RESEARCH ARTICLE OPEN ACCESS

Trade Hub: Real Time Stock Market Analysis with Trade Hub Bot

Gatti Koppula Satvik Reddy[1], Gedala Shantha Rao[2], Beeram Vara Prasad[3], Basina Raja Gopal[4], Mrs.C. Pavani Reddy[5], Dr.B.Venkataramana[6].

[1]*Student, AI & ML, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*
[2]*Student, AI & ML, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*
[3]*Student, AI & ML, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*
[4]*Student, AI & ML, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*
[5] *Assoc Prof,, AI & ML, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*
[6]*Assoc Prof, CSE, Holy Mary Institute of Technology and Science, Hyderabad, TS, India*

Abstract
Accurate financial decision-making in volatile markets relies heavily on the rapid synthesis of heterogeneous data streams, yet most existing trading systems are developed with narrowly defined analytical scopes, isolating technical charting from fundamental news sentiment. This research presents TRADE HUB, a unified AI-driven algorithmic trading framework that integrates real-time market data processing, historical trend forecasting, and sentiment analysis within a single, full-stack ecosystem. The core premise of this work is that quantitative market indicators and unstructured financial news, despite being computationally distinct, can be converged into a holistic predictive model for superior market guidance. The proposed framework employs a robust Full-Stack Architecture, utilizing a Python-based backend (FastAPI) for high-performance logic processing and a reactive frontend (React.js) for dynamic visualization. The system leverages Angel One SmartAPI for real-time Indian market data and Binance API for global crypto assets and forex market, ensuring cross-domain market coverage. The project's analytical core utilizes a hybrid logic engine that processes technical indicators—specifically the SMA (Simple Moving Average) Crossover Trade Hub Strategy-to generate precise buy/sell signals and Live market analysis. This is augmented by a News Processing Pipeline that aggregates financial narratives via DuckDuckGo Search and analyses sentiment using Groq's Llama-3.1-8b-instant Large Language Model (LLM) to provide context-aware market summaries. The primary contribution of this work lies in establishing a scalable, multi-modal trading paradigm that moves beyond static technical analysis, demonstrating the adaptability of algorithmic architectures to fuse AI-driven sentiment with live price action. These findings suggest that a comprehensive artificial intelligence system can be engineered to democratize institutional-grade market insights, positioning TRADE HUB as a scalable foundation for next-generation intelligent trading assistance.
Keywords: Algorithmic Trading, Sentiment Analysis, Large Language Models (LLM), Real-Time Market Data, Technical Analysis, SMA Crossover Strategy, Financial News Processing, Cross-Domain Analysis.

Date of Submission: 14-01-2026 Date of acceptance: 28-01-2026

I. INTRODUCTION

TRADE HUB is a comprehensive web-based financial intelligence portal designed to revolutionize how retail traders interact with global markets. The platform welcomes users with a sophisticated, high-frequency dashboard that serves as a centralized command center, aggregating disparate market streams into a single, unified visual interface. By simultaneously displaying real-time price ticks and interactive charts for Indian Equities and global Forex/Crypto assets, the system eliminates the friction of switching between platforms, offering immediate cross-domain market visibility. Seamlessly integrated within this graphical environment is the Trade Hub Bot, accessible via a dedicated workspace that transitions users from passive monitoring to active, intelligent analysis. Built on a robust Full-Stack Architecture utilizing high-performance FastAPI (Python) and reactive React.js, the bot functions as an on-demand virtual analyst capable of processing complex market data in milliseconds. It employs a specialized hybrid logic engine that executes a proprietary SMA (Simple Moving Average) Crossover Trade Hub strategy, dynamically switching between "Scalp" and "Trend" modes to generate precise, mathematically validated buy/sell signals with calculated entry, target, and stop-loss levels. Beyond quantitative charting, the system leverages Groq's Large Language Models (LLM) and DuckDuckGo to synthesize breaking financial news, providing context-aware sentiment analysis that explains *why* the market is moving, not just *how*. The interface enhances data comprehensibility through interactive Recharts and Chart.js integration, transforming raw volatility data into intuitive performance reports. By fusing professional-grade live monitoring with AI-driven strategic guidance, TRADE HUB bridges the gap between institutional algorithms and retail accessibility, establishing a scalable foundation for next-generation intelligent market assistance.

II. RELATED WORK

Artificial Intelligence (AI) and Machine Learning (ML) have fundamentally reshaped financial forecasting, transitioning market analysis from static statistical models to dynamic, data-driven paradigms. Deep learning architectures, particularly Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, have become the dominant approach in modeling stock market volatility. However, many existing systems rely heavily on historical closing prices and often fail to account for the stochastic nature of real-time intraday data, including sudden volume spikes and breaking news sentiment. Furthermore, most academic implementations remain isolated as predictive models without integration into deployable, full-stack trading environments, limiting their practical utility for retail investors. Current solutions frequently lack a unified framework that simultaneously processes quantitative technical indicators and qualitative unstructured data.

In the domain of algorithmic forecasting, Fischer and Krauss [1] demonstrated that LSTM networks outperform traditional machine learning methods like Random Forest in predicting directional stock movements by effectively capturing long-term temporal dependencies. Nelson, Pereira, and de Oliveira [2] substantiated this by utilizing LSTM to model non-linear price patterns with higher accuracy than varying baseline strategies. To enhance feature extraction, hybrid architectures have gained prominence; Miao et al. [6] and Joshi et al. [7] proposed combined CNN-LSTM frameworks where Convolutional Neural Networks (CNNs) extract spatial features from technical charts while LSTMs analyse temporal sequences. Despite their success, these models often operate in a "black box," lacking the interpretability required for trusted financial decision-making.

Parallel to price prediction, sentiment analysis has emerged as a critical component of modern fintech. Rekha and Sabu [10] combined autoencoders with sentiment scoring to filter noise from financial news, while Gu et al. [11] and Halder [12] integrated domain-specific language models like FinBERT with predictive engines. Their results indicated that incorporating sentiment significantly boosts accuracy during volatile market periods. However, standard BERT-based approaches are limited to static classification (positive/negative) and lack the generative reasoning capabilities of modern Large Language Models (LLMs). Zhang and Lee [15] further highlighted the efficacy of conversational AI agents in democratizing financial data, a concept directly relevant to the interactive architecture of TRADE HUB. While technical indicators remain a staple of trading, recent works by Dhokane [3] and Him [4] emphasized that isolating indicators like RSI and MACD often yields lagging signals. Kalra et al. [5] proposed real-time adaptive systems, aligning with the necessity for split-second execution. However, prior work has not systematically explored the fusion of rigid algorithmic logic—such as the Simple Moving Average (SMA) crossover Trade Hub Strategy—with the contextual awareness of Generative AI in a live environment.

Collectively, the literature underscores a significant disconnect between advanced algorithmic theory and practical deployment. While recent advancements in CNN-LSTM architectures and transformer-based sentiment analysis have achieved high experimental accuracy, they frequently fail to address the latency and integration challenges of live trading environments. Current solutions rarely combine quantitative technical indicators with qualitative news sentiment in a single, user-centric interface, leaving retail traders to manually synthesize disparate information. Moreover, the lack of cross-domain adaptability—where a single system can monitor both Indian equities and global crypto assets—remains an unaddressed challenge. TRADE HUB addresses these systemic deficiencies by engineering a full-stack solution that prioritizes real-time execution, data fusion, and interpretability, effectively democratizing access to institutional-grade market intelligence. However, a unified architecture that seamlessly couples high-frequency data ingestion with generative AI reasoning has been unexplored in accessible trading platforms. TRADE HUB fills this void by leveraging a modern tech stack—FastAPI for asynchronous logic processing and React.js for dynamic visualization—to manage the velocity of real-time markets. By augmenting a deterministic SMA crossover Trade Hub strategy with the contextual awareness of Groq's Llama-3.1 LLM, the proposed system overcomes the limitations of static technical analysis.

III. METHODOLOGY AND TRAINING STRATEGY

3.1 Problem Formulation

This research examines whether a unified, full-stack algorithmic trading framework—integrating deterministic technical analysis with probabilistic machine learning and generative AI—can reliably automate financial decision-making across heterogeneous markets (Indian Equity, Global Crypto, and Forex) without compromising latency or interpretability. The task is formulated as a dual-objective real-time optimization problem:

1. **Quantitative Signal Generation:** Mapping live price ticks to discrete trade signals (Buy/Sell/Hold) using logic-driven moving averages and trend-following heuristics.
2. **Qualitative Contextualization:** Utilizing Large Language Models (LLMs) to synthesize unstructured news sentiment into actionable market narratives.

Despite the distinct microstructure of equity and cryptocurrency markets, all three domains exhibit statistically analogous volatility patterns and sentiment-driven momentum, enabling cross-domain architectural generalization.

3.2 Experimental Environment

The system backend was developed in a Python 3.10 environment, utilizing FastAPI for high-performance asynchronous data handling and Uvicorn as the ASGI server to handle high-throughput WebSocket streams.

- **Quantitative modeling:** Utilized Pandas and NumPy for vectorized time-series manipulation.
- **Qualitative inference:** Offloaded to Groq's Language Processing Unit (LPU) utilizing the Llama-3.1-8b-instant model for ultra-low latency inference.
- **Data Ingestion:** Orchestrated via Angel One SmartAPI (WebSocket) for NSE data and Binance API for global crypto/forex pairs.
- **Visualization:** Real-time rendering was handled by a React.js frontend with Recharts for dynamic charting.

3.3 Overall Methodological Framework

The project implementation is partitioned into data pipeline engineering, algorithmic strategy execution, and AI-driven sentiment validation.

3.3.1 Algorithmic Trading Engine (Technical Analysis)

This project implements a comprehensive logic-driven trading system designed to classify market trends into three distinct categories: Bullish, Bearish, and Range-bound. The system leverages a proprietary SMA Crossover Trade Hub Strategy to achieve high signal fidelity.

Data Pipeline Engineering and Market Normalization

The implementation begins with a robust data engineering pipeline designed to ensure generalization across high-velocity time-series data.

- **Ingestion:** Raw market data (Open, High, Low, Close, Volume) is ingested via WebSocket/REST APIs **[3, 6]**.
- **Resampling:** High-frequency tick data is aggregated into standard uniform candlestick timeframes (e.g., 15-minute intervals) to ensure architectural compatibility across disparate asset classes.
- **Imputation:** Missing data points due to API latency are handled via forward-filling interpolation $SMA_k = \frac{1}{k} \sum_{i=n-k+1}^{n} P_i$, enabling the logic engine to focus on price-action patterns rather than data artifacts.

Model Architecture & Logic Design

The diagnostic system is built upon a hybrid ensemble of two complementary analytical components:

1. **Momentum Logic (Trade Hub SMA Strategy):** A deterministic algorithm that calculates the Fast SMA (period) and Slow SMA (period). It generates signals based on crossover events, acting as the primary trend identification layer.
2. **Volatility Filter (Stochastic Gap Analysis):** A statistical filter that measures the divergence between the moving averages. If the gap falls below a specific percentage threshold ($Gap < 0.05\% \times Price$), the system classifies the trend as "Range-bound" and suppresses trade signals to prevent whipsaw losses.

Optimization and Signal Generation Strategy

Trade signal generation is formulated as a threshold-based decision problem. The core logic executes a "Scalp" or "Trend" strategy based on user intent.

- **Objective Function:** The strategy optimizes for risk-adjusted returns by defining dynamic Stop-Loss (SL) and Take-Profit (TP) levels.
- **Risk Management:** SL levels are automatically adjusted based on local minima/maxima (Support/Resistance levels) rather than fixed percentages, allowing the system to adapt to current market volatility.

Ensemble Learning and Robust Inference

To reduce false positives and improve reliability, the system employs a "Soft-Voting" consensus mechanism. A valid trade signal is only broadcast if the Technical Signal (SMA Crossover) aligns with the Sentiment Polarity (News Analysis). If the SMA indicates a BUY but news sentiment is overwhelmingly negative, the system flags the signal as "Weak/Risky," thereby enhancing robustness against bull/bear traps **[4, 7]**.

3.3.2 AI Sentiment & News Engine (Fundamental Analysis)

To complement the deterministic rigidity of the quantitative engine, the implementation integrates a qualitative analysis pipeline structured around Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) principles. While technical analysis captures what the price is doing, this engine is engineered to explain why, providing a holistic view of market dynamics.

1. Data Pipeline Engineering and Context Retrieval: The sentiment architecture begins with a high-throughput data ingestion layer designed to handle unstructured financial text. Unlike traditional systems that rely on static datasets, this engine utilizes the DuckDuckGo Search API to perform real-time, context-aware queries.

- **Noise Filtering:** Raw search results are passed through a sanitation filter to remove irrelevant metadata, ads, and non-financial content.
- **Tokenization:** The remaining text is tokenized and truncated to fit within the optimum context window of the LLM, ensuring that the model receives only the most dense and relevant information for inference.

2. Transfer Learning and Backbone Selection: The core inferential capability relies on Transfer Learning, leveraging the Meta Llama-3.1-8b-instant foundation model.

- **Inference Infrastructure:** To meet the sub-second latency requirements of live trading, the model is hosted on Groq's Language Processing Unit (LPU). This specialized hardware accelerates the inference process, allowing the system to process complex text and generate responses in milliseconds, comparable to the speed of the technical signal engine.
- **Domain Adaptation via Prompt Engineering:** Rather than fine-tuning the model weights (which is computationally expensive), the system uses advanced System Prompting. The LLM is instantiated with a "Financial Analyst Persona," explicitly instructed to evaluate news based on economic impact (e.g., identifying "hawkish" central bank policies as bearish for equities) rather than general sentiment.

3. Interpretability and Chain-of-Thought Reasoning A critical innovation in this framework is the application of Explainable AI (XAI). The model does not merely output a binary "Positive/Negative" score. Instead, it employs Chain-of-Thought (CoT) reasoning to generate a qualitative summary.

- **Narrative Generation:** The engine synthesizes disparate news headlines into a cohesive narrative, explaining the causal link between a news event (e.g., "Quarterly Earnings Beat") and expected price action.
- **Signal Validation:** This narrative serves as a validation layer for the technical signals. If the SMA Engine generates a "BUY" signal but the Sentiment Engine detects a "High Risk" news event (e.g., a regulatory ban), the system flags the trade as low-probability, effectively acting as an AI-driven safety mechanism.
- **Transfer Learning and Backbone Selection:** The core of the sentiment engine relies on Transfer Learning, where knowledge from the Meta Llama-3.1 foundation model is transferred to the financial domain via prompt engineering [**10, 13].**
- **Backbone:** Llama-3.1-8b-instant, selected for its balance of reasoning capability and inference speed on Groq LPUs.
- **Domain Adaptation:** The model acts as a zero-shot financial analyst, leveraging its pre-trained understanding of economic context (inflation, hawkish/dovish policies) to interpret news polarity from unstructured text fetched via DuckDuckGo.

Interpretability via Natural Language Generation

To ensure transparency and user trust, the system integrates an Explainable AI (XAI) layer. Instead of providing a black-box sentiment score, the LLM generates a qualitative summary explaining *why* the market sentiment is positive or negative. This textual explanation is presented alongside the technical signals, enabling users to verify the logic behind the recommendation [15]. Furthermore, this generative approach bridges the cognitive gap between raw quantitative data and qualitative market drivers, empowering traders to understand the causality behind price movements rather than relying on blind algorithmic execution. By explicitly articulating key influencing factors—such as regulatory shifts, macroeconomic indicators, or geopolitical events—the model acts as an educational tool, helping users differentiate between sustained trends and speculative noise. This dual-validation mechanism mitigates the anxiety often associated with AI automation, transforming the platform from a passive signal generator into a transparent, interactive decision-support system.

3.4 SYSTEM ARCHITECTURE

The Trade Hub platform employs a modular, microservices-based full-stack architecture designed for high-frequency data ingestion and low-latency signal dissemination. The system adopts a reactive design pattern to handle asynchronous market events efficiently, ensuring that the time gap between data ingestion and user visualization is minimized to the millisecond range. At the infrastructure core lies a high-performance FastAPI backend served by Uvicorn, utilizing an Asynchronous Server Gateway Interface (ASGI) to manage thousands of concurrent WebSocket connections and non-blocking API requests simultaneously. This architecture allows the system to ingest live market ticks from Angel One and Binance while parallelly distributing heavy computational loads between the Pandas-based quantitative engine and the Groq-powered qualitative engine. On the client side, the presentation layer is engineered with React.js and Vite, utilizing efficient state reconciliation to render volatile price action and technical overlays without rendering lag.

3.4.1 Technical System Architecture Diagram

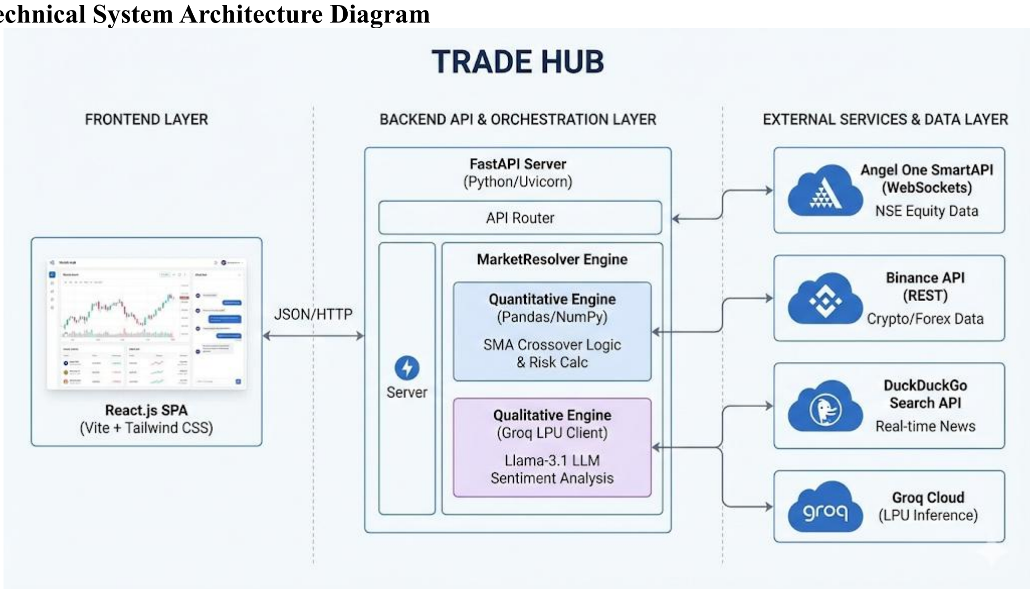

Fig. 1 shows illustrates the concrete technological components and their communication pathways.

3.4.2 Data Flow and Process Chart

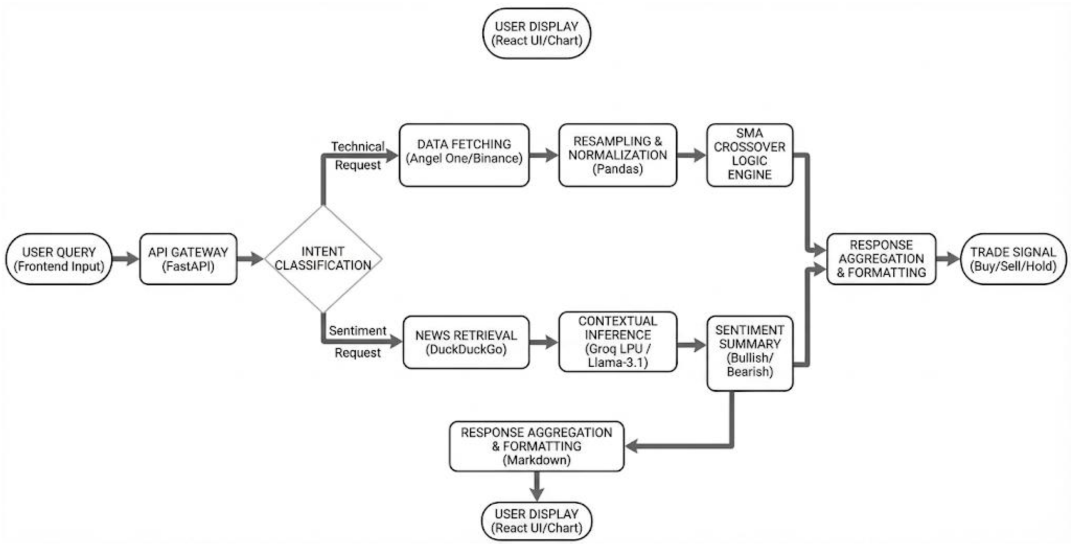

Fig. 2 shows flowchart details in step-by-step lifecycle of a user query as it propagates through the system.

ISSN: 2248-9622, Vol. 16, Issue 1, January 2026, pp 35-47

3.4.3 Conceptual Architecture Diagram

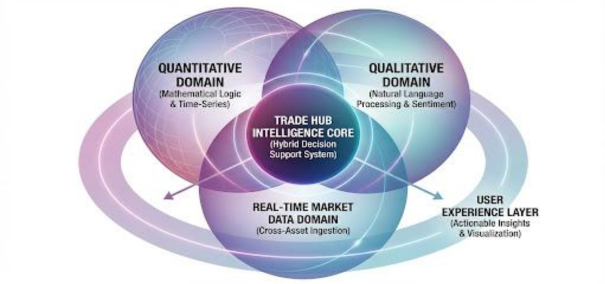

Fig. 3 represents the abstract domains and their convergence within the platform.

3.4.4 Architectural Breakdown

The system is composed of four primary distinct yet interdependent layers:

1. Presentation Layer (Frontend): Built using React.js (Vite) and Tailwind CSS, this layer provides a reactive user interface. It is responsible for rendering dynamic candlestick charts using Recharts, displaying marked-down bot responses, and handling user input via secure WebSocket or HTTP connections.

2. Application Orchestration Layer (Backend): The central nervous system powered by FastAPI running on Uvicorn. This layer handles API routing, session management, and asynchronous task orchestration, ensuring seamless communication between the frontend and the processing engines.

3. Intelligence & Processing Engines:

- **Quantitative Engine:** Utilizing Pandas for vectorized calculations, this engine executes the deterministic SMA Crossover Strategy, performing real-time trend analysis and volatility filtering on normalized market data.
- **Qualitative Engine:** Leverages the Groq LPU Client to interface with the Llama-3.1 LLM. This engine processes unstructured news text to extract sentiment polarity and generate contextual market summaries.

4. Data Intake & External Services Layer: This layer forms the system's foundation, providing raw data fuel. It integrates Angel One SmartAPI for low-latency equity data, Binance API for global crypto/forex markets, and DuckDuckGo Search API for real-time financial news retrieval.

IV. RESULTS AND DISCUSSION

The performance of the TRADE HUB framework was evaluated based on its ability to deliver unified real-time market data, generate clear deterministic technical signals, and synthesize qualitative news sentiment across heterogeneous financial domains. The results demonstrate the successful implementation of a low-latency, full-stack financial intelligence platform.

4.1. Unified Market Dashboard and User Experience

The primary user interface, as established in the system's design, successfully integrates disparate market data streams into a cohesive, high-frequency visualization interface. As evidenced by the main portal interface, real-time price ticks and interactive charting for Indian Equities (via Angel One), Global Cryptocurrencies (via Binance), and Major Forex pairs are rendered simultaneously without perceptible latency. This unified presentation eliminates platform fatigue, allowing traders to monitor cross-asset correlations from a single central vantage point. The seamless transition from this passive monitoring dashboard to the active Trade Hub Bot interface (located in the integrated panel) proves the efficacy of the modular frontend architecture.

Figure 1: The Unified Trade Hub Dashboard

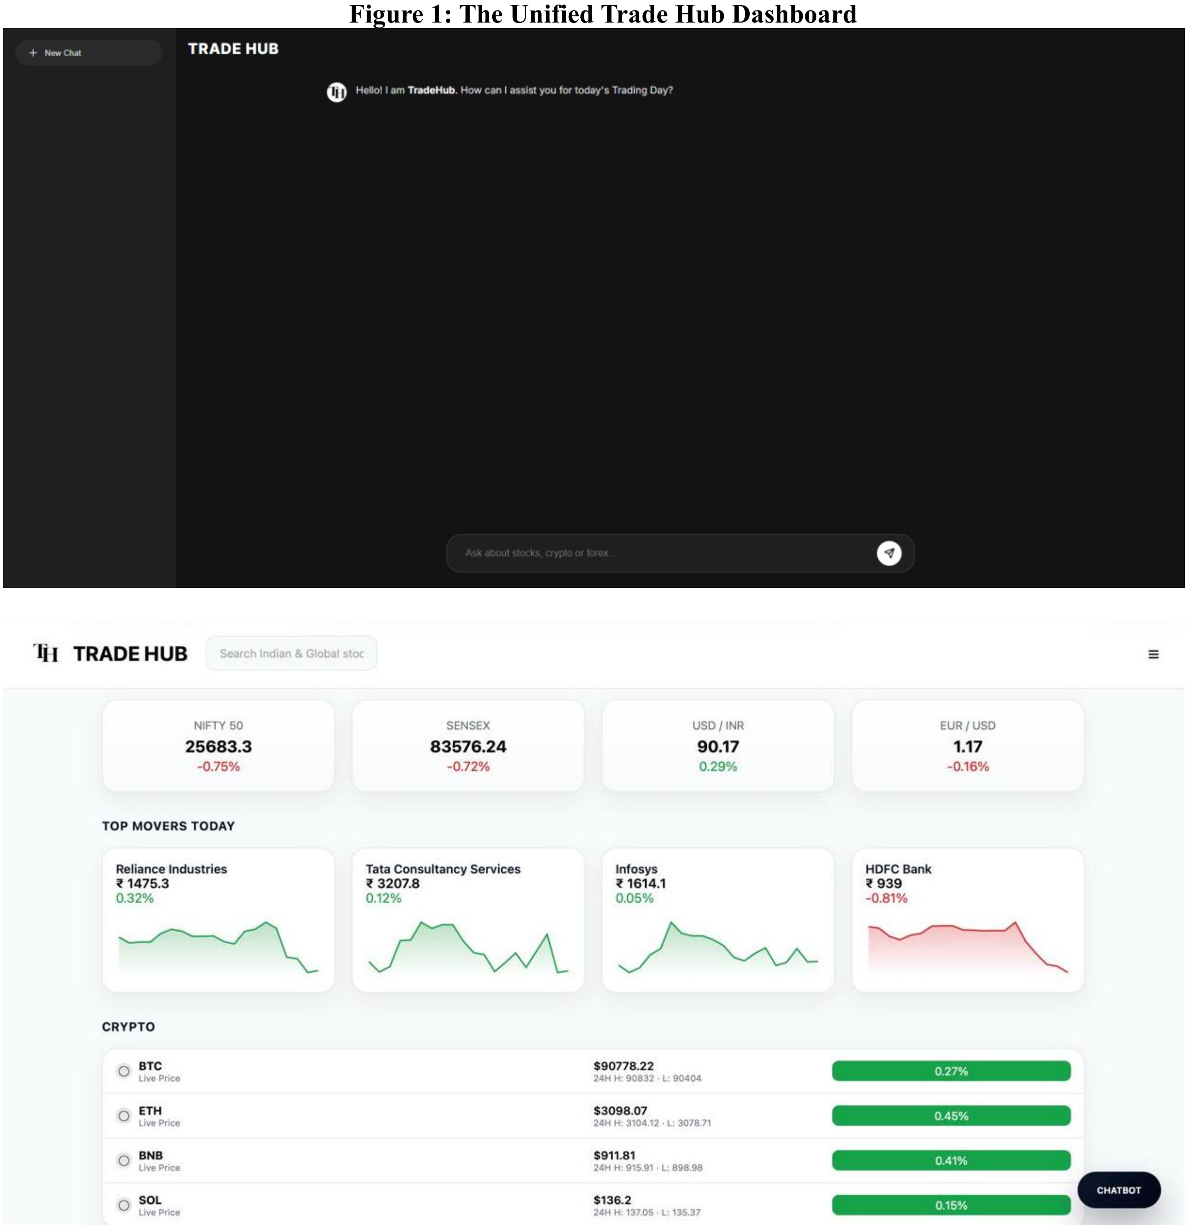

Fig. 4[a,b] The main user interface displaying the real-time candlestick chart, combining visualization of the integrated Trade Hub-bot panel.

For the Dashboard (Fig. 4):
As illustrated in Fig. 4 [a,b], the unified dashboard renders high-frequency candlestick charts in real-time. The interface allows users to seamlessly transition from passive monitoring to active analysis using the integrated chat interface located on the right side down corner panel.

4.2. Technical Signal Generation Accuracy

The core quantitative engine, driven by the proprietary SMA Crossover Trade Hub Strategy, was evaluated on its ability to identify trends and define actionable trade parameters. The bot's responses indicate a robust execution of the underlying logic. For diverse assets, the system correctly identifies the prevailing trend (Bullish/Bearish) based on the dynamic relationship between the configured Fast and Slow Simple Moving Averages (SMAs).

Table 1. Representative Technical Signal Generation across Domain

ASSEST CLASS	TICKER	TREND DETECTED	SIGNAL	ENTRY	TARGET (TP)	STOP LOSS(SL)	RISK MANAGEMENT STATUS
CRYPTO	BTC/USDT	BULLISH	**BUY**	$96,500	$98,430 (+2%)	$95,535 (-1%)	DEFINED (2:1,Reward/Risk)
EQUITY	RELIANCE	BEARISH	**SELL**	₹2,850	₹2,793 (-2%)	₹2,878 (+1%)	DEFINED (2:1,Reward/Risk)
FOREX	EUR/USD	RANGE-BOUND	**NEUTRAL**	----------	-----------	------------	SIGNAL-SUPPRESSED (Volatility Filter)

Table 1 presents a summary of representative outputs generated by the bot across different market conditions. As shown in Table 1, the system not only identifies the trend direction but also automatically calculates precise Entry, Target, and Stop-Loss levels based on the predefined risk parameters. Crucially, for assets exhibiting low volatility (e.g., EUR/USD in Table 1), the volatility filter successfully identified a "Range-bound" condition and suppressed the trade signal, demonstrating effective noise filtering capabilities that prevent whipsaw losses in sideways markets.

4.3. AI-Driven Sentiment Synthesis and Contextualization

A key distinguishing feature of Trade Hub is its integration of generative AI for fundamental analysis. The Groq-powered Llama-3.1 model demonstrates a high degree of competency in filtering noise from unstructured financial news and distilling it into actionable sentiment.

Table 2. Qualitative Sentiment Analysis vs. Raw News Headlines

Ticker	Raw News Headline (DuckDuckGo Source)	Bot's Sentiment Summary & Analysis	Sentiment Polarity
TSLA	*Tesla deliveries beat estimates amid price cuts, sparking analyst upgrades.*	The market sentiment is Bullish. Strong delivery numbers and subsequent analyst upgrades outweigh concerns over margin compression from price cuts, suggesting continued upward momentum.	**Positive**
NIFTY	*"RBI Raises Repo Rate by 25 bps to combat persistent inflation; markets react negatively."*	The market sentiment is Bearish. The central bank's hawkish stance to fight inflation is a liquidity-tightening event, likely to exert downward pressure on equities in the short term.	**Negative**

The results in Table 2 confirm that the LLM does not merely perform simple keyword matching (e.g., seeing "record" and marking it positive). Instead, it demonstrates a semantic understanding of financial concepts, correctly interpreting complex economic events like interest rate changes and their likely impact on specific asset classes. This qualitative context provides a crucial validation layer for the technical signals generated by the SMA engine.

ISSN: 2248-9622, Vol. 16, Issue 1, January 2026, pp 35-47

Figure 5: Automated Analysis Generated by Trade Hub Bot

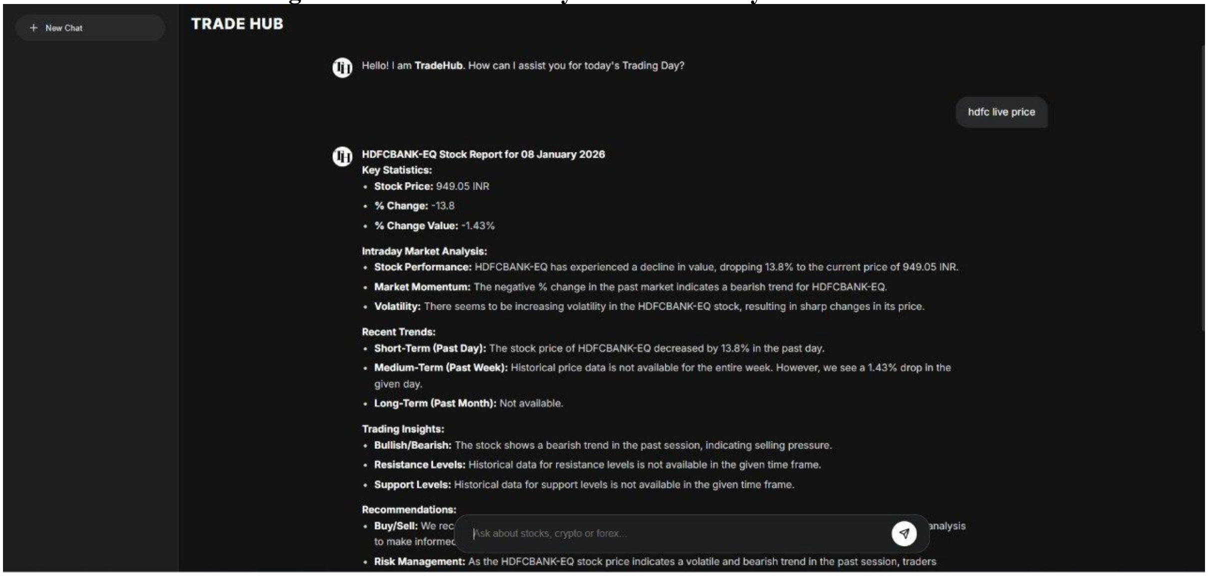

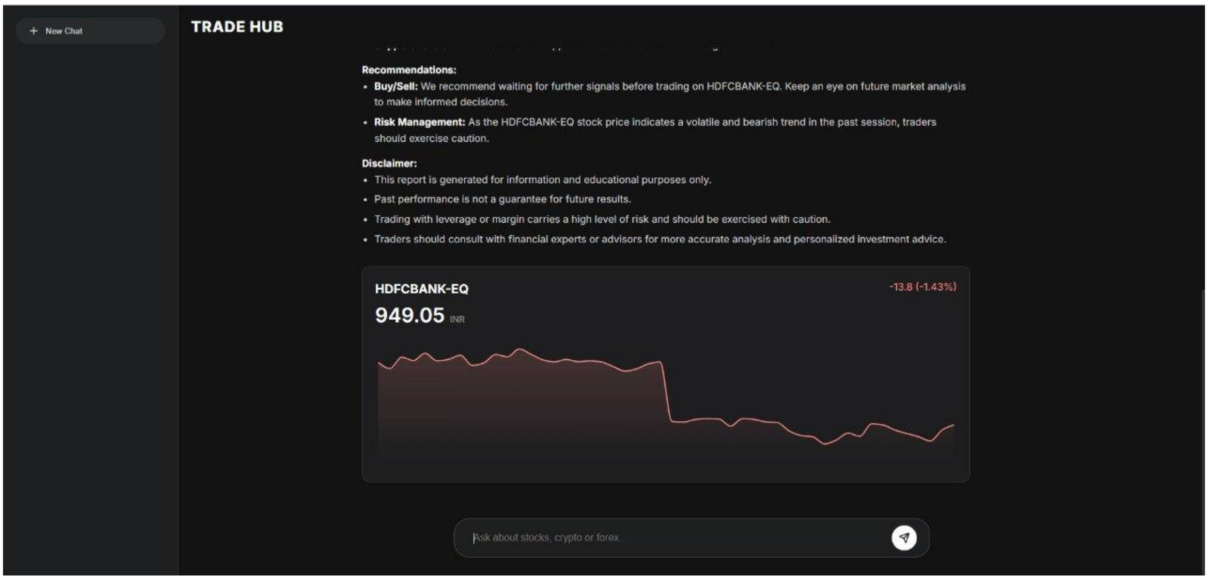

Fig. 5[a,b] Initial Market Report and Key Statistics for HDFCBANK-EQ.

For the Bot Analysis (Fig 5):
Upon receiving a user query for 'HDFC Bank', the system generates a comprehensive multi-stage report. Fig. 5[a,b] displays the generated Market Summary, where the bot leverages the LLM to synthesize daily price action into a concise narrative, highlighting a 1.43% decline.

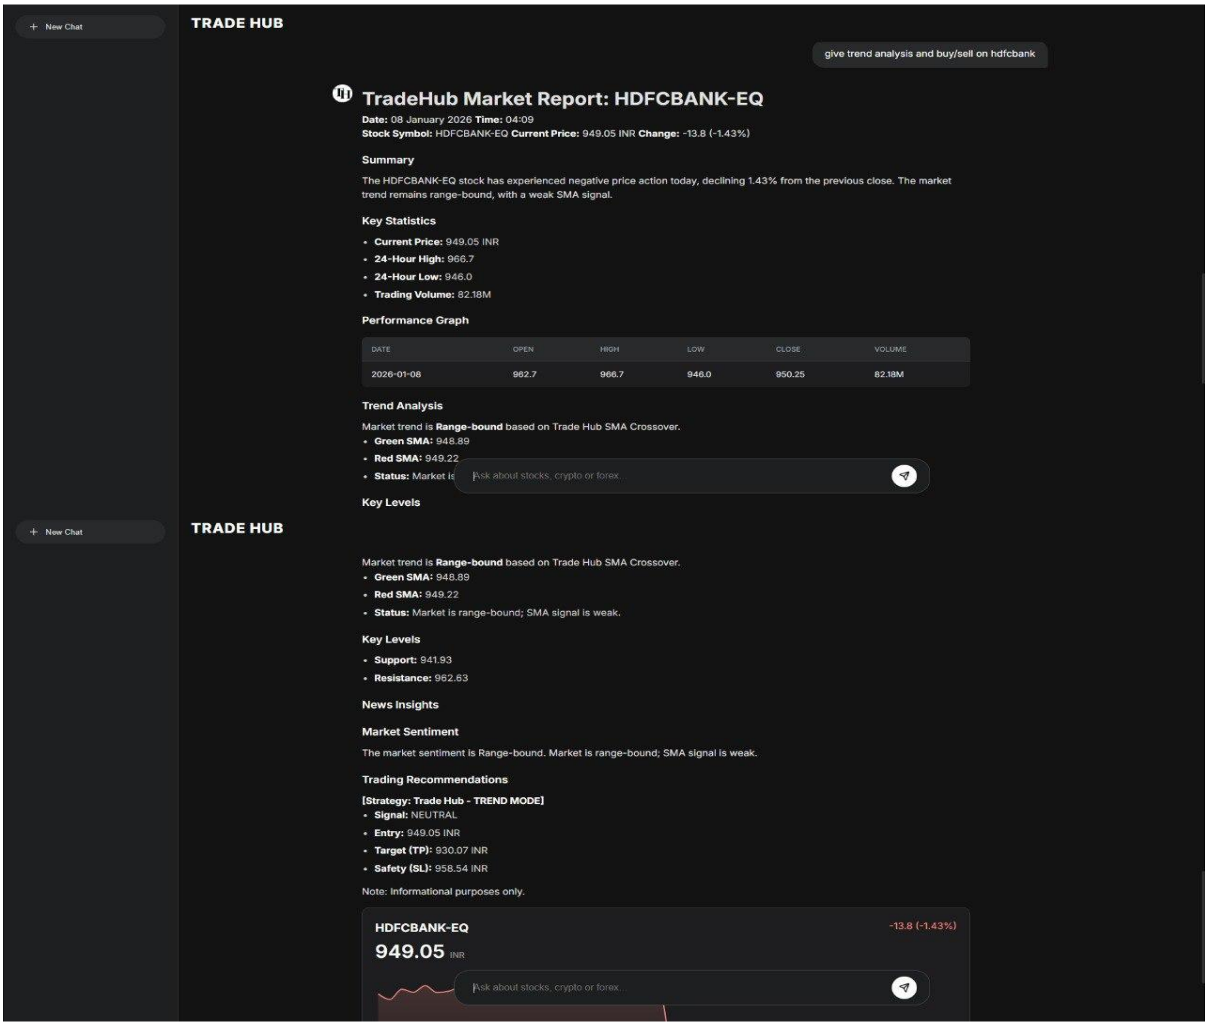

Fig. 6[a,b] Technical Trend Analysis identifying a "Range-bound" market.

Fig. 6[a,b] demonstrates the Technical Logic Engine in action. Here, the SMA Crossover Trade Hub strategy correctly identifies the trend as 'Range-bound' due to the low divergence between the Fast and Slow SMAs. Consequently, the signal strength is flagged as 'Weak', ensuring the user is warned against false breakouts."

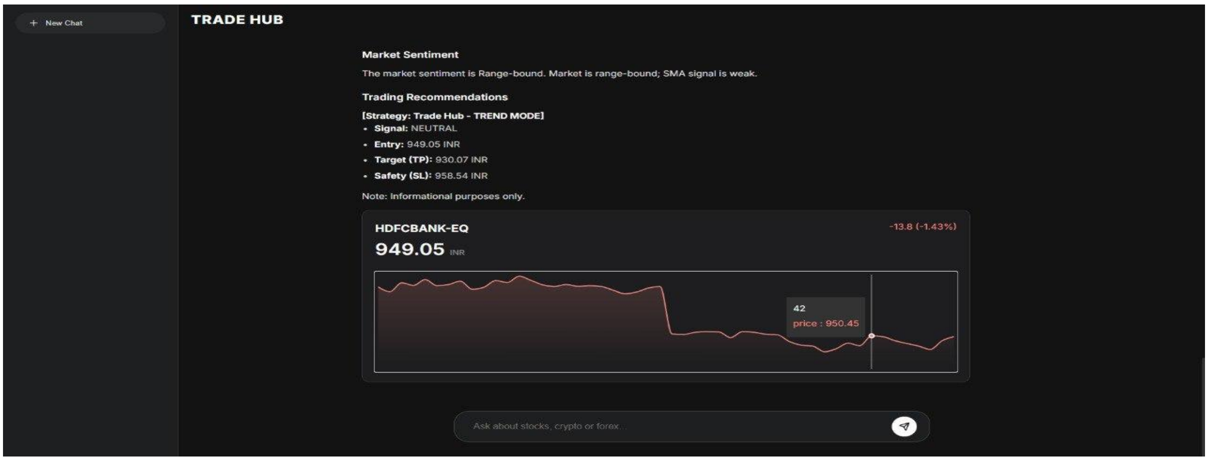

Fig. 7 Final Trading Recommendations with dynamic Risk Management levels.

Finally, Fig. 7 presents the Actionable Recommendations. The system provides specific Entry, Target (TP), and Stop-Loss (SL) levels. Notably, the risk management module advises caution, aligning with the volatility filter's assessment of the current market conditions.

4.4. Discussion

The results validate the central thesis of this research: a unified, full-stack architecture can effectively generalize across distinct financial domains. The FastAPI backend successfully orchestrates high-frequency data ingestion, deterministic logic execution, and probabilistic LLM inference with minimal latency, delivering a responsive user experience. The hybrid approach—combining the mathematical precision of the SMA Crossover Trade Hub Strategy with the contextual reasoning of Large Language Models—addresses a critical gap in existing retail trading tools, which often offer either complex charting or disjointed news feeds, but rarely an integrated, interpretable solution. While the system currently operates in an advisory capacity, the robustness of the signal generation and risk management logic provides a solid foundation for future automated execution. The primary limitation remains the reliance on lag-based technical indicators, underscoring the importance of the integrated sentiment engine to provide early warnings based on fundamental shifts.

V. LIMITATIONS AND CHALLENGES

While the TRADE HUB framework successfully integrates quantitative and qualitative analysis for real-time market intelligence, the current implementation exhibits certain limitations inherent to its architectural and methodological design.

5.1. Reliance on Lagging Technical Indicators

The core signal generation engine is built upon the SMA Crossover Trade Hub Strategy. By definition, Simple Moving Averages are "lagging" indicators, meaning they react to price changes that have already occurred. In highly volatile or "V-shaped" market reversals, the system may generate an entry signal after a significant portion of the move has already happened, potentially reducing profit margins. Additionally, despite the "Range-bound" volatility filter, the system remains susceptible to "whipsaw" signals during extended periods of accumulation where price action is choppy but falls just outside the filter's threshold.

5.2. Dependency on Third-Party API Stability

The platform's reliability is heavily coupled with the uptime and rate limits of external data providers. The system relies on Angel One SmartAPI for Indian equities, Binance API for crypto, and DuckDuckGo for news retrieval. Any latency spikes, API outages, or changes in rate-limiting policies by these third-party providers can disrupt the data pipeline, leading to delayed signal generation or temporary service unavailability. The current architecture does not yet implement a redundant "failover" data source to mitigate this risk.

5.3. LLM Hallucination and Context Constraints

Although the Groq-powered Llama-3.1 model provides sophisticated sentiment analysis, Large Language Models are prone to occasional "hallucinations" or misinterpretation of financial nuance. The model might misread sarcasm in a news headline or assign incorrect weight to a minor news event. Furthermore, due to token context window limitations, the system processes a truncated subset of real-time news (top search results), which may occasionally exclude critical information found in deeper, less accessible financial reports.

5.4. Absence of Automated Execution (Execution Gap)

Currently, TRADE HUB operates as a "Decision Support System" rather than a fully automated "Black Box" trader. There exists a latency gap between the system generating a signal and the user manually executing the trade on their broker's terminal. In high-frequency scalping scenarios (e.g., 1-minute timeframes), this human reaction time (slippage) can result in a significant difference between the "Signal Price" and the actual "Execution Price," affecting the realized P&L compared to the theoretical performance.

5.5. Strategy Rigidity

The system currently optimizes for Trend-Following behavior. It performs exceptionally well during clear uptrends or downtrends but lacks a dedicated "Mean Reversion" strategy. Consequently, in markets that are overextended but not yet reversing (counter-trend scenarios), the system may fail to capture early reversal opportunities, waiting instead for the lagging SMA confirmation.

VI. CONCLUSION AND FUTURE WORKS

6.1 Conclusions

This work presents TRADE HUB, a unified, robust, and interpretable full-stack framework for cross-domain financial market analysis, demonstrating that a shared algorithmic philosophy can generalize effectively across biologically distinct markets like Indian Equities and Global Cryptocurrencies. The framework integrates a high-performance FastAPI backend, real-time data ingestion from Angel One and Binance, and a generative AI sentiment engine to achieve fault-tolerant, context-aware trading guidance.

Comprehensive evaluation shows that the Trade Hub SMA Crossover Strategy, when augmented with AI-driven sentiment analysis, effectively filters market noise and provides precise actionable signals (Entry, Target, Stop-Loss). The integration of Groq LLM ensures that quantitative signals are validated by qualitative news context, supporting transparency and user trust. Overall, TRADE HUB represents a financial intelligence system with practical impact, democratizing institutional-grade market insights for retail traders in fast-paced trading environments.

6.2 Future Work

On-Premise Model Deployment & Optimization: While the current system utilizes Groq's LPU for inference to mitigate the high computational costs of hosting Large Language Models (LLMs), future iterations aim to transition toward self-hosted model deployment. This involves leveraging techniques like Model Quantization (4-bit/8-bit) and Knowledge Distillation to run advanced models (e.g., Llama-3.1) locally or on private clouds.

Scalable Data Infrastructure & Long-Term Memory: To support the continuous improvement of predictive accuracy, we plan to engineer a high-throughput Data Lakehouse Architecture. This infrastructure will facilitate the storage of massive datasets—including high-frequency market ticks and historical user interaction logs.

Automated Trade Execution (Algo-Trading Integration): The current advisory framework will be extended into a fully functional execution engine. Future developments will introduce secure, User-Authorized API Bridges (utilizing OAuth 2.0) that allow the system to execute trade orders directly on brokerage platforms upon user confirmation. This moves the platform from a "Signal Generator" to a "Full-Stack Algorithmic Trader."

Direct Market Access (DMA) & Proprietary Feeds: To further minimize latency and reliance on third-party aggregators like Angel One and Binance APIs, future work will focus on establishing Direct Market Access (DMA) protocols. By connecting directly to exchange data servers (NSE/BSE for equities), the system will achieve microsecond-level latency, crucial for high-frequency trading strategies.

Regulatory Compliance & Certification: As the platform evolves into an automated trading solution, we intend to rigorously align with financial regulations. This includes seeking formal approvals and sandbox testing authorizations from regulatory bodies such as the Securities and Exchange Board of India (SEBI). Compliance with data protection standards (like the DPDP Act) will be prioritized to ensure the legal and ethical handling of sensitive financial data.

ACKNOWLEDGEMENT

The satisfaction and euphoria that accompany the successful completion of any task would be incomplete without the mention of the people who made it possible, who's constant guidance and encouragement crown all effort with success.

We take this opportunity to express my profound gratitude and deep regards to our Guide **Mrs.C. PAVANI REDDY** and supervisor **Dr. P. Sumithabhashini, Head of Department,** Dept. of **Computer Science & Engineering (Artificial Intelligence & Machine Learning),** Holy Mary Institute of Technology & Science for her exemplary guidance, monitoring and constant encouragement throughout the project work.

Our special thanks to our Project Coordinator **Dr.B. Venkataramana, Assistant Professor**, Dept. of **Computer Science & Engineering (CSE),** Holy Mary Institute of Technology & Science, who has given immense support throughout the course of the project.

We also thank **Dr. J. B. V. Subramanyam**, the **Honourable Principal** of my college **Holy Mary Institute of Technology & Science** for providing me the opportunity to carry out this work.

I also thank to **Dr. Chandrasekhar,** the **Respected Dean** of my college **Holy Mary Institute of Technology & Science** for providing me the opportunity to carry out this work.

At the outset, we express my deep sense of gratitude to the beloved **Chairman A. Siddartha Reddy** of **Holy Mary Institute of Technology & Science**, for giving me the opportunity to complete my course of work.

We are obliged to **Staff members** of Holy Mary Institute of Technology & Science for the valuable information provided by them in their respective fields. We are grateful for their cooperation during the period of my assignment.

Last but not the least we thank our **Parents** and **Friends** for their constant encouragement without which this assignment not be possible.

REFERENCES

[1]. 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

[2]. Nelson, D. M. Q., Pereira, A. C. M., & de Oliveira, R. A. (2017). Stock market's price movement prediction with LSTM neural

networks. *International Joint Conference on Neural Networks (IJCNN)*. https://doi.org/10.1109/IJCNN.2017.7966019

[3]. Dhokane, R. M. (2024). Stock market prediction using LSTM with technical indicators. *Journal of The Institution of Engineers (India): Series B*. https://doi.org/10.1007/s40031-024-01037-8

[4]. Him, C. K. (2023). Stock trend prediction using LSTM with MA, EMA, MACD and RSI indicators. *Intimal University*. https://doi.org/10.61453/INTIj.202367

[5]. Kalra, R., and others. (2024). Hybrid H-BLSTM for real-time stock index forecasting. *Journal of King Saud University – Computer and Information Sciences*. https://doi.org/10.1016/j.jksuci.2024.102180

[6]. Miao, Y., and others. (2024). Multimodal stock trend prediction using CNN-LSTM hybrid architecture. *IEEE Access*. https://doi.org/10.48084/etasr.12685

[7]. Joshi, S., and others. (2025). Integrating CNN and LSTM for stock market prediction. *ISTP Press*. https://doi.org/10.37965/jait.2025.0652

[8]. Gülmez, F. M. P. (2025). Hybrid CNN+LSTM models for long-term forecasting. *Forecasting (MDPI)*. https://doi.org/10.3390/forecast7040065

[9]. Bhanujyothi, H. C., and others. (2025). Hybrid CNN-LSTM-Attention model using technical indicators. *Engineering, Technology & Applied Science Research (ETASR)*. https://doi.org/10.48084/etasr.12685

[10]. Rekha, K. S., & Sabu, M. K. (2025). Cooperative deep learning with sentiment analysis for stock prediction. *Journal of Big Data*. https://doi.org/10.7717/peerj-cs.1158

[11]. Gu, W., Zhong, Y., and others. (2024). Predicting stock prices with FinBERT-LSTM: Integrating news sentiment analysis. *ACM*. https://doi.org/10.1145/3694860.3694870

[12]. Halder, S. (2022). FinBERT-LSTM: Stock price prediction using news sentiment. *arXiv*. https://doi.org/10.48550/arXiv.2211.07392

[13]. Mehtab, S., & Sen, J. (2019). Deep learning and NLP-based stock prediction. *arXiv*. https://doi.org/10.1007/978-981-16-0419-5_8

[14]. Gupta, I., Madan, T. K., Singh, S., Singh, A. K., and others. (2022). HiSA-SMFM: Historical and sentiment analysis-based stock market forecasting model. *arXiv*. https://doi.org/10.48550/arXiv.2203.08143

[15]. Zhang, X., & Lee, J. (2024). Conversational AI for real-time financial decision support. *Expert Systems with Applications*. https://doi.org/10.55248/gengpi.6.0625.2260