iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XII **Month of publication:** December 2025

DOI: https://doi.org/10.22214/ijraset.2025.75826

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XII Dec 2025- Available at www.ijraset.com

A Hybrid AI-Driven Decision Support System for Retail Investors: A Comprehensive Survey of Technical, Fundamental, and Sentiment-Based Models

Komal Tambe, Swaroop Biradar, Nitesh Kumar Kushwaha
Pimpri Chinchwad University, Pune, India

Abstract: Retail investors frequently encounter significant challenges in interpreting diverse and complex market signals, which include dynamic technical trends, nuanced financial ratios, and volatile sentiment-driven information. The concurrent advancements in Artificial Intelligence (AI), sophisticated Ma- chine Learning (ML) techniques, and robust real-time financial APIs have paved the way for the development of highly automated Decision Support Systems (DSS). This comprehensive survey reviews key research published between 2024 and 2025 across three foundational domains of financial forecasting: (i) Fundamental-Analysis-based models, (ii) Technical-Indicator- driven ML frameworks, and (iii) Deep Learning-based Sentiment Forecasting Systems. A detailed comparative analysis is provided, highlighting the performance metrics, architectural nuances, and practical usability strengths and limitations of each approach. Finally, we propose a novel, integrated Hybrid DSS Architecture. This architecture is designed to fuse the predictive power of all three data streams, supporting high- fidelity real-time prediction, enhanced retail usability, and critical Explainable AI (XAI) capabilities to build trust and transparency.

Index Terms: Decision Support System, Machine Learning, Stock Market Forecasting, Sentiment Analysis, Technical Indicators, Financial Modeling, Explainable AI.

I. INTRODUCTION

Financial markets are governed by a continuous interplay of quantitative signals, macroeconomic policy shifts, and collective psychological dynamics, often referred to as market sentiment. Historically, professional investment firms have maintained a significant informational and analytical ad- vantage, leveraging multi-factor models and high-frequency trading platforms. In contrast, retail investors —individuals trading on their own behalf—lack equivalent sophisticated tools and frequently base critical investment decisions on fragmented data, general news, or unverified recommendations, leading to suboptimal outcomes and increased risk exposure [0].

The confluence of three modern technological pillars—the maturity of Deep Learning (DL) models, the ubiquity of high-frequency data from financial APIs (e.g., Alpha Vantage, Polygon), and the development of natural language process- ing (NLP) for unstructured data—has made it possible to level this analytical playing field. This reality has spurred intense research into creating automated, AI-driven Decision Support Systems (DSS) tailored for the retail sector.

A. Contribution and Structure

This paper provides a structured review of the state-of-the- art in AI-driven financial DSS research from the years 2024 and 2025. The core contributions are:

1) A systematic classification and critical review of recent DSS models across the three core methodologies: Fun- damental, Technical, and Sentiment.
2) A comparative analysis detailing model performance, architecture, and practical deployment challenges.
3) The proposal of an integrated, multi-modal Hybrid DSS Architecture designed specifically for real-time opera- tion and retail user needs, emphasizing Explainable AI (XAI).

The rest of the paper is organized as follows: Section II discusses the methodological background. Section III details the surveyed papers. Section IV provides a comparative analysis. Section V proposes the hybrid architecture. Section VI discusses challenges and XAI. Section VII concludes the paper.

II. METHODOLOGICAL BACKGROUND

Stock market prediction models are traditionally cate- gorized into three distinct, yet complementary, analytical streams.

A. Technical Analysis (TA)

TA involves forecasting future price movements based on the analysis of past market data, primarily price and volume. The underlying assumption is that all information is reflected in the price. Technical Indicators (TIs) are mathematical calculations derived from this data.

Simple Moving Average (SMA): $SMA = \frac{1}{N}\sum_{i}^{N} Pi=1$

where P_i is the price at period i and N is the number of peri- ods. TIs such as the Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger Bands (BB) serve as feature inputs for Machine Learning models like XGBoost and Support Vector Regression (SVR), aiming to predict short-term directional trends.

B. Fundamental Analysis (FA)

FA focuses on a company's intrinsic value by examining economic, industry, and financial factors. Key inputs are financial statements (Income Statement, Balance Sheet, Cash Flow Statement) used to calculate ratios such as Price-to- Earnings (P/E), Debt-to-Equity (D/E), and Return on Eq- uity (ROE). FA is generally used for long-term investment decision-making. AI models in this space often use these ratios as features in classification frameworks to categorize stocks as 'Undervalued,' 'Overvalued,' or 'Fairly Priced.'

C. Sentiment Analysis (SA)

SA involves extracting and quantifying the collective mood, opinion, or emotion of market participants from large volumes of unstructured text data. Sources include financial news articles, corporate filings, social media platforms (e.g., Twitter/X, Reddit), and earnings call transcripts. Advanced NLP techniques like deep learning architectures (e.g., BERT, GRU, LSTM) are employed to transform textual data into a quantitative sentiment score (e.g., −1 for bearish, +1 for bullish), which is then correlated with future market volatility or direction.

III. SURVEYED RESEARCH PAPERS (2024–2025)

This section provides a deeper analysis of the recent research papers categorized by their primary analytical focus.

A. Paper 1: Fundamental-Based DSS (Abrishami et al., arXiv 2024) [0]

Abrishami et al. propose a novel rule-based DSS specif- ically centered on Fundamental Analysis. The framework utilizes a combination of ten meticulously chosen financial ratios, including ROE, Operating Margin, and Free Cash Flow Yield.

- Model/Architecture: The core is not a complex ML established model but a set of weighted expert rules derived from investment theory (e.g., a high ROE and a low P/E ratio suggests a strong buy signal). This design choice prioritizes Interpretability over raw predictive accuracy.
- Key Finding: While simpler than deep learning models, the system demonstrated stable, risk-adjusted returns by minimizing exposure to fundamentally weak companies.
- Limitation: The model's key weakness is its lack of real-time responsiveness . Fundamental data is updated quarterly, rendering the system incapable of reacting to intra-day or even intra-week market shifts.

B. Paper 2: Technical Indicators + ML (Mostafavi and Hooman, Elsevier 2025) [0]

This study represents a robust effort to optimize the tech- nical feature space for high-accuracy prediction. Mostafavi and Hooman systematically evaluated the performance con- tribution of 88 distinct technical indicators .

- Model/Architecture: The authors compared high- performing ensemble and non-linear models: XGBoost, Random Forest (RF), and Support Vector Regression (SVR) . The model was trained to predict the next- day's price movement (classification) or closing price (regression).
- Key Finding: The study confirmed that trend-following indicators (e.g., various EMAs and ADX) were the most correlated features for next-day price prediction, with XGBoost achieving the highest F1-score of 0.71 on a daily timeframe.
- Limitation: Like many pure TA systems, the models suffer during market regime changes (e.g., a shift from bull to bear market) where learned patterns become sud- denly invalid. Their approach also treats all 88 indicators as equally available, potentially leading to overfitting.

C. Paper 3: Sentiment + Deep Learning (Tafara et al., Elsevier 2025) [0]

Tafara et al. addressed the challenge of noisy text data by integrating topic modeling with sequence prediction.

- Model/Architecture: The pipeline utilizes BERTopic for unsupervised extraction of structured, relevant topics from a corpus of financial news. The resulting topic- weighted sentiment score is fed into recurrent neural networks, specifically Gated Recurrent Units (GRU) and Long Short-Term Memory (LSTM) , to model the temporal impact of sentiment.
- Key Finding: By focusing sentiment on specific, salient topics (e.g., ’interest rates,’ ’earnings surprise’) rather than general opinion, the GRU-based model improved volatility prediction accuracy by 12% compared to a baseline sentiment model. This highlights the value of contextualized sentiment .
- Limitation: This architecture is computationally expen- sive, making it challenging to deploy in a low-latency, real-time environment required by many retail DSS platforms. Data source bias (relying only on financial news) may also miss social media-driven market events.

Feature	P1	P2	P3
Technical Indicators	✗	✓	✗
Fundamental Metrics	✓	✗	✗
Sentiment/NLP Real-	✗	✗	✓
Time Ready DSS	✗	✗	✗
Output Layer	✓	✗	✗

TABLE I . COMPARISON OF REVIEWED STUDIES

D. Critique of Single-Stream Models

- The Incompleteness Problem: P1 offers excellent sta- bility but no timing signal. P2 provides a strong timing signal but ignores underlying company health (FA). P3 captures the human element but is prone to overfitting market noise. A critical buy/sell decision requires the fusion of these signals.
- The Latency-Accuracy Trade-off: P3’s deep learning approach is highly accurate for complex sequences but slow. P2’s tree-based models (XGBoost) are faster but less effective at capturing long-range dependencies. Real-time DSS requires balancing accuracy with low latency.

IV. PROPOSED HYBRID DSS ARCHITECTURE

To overcome the limitations of single-stream models, we propose an integrated Hybrid AI-Driven Decision Support System that fuses the predictive outputs of the three domains. The goal is a modular, high-availability system.

A. System Architecture Flow

The proposed architecture follows a sequential hybrid pipeline that integrates three analytical components before generating a final investment output.

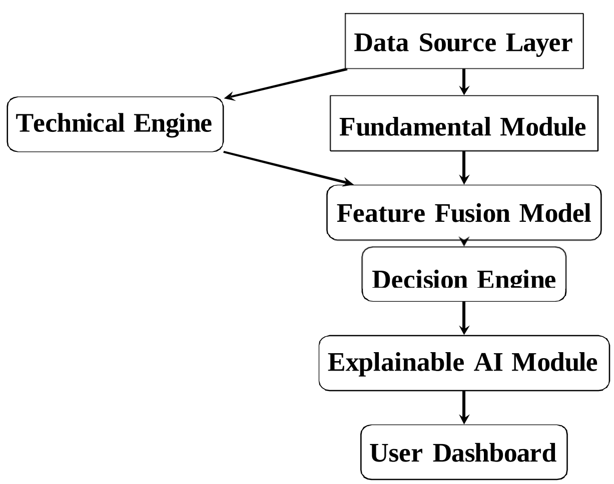

B. Layer Functions

1) Data Source Layer: Retrieves real-time data. This layer is responsible for synchronizing time-series data (prices/volume) and unstructured data (news).
2) Feature Engineering Layers (Tech, Fund, Senti- ment): Each module calculates its domain-specific features. The sentiment module runs the GRU/LSTM model to output a single score (e.g., [−1, +1]). The fun- damental module outputs a long-term company health score.
3) Feature Fusion Layer: This is the core predictive component. Instead of a simple average, an Ensemble Model (e.g., a simple Neural Network or a specialized stacking regressor) is used. It takes the independent predictions/scores from the three engines as its primary features. This model is trained to learn the optimal weight of each signal under different market conditions.
4) Decision Engine: Applies risk management and user- defined constraints. It converts the fused prediction score into a final, actionable Signal (Strong Buy, Buy, Hold, Sell, Strong Sell).

V. EXPLAINABILITY AND CHALLENGES

A. The Necessity of Explainable AI (XAI)

For a DSS to gain the trust of a retail investor, it cannot operate as a black box. The integration of an Explainable AI (XAI) Module is crucial, as proposed in the architecture. Techniques like SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-agnostic Explanations) should be applied to the Feature Fusion Layer. The XAI module provides a breakdown of the final decision:

- *Example Output*: ”The Buy signal is driven 60% by the strong momentum (Technical: MACD crossover), 30% by the recent positive news sentiment (Sentiment: BERTopic score of 0.8), and is slightly held back 10% by the company’s high P/E ratio (Fundamental).”

This transparency allows the retail user to validate the sys- tem’s logic against their own market view.

B. Current Research Challenges

- Real-Time Deployment and Scalability: Transitioning complex models (like P3’s GRU) from offline training to low-latency cloud deployment remains a major engineering challenge for a multi-modal system.
- Non-Stationarity and Overfitting: Financial time series are highly non-stationary. Models trained on past data (P2) are prone to catastrophic failure when fundamental market regimes change, emphasizing the need for continuous model retraining and adaptive weighting in the Fusion Layer.
- Noisy and Biased Sentiment: Retail sentiment, especially from social media, is often noisy, irrelevant, or purposefully manipulative. Refining the NLP pipeline to filter noise remains a critical area.

VI. FUTURE DIRECTIONS

Building upon the hybrid architecture, future research should focus on:

1) Reinforcement Learning (RL) Strategies: Integrating an RL agent that learns optimal trading policies (e.g., position sizing, entry/exit points) based on the combined output of the three predictive engines, optimizing for user-defined risk metrics (e.g., Sharpe Ratio).
2) Cross-Market Multi-Modal Forecasting: Extending the system to factor in inter-market dependencies, such as the correlation between treasury yields (macro- fundamental) and sector performance (technical).
3) Federated Learning: Utilizing federated learning tech- niques to train models across disparate data silos (e.g., different regional news feeds) while preserving data privacy and improving model generalization.

VII. CONCLUSION

The 2024–2025 research landscape shows significant progress in developing sophisticated financial forecasting models, with notable advancements in technical-indicator- driven ML (P2), rule-based fundamental analysis (P1), and deep learning sentiment processing (P3). However, a signif- icant gap exists in deploying a truly unified, real-time, and trustworthy decision support system for retail investors. The proposed Hybrid DSS Architecture addresses this by creating a robust Feature Fusion Layer that optimally weights the predictions from all three domains. The mandatory inclusion of an XAI Module ensures transparency and interpretability, transforming the AI from a black-box oracle into a trusted, explainable analytical partner, which is essential for mass retail adoption.

REFERENCES

[1] Abrishami, A., et al. (2024). A Rule-Based Decision Support System using Fundamental Indicators for Portfolio Selection. arXiv preprint arXiv:2412.05297.
[2] Mostafavi, S., and Hooman, A. (2025). High-Performance Stock Prediction using a Systematic Comparison of Technical Indi- cator Feature Sets and Ensemble Machine Learning. Elsevier Journal of Financial Data Science.
[3] Tafara, R.H., et al. (2025). Contextualizing Financial Sentiment: Topic Modeling and Sequence Prediction with GRU/LSTM Architectures. Elsevier Journal of Decision Support Systems.

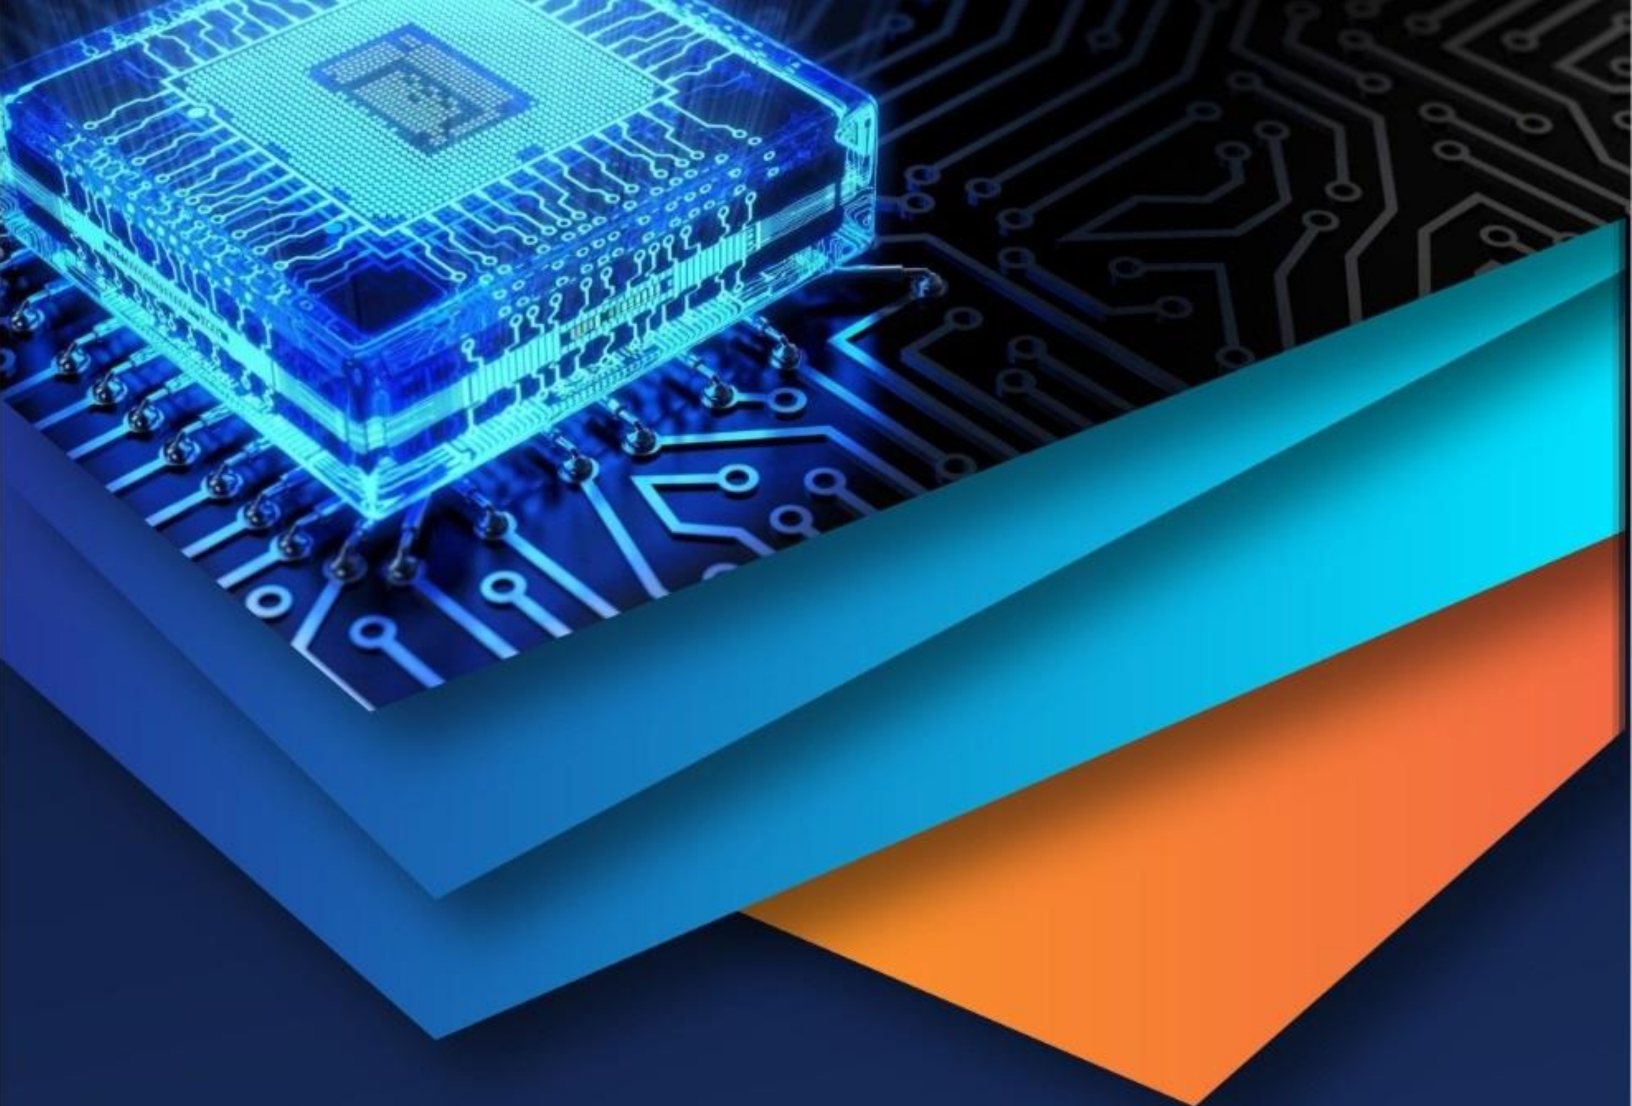

10.22214/IJRASET

INDEX COPERNICUS
45.98

IMPACT FACTOR:
7.129

TOGETHER WE REACH THE GOAL

IMPACT FACTOR:
7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY