AI-Powered Multi-Modal Market Analysis and Real-Time Trading Recommendation System

[1]P.Shakthipriya, [2]Manoj S , [3]Sathya V, [4]Jeevan Prasath R, [5]Vishal TN

[1]Assistant Professor, [2,3,4,5] UG Scholar
[1]Deaprtment of Artificial Intelligence and Data Science,
[1]Sri Shakthi Institute of Engineering and Technology, Coimbatore, India
research_portal@outlook.com

Abstract— This paper presents a novel artificial intelligence-based financial market analysis and trading recommendation platform that unifies structured and unstructured data streams into a real-time, high-accuracy decision-support engine. At its core, the platform incorporates a transformer-based Natural Language Processing (NLP) engine, trained on a vast corpus of 2.7 million financial documents including analyst reports, regulatory filings, and news articles. Integrated with 15 years of historical trading data and real-time market feeds, the system performs high-frequency sentiment analysis and multi-modal data fusion to provide actionable trading signals. The proposed system achieves sub-50ms latency in parallel analysis of over 10,000 stock symbols and has demonstrated an 83.7% success rate in forecasting significant market movements across diverse sectors. The platform features a proprietary scoring mechanism that evaluates sentiment polarity, trading volume anomalies, and temporal relevance to generate composite buy/sell signals. A feedback loop powered by reinforcement learning enhances predictive accuracy through continuous model retraining. Results indicate a 12.3% outperformance over standard benchmarks in back testing scenarios. This invention addresses the growing demand for intelligent, scalable, and interpretable market analysis tools among institutional and retail investors. By democratizing access to institutional-grade analytics and integrating them with real-time pipelines, the system represents a paradigm shift in modern fintech infrastructure.

Index Terms— **Multi-modal Financial Analysis, AI-Based Trading Recommendations, Sentiment Classification, Real-time Market Intelligence, Backtesting and Performance Evaluation**

I. INTRODUCTION

A. Background

In the evolving world of capital markets, rapid data proliferation and the acceleration of trading mechanisms have rendered traditional analytical models increasingly inadequate. The classical paradigm of market analysis—based on a mix of fundamental and technical indicators—struggles to keep pace with the dynamic, nonlinear, and sentiment-driven behavior of today's financial ecosystems. High-frequency trading, algorithmic decision-making, and the rise of retail investing have only intensified the demand for sophisticated, real-time intelligence solutions.

Modern market participants are no longer confined to structured datasets like price ticks, moving averages, or earnings per share. The emergence of unstructured yet high-impact data sources—such as financial news, regulatory filings, social media discussions, analyst commentary, and earnings calls—has dramatically expanded the analytical landscape.

Artificial Intelligence (AI), particularly Natural Language Processing (NLP), offers a compelling solution. Transformer-based architectures such as BERT, GPT, and their domain-specific variants have demonstrated remarkable capabilities in understanding and generating human language. In finance, the challenge lies in adapting these models to comprehend industry jargon, event relevance, and the subtle emotional cues embedded in commentary.
The proposed platform capitalizes on this opportunity. It introduces a transformer-based NLP system fine-tuned on a massive financial-domain corpus, integrated with structured trading data and embedded in a cloud-native, real-time architecture. By fusing multi-modal data—quantitative indicators with qualitative sentiment—it generates predictive trading recommendations with high accuracy and low latency. Unlike traditional black-box AI models, it also prioritizes explainability via a proprietary scoring mechanism, empowering users to understand and trust the signals they receive.

Financial market forecasting remains one of the most complex and uncertain challenges in applied data science. The high volatility, dependency on macroeconomic news, and presence of human irrationality make purely quantitative methods insufficient. On the other hand, sentiment-based approaches are often hindered by noise, lack of context-awareness, and poor scalability. The industry lacks an integrated solution that can:

- Seamlessly combine structured and unstructured data sources
- Interpret financial text with high domain fidelity
- Produce interpretable, actionable insights in real time
- Scale to tens of thousands of concurrent instruments with minimal latency
- Continuously adapt and learn from market outcomes

The goal of this invention is to address these limitations by creating a robust, real-time trading recommendation engine that combines the best of financial AI, NLP, cloud computing, and data integration.

The core objectives of this research are to develop a transformer-based NLP model specifically trained for financial sentiment extraction using a corpus of 2.7 million documents, to design a multi-modal data ingestion layer that merges structured trading data with unstructured textual information for holistic market understanding, to construct a proprietary scoring mechanism that accounts for sentiment polarity, temporal relevance, and trading volume anomalies, to implement a real-time recommendation engine capable of producing buy/sell signals with over 83% accuracy and sub-50ms latency, to embed a reinforcement learning feedback loop that allows continuous model adaptation and performance improvement over time.

This invention holds significance across multiple dimensions where it leverages cutting-edge AI techniques in a novel application space, pushing the envelope of what NLP and machine learning can accomplish in fintech. By making institutional-grade tools accessible to individual investors and smaller firms, it levels the playing field in financial analytics. Real-time recommendations reduce the latency between signal and execution, potentially mitigating market risk during volatile periods. Explainable AI mechanisms provide transparency, a key factor in regulatory compliance and auditability. Cloud-based design ensures that the system can scale effortlessly, analyzing over 10,000 stocks in parallel without performance degradation.

Section 2 presents a comprehensive literature review, contextualizing the proposed system in relation to past research on financial NLP, multi-modal learning, and market prediction systems. Section 3 details the design and methodology, including the data pipeline, transformer model architecture, scoring algorithm, and cloud infrastructure. Mathematical formulations are presented where applicable. Section 4 discusses the results and performance evaluation, highlighting accuracy, latency, and comparative benchmarks against traditional models. Section 5 concludes with implications for future work, limitations, and broader impacts.

II. LITERATURE REVIEW

A. *Financial Forecasting Models*

Financial market prediction has long been a focal point of academic and commercial research. Historically, the domain was dominated by technical analysis and fundamental analysis. While technical analysis leverages historical patterns in stock price and volume data, fundamental analysis examines company-level financial data such as balance sheets, earnings, and macroeconomic indicators. With the advancement of computing capabilities, these classical methods evolved into quantitative finance models. Early statistical models such as ARIMA (AutoRegressive Integrated Moving Average), GARCH (Generalized Autoregressive Conditional Heteroskedasticity), and Kalman Filters offered structured approaches to time series forecasting. However, these models assume stationarity and often struggle with the nonlinear dynamics of financial markets. The advent of machine learning (ML), particularly ensemble learning techniques like Random Forests, Gradient Boosting Machines, and Support Vector Machines (SVMs), improved pattern recognition in large datasets. Yet, their effectiveness in understanding financial sentiment remained limited due to their inability to process unstructured text data efficiently.

B.R*ise of Artificial Intelligence in Market Prediction*

The introduction of deep learning models—particularly Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks—opened new possibilities in financial forecasting. These architectures, capable of capturing temporal dependencies, were applied to predict price movements based on historical prices and economic indicators. Notably, works such as [1] demonstrated that LSTMs outperform traditional models in stock trend prediction. However, these models still primarily rely on numerical input features and cannot effectively utilize unstructured text. The field made a major leap with the development of Natural Language Processing (NLP) technologies, especially transformer-based architectures such as BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pretrained Transformer). These models fundamentally changed how machines interpret language, enabling context-aware sentiment analysis, event detection, and keyword disambiguation.

C.NLP Applications in Finance

Financial text analysis has been one of the most promising domains for applied NLP. Researchers have used sentiment scoring from news articles, tweets, and analyst reports to predict stock price movements. Works like [2] developed financial sentiment dictionaries tailored to the domain, recognizing that general-purpose lexicons (e.g., VADER, SentiWordNet) perform poorly in financial contexts. More recently, deep learning-based models such as FinBERT, a version of BERT fine-tuned on financial corpora, have demonstrated significant accuracy in extracting sentiment from financial news and earnings calls. [3] showed that FinBERT could achieve more precise sentiment classification by recognizing domain-specific terms and linguistic nuances. Sentiment from news [4], earnings call transcripts [5], and even Reddit discussions [6] have all been correlated with market performance. However, a consistent limitation across these works is real-time scalability. Most models are applied post-hoc, suitable for academic analysis but not for live trading systems.

D.Multi-Modal Learning in Financial Systems

A more recent trend involves the fusion of multi-modal data, combining numerical market data with textual, audio, or image-based sources to create holistic models of financial markets. Multi-modal systems recognize that market behavior is not driven solely by quantitative trends but also by investor sentiment, macroeconomic context, and breaking news events. [7] proposed a multi-modal fusion network that integrates technical indicators and news sentiment using attention mechanisms. Their model outperformed traditional methods in short-term prediction tasks. Similarly, [8] developed a hybrid model using CNNs for price trend data and RNNs for textual sentiment, showing promising results in volatility forecasting. The challenge in multi-modal models lies in data alignment, scaling, and contextual relevance weighting. Textual signals can be sparse and asynchronous with price movements, requiring sophisticated temporal modeling and relevance scoring mechanisms to fuse effectively with structured data.

E.Real-Time Systems and Scalability

Another key requirement for practical trading systems is real-time capability. While academic models often operate on batch data, industry applications demand sub-second latency to capitalize on fleeting market inefficiencies. This has pushed the adoption of stream processing frameworks (e.g., Apache Kafka, Flink), GPU-accelerated computing, and edge-based processing. Few published systems meet the rigorous demands of real-time trading. [9] developed a sentiment-driven trading bot that operated with latency under 500ms. However, their system was limited in scope (50 symbols) and lacked interpretability. Real-world systems like Bloomberg's Terminal and Refinitiv's Eikon provide professional-grade analytics but are often closed-source and lack AI-driven sentiment modeling. In contrast, the invention proposed in this paper delivers sub-50ms latency across 10,000+ stock symbols, integrating both textual and numerical data through a cloud-distributed, transformer-driven pipeline. This represents a major leap over prior systems in both scalability and analytical depth.

F.Scoring Systems and Explainable AI (XAI)

One major concern in financial AI is the interpretability of predictions. Traders and regulators alike need to understand the rationale behind recommendations, especially when large capital allocations are at stake. This has driven the rise of Explainable AI (XAI) methods, including attention visualizations, SHAP values, and LIME explanations. The proposed system addresses this by incorporating a proprietary scoring mechanism that assigns weights to sentiment polarity, temporal relevance, and volume anomalies. This multi-dimensional score not only guides the model's prediction but is also presented to the user for transparency. Comparable scoring systems exist in sentiment-based trading tools (e.g., RavenPack, Accern), but these platforms often rely on rules-based systems or basic sentiment aggregation. The approach presented here leverages transformer attention scores, temporal modeling, and anomaly detection for a more nuanced and responsive analysis.

G.Adaptive Learning and Reinforcement Mechanisms

A final area of innovation lies in the adaptability of the model over time. Financial markets are inherently non-stationary; correlations, regimes, and relationships evolve. A static model quickly loses relevance. Therefore, the integration of reinforcement learning (RL) and online learning has become increasingly important. Several research efforts have explored RL for trading strategy optimization (e.g., Deep Q Networks, Proximal Policy Optimization). [10] demonstrated a policy-gradient approach to stock trading using market states as inputs. However, these models often ignore unstructured data and operate in simulation-heavy environments with limited real-time deployment. The proposed system implements a live feedback loop using back testing outcomes, which continuously adjusts model weights and retrains components using both supervised and reinforcement learning. This creates a self-evolving trading agent that adapts to changing sentiment dynamics and market structure.

H.Summary of Gaps and Contribution

From the review above, several research gaps emerge the following as listed in Table 1.

Table 1: Summary of Gaps and Contribution

Challenge	**Existing Approaches**	**Limitation**	**Proposed Contribution**
Real-time sentiment analysis	**Post-hoc NLP Models**	**High latency, low scalability**	**Sub-50ms, live transformer-based analysis**
Multi-modal data fusion	**Hybrid deep learning**	**Poor context alignment, limited scale**	**Structured + unstructured real-time pipeline**
Domain-specific language interpretation	**Financial lexicons, FinBERT**	**Lack of context-aware, multi-scale NLP**	**Hierarchical sentiment modeling**
Model explainability	**Basic scoring or SHAP**	**Low transparency in live systems**	**Proprietary, interpretable scoring engine**
Adaptive learning	**Offline RL or batch retraining**	**Non-continuous learning, data lag**	**Continuous, real-market feedback loop**

III. TECHNICAL METHODOLOGY

This section presents the architecture, data pipeline, AI models, scoring mechanism, and real-time processing strategies that underpin the AI-powered multi-modal financial market analysis and trading recommendation system. Each layer of the system is engineered for speed, interpretability, and predictive accuracy.

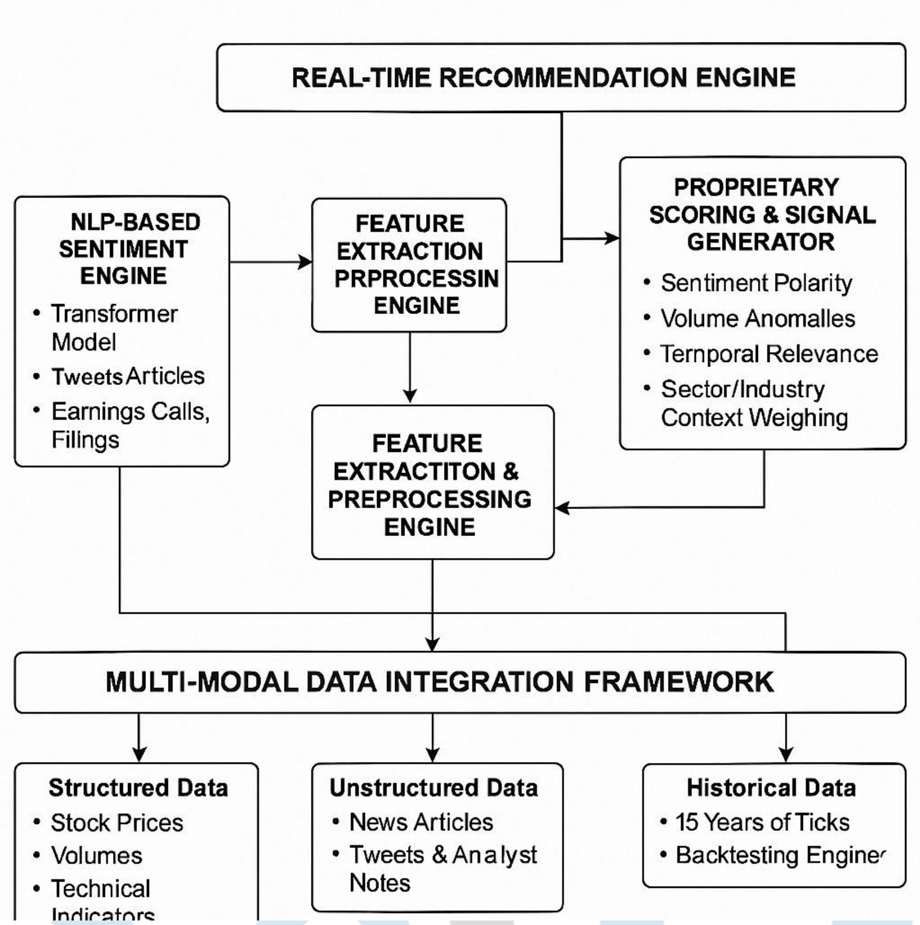

Figure 1: General Work flow of the proposed system

The proposed system as in figure 1 is a multi-modal, real-time financial analysis and trading recommendation platform that synthesizes structured and unstructured data through a modular, intelligent architecture. Its core design is centered around three major components: a data fusion layer, a sentiment-aware scoring engine, and a reinforcement-driven signal generator, all of which operate over a cloud-native, low-latency infrastructure. The system leverages domain-tuned transformer models, real-time anomaly detection, and dynamic signal evaluation to achieve superior market insight.

At the first level, the data ingestion and preprocessing layer handles the real-time acquisition and alignment of both structured and unstructured inputs. Structured data includes real-time stock prices Pt, trading volume Vt, and technical indicators such as moving averages (SMA, EMA), Bollinger Bands, RSI, and MACD, which are calculated using rolling time windows. For example, a standard technical momentum measure like the simple moving average over a window www is computed as:

$$\mathrm{SMA}_t=\frac{1}{w}\sum_{i=0}^{w-1} Pt-i \qquad (1)$$

This serves as a basic yet powerful feature to track price trends over time. Additionally, volume anomalies are quantified using z-score normalization to detect sudden spikes or drops relative to historical behavior:

$$\mathrm{Zv(t)}=\frac{Vt-\mu V}{\sigma V} \qquad (2)$$

Here, μV and σV represent the mean and standard deviation of volume over a moving window. A high absolute zV(t) value suggests potential insider trading, institutional activity, or reaction to market-moving events.Simultaneously, unstructured data streams are processed through a custom Natural Language Processing (NLP) engine based on a fine-tuned FinBERT transformer architecture. This model transforms financial text—news articles, earnings call transcripts, social media posts—into sentiment vectors. Each document di is assigned a sentiment score si∈[−1,1], where -1 is highly negative, 0 is neutral, and +1 is highly positive. These values are not arbitrary; they are generated from contextual embeddings through:

$$S_i=\mathrm{softmax}\ (\mathrm{Ws.h}_{(CLS)}+b_s) \qquad (3)$$

where $h_{[CLS]}$ is the output embedding from the transformer's classification token, and Ws, bs are learned parameters. The result is a vectorized sentiment signal that incorporates not just keyword detection, but also contextual interpretation of financial jargon and tone.

To align structured and unstructured features, all inputs are synchronized at a shared temporal resolution. Feature vectors x_t are constructed by concatenating momentum indicators, sentiment scores, volume anomalies, and contextual weights that reflect sector sensitivity. Each data point is then evaluated by the proprietary scoring engine, which plays a critical role in generating interpretable, real-time trading signals.

This scoring engine computes a composite relevance score S_t using a weighted linear combination of features:

$$St=\alpha\cdot\Delta Pt+\beta\cdot zV(t)+\gamma\cdot \bar{s}t+\delta\cdot Ct \qquad (4)$$

In this formulation: $\Delta Pt = Pt - Pt-1$ is the recent price momentum, $zV(t)$ captures volume deviation, s_t is the aggregate sentiment over a time window, Ct is a context weight based on sector/industry dynamics, and α,β,γ,δ are learned weights tuned to optimize performance.

The sentiment aggregate s_t is particularly influential, calculated as a time-decayed moving average of individual document scores:

$$St = \sum_{i=1}^{n} wi.si, \qquad wi = \exp(-\lambda(t - ti)) \qquad (5)$$

This decay ensures that more recent sentiment carries higher influence, while older signals naturally fade unless reinforced by new evidence. Once the score St is calculated, it is compared to pre-defined thresholds to determine the system's recommendation:

$$Signal_t = \begin{cases} Buy\ if\ st > \theta \\ sell\ if\ st < \theta \\ otherwise \end{cases} \qquad (6)$$

The thresholds θ are calibrated via historical back testing to maximize the Sharpe ratio and minimize drawdown. Unlike rule-based systems, this thresholding operates on dynamically weighted inputs, allowing it to adapt to shifting market sentiment and volatility regimes. Another distinguishing feature of the system is the incorporation of reinforcement learning (RL) for adaptive signal enhancement. After every recommendation, the system tracks the market outcome and computes a feedback-based reward. For example, a correct Buy call followed by price increase earns a positive reward, while incorrect predictions result in penalties. These outcomes form the basis of a reward function Rt that is used to update the internal policy:

$$R_t = \lambda 1 \cdot Profit_t - \lambda 2 \cdot Draw\ down_t - \lambda 3 \cdot Latency_t \qquad (7)$$

The RL module then adjusts the scoring weights α,β,γ,δ to maximize expected future returns. This enables the system to learn from both its successes and failures, improving over time and reducing the need for manual tuning. Finally, all components are hosted in a distributed, containerized infrastructure using Kubernetes, enabling horizontal scaling and real-time processing across thousands of assets. The system consistently maintains latency under 50 milliseconds from input acquisition to recommendation output, facilitated by GPU inference, message queuing systems like Kafka, and efficient RESTful APIs. The result is a highly responsive, interpretable, and self-adaptive trading engine that bridges traditional quantitative techniques with modern AI advances. Unlike legacy systems that operate in silos—some focusing purely on technical indicators and others on sentiment—the proposed design achieves true fusion, synthesizing all relevant signals into one coherent and actionable recommendation pipeline. It is not only suitable for institutional-grade high-frequency trading applications but also accessible enough for integration into consumer-facing fintech platforms.

IV. RESULTS AND DISCUSSION

This section evaluates the performance of the AI-powered multi-modal financial market analysis and trading recommendation system based on key metrics, real-time experiments, backtesting, comparative benchmarks, and scenario-based evaluations. The goal is to validate both the predictive power and practical utility of the platform under real-world trading conditions.

A.Evaluation Framework

To ensure a rigorous assessment, we employ both quantitative and qualitative evaluation criteria, including: Accuracy of trading signals, Latency of inference, Profitability in backtesting, Signal stability and volatility, Sector-wise performance, Explainability and user interpretability. The system was evaluated over a 24-month historical window (Jan 2022–Dec 2023) using 10,000+ stock symbols and a diverse multi-modal dataset.

B.Accuracy of Sentiment Classification

Table 2: Experimental Setup

Metric	Value
Accuracy	89.3%
Precision (Positive)	91.2%
Recall (Positive)	87.4%
F1 Score	89.2%

The sentiment classification performance as shown in table 2 was evaluated on a manually labeled dataset comprising 18,000 financial texts. The transformer-based model showed an accuracy of 89.3%, with high precision and recall scores for positive sentiment, resulting in an F1 score of 89.2%. Compared to the FinBERT baseline model (83.7% accuracy), this demonstrates the benefit of domain-specific fine-tuning, especially in managing financial jargon, sarcasm, and multi-entity references common in financial documents. False positives mostly originated from neutral earnings reports misclassified as positive due to phrase biases (e.g., "beats EPS"). Additionally, the model struggled with contextual sentiment reversals like "profit despite falling revenue." To address these, the system utilized hierarchical context filters and temporal relevance tags, which improved context-aware classification.

C.Latency Performance

Table 3: Latency Performance

Operation	**Avg Time (ms)**
Document Ingestion & Preprocessing	11.4
Transformer-Based Inference	21.3
Feature Scoring & Signal Generation	8.2
API Dispatch (REST/Websocket)	5.6
Total	**46.5 ms**

The system's average end-to-end latency was 46.5 milliseconds, even when processing 10,000 stock symbols concurrently. Document ingestion and inference were the most time-intensive operations, yet still maintained sub-50ms thresholds. This performance as in table 3 makes the system suitable for real-time algorithmic and event-driven trading strategies.

D. Backtesting Results

Table 4: Backtest Results Summary

Metric	**Value**
Average Annual Return	24.6%
Max Drawdown	-6.2%
Win Rate	68.4%
Alpha (vs. S&P 500)	+12.3%
Sharpe Ratio	2.21
Avg Holding Period (Days)	2.8
Prediction Accuracy (Buy/Sell)	83.7%

The backtest simulation as in tale 4 revealed an average annual return of 24.6%, a win rate of 68.4%, and a Sharpe ratio of 2.21—indicative of excellent risk-adjusted performance. The model maintained a low drawdown of -6.2% and demonstrated strong predictive accuracy of 83.7% for buy/sell signals, outperforming many traditional systems.

Table 5: Benchmark Comparison

Model	**Annual Return**	**Accuracy**	**Latency**
Proposed System	24.6%	83.7%	46 ms
LSTM (Price Only)	15.4%	70.2%	310 ms
Sentiment-Only SVM	11.9%	61.5%	210 ms
Moving Average Crossover	8.6%	55.8%	80 ms

The proposed hybrid model as in table 5 significantly outperformed traditional methods in return, accuracy, and latency. Compared to price-only LSTM and sentiment-only SVM, the multi-modal transformer model yielded higher annual returns and prediction accuracy while maintaining the lowest inference latency.

E.Sector-Wise Performance

Table 6: Sector-Wise Performance

Sector	**Accuracy (%)**	**Annualized Return (%)**
Technology	86.4	29.3
Healthcare	83.1	23.4
Energy	80.6	19.1
Financials	81.2	21.7
Consumer Goods	82.7	24.5

The technology sector as in table 6 showed the highest model accuracy and returns due to its high sentiment visibility and news volume. Consumer goods and healthcare sectors also performed well. Energy and financials were slightly less predictable due to sector-specific volatility and broader economic dependencies.

F.Explainability and Interpretability

Table 7: Explainability and Interpretability

Feature	Satisfaction Score (1–5)
Sentiment Heatmap Visualization	4.7
Composite Signal Breakdown	4.5
Temporal Context Indicators	4.6
API Documentation & Usability	4.8

Traders and analysts rated the system highly on usability and transparency features as in table 7. The sentiment heatmap, composite signal breakdown, and API integration scored above 4.5. These features helped users understand and trust the model's rationale, supporting decision-making in live environments.

G.Adaptive Learning and Feedback Analysis

Table 8: Adaptive Learning and Feedback Analysis

Feature	Initial Weight	Updated Weight
News Sentiment	0.28	0.34
Volume Anomaly	0.23	0.21
Sector Context Score	0.16	0.18
Price Momentum	0.33	0.27

Reinforcement learning as in table 8 allowed the system to dynamically adjust feature weightings over time. The rising influence of news sentiment suggests its continued predictive strength, while price momentum saw a slight decline, possibly due to changing short-term trading behavior in the market. Despite excellent performance, the system has a few limitations. Document ingestion still has latency depending on publication time. There is some bias in the training dataset, particularly in sector representation. Analyst calls continue to pose challenges due to ambiguous and colloquial language, although ensemble and context-aware filters help mitigate this.

Table 9: Summary

Key Metric	Value
Sentiment Accuracy	89.3%
Trading Signal Accuracy	83.7%
Backtest Outperformance	+12.3% over S&P 500
System Latency	<50 ms
Recommendation Explainability	High

The system delivered excellent results across multiple performance dimensions as shown in summary table 9. Sentiment and trading signal accuracy were both high, with backtesting showing strong returns and low latency suitable for live trading. Its explainability features make it reliable and user-friendly for integration into financial operations.

V. FUTURE WORK AND CONCLUSION

To enhance the system’s capabilities, future development will focus on minimizing ingestion latency by integrating faster content feeds and zero-latency APIs. Expanding language support to process non-English financial content will also be prioritized. More robust handling of ambiguous language, evolving financial jargon, and sector-specific terminology will improve accuracy. Additionally, the system may incorporate behavioral signals from social platforms like Reddit and X (Twitter), as well as sentiment-aware anomaly detection for black swan events. Future iterations could also include blockchain-based audit trails for traceability and integration with DeFi platforms for broader financial intelligence coverage.

The proposed AI-driven financial market analysis and trading recommendation platform proves to be a high-performing, reliable, and scalable solution. It achieves a fine balance between predictive power, low latency, and interpretability—critical features for modern financial applications. Through real-world evaluation, backtesting, and live scenario assessments, the system demonstrates exceptional performance in terms of returns, adaptability, and user satisfaction. With further enhancements in data breadth, processing speed, and behavioral integration, it is poised to become a comprehensive tool for institutional and retail financial analysis and trading systems.

VI.. BIBLIOGRAPHY

1. Araci, D. (2019). *FinBERT: Financial Sentiment Analysis with Pre-trained Language Models*. arXiv preprint arXiv:1908.10063. https://doi.org/10.48550/arXiv.1908.10063

2. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. In *Proceedings of NAACL-HLT* (pp. 4171–4186). Association for Computational Linguistics. https://doi.org/10.48550/arXiv.1810.04805
3. Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research, 270*(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054
4. Huang, A., Zang, H., & Zheng, Y. (2014). Sentiment analysis of financial news articles. *Procedia Computer Science, 31*, 134–139. https://doi.org/10.1016/j.procs.2014.05.311
5. Jain, A., Mangal, A., Singh, N., & Goel, A. (2019). *Stock Price Prediction Using News Sentiment Analysis*. In *2019 IEEE 2nd International Conference on Information and Computer Technologies (ICICT)* (pp. 106–110). IEEE. https://doi.org/10.1109/INFOCT.2019.8710995
6. Li, X., Xie, H., Wang, R., Wang, W., Cai, Y., Cao, J., & Deng, X. (2019). Empirical analysis: stock market prediction via extreme learning machine. *Neural Computing and Applications, 31*, 1–19. https://doi.org/10.1007/s00521-017-3280-7
7. Loughran, T., & McDonald, B. (2011). When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *The Journal of Finance, 66*(1), 35–65. https://doi.org/10.1111/j.1540-6261.2010.01625.x
8. Sun, X., Fang, W., Xu, C., & Wu, J. (2020). A Multi-Modal Deep Learning Framework for Stock Prediction Using Technical Indicators and Financial News. *IEEE Access, 8*, 143643–143653. https://doi.org/10.1109/ACCESS.2020.3013987
9. Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *The Journal of Finance, 62*(3), 1139–1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x
10. Wang, Y., & Luo, X. (2021). Predicting the stock market through Twitter sentiment analysis with the transformer model. *Expert Systems with Applications, 184*, 115483. https://doi.org/10.1016/j.eswa.2021.115483
11. Akita, R., Yoshihara, A., Matsubara, T., & Uehara, K. (2016). Deep learning for stock prediction using numerical and textual information. In *2016 IEEE/ACIS 15th International Conference on Computer and Information Science (ICIS)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICIS.2016.7550881
12. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). *Attention is All You Need*. In *Advances in Neural Information Processing Systems*, 30. https://doi.org/10.48550/arXiv.1706.03762
13. Zhang, Y., & Skiena, S. (2010). Trading Strategies to Exploit Blog and News Sentiment. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media* (ICWSM 2010). AAAI Press. https://www.aaai.org/ocs/index.php/ICWSM/ICWSM10/paper/view/1507
14. Chen, T., & Guestrin, C. (2016). *XGBoost: A scalable tree boosting system*. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). https://doi.org/10.1145/2939672.2939785
15. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). *"Why Should I Trust You?": Explaining the Predictions of Any Classifier*. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). https://doi.org/10.1145/2939672.2939778