ISSN 2348-1196 (print)
International Journal of Computer Science and Information Technology Research ISSN 2348-120X (online)
Vol. 14, Issue 2, pp: (23-30), Month: April – June 2026, Available at: www.researchpublish.com

# AI-Driven Portfolio Optimization: Integrating Sentiment Analysis, Reinforcement Learning, and Personalised Advisory

Sana Perween[1], Aayush Rathi[2], Tanush Vora[3]

[1,2,3]Department of Electronics and Telecommunication Engineering, Thakur College of Engineering and Technology, Mumbai, India

DOI: https://doi.org/10.5281/zenodo.19627138

Published Date: 17-April-2026

*Abstract:* **Traditional portfolio optimization methods, such as Mean-Variance Optimization (MVO) and the Capital Asset Pricing Model (CAPM), rely on static historical data and linear assumptions about asset returns and correlations, making them ineffective during volatile market conditions, when relationships between assets break down and historical patterns fail to generalise. These classical approaches also cannot incorporate unstructured data sources such as news sentiment, social media discourse, or macroeconomic indicators, leaving valuable market signals untapped. This paper presents FinanceWiz, an integrated AI-powered portfolio optimization platform that combines LSTM-based return prediction, sentiment analysis from news and social media, reinforcement learning for dynamic rebalancing, and a personalised AI investment coach. The system was evaluated on 10 years of multi-asset data across equities, forex, commodities, and cryptocurrencies. The AI advisory module showed that it could respond to different situations. Our method works when the market is very volatile and gives better risk adjusted returns than Mean Variance Optimizer (MVO) baseline.**

*Keywords:* **Portfolio Optimization, Artificial Intelligence, Machine Learning, Reinforcement Learning, Risk Management, Sentiment Analysis, Investment Advisory.**

## I. INTRODUCTION

The optimization of portfolios continues to constitute one of the most crucial issues of modern finance. Indeed, the issue here is one of how best to allocate money among the different kinds of financial instruments to obtain maximum gains for the minimum risk. Although the goal itself seems simple enough, its practical application proves to be quite complicated. Ever since Harry Markowitz came up with the concept of mean-variance optimization back in 1952, mathematical modelling has played an important part in defining the investment strategies of both institutional investors and private individuals.[1] These methods of investment planning served as the basis of many institutions' activities and were widely used by people investing on their own initiative for decades now. Nevertheless, these models rely upon several simplified assumptions, such as normal distribution of returns, linear and constant relations between investments, and the inclusion of all information in price behavior. Financial markets show such qualities as fat tails, volatility clustering, regime changes, and nonlinearity, all of which classical models find difficult to incorporate. According to the Chartered Financial Analyst Institute (CFA) in 2024, current asset managers work with data exceeding 50 million pieces per day, including price and volume data, social media sentiment, news, macroeconomic indicators, and Environmenmental, Social and Governance (ESG) indicators.[2] The sheer high dimensional and heterogeneous nature of such data make conventional statistical techniques ineffective for timely investment decision making.

Artificial Intelligence (AI) and Machine Learning (ML) have emerged as powerful tools to address these challenges. Advanced models have been applied to overcome these challenges by modelling non-linearity in asset return processes, adapting dynamic changes in the financial environment, and extracting relevant information from unstructured data, which would continually update portfolio weightings according to the changing signals. Himanshu Chaudhry et al. in 2025 proposed a risk adjusted deep reinforcement learning approach that resulted in better annualized Sharpe ratio than benchmarked portfolios in multiple global equity markets.[3] Moreover, Ashrafzadeh et al. found that stock clustering combined with deep learning models enhanced return prediction performance.[4]

In this paper, we introduce FinanceWiz, a comprehensive AI-based platform for portfolio optimization platform that addresses these challenges. The key contributions of the FinanceWiz platform are:

- An integrated prediction-optimization system that incorporates return prediction and learning optimization.
- Allocation of portfolio using behaviour.
- Risk Assessment based on appetite.
- Sentiment analysis in real-time affecting return predictions and allocation decisions.
- AI coach providing contextual, non-prescriptive financial guidance.

## II. PROBLEM DEFINITION

The nature of classical approaches to portfolio optimization such as Mean-Variance Optimization (MVO) and the Capital Asset Pricing Model (CAPM), is backward-looking because they are completely based on historical data in estimating future performance. During times of relative stability, this approach can approximate reality. However, during transitions in the market, this approach falls short. Classical approaches, such as CAPM and MVO, make an assumption of linearity and stationarity of relationships between securities. The degree of correlation between two securities or a stock and a commodity is seen as a static number that can be predicted using historical data and then used with confidence in making future decisions. In times of market turmoil, securities that typically have no relationship end up being very much correlated in what is called correlation breakdown. Portfolio optimizations built with the classical MVO approach suffered huge drawdowns in the market crash in March 2020, caused by the COVID-19 pandemic that, in most cases, were higher than equally-weighted benchmark strategies. Another critical limitation of classical models lies in their inability to process and incorporate the vast and rapidly expanding volume of unstructured financial data. Modern financial markets generate information at an unprecedented scale, including news articles, earnings call transcripts, social media sentiment, and macroeconomic reports. Traditional mathematical frameworks cannot incorporate the huge volume of unstructured data existing in the financial market because of the inherent limitations of mathematics.

However, even the current portfolio optimization systems that make use of AI technology face many issues, which have been discussed in this paper. Firstly, there are very few studies in the area of AI that combine return prediction and portfolio optimization. Secondly, all the current systems treat all users the same way, irrespective of the fact that they vary widely in terms of their risk tolerance, investment period, age, financial position, experience level, and purpose of investment. Thirdly, all current systems operate on a black box basis without justifying the decision-making process. These limitations highlight the need for a more comprehensive and integrated approach to portfolio optimization, one that not only leverages the predictive power of AI but also ensures adaptability, personalization, and transparency. These issues are briefly explained in Table I below.

**TABLE I: COMPARATIVE ANALYSIS OF PORTFOLIO OPTIMIZATION APPROACHES**

| Parameter | Traditional Methods | Modern Methods |
|---|---|---|
| Data Sources | Historical Returns, statistics | High frequency data, text, sentiment, imagery |
| Assumptions | Linear, static correlations | Non-linear, adaptive, data-driven |
| Market Handling | Stable, gradual changes | Volatile, dynamic, fragmented |
| Optimization Process | Analytical solutions, closed-form | Learning algorithms, iterative search |
| Flexibility | Low (fixed rules) | High (robust to noise) |
| Performance in Volatility | Poor/average | Often superior |
| Transparency | High | Lower |

## III. LITERATURE REVIEW

### *A. Foundations of Portfolio Theory*

One of the earliest contributions to finance is Markowitz's work on Mean-Variance Optimization (MVO), which is credited as the birth of Modern Portfolio Theory (MPT). MPT proved that a diversified portfolio can achieve the lowest risk while maintaining maximum return by investing in assets that have a low correlation. MPT resulted in the creation of the "efficient Frontier" graph that illustrates the maximum return for the lowest risk associated with the investments. You can reduce your overall risk through investment diversification while still getting returns.[1][5]

Another significant contribution to portfolio theory is CAPM, which was developed in the 1960s by Sharpe, Lintner, and Mossin. This theory expands upon the Modern Portfolio Theory to focus on the relationship between the systematic risk of the asset and its expected return using the beta coefficient (β)[5][6].

The CAPM equation is:

$$E(R_i) = R_f + \beta_i[\, E(R_m) - R_f \,]$$

Where:

- $E(R_i)$ = expected return of asset i
- $R_f$ = risk-free rate
- $\beta_i$ = beta of the asset (its sensitivity to the market portfolio)
- $E(R_m)$ = expected return of the market portfolio
- $E(R_m) - R_f =$ market risk premium

As per CAPM theory:

- A stock with $\beta = 1.0$ is expected to move in line with the market.
- If $\beta > 1$, the asset is more volatile than the market.
- If $\beta < 1$, it is less volatile.
- CAPM suggests that higher beta (higher market risk) should be rewarded with higher expected returns.

Beta coefficient (β) is derived based on historical performance, but it cannot be used to predict future risk. Other aspects influencing the return are overlooked, like the size effect, value effect, and sentiment effect.

### *B. AI and Machine Learning in Portfolio Management*

Sutiene et al. (2024) present a thorough literature review on the impact of AI in portfolio management at planning, implementation, and feedback levels. The study finds that there are three major applications where AI provides considerable benefits, namely, (1) high-dimensional forecasting with dimensionality reduction techniques such as Principal Component Analysis (PCA) and Partial Least Square (PLS), (2) time-series forecasting with deep learning models such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), and (3) correlation and network analysis through clustering techniques and graph-based methodologies [2].

Choudhary et al. (2025) developed a risk-adjusted deep reinforcement learning (RA-DRL) framework that exploits the capabilities of three different reward functions, namely, log returns, differential Sharpe ratio, and maximum drawdown. Based on empirical findings in four world stock markets, i.e., Sensex, Dow, Taiwan Stock Exchange (TWSE), and Integrated Budget and Expenditure System (IBEX), the study proves that the suggested RA-DRL framework yields cumulative returns of 124.83% in the Sensex index compared to base Deep Reinforcement Learning (DRL) agents and conventional benchmarks [3]. Jiang et al. (2024) also proved that deep reinforcement learning frameworks yield better performances than classical portfolio allocation frameworks under various market circumstances and especially in volatile conditions.[7]

### *C. Sentiment Analysis for Portfolio Optimization*

According to Muthivhi and van Zyl (2022), it was discovered that when sentiment analysis is incorporated into asset price prediction for portfolio investments, it results in improved portfolio performance. The revenue of the proposed sentiment portfolios was significantly higher than without the use of sentiment analysis. In addition, they observed that the proposed portfolios were also more volatile, which made them less appealing to risk-averse investors.[8] In the next advancement on this subject, Mantshimuli (2025) came up with the idea of a sentiment-aware portfolio optimization model that incorporates Conditional Value-at-Risk (CvaR) diversification, Deep Deterministic Policy Gradient (DDPG) and Financial Bidirectional Encoder Representations from Transformers (FinBERT) sentiment analysis. This model performed exceptionally well, with an annual return of 34.57% and Sharpe Ratio of 1.05, significantly exceeding other benchmark models, such as Most Diversified Portfolio and Minimum CVaR portfolios.[9] Nyakurukwa and Seetharam (2025) suggested a portfolio optimization method based on online sentiment connectedness. They created portfolios by taking into consideration the daily news sentiments as well as the social media sentiments of firms using minimum directional connectedness approach based on Time Varying Parameter Vector Auto-regression (TVP-VAR), where their results showed that minimum sentiment connectedness portfolio outperforms other traditional approaches in the case of cumulative returns especially the negative news sentiment portfolios [10].

## IV. RESEARCH GAPS

The following limitations have been observed in existing models and provide motivation for this research work. Many current research works consider return prediction, optimization, or sentiments individually, not combining all of these aspects together. Integration of various AI technologies within existing systems is not complete. While Choudhary et al. [3] employ reinforcement learning technology for stock portfolio optimization, there is no consideration of the aspect of sentiments. While Muthivhi and van Zyl [8] show how sentiment is used for prediction, they apply the classical mean-variance optimization method, and not adaptive techniques.

None of the existing systems has the functionality to offer user guidance on a personalised basis or provide explainable AI Coaching. As highlighted by Cho et al. [3], large language models are capable of altering risk profiles based on different personas; however, none of the existing solutions uses this approach for actionable portfolio guidance with explainable reasoning. The lack of an AI Coach explaining the rationale behind allocation makes it difficult to instil trust and confidence among users. The risk assessment is either missing or static and does not consider personal factors like age, income and experience. Existing AI Portfolio solutions consider only equities and ignore diversification through assets like commodities, forex, and cryptocurrencies.

**TABLE II: COMPARATIVE ANALYSIS OF LITERATURE REVIEW**

| **Ref.** | **Method** | **Core Technique** | **Key Limitations** |
|---|---|---|---|
| [1] | Risk-Adjusted DRL (RA-DRL) | PPO and CNN | No sentiment integration |
| [5] | Mean – Variance Optimization | Quadratic Programming | No market adaptation |
| [6] | Capital Asset Pricing Model (CAPM) | Linear Regression | Ignores non-market risks |
| [7] | Sentiment LSTM and MVO | LSTM and Mean-Variance | Increases volatility |
| [10] | Minimum Sentiment Connectedness | TVP-VAR and Sentiment Analysis | Limited asset classes |
| Proposed | FinanceWiz | LSTM, RL and LLM Coach | Higher inference time |

These limitations provide the impetus for the development of FinanceWiz, which would tackle all the above-stated issues using a combination of LSTM prediction, Optimization, Sentiment Analysis, and LLM-based AI Coach.

## V. METHODOLOGY

The overall structure of the system comprises six modules that work together to make the solution possible: (1) data acquisition and processing, (2) risk profiling of users, (3) sentiment analysis, (4) return prediction based on LSTM networks, (5) optimization of portfolios via reinforcement learning algorithms, as well as classic mean-variance optimization technique, and (6) AI coach based on LLM technologies. x.

### *A. System Architecture Overview*

The FinanceWiz platform operate on a modular, four-level architecture.

Layer 1: Data Layer – This level will be responsible for the gathering and storage of market, news, and user data.

Layer 2: Prediction Layer - Here, there are two models – LSTM-based Return Predictor and Sentiment Analysis model.

Layer 3: Optimization Layer – Classical Mean Variance Optimizer and Reinforcement Learning Agent will be stored here.

Layer 4: Advisory Layer - The LLM-based AI Coach will operate on this level.

### *B. Data Collection and Preprocessing*

Data is sourced from the following sources:

- Market Data: This comprises daily stock prices for equities, commodities, forex, and cryptocurrencies extracted via Yahoo Finance API. .
- Alternative Data: Sentiment analysis data comprises news headlines and articles from financial news sites. In addition, the system incorporates sentiment analysis data from social media platforms such as Twitter and Reddit.
- User Data: This comprises financial objectives and risks preferences obtained through a questionnaire.
- Data Preprocessing: The unprocessed data is processed using feature engineering to extract meaningful attributes. Feature engineering involves transforming raw price data into technical indicators, which include the Relative Strength Index (RSI), Exponential Moving Average (EMA) and Moving Average Convergence Divergence (MACD).

Relative Strength Index (RSI) is a momentum oscillator used to measure the velocity and degree of change in price changes on a scale of 0 to 100. An RSI reading of more than 70 usually indicates that the stock is overbought, and hence, it is due for a price correction whereas a reading less than 30 implies that the stock might be oversold, and thus it will eventually move upwards. Exponential Moving Average (EMA) is one form of moving average that gives a higher priority to the latest data compared to other data. Moving Average Convergence Divergence (MACD) is a trend-following momentum indicator that shows the relationship between two moving averages of an asset's price.

### *C. User Risk Assessment and Profiling*

The algorithm assesses the user's attributes from five different angles to define their adequate risk profile: age, yearly income, experience level (novice, intermediate, professional), purpose of investments (capital protection, earnings, growth, aggressive growth), and individual attitude towards risks (conservative, balanced, aggressive). On the basis of the analysis performed, the user is categorized into one of the three groups.

### *D. Sentiment Analysis*

The latest news headlines, along with related articles concerning each asset are retrieved from news feeds. The process of pre-processing involves splitting the article into sentences using the sentence tokenizer and identifying entities based on case-insensitive regular expressions that detect companies' names and ticker symbols. Sentiment analysis is performed by means of feeding the texts into the transformer library and receiving a sentiment value $\sigma \in [-1,1]$ for each text component, where $-1$ means strongly negative, 0 corresponds to neutral sentiment, and +1 stands for strongly positive sentiment. The model's output includes three probabilities: $P_{pos}$, $P_{neu}$, and $P_{neg}$, with the final score calculated as $\sigma = P_{pos} - P_{neg}$.[9]

News headlines together with metadata such as the source, timestamp, and related symbols are compiled in the news layer.

### *E. Return Prediction using LSTM Networks*

Long Short-Term Memory (LSTM) networks are employed for predicting asset returns due to their ability to capture long-term dependencies in sequential financial data.

The LSTM cell at time $t$ is defined by the following equations:

$$
\begin{aligned}
f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\
i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\
\tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\
C_t &= f_t * C_{t-1} + i_t * \tilde{C}_t \\
o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\
h_t &= o_t * \tanh(C_t)
\end{aligned}
$$

where $f_t$, $i_t$, and $o_t$ represent the forget, input, and output gates, respectively, $C_t$ is the cell state, and $h_t$ is the hidden state. The input vector $x_t$ includes historical prices, technical indicators (RSI, EMA, MACD), and sentiment scores $\bar{\sigma}_{i,t}$ .[11] The model is trained on 80% of historical data (training set), validated on 10%, and tested on the remaining 10%.

### *F. AI Coach for Personalised Guidance*

The AI Coach's response is based on user trading patterns, asset allocation, risk appetite, and financial goals. The Personalised AI Coach module uses a large language model to give users non-prescriptive financial advice. When a user makes a request, the system retrieves necessary information from three sources: (1) the user's existing investment portfolio and risk tolerance from the database, (2) the latest market data and sentiment analysis of investments in the portfolio, and (3) financial guidelines defined in advance.[12] These inputs are then combined into a structured prompt, which is passed to the LLM for inference. The model processes this multi-source contextual information to generate coherent, relevant, and user-specific outputs. Importantly, the AI Coach is designed to provide non-prescriptive guidance, meaning that it does not issue direct buy or sell commands. Instead, it explains potential strategies, highlights risks and trade-offs, and provides reasoning behind different allocation choices. This approach encourages informed decision-making while maintaining user autonomy. The communication style of the AI Coach can also be adapted based on the user's level of financial literacy, offering simplified explanations for novice investors and more technical insights for experienced users. The AI Coach addresses one of the key limitations of existing portfolio optimization systems—the lack of user-centric and interpretable decision support.

## VI. RESULT AND DISCUSSION

Results from the FinanceWiz experiment will be discussed here. These include comparing the performance to the benchmarks, understanding the individual contributions via ablation studies and significance tests. .

### *A. Experimental Setup*

Experiments were run on the test data set of the last ten years. This time horizon covers various economic cycles including the stable phase after 2016, the market crash caused by COVID-19 and the subsequent recovery between 2020 and 2021, the inflationary upsurge and bear market in 2022 and finally the recovery in 2023 and onwards. Four benchmark models were designed for comparison and include (1) the Equal Weighted Portfolio (EWP) model, where equal allocation was carried out across all securities every month; (2) the MVO strategy that employs classical Markowitz portfolio optimization; (3) LSTM-only portfolio, which utilizes the predicted returns via LSTM with MVO weights but no sentiment or reinforcement learning component. Finally, (4) the Buy and Hold (B&H) strategy where an initially balanced EWP is not rebalanced further.

The system utilises three categories of data, summarised in Table III.

**Table III: Dataset Summary**

| Category | Size | Purpose |
|---|---|---|
| Market Data | 15 Asset Classes | Price & Volume history |
| News Data | 10,000 Articles | Sentiment Analysis |
| User Profiles | 1000 Profiles | Risk Assessment |

### *B. Comparative Performance Results*

Table IV presents the comparative performance of FinanceWiz against the baseline strategies.

**TABLE IV: COMPARATIVE PORTFOLIO PERFORMANCE**

| Strategy | Annual Return (%) | Sharpe Ratio | Max Drawdown (%) |
|---|---|---|---|
| Buy and Hold (B&H) | 10.4 | 0.68 | -22.3 |
| Equal Weight Portfolio(EWP) | 11.8 | 0.82 | -19.4 |
| Mean-Variance (MVO) | 13.5 | 0.95 | -16.2 |
| LSTM-only (No Sentiment analysis) | 16.2 | 1.12 | -11.5 |
| FinanceWiz (Ours) | 19.2 | 1.34 | -8.7 |

FinanceWiz had the best annualized return of 19.2%, bettered the Sharpe ratio of 1.34 and minimized maximum drawdown to 8.7%.

### *C. AI Coach and User Engagement*

User study was conducted by interacting with AI Coach among 30 users (students and faculty) for two weeks. Users subsequently provided feedback via a questionnaire. Major Findings:

**TABLE V: SURVEY FINDINGS**

| Metric | Result |
|---|---|
| Average session duration | 10 minutes |
| User satisfaction (1-5 scale) | 4.3/5 |
| Perceived usefulness (1-5 scale) | 4.2/5 |
| Reduction in follow-up queries | 65% (compared to no-AI) |

The AI coach managed to clarify and explain reasons for portfolio allocation thereby solving the "black-box" issue often attributed to Artificial Intelligence used in finance.

### *D. Limitations*

There are certain limitations even in light of good results of our work. Though 10 years of data comprises many market cycles, the data doesn't include any extreme events like the global financial crisis or the dot-com bubble and hence the results might not be extrapolated in case of such events. Sentiment analysis has been done by analyzing only a few key sources that could have been helpful in giving extra inputs. AI Coach utilizes API which makes it dependent on external third-party services with associated costs. Finally, the application has been tested for Indian and USA stock markets, applicability to other emerging markets with different microstructures requires validation.

The results proved significant. By incorporating the analysis of sentiment, prediction of returns, and reinforcement learning optimization, we obtain better risk-adjusted returns than using conventional techniques and individual AI systems. The results of our ablation experiment confirm the importance of each element. By incorporating sentiment analysis, it proves the behavioural finance approach that the sentiments of investors extracted from news articles can be used for predictions beyond prices.

## VII. CONCLUSION

This paper introduces the concept of FinanceWiz, an AI-powered portfolio optimization system that aims to address the shortcomings of traditional models like MVO and CAPM. The system is built on five core components, namely, risk assessment by users, sentiment analysis, LSTM-based prediction of returns, hybrid portfolio optimization, and the use of LLM-based AI coaches.FinanceWiz has been trained for 10 years of data history and has resulted in an annualized return rate of 19.2% and a Sharpe ratio of 1.34, which outperforms both the equal weight baseline and the MVO model. The maximum drawdown has been lowered to 8.7%, marking a 55% improvement from the equal-weight strategy.

Building on the current work, several directions for future research and development are identified. Firstly, integrating the system with various trading platforms for automated trading. Integration with brokerage websites like Zerodha, Upstox, Angel One (for India) or Alpaca, Interactive Brokers (for international securities) will allow automated trade execution by the system which will make it easier for us to get investment data as well as link the current portfolio with our website for better understanding of the current Portfolio. In this case we will need additional modules for portfolio monitoring, placing orders, positions and risk management while executing trades. Secondly, integration with banks for deposits and withdrawals. This involves integrating payment gateways and other goal-based saving features. Thirdly, integration of the developed AI models with financial strategies. Although our current framework uses DRL for optimization, the integration of proven financial strategies may improve results and interpretation. Fourthly, allocating assets and diversifying across various assets properly with the right asset names and symbols. This combination of aspects puts our FinanceWiz website on a realistic path towards a fully automated intelligent wealth management system.

## REFERENCES

[1] H. Markowitz, "Portfolio selection," The Journal of Finance, vol. 7, no. 1, pp. 77-91, Mar. 1952.

[2] K. Sutiene et al., "Enhancing portfolio management using artificial intelligence: Literature review," Frontiers in Artificial Intelligence, vol. 7, p. 1371502, Apr. 2024.

[3] H. Choudhary, A. Orra, K. Sahoo, and M. Thakur, "Risk-adjusted deep reinforcement learning for portfolio optimization: A multi-reward approach," International Journal of Computational Intelligence Systems, vol. 18, no. 1, p. 126, 2025.

[4] M. Ashrafzadeh, M. Sadrani, and S. H. Zolfani, "Deep learning and machine learning models for portfolio optimization: Enhancing return prediction with stock clustering," Results in Engineering, p. 106263, 2025.

[5] J. E. Beasley, "Portfolio optimization: Models and solution approaches," in INFORMS TutORials in Operations Research, 2013, pp. 201-221.

[6] W. F. Sharpe, "Capital asset prices: A theory of market equilibrium under conditions of risk," The Journal of Finance, vol. 19, no. 3, pp. 425-442, Sep. 1964.

[7] Y. Jiang, J. Olmo, and M. Atwi, "Deep reinforcement learning for portfolio selection," Global Finance Journal, vol. 62, p. 101016, 2024.

[8] M. Muthivhi and T. L. van Zyl, "Fusion of sentiment and asset price predictions for portfolio optimization," in Proc. 25th Int. Conf. Information Fusion (FUSION), IEEE, 2022, pp. 1-8.

[9] L. A. Mantshimuli, "Sentiment-aware portfolio optimization: CVaR-based diversification with deep reinforcement learning," IEEE Access, vol. 13, pp. 182704-182718, 2025.

[10] K. Nyakurukwa and Y. Seetharam, "Leveraging sentiment connectedness for portfolio optimization: A minimum connectedness approach," Review of Behavioral Finance, vol. 17, no. 5, pp. 908-917, 2025.

[11] A. G. Olanrewaju, "Artificial intelligence in financial markets: Optimizing risk management, portfolio allocation, and algorithmic trading," International Journal of Research Publication and Reviews, vol. 6, no. 3, pp. 8855-8870, Mar. 2025.

[12] A. Voronina, O. Romanko, R. Cao, R. H. Kwon, and R. Mendoza-Arriaga, "Generative AI-enhanced sector-based investment portfolio construction," arXiv preprint arXiv:2512.24526, 2025.