



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 4, April 2026

A Multi-Agent Machine Learning Approach to Automated Trading using Sentiment and Technical Analysis

Shadab Shamsheer Alam

Department of Data Science, S.I.E.S. College of Arts, Science & Commerce (Autonomous), Sion (W), Mumbai,
Maharashtra, India¹

ABSTRACT: Cryptocurrency markets are characterized by extreme price volatility and continuous 24/7 operation, creating significant challenges for human traders attempting to make consistent and informed decisions. This paper proposes a multi-agent artificial intelligence framework for automated Bitcoin trading that simultaneously leverages three independent analytical perspectives. The first agent employs an XGBoost gradient boosting classifier and a Long Short-Term Memory (LSTM) neural network to process over 200 engineered market indicators, including RSI, MACD, and Bollinger Bands, for short-term directional price forecasting. The second agent focuses on candlestick and chart pattern identification through a hybrid approach combining a rule-based detection engine covering more than 30 established formations with a Transformer encoder network, achieving a classification accuracy of 47%. The third agent assesses prevailing market sentiment by analyzing financial news headlines via the Finnhub API and applying the VADER lexicon-based scoring tool. Each agent independently generates a directional signal accompanied by a confidence score. These signals are consolidated by a central Brain Decision Engine through an adaptive weighted voting mechanism. The Brain dynamically adjusts each agent's contribution based on historical prediction accuracy and incorporates a progressive neural blending component that refines weight allocation as trade history accumulates. A multi-layered risk management framework provides ATR-based stop-losses, portfolio drawdown limits, and confidence-threshold filtering. A Flask and WebSocket-powered web dashboard delivers real-time visualization of system activity. Paper trading evaluation with a virtual \$10,000 capital base confirms that multi-agent consensus correlates positively with trade quality, and the adaptive weighting approach shows measurable improvement in signal reliability across different market regimes.

KEYWORDS: Multi-agent systems, Bitcoin trading automation, XGBoost, LSTM, Transformer encoder, VADER sentiment analysis, candlestick pattern recognition, ensemble decision-making, adaptive weighted voting, risk management.

I. INTRODUCTION

The global cryptocurrency market has undergone remarkable growth over the past decade, with Bitcoin consistently ranking as the most actively traded digital asset by market capitalization and daily volume. In contrast to traditional stock exchanges, which operate within predefined trading windows, cryptocurrency platforms remain active at all times, creating a continuous and unforgiving environment for market participants. Price behavior in this domain is shaped by a diverse range of drivers, spanning quantitative chart-based signals, macroeconomic developments, geopolitical events, regulatory announcements, and the collective sentiment expressed across social media platforms and financial news channels. The convergence of these factors renders consistent manual trading exceptionally difficult.

Conventional algorithmic approaches have sought to automate trading decisions by parsing historical price data through established mathematical indicators such as the Relative Strength Index (RSI), Moving Average Convergence Divergence (MACD), and Bollinger Bands. Although these tools are effective at identifying momentum shifts, trend continuations, and extreme market conditions, they are fundamentally limited to information contained within price and volume data alone. A purely technical system will fail to anticipate sudden market dislocations triggered by regulatory news or high-profile social media commentary. Conversely, a sentiment-only

system may generate signals without awareness of established support or resistance boundaries or deteriorating momentum conditions.

This gap motivates a multi-agent design philosophy, wherein distinct analytical modules operate independently across different data domains and their outputs are subsequently integrated by a coordinating layer. The present research implements three such agents: a Technical Analysis Agent employing XGBoost and LSTM models trained on over 200 derived market features; a Pattern Recognition Agent combining deterministic rule-based candlestick detection with a Transformer encoder for learned sequential classification; and a Sentiment Analysis Agent that processes real-time financial news through the VADER scoring framework.

A central Brain Decision Engine aggregates signals from all three agents using a weighted voting scheme with adaptive capabilities. The Brain operates initially in a deterministic mode with predefined weights, transitioning progressively toward a neural blending mode that learns optimal agent contributions from completed trade outcomes. This design ensures system stability during the early operational phase while permitting data-driven refinement as trading history accumulates. Risk control is enforced through ATR-calibrated stop-loss placement, reward-to-risk constrained take-profit targets, portfolio-level drawdown monitoring, and a minimum confidence threshold for trade execution.

1.1 Research Objectives

- Develop and rigorously benchmark three ML models—XGBoost, LSTM, and Transformer encoder—for directional Bitcoin price prediction across three market classes.
- Implement a multi-agent coordination framework wherein independent analytical agents process distinct data modalities and forward outputs to a centralized Brain Decision Engine.
- Evaluate ensemble signal quality using a multi-metric framework assessing agent agreement levels, confidence ranges, and trade execution outcomes.
- Identify regime-specific agent strengths and demonstrate adaptive weight allocation that evolves from live trade performance.
- Operationalise a Flask and WebSocket-powered real-time dashboard supporting interpretable visualization of system activity and trade rationale.

1.2 Paper Organisation

The remainder of this paper is organised as follows. Section 2 surveys relevant prior literature. Section 3 describes the dataset and formalises the problem statement. Section 4 details the methodology including all data preprocessing and model architectures. Section 5 presents the system architecture and implementation design. Section 6 reports experimental results and comparative analysis. Section 7 provides a critical discussion of findings. Section 8 concludes the paper and outlines directions for future research.

II. LITERATURE REVIEW

Research into machine learning-based financial forecasting spans multiple disciplines, drawing from time-series analysis, natural language processing, and multi-agent systems. This section reviews key contributions across four areas relevant to the proposed framework.

2.1 Machine Learning for Technical Analysis

Supervised learning methods applied to price-derived features have demonstrated consistent improvements over random baselines in market prediction tasks. Patel et al. [6] explored a preprocessing strategy that converted raw indicator values into trend-directional representations before classification, finding that this transformation substantially improved performance across a range of classifiers including support vector machines, artificial neural networks, random forest, and naive Bayes when applied to equity index prediction. The XGBoost algorithm, introduced by Chen and Guestrin [1], extended gradient-boosted trees with explicit regularization, parallel computation, and sparse-aware splitting, establishing itself as a dominant approach in structured data competitions including financial forecasting benchmarks. McNally, Roche, and Caton [7] provided one of the earliest direct comparisons of standard recurrent architectures against LSTM models for Bitcoin price forecasting, concluding that LSTM networks captured longer temporal dependencies more reliably. Hochreiter and Schmidhuber [2] originally

proposed the LSTM architecture to resolve the vanishing gradient problem in sequential modeling, and it has since become a foundational tool for financial time-series tasks.

2.2 Sentiment Analysis in Financial Markets

The relationship between collective expressed opinion and subsequent price behavior has attracted sustained academic interest. Bollen, Mao, and Zeng [5] demonstrated that Twitter-derived mood dimensions, particularly measures of social calmness, carried statistically significant predictive information for Dow Jones Industrial Average movements when integrated alongside traditional economic indicators. Domain adaptation of pre-trained language models for financial text was advanced by Araci [8], whose FinBERT model achieved substantially higher accuracy than general-purpose sentiment classifiers on earnings call transcripts and financial news. For the cryptocurrency context, where communication is typically informal and sentiment-laden, Hutto and Gilbert [4] showed that VADER provides reliable polarity scoring without the computational overhead of transformer-based models, making it well-suited to real-time operational requirements.

2.3 Multi-Agent and Ensemble Approaches

Agent-based modeling in finance was explored by Lux and Marchesi [9], who showed that markets populated by heterogeneous agents following fundamentalist and chartist strategies spontaneously reproduced empirically observed properties such as volatility clustering and fat-tailed return distributions. In the machine learning literature, ensemble diversity has been consistently identified as a driver of performance gains, with diverse models compensating for each other's weaknesses and reducing generalization error. The Transformer architecture introduced by Vaswani et al. [3], built on scaled dot-product self-attention, extended beyond natural language tasks into time-series classification, where position-independent weighting offers advantages over fixed-window convolutions and recurrent bottlenecks. Jiang et al. [11] further explored reinforcement learning for portfolio management, highlighting the potential for policy-based optimization in financial decision-making.

2.4 Research Gap

The reviewed literature confirms that each of the three analytical domains—technical indicators, sentiment signals, and pattern recognition—contributes independent predictive value. However, existing implementations typically treat these domains in isolation. Fully integrated multi-agent frameworks that coordinate across fundamentally different signal sources, with adaptive weighting that evolves from trade outcomes, remain underexplored. Furthermore, existing systems rarely incorporate real-time dashboards that provide interpretable, confidence-weighted reasoning traces alongside live trading signals. This research addresses all three dimensions in an integrated manner.

III. DATASET DESCRIPTION AND PROBLEM FORMULATION

3.1 Dataset Characteristics

This study employs historical and real-time OHLCV (Open, High, Low, Close, Volume) data for the BTC/USD trading pair sourced through the yfinance library, with supplementary financial news data retrieved via the Finnhub API. The OHLCV dataset spans multi-year historical hourly candlestick data, providing a rich sequential record of price dynamics across diverse market regimes including sustained uptrends, sharp reversals, sideways consolidation, and high-volatility episodes. A comprehensive feature engineering stage derives over 200 quantitative indicators from this raw data.

Table 1: Feature Category Summary and Descriptions

Category	Feature Names	Type	Role in Prediction
Momentum Oscillators	RSI, Stochastic, Williams %R, CCI	Numeric	Identify overbought/oversold market conditions
Trend-Following	MACD histogram, EMA crossovers, ADX	Numeric	Detect directional trend strength and momentum
Volatility Measures	Bollinger Band width, ATR, Historical Vol.	Numeric	Capture price dispersion and range expansion
Volume Signals	OBV, Volume change rate, VWAP deviation	Numeric	Confirm price moves with participation data
Pattern Features	Doji, Hammer, Engulfing, Morning Star, etc.	Mixed	Encode candlestick formation signals
Sentiment Index	VADER compound score (Finnhub headlines)	Numeric	Capture real-time market narrative direction
Target Variable	Direction: Bullish / Neutral / Bearish	3-class	Forward return over 6-hour prediction horizon

A critical characteristic of the target variable construction is its three-class formulation. Labels are derived from the forward price return over a six-hour horizon and thresholded to produce an approximately balanced distribution with a random chance baseline of 33.3 percent. This balanced construction avoids the systematic bias that would arise from predicting a perpetually upward-trending asset using a binary label without return-magnitude thresholding.

3.2 Formal Problem Statement

Let $D = \{(x_i, y_i) \mid i = 1, \dots, N\}$ represent the labelled dataset of N BTC/USD hourly records, where each feature vector x_i encodes over 200 engineered market indicators and y_i in $\{\text{Bullish, Neutral, Bearish}\}$ denotes the three-class directional outcome. The learning objective is to identify a mapping f that maximises directional classification accuracy across unseen market windows while maintaining calibrated confidence estimates suitable for position-sizing decisions by the Brain Decision Engine.

IV. METHODOLOGY

The proposed system is structured as a pipeline in which independent analytical agents process distinct data modalities and forward their outputs to a centralized decision layer. The research adopts a modular architecture philosophy, wherein each analytical component is independently trained, evaluated, and integrated through a well-defined signal interface.

4.1 Data Preparation Pipeline

4.1.1 Feature Engineering

A comprehensive feature engineering stage derives over 200 quantitative indicators from raw OHLCV data, encompassing momentum oscillators such as RSI, trend-following metrics including MACD histogram and exponential moving average crossovers, volatility measures including Bollinger Band width and ATR, and volume-based signals. All feature values are standardized through z-score normalization. The scaler was fitted exclusively on training data and applied to the test set to prevent data leakage.

4.1.2 Label Construction

Ground truth labels are constructed as a three-class variable—bullish, neutral, and bearish—derived from the forward price return over a six-hour horizon. Positive returns exceeding a defined threshold are labelled bullish,

returns below the negative threshold are labelled bearish, and intermediate returns are labelled neutral. This construction yields an approximately balanced distribution with a random chance baseline of 33.3 percent.

4.1.3 Train-Test Split

An 80/20 chronological train-test split is applied, preserving temporal ordering to prevent look-ahead bias. The training partition is used exclusively for model fitting and hyperparameter optimization; the test partition is held out for final performance evaluation. No data augmentation or synthetic oversampling is applied, as the three-class label construction naturally produces a balanced distribution.

4.2 Technical Analysis Agent

This agent integrates two independently trained models. The first is an XGBoost classifier configured with 100 boosting rounds and a maximum tree depth of 6, operating on 13 selected features including RSI, MACD, Bollinger Band relative position, and volume change rate, achieving 42.2 percent test accuracy. The second is a two-layer LSTM network with 64 hidden units per layer, receiving a 20-timestep sliding window of the same 13 features as sequential input and attaining 39.7 percent test accuracy. Directional agreement between models elevates the combined confidence score, while disagreement produces a split signal with proportionally reduced confidence.

4.3 Pattern Recognition Agent

The Pattern Agent operates through two complementary mechanisms. The first is a deterministic rule engine evaluating each candlestick sequence against a library of over 30 established formations, including Doji, Hammer, Bullish and Bearish Engulfing, Morning Star, and Three White Soldiers, alongside pivot-point-based support and resistance estimation. The second is a Transformer encoder comprising two attention layers with four heads each and 128-dimensional feed-forward sublayers, processing 20-step OHLCV sequences projected into 64-dimensional embeddings with sinusoidal positional encodings, achieving standalone classification accuracy of 47.0 percent.

4.4 Sentiment Analysis Agent

Real-time Bitcoin-related news headlines are retrieved through the Finnhub API and scored using the VADER compound polarity metric. Individual scores are aggregated into a composite index from -90 to +90, classified as bullish (above +20), bearish (below -20), or neutral. The agent operates on a six-hour refresh cycle aligned with the other agents' prediction horizon, and a temporal decay function reduces the weight of older articles to ensure recency of the sentiment signal.

4.5 Brain Decision Engine

The Brain accepts the directional classification and confidence score from each agent and applies weighted voting with initial weights of 0.45 (Technical), 0.25 (Pattern), and 0.30 (Sentiment). Inter-agent agreement directly modulates the output confidence value. A minimum confidence threshold of 30 percent is enforced before any trade order is generated. A progressive neural blending mechanism scales from zero to 30 percent influence as trade history grows, and agent weights are updated after each completed trade through exponential moving averages applied to individual prediction accuracy records.

4.6 Risk Management

Position sizing is capped at 25 percent of available capital and scaled by the Brain's confidence output. Each trade receives an ATR-based stop-loss placed 1.5 times the current ATR below entry and a take-profit at twice the risk distance, enforcing a 2:1 reward-to-risk ratio. A portfolio-level drawdown monitor suspends all trading activity when cumulative losses reach 5 percent of initial capital, and a consecutive-loss cooling mechanism introduces mandatory waiting periods following adverse streaks.

V. SYSTEM ARCHITECTURE AND IMPLEMENTATION

The prediction system is structured as a four-layer pipeline. The data flow from raw market data through multi-agent analysis, ensemble decision-making, and real-time visualization is described in the following subsections.

5.1 Four-Layer Architecture

Layer 1 (Data Ingestion): Real-time and historical OHLCV data is sourced via yfinance and financial news headlines via the Finnhub API, feeding a continuous data stream to all three analytical agents. Layer 2 (Agent Processing): The Technical, Pattern, and Sentiment agents independently process their respective data modalities and emit directional signals with confidence scores on a unified five-minute cycle. Layer 3 (Brain Aggregation): The Brain Decision Engine applies adaptive weighted voting, confidence thresholding, and risk management filters before emitting a trade decision. Layer 4 (Dashboard Output): The Flask and WebSocket-powered dashboard delivers real-time visualization of agent signals, Brain decisions, open positions, and cumulative performance.

5.2 Technology Stack

Table 2: Technology Stack Summary

Component	Technology	Version	Purpose
Programming Language	Python	3.10+	Core implementation language
Data Ingestion	yfinance, Finnhub API	Latest	OHLCV and news headline retrieval
ML Algorithms	scikit-learn, XGBoost	1.4 / 2.0	Model training, preprocessing, evaluation
Deep Learning	PyTorch / TensorFlow	2.x	LSTM and Transformer encoder training
NLP / Sentiment	VADER (nlTK)	3.3+	Lexicon-based polarity scoring
Web Backend	Flask	3.x	REST API for trade signal delivery
Real-time Comms	WebSocket	---	Sub-second dashboard data push
Visualisation	Plotly, Matplotlib	5.x / 3.x	Interactive charts and performance plots

5.3 Dashboard Design

The web dashboard encapsulates the full inference pipeline within a single-page interface. On each five-minute cycle, the dashboard updates live candlestick charts, individual agent signal breakdowns with confidence meters, the Brain's consolidated decision with its reasoning trace, open position summaries, running profit and loss figures, and a time-stamped activity log. The interpretable output structure supports continuous monitoring and manual override capability.

VI. EXPERIMENTAL RESULTS AND ANALYSIS

6.1 Individual Model Performance

All models were trained on historical BTC/USD hourly data with an 80/20 chronological train-test split. The three-class classification baseline stands at 33.3 percent. Table 3 presents the performance summary. All models exceeded the random baseline by 6.4 to 13.7 percentage points, confirming that each agent extracts genuine predictive signal from its respective data domain.

Table 3: Individual Model Performance on BTC/USD Test Set

Model	Agent	Accuracy	Baseline+	Input Type
XGBoost	Technical	42.2%	+8.9 pts	13 engineered indicators
LSTM	Technical	39.7%	+6.4 pts	20-timestep sequence window
Transformer	Pattern	47.0%	+13.7 pts	20-step OHLCV sequences
VADER	Sentiment	Corr.-based	N/A	Financial news headlines

The Transformer encoder achieved the highest individual accuracy at 47.0 percent, attributable to the self-attention mechanism’s capacity to detect formations of variable temporal length. XGBoost contributed reliable directional predictions at 42.2 percent with particular strength in bullish signal identification, while the LSTM model’s 39.7 percent accuracy provided complementary sequential information.

6.2 Ensemble Signal Quality

The most significant empirical finding is that signal quality improves substantially when multiple agents converge on the same directional prediction. Table 4 presents the relationship between agent agreement level and the corresponding confidence range, position allocation, and execution decision.

Table 4: Signal Quality by Agent Agreement Level

Agreement Level	Confidence Range	Position Size	Execution Decision
All 3 aligned	65–80%	Full size	Trade executed
2 of 3 aligned	35–55%	Reduced	Trade executed with caution
1 of 3 aligned	Below 35%	None	Trade skipped
All conflicting	Below 20%	None	No action taken

When all three agents aligned, confidence ranged from 65 to 80 percent and corresponded to the most reliable trade outcomes. Partial agreement produced intermediate confidence and reduced allocations. Full disagreement suppressed all trading activity, successfully preventing execution under genuinely ambiguous market conditions.

6.3 Agent Behavior Across Market Regimes

Performance analysis revealed that each agent held a comparative advantage in specific market regimes. Table 5 summarizes these regime-specific strengths.

Table 5: Agent Strengths by Market Regime

Market Condition	Best Agent	Observation
Sustained uptrend	Technical	Momentum indicators matched bullish direction
Sharp reversal	Pattern	Reversal candles detected before price turned
News-driven spike	Sentiment	Headlines captured before price agents reacted
Sideways range	Pattern	Support/resistance zones identified accurately
High volatility	Technical	ATR and Bollinger Bands reacted to expanding range

The Brain's adaptive weight mechanism responded by progressively elevating the weight of the most accurate agent for prevailing conditions. The Technical Agent received increased influence during directional trending phases, while the Pattern Agent's weight grew during consolidation and range-bound periods.

6.4 Paper Trading Deployment

The system was deployed with \$10,000 virtual capital, executing at five-minute intervals subject to confidence thresholds and risk constraints. Table 6 presents the full operational configuration.

Table 6: Deployment Configuration and Operational Metrics

Parameter	Value
Starting capital	\$10,000 (virtual)
Tick interval	5 minutes
Stop-loss method	1.5× ATR from entry
Take-profit target	2:1 reward-risk ratio
Max. position size	25% of capital
Drawdown pause at	5% of initial capital
Decision latency	Under 2 seconds
Dashboard update	Under 1 second

The risk management layer prevented significant capital erosion during consecutive losing signal periods. The drawdown mechanism suspended operations before losses could compound, and confidence filtering reduced total trade count while improving the proportion of high-quality entries. Decision latency remained under two seconds for all observed cycles.

6.5 Key Observations

- The Transformer encoder achieved the highest individual model accuracy (47.0%), confirming that self-attention mechanisms are well-suited to variable-length pattern recognition in OHLCV sequences.
- Multi-agent consensus (all 3 aligned) produced confidence scores of 65–80% and yielded the highest-quality trade outcomes, validating the core ensemble hypothesis.
- Adaptive weight reallocation measurably improved signal reliability across changing market regimes, with the Brain correctly upweighting the best-performing agent for prevailing conditions.
- The ATR-based risk framework successfully prevented drawdown compounding; the 5% drawdown circuit breaker triggered appropriately during adverse signal streaks.
- Dashboard decision latency consistently remained below 2 seconds, confirming operational viability for live deployment.

VII. DISCUSSION

7.1 Interpretation of Results

The experimental results collectively confirm three central propositions. First, ensemble coordination across fundamentally different data domains offers a material advantage over any single analytical perspective. The Transformer pattern encoder, XGBoost technical classifier, and VADER sentiment scorer each operate on orthogonal information sources; their combination reduces the risk of systematic blind spots that would afflict a single-domain system. Second, adaptive weighting is essential for maintaining performance across market regime transitions. Static weight assignments that perform well in trending markets degrade during consolidation, and vice versa. The progressive neural blending mechanism's ability to rebalance agent contributions in response to live

performance directly addresses this limitation. Third, the multi-layered risk management framework is not merely supplementary but operationally essential.

7.2 Trading and System Implications

The regime-specific agent behavior observed in Table 5 yields actionable architectural insights. The Technical Agent's superior performance during trending phases confirms that momentum-based indicators retain predictive value under directional market conditions, supporting their continued inclusion in the feature set. The Pattern Agent's strength at reversals validates the inclusion of candlestick pattern recognition as a distinct analytical perspective rather than a redundant subset of technical analysis. The Sentiment Agent's advantage during news-driven events highlights the importance of real-time textual data processing in cryptocurrency markets, where regulatory announcements and social media events drive disproportionate price movements.

7.3 Limitations

Several limitations of the present study merit acknowledgement. The paper trading evaluation, while providing a controlled deployment environment, does not capture the full complexity of live market execution including slippage, liquidity constraints, and exchange latency. The VADER sentiment scorer, while computationally efficient and well-suited to real-time operation, does not capture domain-specific nuance in cryptocurrency-related financial text as effectively as fine-tuned transformer models such as FinBERT. Additionally, the system was evaluated on a single asset (BTC/USD); cross-asset generalization to altcoins or traditional financial instruments has not been validated.

VIII. CONCLUSION AND FUTURE DIRECTIONS

8.1 Conclusion

This paper introduced a multi-agent artificial intelligence framework for automated Bitcoin trading, integrating three independently operating analytical agents within a coordinated decision architecture. The Technical Agent combined XGBoost and LSTM models trained on over 200 market-derived features to assess price momentum and trend conditions. The Pattern Agent paired a comprehensive rule-based candlestick detection library with a Transformer encoder to recognize both established and learned sequential price formations. The Sentiment Agent applied VADER-based scoring to real-time financial news, providing a complementary qualitative dimension absent from price-only analyses.

The central Brain Decision Engine fused agent outputs through adaptive weighted voting, progressively incorporating a neural blending component as operational experience accumulated. Empirical results confirmed that inter-agent consensus correlates positively with signal quality and that the adaptive weight mechanism successfully prioritized the most accurate agent for the prevailing market condition. The multi-layered risk management framework, encompassing ATR-based stops, drawdown limits, confidence filtering, and position sizing constraints, provided robust capital protection throughout the paper trading evaluation.

8.2 Future Research Directions

1. Real-Time On-Chain Analytics Integration: Incorporating blockchain-derived metrics (active addresses, exchange inflows, miner behavior) as a fourth analytical agent to extend the framework beyond price and text data.
2. Rigorous Historical Backtesting: Developing a full backtesting framework with realistic transaction cost modeling, slippage simulation, and regime-labelled performance attribution.
3. Multi-Asset Extension: Extending the framework to multiple cryptocurrency pairs and investigating whether agent weight profiles generalize across assets or require per-asset calibration.
4. Deep Reinforcement Learning for Brain Optimization: Replacing the adaptive weighted voting mechanism with a policy-gradient reinforcement learning agent that directly optimizes risk-adjusted returns from historical trade sequences.
5. FinBERT Sentiment Upgrade: Replacing VADER with a fine-tuned FinBERT or CryptoBERT model to improve sentiment accuracy on domain-specific financial and social media text.

6. Federated Learning for Privacy Preservation: Exploring multi-exchange collaborative model training without centralising sensitive order flow data, enabling more robust generalization while maintaining exchange confidentiality.
7. Temporal Model Drift Monitoring: Implementing automated retraining pipelines with drift detection metrics to maintain prediction quality as cryptocurrency market microstructure evolves.

REFERENCES

1. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, San Francisco, CA, USA, 2016, pp. 785–794. <https://doi.org/10.1145/2939672.2939785>
2. S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
3. A. Vaswani et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998–6008.
4. C. J. Hutto and E. Gilbert, "VADER: A parsimonious rule-based model for sentiment analysis of social media text," in Proc. 8th Int. AAAI Conf. Weblogs and Social Media, Ann Arbor, MI, USA, 2014, pp. 216–225.
5. J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," *Journal of Computational Science*, vol. 2, no. 1, pp. 1–8, 2011.
6. J. Patel, S. Shah, P. Thakkar, and K. Kotecha, "Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques," *Expert Systems with Applications*, vol. 42, no. 1, pp. 259–268, 2015.
7. S. McNally, J. Roche, and S. Caton, "Predicting the price of Bitcoin using machine learning," in Proc. 26th Euromicro Int. Conf. Parallel, Distributed and Network-Based Processing, Cambridge, UK, 2018, pp. 339–343.
8. D. Araci, "FinBERT: Financial sentiment analysis with pre-trained language models," arXiv preprint arXiv:1908.10063, 2019.
9. T. Lux and M. Marchesi, "Scaling and criticality in a stochastic multi-agent model of a financial market," *Nature*, vol. 397, no. 6719, pp. 498–500, 1999.
10. K. Chen, Y. Zhou, and F. Dai, "A LSTM-based method for stock returns prediction: A case study of China stock market," in Proc. IEEE Int. Conf. Big Data, Santa Clara, CA, USA, 2015, pp. 2823–2824.
11. Z. Jiang, D. Xu, and J. Liang, "A deep reinforcement learning framework for the financial portfolio management problem," arXiv preprint arXiv:1706.10059, 2017.
12. B. Weng et al., "Predicting short-term stock prices using ensemble methods and online data sources," *Expert Systems with Applications*, vol. 112, pp. 258–273, 2018.
13. M. R. Hassan and B. Nath, "Stock market forecasting using hidden Markov model: A new approach," in Proc. 5th Int. Conf. Intelligent Systems Design and Applications, Wroclaw, Poland, 2005, pp. 192–196.

About Author

Shadab Shamsheer Alam is a postgraduate candidate in the Data Science programme at S.I.E.S. College of Arts, Science & Commerce (Autonomous), Mumbai, India. His research interests encompass machine learning, deep learning, natural language processing, financial analytics, and algorithmic trading systems. This paper presents his original research contribution in the domain of automated trading, focusing on a multi-agent machine learning framework that integrates sentiment analysis from textual data with technical indicators for improved decision-making in financial markets.

Dr. Abuzar Ansari is Professor and Head of the Department of Data Science at S.I.E.S. College of Arts, Science & Commerce (Autonomous), Mumbai. His research domain spans applied machine learning, educational data mining, and natural language processing. He serves as Principal Investigator for multiple funded data science research projects within the institution.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details