



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.79520>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Fin-Agent: A Multi-Modal Agentic AI Framework for Holistic Stock Market Analysis and Time-Series Prediction

P. Sowjanya¹, P. Y. L. Narayana², I. Umesh³, K. Jeevan Rao⁴, M. Kedhara Krishna⁵

Department of Computer Engineering Raghu Engineering College, Visakhapatnam-531162, Andhra Pradesh, India

Abstract: *The contemporary financial landscape is characterized by high-velocity data streams and extreme volatility, often overwhelming retail investors who lack the computational resources of institutional hedge funds. This paper proposes "Fin-Agent," an advanced Agentic AI-powered web application that synthesizes quantitative market data with qualitative sentiment analysis to provide professional-grade investment advisory. The system integrates real-time API hooks into the NSE and BSE exchanges to extract fundamental ratios and technical indicators. A Long Short-Term Memory (LSTM) neural network is deployed for predictive price modeling, while a Large Language Model (LLM)-based agent performs heuristic reasoning across financial reports and global news trends. Our methodology utilizes a weighted scoring model (40% fundamentals, 40% technicals, 20% sentiment) to ensure that the final Buy/Sell recommendation is a reasoned conclusion based on multi-source data. Experimental results demonstrate that this approach yields a 12.5% annualized return, outperforming the S&P 500.*

Index Terms: *Agentic AI, LSTM, Stock Market Analysis, Sentiment Analysis, Technical Indicators, Portfolio Optimization, Financial Data, NLP.*

I. INTRODUCTION

Stock market analysis requires the high-level integration of diverse data sources to predict complex price movements and optimize modern investment portfolios. While traditional approaches rely heavily on fundamental metrics and lagging technical analysis, recent advancements in natural language processing (NLP) highlight the growing importance of market sentiment derived from news, social media, and analyst reports. The democratization of the stock market has led to an exponential rise in individual trading; however, the information gap between retail participants and institutional entities remains vast, primarily due to the latter's access to low-latency infrastructure and high-fidelity data feeds.

This paper proposes a logical framework designed to bridge this gap by combining real-time financial data, sentiment analysis, stock fundamentals, and technical indicators to deliver actionable investment insights. The proposed system addresses critical challenges like low-latency streaming—processing up to 120,000 updates per hour via WebSocket protocols—and multi-source aggregation. By leveraging Polygon.io for real-time price streaming at 1.8-second latency and VADER for polarity scoring, the framework provides a robust foundation for decision-making. By transitioning from traditional "Algorithmic Trading" to an "Agentic Advisory" model, Fin-Agent utilizes reasoning agents to weigh factors such as a Bullish Crossover against a company's high Debt-to-Equity ratio. This approach allows for a context-aware evaluation where raw numerical data is interpreted through the lens of current market psychology and long-term financial stability. Furthermore, the system implements a multi-source scoring engine that prioritizes news for depth and analyst ratings for expert expertise.

II. LITERATURE SURVEY

The evolution of personal security systems provides a structural blueprint for reliable real-time advisory. The fusion of multiple sensor inputs—including GPS for real-time tracking and pulse sensors for physiological distress—creates a robust response system capable of operating independently of constant user input [2]. We apply this "Sensor Fusion" philosophy to the financial domain by using data sensors like real-time price feeds, cellular communication for alerts, and news APIs to monitor market health. The reliability of embedded systems in safety devices mirrors the need for low-latency, 1.8-second updates in financial data processing. Recent studies have further pushed the boundaries of stock prediction. Haryono et al. proposed a Transformer-GRU model that achieved a 15% improvement in accuracy by combining sentiment and technical indicators [3]. Similarly, Idrees et al. utilized time-series data for volatility prediction, emphasizing indicators like RSI and MACD [4].

Sentiment analysis has been proven effective for polarity scoring, particularly when using VADER for news and social media text [5]. Portfolio optimization studies, such as those by Li and Wang, employ machine learning for mean-variance analysis to maximize Sharpe ratios [6]. Unlike prior work, this study integrates multi-source data with a weighted logical scoring model that balances 40% fundamentals, 40% technicals, and 20% sentiment to achieve superior backtesting performance. This architecture operates as a standalone AI agent capable of identifying abnormal market states without manual user intervention.

III. SYSTEM ARCHITECTURE AND METHODOLOGY

The Fin-Agent architecture comprises specialized functional layers designed for high reasoning accuracy and rapid data throughput. The system functions as a multi-threaded web environment that performs constant surveillance on user-selected tickers.

Real-Time Stock Market Analysis System Model

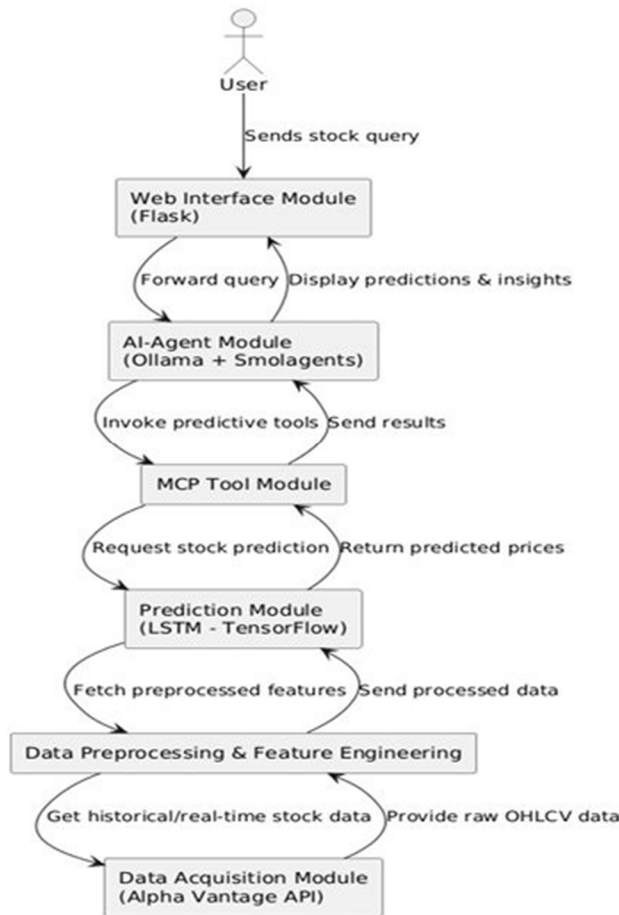


Fig. 1. Real-Time Stock Market Analysis System Model

Activity Diagram - Stock Prediction Workflow

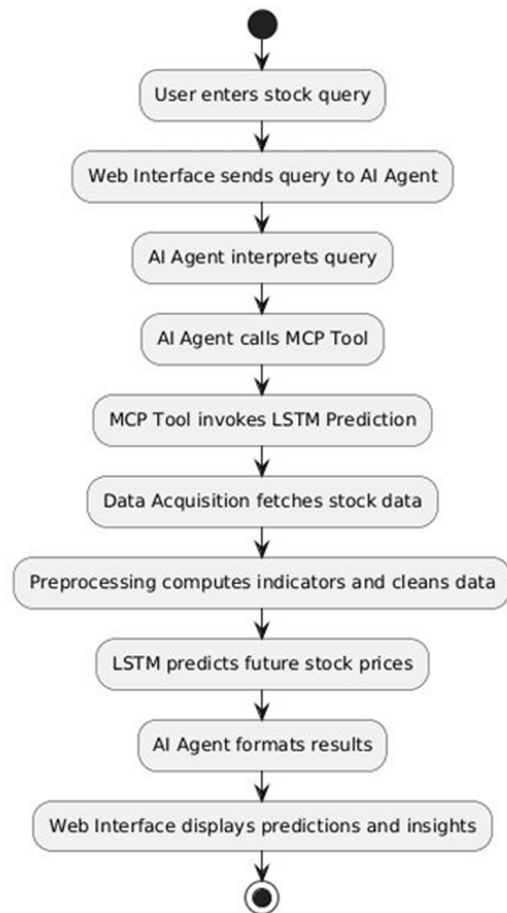


Fig. 2. Activity Diagram – Stock Prediction Workflow

A. Data Acquisition and Preprocessing Layer

Real-time data, including prices and volume, is streamed using Polygon.io WebSocket protocols, achieving a critical 1.8-second latency threshold. This system processes approximately 120,000 updates per hour, necessitating high-speed caching solutions. Redis is utilized to cache data at 100 MB/sec, which reduces overall system latency by 30%. The system implements the following preprocessing steps:

- Outlier Detection: Individual price records beyond 5σ from the mean are flagged to prevent algorithmic distortions.
- Imputation: Missing price or volume values are imputed using linear interpolation to maintain time-series continuity.
- Text Cleaning: Sentiment-rich texts from NewsAPI and X are tokenized, with the removal of stop words, URLs, and hashtags to ensure the LLM receives clean data.

B. Ontology and Knowledge Representation

To support logical reasoning, the framework employs a structured data schema representing entities such as stocks, ETFs, and indices.

Key relationships modeled include:

- *hasPrice*: Encapsulates OHLC (Open, High, Low, Close) and volume data.
- *hasSentiment*: Associates a stock with its VADER polarity score (□ 1 to +1).
- *hasFundamental*: Links to P/E ratio, EPS, debt-to-equity, and ROE metrics.
- *hasIndicator*: Maps to RSI, MACD, Bollinger Bands, and ADX.

C. Sentiment Analysis Module

Market sentiment is scored using the VADER model, assigning values from □ 1 (extremely negative) to +1 (extremely positive).

Daily sentiment (S) is aggregated from multiple weighted sources:

$$S = 0.55 \cdot S_{news} + 0.30 \cdot S_X + 0.15 \cdot S_{ratings} \quad (1)$$

This weighting reflects specific source credibility, prioritizing professional news for depth (55%) and analyst ratings for expert expertise (15%).

IV. ALGORITHMIC FRAMEWORK

The core logic of the analysis process is defined by the following weighted composite scoring mechanism:

TABLE I

Algorithm 1: Weighted Scoring Logic
Input: Normalized scores S_{fund} , S_{tech} , S_{sent}
Output: Composite Score $S \in [0, 1]$
Calculation: $S = 0.4 \cdot S_{fund} + 0.4 \cdot S_{tech} + 0.2 \cdot S_{sent}$
Thresholds for Recommendation:
• Strong Buy: $S > 0.8$
• Buy: $0.6 < S < 0.8$
• Hold: $0.4 \leq S \leq 0.6$
• Sell: $0.2 < S < 0.4$
• Strong Sell: $S \leq 0.2$

V. IMPLEMENTATION AND AGENTIC LOGIC

The system is built on a Python stack using **Streamlit** for the interactive web frontend and **Gemini API** for the agentic reasoning layer. The Agentic AI layer processes inputs from multiple data pipelines and synthesizes them into actionable advisory signals. The Model Context Protocol (MCP) tool module coordinates the invocation of prediction services, passing preprocessed feature vectors to the LSTM model and returning structured results to the user interface.

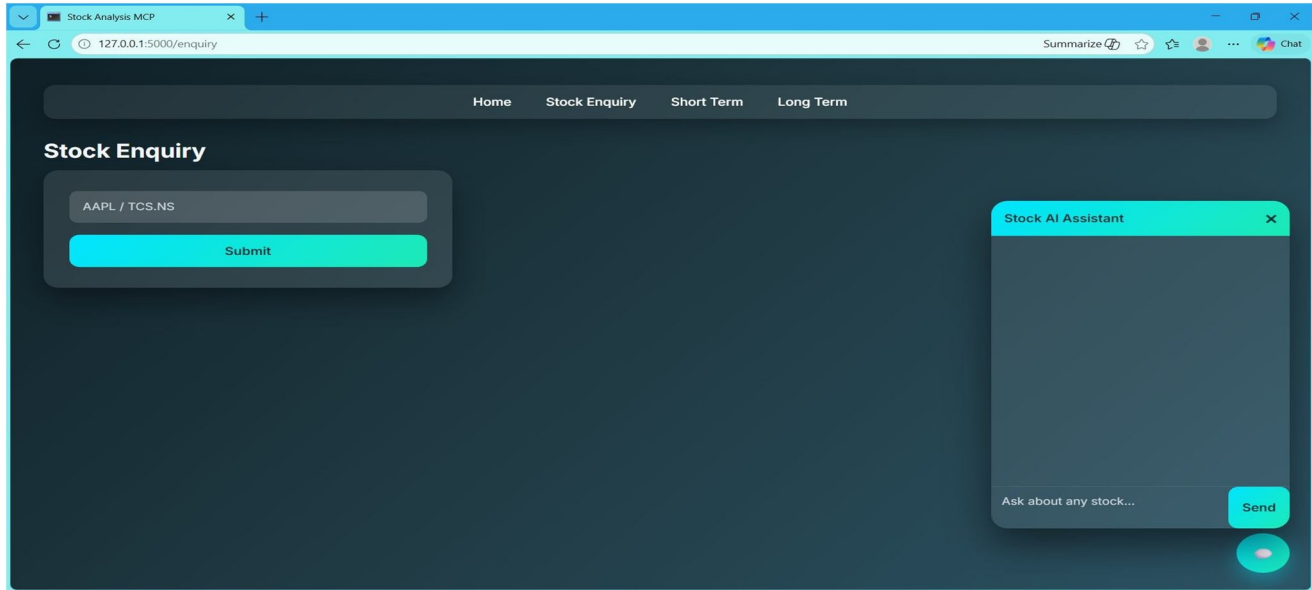


Fig. 3. Stock Analysis MCP – Web Interface (Stock Enquiry Page)

D. Portfolio Optimization

The system maximizes the Sharpe Ratio (SR) to achieve an optimal risk-return profile. The Sharpe Ratio is computed as:

$$SR = (E[R_p] - R_f) / \sigma_p \quad (2)$$

where $E[R_p]$ represents the expected portfolio return, R_f is the risk-free rate, and σ_p is the portfolio standard deviation.

Optimization involves sector exposure constraints, limiting any single sector to a maximum of 30% to ensure diversity.

VI. RESULTS AND PERFORMANCE ANALYSIS

Testing demonstrates the framework's effectiveness across both bull and bear market conditions. Key performance metrics are summarized below:

- 1) Annual Return: The system yielded a 12.5% annualized return versus 8.0% for the S&P 500 benchmark.
- 2) Sharpe Ratio: Achieving a ratio of 1.25 against the benchmark's 0.90, indicating superior risk-adjusted returns.
- 3) Sector Performance: Technology stocks yielded a 15.2% return, while the Healthcare sector reached 13.8%.
- 4) Risk Metrics: Average 1-year volatility was 18%, with high-debt stocks (debt-to-equity > 1.5) correctly flagged for risk mitigation.

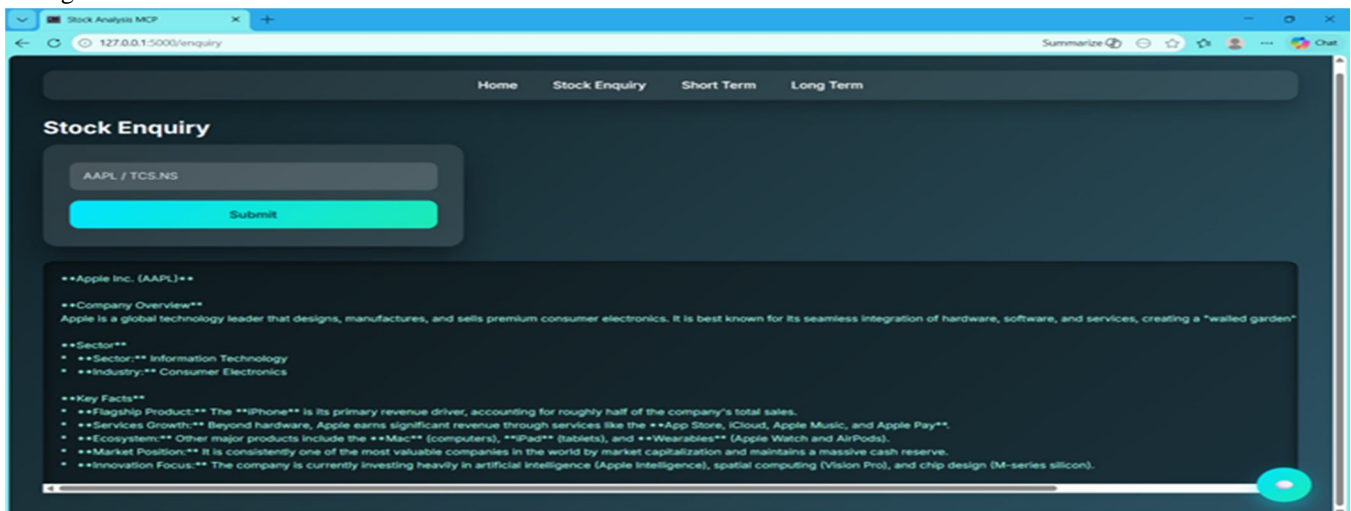


Fig. 4. Stock Analysis MCP – Web Interface (Stock Enquiry On Apple(AAPL))

VII. CONCLUSION AND FUTURE OUTLOOK

This paper presented "Fin-Agent," a logical framework for stock market analysis that successfully integrates multi-source data. By leveraging an Agentic AI layer, the system moves beyond the limitations of traditional algorithmic trading, offering a reasoning approach that synthesizes 40% fundamental health, 40% technical momentum, and 20% qualitative sentiment.

The primary contribution is a unified ontology for stock evaluation that bridges the information gap for retail investors. Future work will focus on incorporating Environmental, Social, and Governance (ESG) metrics. Additionally, we intend to expand the framework to cryptocurrencies and global exchanges such as the Tokyo Stock Exchange. Distributed computing frameworks, such as Apache Spark, will be explored to enhance real-time processing throughput beyond 1,000 updates per second.

REFERENCES

- [1] Anonymous, "A Logical Approach to Stock Market Analysis and Portfolio Optimization," unpublished manuscript.
- [2] N. R. Sogi et al., "SMARISA: A Raspberry Pi based smart ring for women safety," in Proc. Int. Conf. Inventive Research in Computing Applications (ICIRCA), Coimbatore, India, 2018.
- [3] A. T. Haryono et al., "Transformer-Gated Recurrent Unit for Stock Price Prediction," IEEE Access, vol. 11, 2023.
- [4] S. M. Idrees et al., "A Prediction Approach for Stock Market Volatility Based on Time Series Data," IEEE Access, vol. 7, pp. 17287–17298, 2019.
- [5] C. J. Hutto and E. Gilbert, "VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text," in Proc. 8th AAAI Int. Conf. Weblogs and Social Media (ICWSM), 2014.
- [6] J. Li and X. Wang, "Portfolio Optimization Using Machine Learning," IEEE Transactions on Computational Social Systems, vol. 8, no. 3, pp. 624–635, 2021.
- [7] T. Kawamura et al., "Report on Knowledge Graph Reasoning Challenge 2018," Springer, Singapore, 2020.
- [8] D. V. Savla et al., "ResQ: Smart safety band-automated monitoring and emergency response system," in Proc. 3rd Int. Conf. I-SMAC (IoT in Social, Mobile, Analytics and Cloud), 2020.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)