect subtle anomalies that may go unnoticed by human analysts.[36]

b) Techniques for detecting specific fraud types (e.g., spoofing)
Specific AI techniques, such as clustering and outlier detection, are used to detect particular types of fraud. For example, clustering can identify groups of trades that exhibit similar suspicious behavior, while outlier detection can highlight trades that deviate significantly from the norm.

3) Behavioral analysis for long-term fraud detection

The fraud detection module also includes behavioral analysis tools that monitor trading behaviors over time. By tracking patterns and trends in trading behavior, the system can identify long-term fraud schemes that may not be apparent in individual trades.

4) Real-time alerts and response mechanisms

When the system detects potential fraud, it generates real-time alerts, allowing immediate investigation and intervention. These alerts are critical for minimizing the impact of fraudulent activities and maintaining market integrity.

3.3 Decision Support System

1) Integration of risk management and fraud detection outputs

The decision support system integrates the risk management and fraud detection module outputs to provide a comprehensive view of the trading environment. By combining these insights, the system can offer more accurate and actionable recommendations to traders and risk managers.

2) Dashboard interface for real-time insights

a) Customizability of the dashboard for different user needs
The decision support system features a customizable dashboard interface that allows users to view the information most relevant to their needs. Traders can configure the dashboard to display key metrics, risk scores, and alerts, enabling them to make quick, informed decisions.

b) Key metrics displayed on the dashboard. The dashboard displays a range of key metrics, including real-time market data, risk scores, fraud alerts, and performance indicators. These metrics provide traders with a clear overview of their trading environment and the potential risks they face.

3) Automated decision-making capabilities

In addition to providing insights, the decision support system can automate specific decision-making processes. For example, it can automatically halt trades that exceed predefined risk thresholds or initiate investigations into suspicious activities. This automation helps reduce traders' cognitive load and promptly ensures appropriate actions are taken.

4) Backtesting and simulation tools for strategy evaluation

The decision support system includes backtesting and simulation tools that allow traders to evaluate their strategies against historical data. By simulating different market conditions, traders can assess the robustness of their approach and make adjustments as needed to improve performance.

4. Evaluation and Results

4.1 Performance Metrics

1) Criteria for evaluating the AI-driven system

a) Accuracy of predictions: The accuracy of the AI-driven system is evaluated based on its ability to predict market movements and detect fraud correctly. This is measured by comparing the system's predictions with actual market outcomes and known instances of fraud.

b) Latency in processing and decision-making: Latency is a critical metric in HFT, where decisions must be made in milliseconds. The system's latency is measured by the time it takes to process data, generate insights, and execute decisions. Lower latency indicates a more responsive and efficient system.

c) Scalability under increasing data loads: Scalability is assessed by the system's ability to maintain performance as the volume of data increases. This is particularly important in HFT, where data volumes can increase. A scalable system can handle larger datasets without significant degradation in speed or accuracy.

2) Comparison of performance metrics with traditional systems

The performance metrics of the AI-driven system are compared with those of traditional risk management and fraud detection systems. This comparison highlights AI's advantages, particularly in speed, accuracy, and adaptability.

3) Importance of continuous monitoring and improvement

Continuous monitoring and improvement are essential to maintaining the system's effectiveness. The system's performance metrics are regularly reviewed, and the AI models are updated to ensure they remain accurate and responsive to market changes.

4.2 Comparison with Traditional Methods

a) Speed of data processing and decision-making
The AI-driven system processes data and makes

decisions significantly faster than traditional systems. This speed is critical in HFT, where delays of even a few milliseconds can lead to missed opportunities or increased risk.

b) Accuracy and reliability of risk assessments and fraud detection
The AI-driven system demonstrates higher accuracy and reliability in risk assessments and fraud detection. This is due to its ability to analyze large datasets, identify complex patterns, and continuously learn from new data.

c) Adaptability of AI models to new data and market conditions
Unlike traditional systems, which often require manual updates, AI models can automatically adapt to new data and change market conditions. This adaptability ensures that the system remains effective even as the market evolves.

d) Benefits of AI over traditional methods in HFT
The AI-driven system offers several benefits over traditional methods, including faster decision-making, more accurate predictions, and greater adaptability. These advantages make it particularly well-suited to the demands of high-frequency trading.

5. Challenges and Future Work

5.1 Challenges in AI Implementation

1) **Data quality and its impact on model accuracy**
One of the primary challenges in implementing AI in HFT is ensuring data quality. Only accurate, complete, noisy data can lead to good model performance and correct predictions. Ensuring high-quality data is essential for the accuracy and reliability of AI models.[37]
2) **Interpretability of AI models in financial contexts**
AI models, particularly those based on deep learning, can be complex and challenging to interpret. This lack of transparency can be a barrier to their adoption in finance, where regulators and stakeholders require clear explanations of how decisions are made.
3) **Regulatory compliance and its implications for AI systems**
Financial markets are heavily regulated, and AI systems must comply with various rules and standards. Ensuring that AI-driven decisions meet regulatory requirements is a significant challenge, particularly as regulations continue to evolve.
4) **Technical challenges in deploying AI at scale**
Deploying AI systems at scale in HFT environments presents several technical challenges, including managing large volumes of data, ensuring low latency, and maintaining system reliability. These challenges must be addressed to realize AI's benefits in HFT fully.[38]

5.2 Future Directions

1) **Exploration of advanced AI techniques (e.g., reinforcement learning)**
Future research could explore the use of advanced AI techniques, such as reinforcement learning, to improve decision-making in HFT. These techniques can enhance the system's ability to adapt to dynamic and uncertain market conditions.
2) **Integration of alternative data sources for enhanced predictions**
Incorporating alternative data sources, such as satellite imagery, weather data, and sentiment analysis, could further enhance the accuracy of AI models in predicting market movements and detecting fraud. These additional data sources could provide new insights and improve the system's robustness.[39]
3) **Improving the interpretability of complex AI models**
Developing methods to make AI models more interpretable without sacrificing performance is a key area for future research. Improved interpretability could increase trust in AI systems and facilitate their adoption in the financial industry.
4) **Future research on the ethical implications of AI in finance**
As AI becomes more prevalent in finance, the ethical implications of its use must be considered. Future research could explore fairness, accountability, and transparency in AI-driven financial systems.

6. Conclusion

6.1 Summary of Findings

1) **Recap of the proposed AI-driven architecture**
This paper presents an AI-driven architecture designed to address the challenges of risk management and fraud detection in high-frequency trading environments. The architecture leverages Amazon SageMaker for AI/ML processing, providing a scalable and efficient solution that operates in real-time.
2) **Key benefits of the system in HFT environments**
The proposed system offers several key benefits in HFT environments, including faster decision-making, more accurate risk assessments, and enhanced fraud detection capabilities. These advantages make the system well-suited to the demands of modern financial markets.
3) **Importance of AI in modern financial markets**
AI is increasingly important in financial markets, offering new ways to manage risk, detect fraud, and make informed trading decisions. The proposed architecture demonstrates AI's potential to improve the efficiency and integrity of high-frequency trading.

6.2 Final Thoughts

1) **Future potential of AI in risk management and fraud detection**
As AI technology evolves, its potential to enhance risk management and fraud detection in HFT will only grow. Continued research and development could lead to even more sophisticated and effective systems.
2) **The evolving role of AI in the financial industry**
AI will likely play an increasingly central role in the financial industry, transforming how markets operate and risks are managed. As AI systems become more advanced and widely adopted, they will profoundly impact the future of finance.[40]

References

[1] R. S. Miller and G. Shorter, "High Frequency Trading: Overview of Recent Developments," Apr. 4, 2016. [Online]. Available: https://digital.library.unt.edu/ark:/67531/metadc847594/m2/1/high_res_d/R44442_2016Apr01.pdf

[2] R. Cooper, J. J. M. Seddon, and B. V. Vliet, "High-Frequency Trading and Conflict in the Financial Markets," SAGE Publishing, vol. 32, no. 3, pp. 270-282, Sep. 2017. [Online]. Available: https://doi.org/10.1057/s41265-016-0031-5

[3] I. Aldridge, "Market Microstructure and the Risks of High-Frequency Trading," RELX Group (Netherlands), Jan. 1, 2013. [Online]. Available: https://doi.org/10.2139/ssrn.2294526

[4] J. Sen, R. Sen, and A. Dutta, "Machine Learning in Finance-Emerging Trends and Challenges," Cornell University, Jan. 1, 2021. [Online]. Available: https://doi.org/10.48550/arXiv.2110

[5] Z. Rouhollahi, "Towards Artificial Intelligence Enabled Financial Crime Detection," Cornell University, Jan. 1, 2021. [Online]. Available: https://doi.org/10.48550/arxiv.2105.10866

[6] K. Kerwin and N. D. Bastian, "Stacked Generalizations in Imbalanced Fraud Data Sets using Resampling Methods," Cornell University, Jan. 1, 2020. [Online]. Available: https://doi.org/10.48550/arXiv.2004

[7] X. Jia and R. Y. Lau, "The Control Strategies for High Frequency Algorithmic Trading," Aug. 1, 2018. [Online]. Available: https://doi.org/10.1109/ccsse.2018.8724810

[8] N. Anantrasirichai and D. Bull, "Artificial intelligence in the creative industries: a review," Springer Science+Business Media, vol. 55, no. 1, pp. 589-656, Jul. 2, 2021. [Online]. Available: https://doi.org/10.1007/s10462-021-10039-7

[9] S. Kazemian, S. Zhao, and G. Penn, "Evaluating Sentiment Analysis Evaluation: A Case Study in Securities Trading," Jan. 1, 2014. [Online]. Available: https://doi.org/10.3115/v1/w14-2620

[10] P. Bedi, S. B. Goyal, and J. Kumar, "Basic Structure on Artificial Intelligence: A Revolution in Risk Management and Compliance," Dec. 3, 2020. [Online]. Available: https://doi.org/10.1109/iciss49785.2020.9315986

[11] P. Das, V. Perrone, N. Ivkin, T. Bansal, Z. Karnin, H. Shen, I. Shcherbatyi, Y. Elor, W. Wu, A. Zolic, T. Lienart, A. Tang, A. Ahmed, J. B. Faddoul, R. Jenatton, F. Winkelmolen, P. Gautier, L. Dirac, A. Perunicic, and L. Rouesnel, "Amazon SageMaker Autopilot: a white box AutoML solution at scale," Cornell University, Jan. 1, 2020. [Online]. Available: https://doi.org/10.48550/arxiv.2012.08483

[12] C. Caglio and S. Mayhew, "Equity trading and the allocation of market data revenue," Elsevier BV, vol. 62, pp. 97-111, Jan. 1, 2016. [Online]. Available: https://doi.org/10.1016/j.jbankfin.2015.10.002

[13] W. Dai, C. Dai, K. R. Choo, C. Cui, D. Zou, and H. Jin, "SDTE: A Secure Blockchain-Based Data Trading Ecosystem," Institute of Electrical and Electronics Engineers, vol. 15, pp. 725-737, Jan. 1, 2020. [Online]. Available: https://doi.org/10.1109/tifs.2019.2928256

[14] J. Wang, H. Zhu, J. Shen, Y. Cao, and B. Liu, "Dual-CLVSA: a Novel Deep Learning Approach to Predict Financial Markets with Sentiment Measurements," Dec. 1, 2021. [Online]. Available: https://doi.org/10.1109/icmla52953.2021.00062

[15] F. Ronaghi, M. Salimibeni, F. Naderkhani, and A. Mohammadi, "COVID19-HPSMP: COVID-19 Adopted Hybrid and Parallel Deep Information Fusion Framework for Stock Price Movement Prediction," Cornell University, Jan. 1, 2021. [Online]. Available: https://doi.org/10.48550/arXiv.2101

[16] S. I. Lee and S. J. Yoo, "Multimodal Deep Learning for Finance: Integrating and Forecasting International Stock Markets," Cornell University, Jan. 1, 2019. [Online]. Available: https://doi.org/10.48550/arxiv.1903.06478

[17] C. Auer, M. Dolfi, A. C. P. L. F. D. Carvalho, C. B. Ramis, and P. Staar, "Delivering Document Conversion as a Cloud Service with High Throughput and Responsiveness," Jul. 1, 2022. [Online]. Available: https://doi.org/10.1109/cloud55607.2022.00060

[18] F. Xhafa, B. Kilic, and P. Krause, "Evaluation of IoT stream processing at edge computing layer for semantic data enrichment," Elsevier BV, vol. 105, pp. 730-736, Apr. 1, 2020. [Online]. Available: https://doi.org/10.1016/j.future.2019.12.031

[19] A. M. Truong, A. Walters, J. Goodsitt, K. Hines, C. B. Bruss, and R. Farivar, "Towards Automated Machine Learning: Evaluation and Comparison of AutoML Approaches and Tools," Nov. 1, 2019. [Online]. Available: https://doi.org/10.1109/ictai.2019.00209

[20] N. Tang, "Big Data Cleaning," Springer Science+Business Media, pp. 13-24, Jan. 1, 2014. [Online]. Available: https://doi.org/10.1007/978-3-319-11116-2_2

[21] P. B. Val, N. Audsley, A. Wellings, I. Gray, and N. Fernández, "Architecting Time-Critical Big-Data Systems," IEEE Computer Society, vol. 2, no. 4, pp. 310-324, Dec. 1, 2016. [Online]. Available: https://doi.org/10.1109/tbdata.2016.2622719

[22] R. Gupta, H. Gupta, U. Nambiar, and M. Mohania, "Efficiently querying archived data using Hadoop," Oct. 26, 2010. [Online]. Available: https://doi.org/10.1145/1871437.1871606

[23] A. Verma, N. Sethi, and N. Jain, "Beyond Hadoop for e-commerce Big Data Analysis through Amazon," Dec. 1, 2018. [Online]. Available: https://doi.org/10.1109/icacat.2018.8933660

[24] Y. Zhao, K. Bandyopadhyay, and C. C. Barnes, "Predictive Maintenance Information Systems: The Underlying Conditions and Technological Aspects," IGI Global, vol. 16, pp. 54-72, Jan. 1, 2020. [Online]. Available: https://doi.org/10.4018/IJEIS.2020040104

[25] A. Ajay and N. Sidnal, "High Frequency Trading with Complex Event Processing," Dec. 1, 2016. [Online]. Available: https://doi.org/10.1109/hipcw.2016.014

[26] C. Karakus, R. Huilgol, F. Wu, A. Subramanian, C. Daniel, D. Çavdar, T. Xu, H. Chen, A. Rahnama, and L. Á. Q. Arias, "Amazon SageMaker Model Parallelism: A General and Flexible Framework for Large Model Training," Cornell University, Jan. 1, 2021. [Online]. Available: https://doi.org/10.48550/arXiv.2111

[27] E. Liberty, Z. Karnin, B. Xiang, L. Rouesnel, B. Coşkun, R. Nallapati, J. Delgado, A. Sadoughi, Y. Astashonok, P. Das, C. Balioglu, S. Chakravarty, M. Jha, P. Gautier, D. Arpin, T. Januschowski, V. Flunkert, Y. Wang, J. Gasthaus, and A. Smola, "Elastic Machine Learning Algorithms in Amazon SageMaker," May 31, 2020. [Online]. Available: https://doi.org/10.1145/3318464.3386126

[28] M. Wajahat, A. Karve, A. Kochut, and A. Gandhi, "MLscale: A machine learning based application-agnostic autoscaler," Elsevier BV, vol. 22, pp. 287-299, Jun. 1, 2019. [Online]. Available: https://doi.org/10.1016/j.suscom.2017.10.003

[29] Y. Gan, M. Liang, S. Dev, D. Lo, and C. Delimitrou, "Sage: practical and scalable ML-driven performance debugging in microservices," Apr. 17, 2021. [Online]. Available: https://doi.org/10.1145/3445814.3446700

[30] P. Ganesh and P. Rakheja, "VLSTM: Very Long Short-Term Memory Networks for High-Frequency Trading," Cornell University, Jan. 1, 2018. [Online]. Available: https://doi.org/10.48550/arxiv.1809.01506

[31] P. Oncharoen and P. Vateekul, "Deep Learning Using Risk-Reward Function for Stock Market Prediction," Dec. 8, 2018. [Online]. Available: https://doi.org/10.1145/3297156.3297173

[32] S. Singh, T. K. Madan, J. Kumar, and A. K. Singh, "Stock Market Forecasting using Machine Learning: Today and Tomorrow," Jul. 1, 2019. [Online]. Available: https://doi.org/10.1109/icicict46008.2019.8993160

[33] A. Alonso-Robisco and J. M. C. Martínez, "Measuring the model risk-adjusted performance of machine learning algorithms in credit default prediction," Springer Nature, vol. 8, no. 1, Jul. 12, 2022. [Online]. Available: https://doi.org/10.1186/s40854-022-00366-1

[34] B. Jin, "An intelligent algorithmic trading based on a risk-return reinforcement learning algorithm," Cornell University, Jan. 1, 2022. [Online]. Available: https://doi.org/10.48550/arxiv.2208.10707

[35] G. K. Kulatilleke, "Challenges and Complexities in Machine Learning based Credit Card Fraud Detection," Cornell University, Jan. 1, 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2208

[36] R. A. L. Torres and M. Ladeira, "A proposal for online analysis and identification of fraudulent financial transactions," Dec. 1, 2020. [Online]. Available: https://doi.org/10.1109/icmla51294.2020.00047

[37] S. Koschnicke, V. Grossmann, C. Starke, and M. Schimmler, "Quality and consistency assurance of quote data for algorithmic trading strategies," Mar. 1, 2014. [Online]. Available: https://doi.org/10.1109/cifer.2014.6924081

[38] J. Bosch, I. Crnković, and H. H. Olsson, "Engineering AI Systems: A Research Agenda," Cornell University, Jan. 1, 2020. [Online]. Available: https://doi.org/10.48550/arxiv.2001.07522

[39] W. Jiang, "Applications of deep learning in stock market prediction: Recent progress," Elsevier BV, vol. 184, p. 115537, Dec. 1, 2021. [Online]. Available: https://doi.org/10.1016/j.eswa.2021.115537

[40] A. Ashta and H. Herrmann, "Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance," Wiley, vol. 30, no. 3, pp. 211-222, May 1, 2021. [Online]. Available: https://doi.org/10.1002/jsc.2404