© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717317

Deepstock : AI Powered Stock Analysis

SHREYASH NIMBARGI[1], VEDANT JOSHI[2], OM HUNDEKARI[3], VEDANT JAGTAP[4]
[1,2,3,4]*Computer Science Engineering, MIT Art, Design & Technology University, Pune, India*

Abstract- The pervasive challenge in modern quantitative finance lies in the failure of traditional models to integrate crucial, non-numerical market factors, leading to poor predictive stability during event-driven volatility. This manuscript details the DeepStock platform, an operative architecture engineered to systematically harmonize structured financial time-series data with highly descriptive, unstructured textual narratives. The system leverages the superior contextual reasoning of the Google Gemini Large Language Model (LLM) to generate robust, quantifiable financial sentiment scores. This sentiment feature vector is fused with traditional technical indicators via an Early Fusion strategy to construct a comprehensive predictive feature set. We hypothesize that models incorporating these hybrid features will yield superior risk-adjusted returns, as validated by the Sharpe Ratio, relative to purely time-series-dependent strategies. The deployment utilizes a Flask/Python framework to provide a free, real-time AI tool focused on analysis and prediction within the NSE and BSE stock markets.

I. INTRODUCTION

1) A. CHALLENGES IN MODERN STOCK PREDICTION AND THE NEED FOR MULTI-MODALITY

Traditional quantitative finance, reliant on historical prices and technical metrics, often exhibits inherent predictive fragility, particularly when confronting sudden, event-driven market volatility.2 This limitation stems from the failure to account for crucial latent factors embedded in market narratives. The central challenge is the synergistic fusion of heterogeneous data: blending continuous time-series data with noisy, unstructured textual content while ensuring precise synchronization and alignment.3 This research addresses this methodological gap by proposing an architecture designed to operationalize this multi-modal data fusion.

2) PROJECT CONTRIBUTION AND HYPOTHESIS

This investigation is based on the hypothesis that a predictive model utilizing feature vectors derived from the structured sentiment analysis of real-time financial news, processed through the Google Gemini LLM, will achieve superior risk-adjusted performance (Sharpe Ratio) compared to traditional technical indicator models. The project contributes a robust operational architecture that performs entity resolution and time-synchronization; it validates the use of Gemini AI for generating quantifiable sentiment features, capitalizing on its superior efficacy over classical methods;5 and it deploys the DeepStock platform, a free, real-time tool focusing on NSE and BSE markets, utilizing Flask/Python for foundational Explainable AI (XAI) insights.7

II. LITERATURE REVIEW

1) A. THE RISE OF TRANSFORMER MODELS IN FINANCIAL TEXT ANALYSIS

The history of NLP in finance, starting with dictionary-based approaches, was limited by its inability to capture semantic nuance and contextual shifts in complex market narratives.6 Current academic inquiry overwhelmingly endorses advanced LLMs for complex financial tasks, valuing them for their deep contextual comprehension rather than simple classification.9 Comparative studies confirm the heightened efficacy of models like the Gemini series in Text-Level Sentiment Analysis (TLSA) against predecessor and rule-based systems.5

2). MULTI-MODAL DATA FUSION IN TIME-SERIES FORECASTING

Empirical research shows that prediction accuracy is significantly enhanced by combining diverse data modalities, including textual news and historical price data.4 The architectural approach in this project aligns conceptually with cascade fusion architectures, where feature layers are sequentially integrated to produce a unified input vector.12 Since the direct

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717317

application of LLMs to time-series forecasting remains preliminary due to the low signal-to-noise ratio,3 extracting a stable, numeric sentiment feature vector (known as Early Fusion) and merging it with traditional numerical data remains the most scientifically validated fusion strategy.4

III. SYSTEM ARCHITECTURE AND IMPLEMENTATION

1) ARCHITECTURAL RATIONALE: PYTHON/FLASK FRAMEWORK

The system utilizes the extensible Python environment, anchored by the Flask micro-framework, which offers an optimal structure for deploying complex AI solutions and handling efficient API interactions.15 The analytical core relies on industry-standard libraries, including Pandas and NumPy, essential for efficient numerical processing and vectorized calculations.

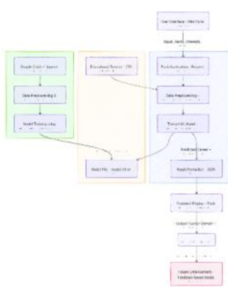

Fig. 1. System architecture of the AI-Powered Stock Analysis

B. Structured Data Acquisition and operational constraints

Historical price and volume data are sourced via the yfinance library, specifically engineered to handle the structured data requirements of the NSE and BSE Stock Screener markets.1 A crucial limitation is the reliance on community-driven APIs, which impose severe rate-limiting constraints (e.g., 1,000 to 2,000 requests per day/hour) on unauthenticated usage.16 This constraint limits the system's immediate scalability and the capacity for large-scale, high-frequency backtesting

C. UNSTRUCTURED DATA ACQUISITION AND ENTITY RESOLUTION

Unstructured data (news headlines and texts) is gathered via Google Search and the Newspaper3k library, an efficient Python tool for extracting structured content from online articles.18 To link this noisy data to the correct financial asset, highly reliable Entity Resolution (ER) is critical. This is achieved using the FuzzyWuzzy library, which employs the Levenshtein distance algorithm to calculate similarity between scraped entity names and pre-defined stock tickers, thereby ensuring data integrity and alignment.

Component Category	Technology Used	Primary Function	Academic Relevance / Research Contribution
Backend & Runtime	Python, Flask, Jinja2	System logic, API handling, web templating	Robust framework for deploying complex AI/LLM solutions and interactive interfaces.[7]
Structured Data	yfinance	Retrieval of historical and real-time stock pricing data	Source of traditional, numerical time-series features. Operational challenges noted due to rate limits.[16]
Unstructured Data	Google Search, Newspaper3k	Web scraping and extraction of news articles/summaries	Generating high-frequency, alternative textual data streams.[18]
Preprocessing	Pandas, NumPy, FuzzyWuzzy	Data handling, numerical calculation, Entity Resolution	Critical step for data cleaning and aligning noisy news entities to financial tickers.[20]
Inference/Analysis	Google Gemini AI	Sentiment scoring, financial analysis, structured insight generation	Leveraging state-of-the-art LLM capabilities for superior financial text classification.[5]

IV. LLM-BASED SENTIMENT FEATURE ENGINEERING

.

1) A. PROMPT ENGINEERING FOR STRUCTURED OUTPUT

To integrate the LLM's textual analysis into a numerical pipeline, the output must be strictly standardized and structured. The methodology employs Structural Formatting and Constraints via prompt engineering, forcing the Gemini model to return machine-readable formats (e.g., JSON) with specific numerical sentiment scores (scaled -1.0 to 1.0), categorical tones, and justifications. This adherence ensures clean data ingestion and mitigates risks associated with non-standardized LLM generation.

2) B. MULTI-MODAL FEATURE VECTOR CONSTRUCTION (EARLY FUSION)

The system utilizes the Early Fusion strategy, processing the textual modality into a numerical vector (F_T) before combining it with the

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717317

structured numerical vector (F_N).4 Individual article scores are aggregated over a temporal window. The final input feature vector ($F_{Combined}$) is constructed by direct concatenation of F_N (lagged returns, volume indicators) and F_T (derived sentiment features). The integrity of this vector relies on the robust chaining of all pre-processing steps (Scraping $\to$ ER $\to$ Structured LLM Output) to maximize the signal-to-noise ratio.

Data Modality	**Source/Library**	**Feature Generation Process**	**Output Feature Format**
Numerical (F_N)	yfinance / Pandas	Calculation of lagged returns, volatility, and volume indicators.	Float vector ($F_N \in \mathbb{R}^k$)
Textual (F_T)	Gemini AI (via Structured Prompting)	**1.** News text scraped. **2.** Prompt Gemini to output structured JSON: {“Sentiment Score”: X, “Tone”: Y}.	Float vector ($F_T \in \mathbb{R}^m$), where m is the dimension of sentiment/tone. [21]
Fusion Mechanism	Python/NumPy	Direct concatenation of F_N and F_T to form a combined feature set $F_{Combined} =$.	Multi-Modal Input Vector ($F_{Combined} \in \mathbb{R}^{k+m}$) [4]

V. EVALUATION FRAMEWORK AND FINANCIAL PERFORMANCE VALIDATION

Beyond basic Accuracy, directional classification success is measured using the Matthews Correlation Coefficient (MCC) and F1-Score.4 MCC is critical for providing a reliable, balanced measure of classification quality in highly imbalanced market datasets.

As technical accuracy often fails to correlate with profitable trading strategies,23 the core validation metric is the Sharpe Ratio (SPR), which measures excess return per unit of total volatility (risk).24 The Sortino Ratio (STR) focuses on downside risk, and Maximum Drawdown (MDD) quantifies worst-case capital erosion.11 Future algorithmic refinement should investigate the integration of Sharpe Ratio-aware loss functions, optimizing the model directly for high risk-adjusted returns.23

Table 3. Proposed Evaluation Metrics for Financial Prediction Systems

Metric Category	**Specific Metric**	**Formula Basis**	**Justification in FinTech Research**
Classification Quality	F1-Score / MCC	Harmonic mean of Precision/Recall; Correlation coefficient.	Measures the true predictive balance and reliability of directional classification in imbalanced market datasets.[4]
Risk-Adjusted Return	Sharpe Ratio (SPR)	(Portfolio Return - Risk-Free Rate) / Volatility (σ)	Essential metric assessing return per unit of total risk. Validates if the prediction translates to effective capital deployment.[23]
Downside Risk	Sortino Ratio (STR)	(Portfolio Return - Risk-Free Rate) / Downside Volatility	Crucial for institutional investors; focuses exclusively on volatility associated with negative returns.[11]
Capital Protection	Maximum Drawdown (MDD)	Max Loss from Peak-to-Trough	Measures the worst-case scenario capital erosion, critical for defining risk tolerance and system stability.[11]

VI. DEPLOYMENT, USER INTERFACE, AND EXPLAINABILITY

The Flask framework provides the platform for demonstrating the LLM’s decision-making interpretability, formalizing the concept of Explainable AI (XAI).26 The DeepStock platform is positioned as a comprehensive, free AI tool that increases user confidence by displaying the entire data chain: scraped news, Gemini-derived structured sentiment (F_T), and the final trading recommendation.8

The current system relies on a local deployment environment (app.run(debug=True)) using SQLite. However, migration to an enterprise-level system requires significant architectural changes. The current synchronous API calls for scraping and LLM inference will not scale effectively. Future development must incorporate asynchronous task handling to ensure responsiveness and manage concurrent data streams under heavy load.7

VII. FUTURE WORK

1. Incorporation of Advanced AI
In order to better understand student interests and increase prediction accuracy, future iterations will make use of deep learning and natural language processing (NLP) models, such as BERT or GPT. Additionally, explainable AI tools will be included to demonstrate the recommendation-making process.

2. Cloud & Mobile Expansion
Create a mobile application using Flutter or React Native, then make it easily accessible by deploying it on cloud platforms like AWS or Heroku. Provide APIs so that educational institutions can incorporate DeepStock AI into their own systems.

3. A More Intelligent Recommendation System
To make the results more precise and industry-relevant, incorporate role-based recommendations, real-time job market data, and psychometric tests.

4. Integration of Institutions
Provide dashboards that allow schools and counselors to monitor student insights while maintaining secure access and data privacy.
5. Scalability & Performance
For seamless deployment and scaling, switch to PostgreSQL databases, use Docker, and handle tasks with Celery + Redis.

6. Optimization & Security
Include automated error handling, input validation, and HTTPS. Boost dependability and response time for high user loads.

7. Constant Research
Collaborate with educators to verify AI forecasts and disseminate results to enhance the system.

VIII. CONCLUSION

The Design and Implementation Of Deepstock Ai, An Ai-Powered Stock Analysis and Prediction Platform Created To Streamline Financial Research And Investment Decision-Making, Were Presented In This Paper. Manual Data Tracking and Subjective Interpretation Are Major Components of Traditional Stock Analysis, Which Frequently Results In Inefficiencies, Missed Insights, And Inconsistent Outcomes.

By Combining Web-Based Visualization Tools, Technical Analysis, And Machine Learning, Deepstock AI Offers an Automated, Data-Driven Solution. Using Ai Models Trained on Actual Market Datasets, The System Examines Past Stock Data, Finds Trends, And Forecasts Future Developments. Users Can Engage with The Platform, View Analysis Results, And Make Well-Informed Trading or Investment Decisions Instantly with a Web Interface Built on Flask.

REFERENCES

[1] DeepStock - Free AI Stock Analysis & Prediction Tool | Real-Time ..., accessed November 4, 2025, https://deepstock.in/

[2] Neural based event-driven stock rally prediction using SEC filings and Twitter data, accessed November 4, 2025, https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1204/reports/custom/15787311.pdf

[3] From Deep Learning to LLMs: A survey of AI in Quantitative Investment - arXiv, accessed November 4, 2025, https://arxiv.org/html/2503.21422v1

[4] Combining Financial Data and News Articles for Stock Price Movement Prediction Using Large Language Models - arXiv, accessed November 4, 2025, https://arxiv.org/html/2411.01368v1

[5] On Assessing the Performance of LLMs for Target-Level Sentiment Analysis in Financial News Headlines - MDPI, accessed November 4, 2025, https://www.mdpi.com/1999-4893/18/1/46

[6] [2412.19245] Sentiment trading with large language models - arXiv, accessed November 4, 2025, https://arxiv.org/abs/2412.19245

[7] Building Scalable AI Solutions with Python and Flask - Shiv Technolabs, accessed November 4, 2025, https://shivlab.com/blog/building-scalable-ai-solutions-python-flask/

[8] Transforming a Financial Model into a Web App with AI (Updated Sept 2025), accessed November 4, 2025, https://www.adventuresincre.com/financial-model-web-app/

[9] [2510.15929] Comparing LLMs for Sentiment Analysis in Financial Market News - arXiv, accessed November 4, 2025, https://www.arxiv.org/abs/2510.15929

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717317

[10] CatMemo at the FinLLM Challenge Task: Fine-Tuning Large Language Models using Data Fusion in Financial Applications - ACL Anthology, accessed November 4, 2025, https://aclanthology.org/2024.finnlp-2.19.pdf

[11] Can LLM-based Financial Investing Strategies Outperform the Market in Long Run? - arXiv, accessed November 4, 2025, https://arxiv.org/html/2505.07078v3

[12] The architecture of the Multi-modal fusion module. The architecture of... | Download Scientific Diagram - ResearchGate, accessed November 4, 2025, https://www.researchgate.net/figure/The-architecture-of-the-Multi-modal-fusion-module-The-architecture-of-the-Multi-modal_fig1_352279961

[13] Free Download Multi-Modal Data Fusion Architecture Plan - Meegle, accessed November 4, 2025, https://www.meegle.com/en_us/advanced-templates/precision_medicine/multi_modal_data_fusion_architecture_plan

[14] Prompt Engineering for Finance 101 | Deloitte US, accessed November 4, 2025, https://www.deloitte.com/us/en/services/consulting/articles/prompt-engineering-for-finance.html

[15] The Advantages of Using Streamlit, Flask, and Django for Building Web Apps - Medium, accessed November 4, 2025, https://medium.com/@akratech/the-advantages-of-using-streamlit-flask-and-django-for-building-web-apps-67f0907861e4

[16] What is the query limit on Yahoo's Finance API? - Stack Overflow, accessed November 4, 2025, https://stackoverflow.com/questions/9346582/what-is-the-query-limit-on-yahoos-finance-api

[17] Why yfinance Keeps Getting Blocked, and What to Use Instead | by Trading Dude - Medium, accessed November 4, 2025, https://medium.com/@trading.dude/why-yfinance-keeps-getting-blocked-and-what-to-use-instead-92d84bb2cc01

[18] Newspaper3k: Article scraping & curation — newspaper 0.0.2 documentation, accessed November 4, 2025, https://newspaper.readthedocs.io/

[19] How To Scrape News Articles with Newspaper3k (Python) - ScraperAPI, accessed November 4, 2025, https://www.scraperapi.com/blog/python-newspaper3k/

[20] Fuzzy String Matching Using FuzzyWuzzy - HKUST Library, accessed November 4, 2025, https://library.hkust.edu.hk/sc/fuzzywuzzy/

[21] Structured output | Gemini API - Google AI for Developers, accessed November 4, 2025, https://ai.google.dev/gemini-api/docs/structured-output

[22] Prompt Engineering Part 4: Architecting Reliable Outputs with Optimization Techniques | by Akanksha Sinha | Medium, accessed November 4, 2025, https://medium.com/@akankshasinha247/prompt-engineering-part-4-architecting-reliable-outputs-with-optimization-techniques-7bfa0b1b98c7

[23] Deep Learning for Stock Performance Prediction: A Sharpe Ratio-Optimized Approach, accessed November 4, 2025, https://www.researchgate.net/publication/390130909_Deep_Learning_for_Stock_Performance_Prediction_A_Sharpe_Ratio-Optimized_Approach

[24] Can LLM-based Financial Investing Strategies Outperform the Market in Long Run? - arXiv, accessed November 4, 2025, https://arxiv.org/html/2505.07078v1

[25] Stock Price Prediction and Portfolio Optimization Based on Mean Variance Model and Random Forest Model - Atlantis Press, accessed November 4, 2025, https://www.atlantis-press.com/article/126008533.pdf

[26] LLMs for Explainable AI: A Comprehensive Survey - arXiv, accessed November 4, 2025, https://arxiv.org/html/2504.00125v1