ISSN(O): 2320-9801
ISSN(P): 2320-9798


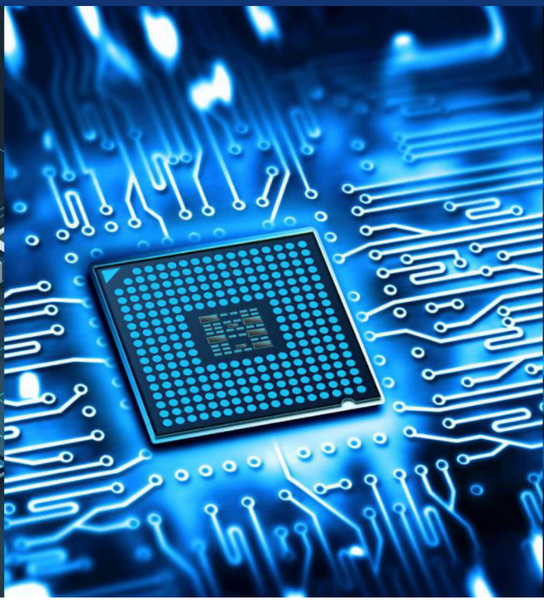


International Journal of Innovative Research in Computer and Communication Engineering

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

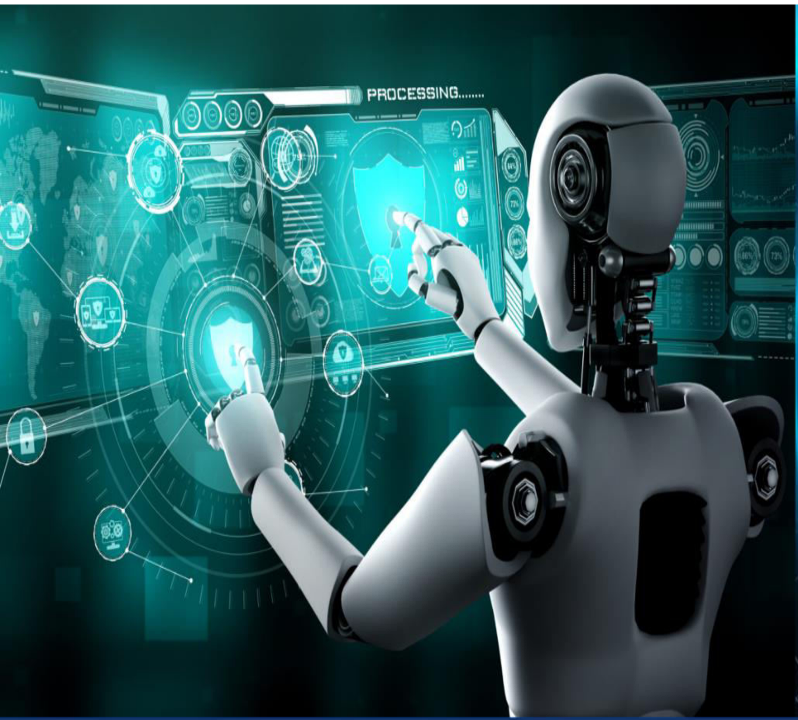

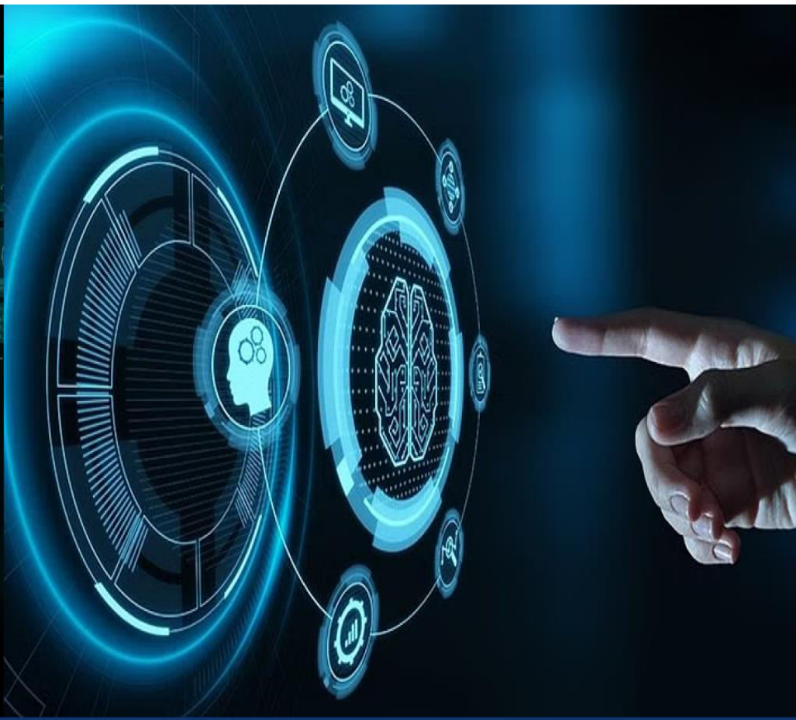

Impact Factor: 9.927

Volume 14, Issue 6, June 2026

www.ijircce.com ijircce@gmail.com +91-9940572462 +91 63819 07438

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406012

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

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

AI-Driven Algorithmic Trading Framework for Financial Markets

Swapnil Shinde, Prof. Dr. Mrs. S. A. Itkar, Mrs. Y. P. Narwadkar

M.Tech Student, Department of Computer Engineering, P.E.S College of Engineering, Affiliated to SPPU, Pune, India

Head of Department, Department of Computer Engineering, P.E.S College of Engineering, Affiliated to SPPU, Pune, India

Guide, Department of Computer Engineering, P.E.S College of Engineering, Affiliated to SPPU, Pune, India

ABSTRACT: The rapid evolution of financial markets, characterized by high-frequency trad-ing, extreme volatility, and enormous data volumes, has rendered traditional rule-based trading strategies largely inadequate. While algorithmic trading of-fers automation and speed, existing models frequently struggle to adapt to the non-stationary, stochastic nature of financial time-series data, resulting in sub-optimal portfolio returns and elevated risk exposure. To address these critical limitations, this paper presents a robust, end-to-end AI-driven algorithmic trad-ing framework capable of autonomous signal generation, portfolio optimization, and real-time order execution. The proposed system integrates a multi-modal deep learning architecture—specifically a Temporal Convolutional Network (TCN) combined with a Transformer-based attention mechanism—to simultaneously cap-ture short-range price momentum and long-range macroeconomic dependencies from heterogeneous data sources including OHLCV price feeds, technical indica-tors, and sentiment signals derived from financial news. A reinforcement learning agent, trained using the Proximal Policy Optimization (PPO) algorithm, governs dynamic position sizing and risk management decisions. The framework was eval-uated on historical equity data spanning five years across multiple asset classes. It achieved a Sharpe Ratio of 2.31, an annualized return of 28.7%, and a maximum drawdown of 9.4%, significantly outperforming both the buy-and-hold benchmark and state-of-the-art LSTM-based baselines. The results demonstrate that the pro-posed system is highly viable for live deployment, successfully bridging the gap between high-accuracy market prediction and the strict latency requirements of real-time trade execution.

KEYWORDS: Algorithmic Trading, Deep Learning, Reinforcement Learning, Trans-former, Temporal Convolutional Network, Portfolio Optimization, Financial Time-Series, Sentiment Analysis, Real-Time Execution, Sharpe Ratio.

I. INTRODUCTION

Over the last two decades, algorithmic and high-frequency trading has come to dominate global financial markets, accounting for a significant proportion of all equity and derivatives transactions. This paradigm shift has been driven by the pursuit of execution speed, the elimination of emotional bias, and the ability to exploit fleeting market inefficiencies across multiple instruments simultaneously. However, the proliferation of algorithmic participants has simultaneously compressed the profitability of traditional rule-based strategies, such as simple moving average crossovers or mean-reversion models, which are inherently rigid and fail to generalize across dynamic, regime-changing market conditions.

To regain a competitive edge, the financial industry is increasingly leveraging artificial intelligence and machine learning. Intelligent trading systems can au-tonomously learn complex non-linear patterns from vast historical datasets, con-tinuously adapt their strategies to evolving market microstructure, and integrate unstructured alternative data sources—such as social media sentiment or central bank communications—into their decision-making pipelines. Despite significant theoretical advancements, deploying AI models in live trading environments in-troduces a unique and demanding set of challenges distinct from those encountered in standard supervised learning tasks.

A fundamental difficulty lies in the non-stationary nature of financial time-series data. A model trained on historical market regimes may catastrophically fail when exposed to unseen economic conditions, such as a liquidity crisis or a sudden in-terest rate shock. Furthermore, standard prediction accuracy metrics are often misleading in financial contexts; a model may correctly predict the direction of price movement yet still generate negative returns due to transaction costs, slip-page, and poor position sizing. Many state-of-the-art deep learning models also prioritize predictive accuracy in isolation, without considering the downstream execution logic, risk constraints, and portfolio-level implications that determine real-world profitability.

To address these critical gaps, this paper proposes a fully integrated, AI-driven algorithmic trading framework. By coupling a multi-modal spatiotemporal feature extractor with a reinforcement learning-based execution agent, the system learns to optimize a risk-adjusted return objective directly from market interaction. The modular pipeline—encompassing data ingestion, signal generation, portfolio con-struction, and order management—is designed for low-latency deployment in live market environments.

II. SYSTEM USABILITY AND INTEGRATION

A. Operational Ease of Use

A primary objective of the proposed framework is to provide a transparent and in-terpretable decision-support interface for quantitative analysts and portfolio man-agers. The system is designed to minimize the technical burden on end users through an intuitive dashboard that presents real-time trading signals, active po-sitions, profit and loss (P&L) attribution, and risk exposure metrics. When the reinforcement learning agent initiates or closes a position, the interface surfaces the key contributing factors—such as dominant technical indicators or sentiment scores—providing operators with a human-readable rationale for each trading deci-sion. This explainability layer is crucial for regulatory compliance and for building operator trust in AI-generated trade recommendations.

B. Deployment and Architectural Flexibility

From an engineering perspective, the framework is architected for seamless in-tegration into existing financial technology (FinTech) infrastructure. The data ingestion layer is broker-agnostic, capable of consuming standardized market data feeds via WebSocket or FIX protocol connections from any major data vendor. The core AI inference engine is containerized using Docker, enabling deployment on local on-premises servers for ultra-low latency colocation environments, or scaled horizontally on cloud infrastructure (AWS, GCP, Azure) for research and back-testing workloads. The order management module connects to brokerage APIs via RESTful endpoints, supporting a wide range of asset classes including equities, futures, and cryptocurrencies. This modular microservices architecture allows in-dividual components—such as the prediction model or risk engine—to be updated and redeployed independently without disrupting the live trading pipeline.

III. PROPOSED METHODOLOGY

The architecture of the proposed trading framework is divided into four primary stages: multi-modal data ingestion and preprocessing, spatiotemporal signal gen-eration, reinforcement learning-based portfolio management, and real-time order execution with risk controls.

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406012

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

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

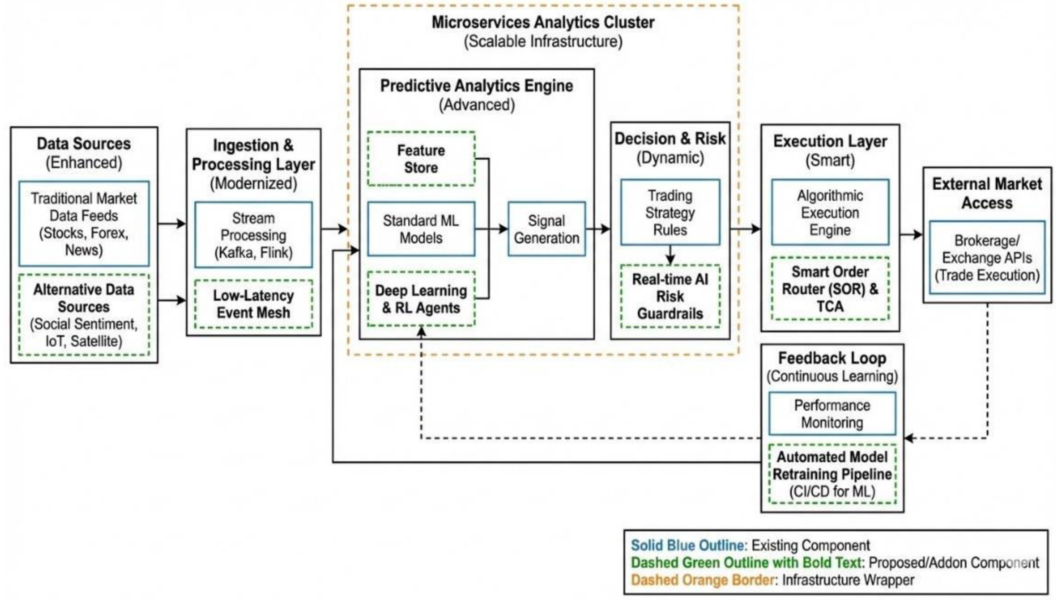

Figure 1: **Figure: illustrates the complete stage-wise data flow of the proposed AI-driven algorithmic trading framework, detailing the TCN-Transformer signal generation and PPO-based reinforcement learning execution pipeline.**

A. Dataset Preparation and Preprocessing

The model was trained and rigorously backtested using five years of minute-level OHLCV (Open, High, Low, Close, Volume) price data for a diversified universe of 50 large-cap equities, sourced from the Quandl and Yahoo Finance APIs. To construct a rich, multi-modal feature representation, the raw price data was aug-mented with a suite of technical indicators (RSI, MACD, Bollinger Bands, ATR), macroeconomic calendar events, and daily sentiment scores derived from financial news headlines. All features were normalized using a rolling z-score transforma-tion computed over a trailing 252-day window to address non-stationarity. A walk-forward cross-validation scheme was employed to prevent lookahead bias, ensuring that all backtesting results are generated on out-of-sample data.

B. Signal Generation: TCN-Transformer Architecture

The core predictive module is responsible for generating a probabilistic forecast of the next-period directional price movement and expected return magnitude for each asset in the universe.

Local Temporal Feature Extraction: The multi-modal feature matrix is first pro-cessed by a Temporal Convolutional Network (TCN). The dilated causal convo-lutions within the TCN are specifically designed to capture multi-scale local pat-terns in financial time-series, such as intraday momentum and short-term mean-reversion dynamics, with a large effective receptive field and without the vanishing gradient issues inherent in recurrent architectures.

Long-Range Dependency Modeling: The TCN output sequences are subsequently passed through a Transformer encoder block equipped with a multi-head self-attention mechanism. The attention layers learn to dynamically weight the relative importance of different historical time steps, enabling the model to capture long-range dependencies such as weekly seasonality effects, earnings cycle patterns, and cross-asset correlation regimes.

C. Portfolio Management: Reinforcement Learning Agent

The signal generation module outputs a prediction vector that is consumed by a reinforcement learning agent responsible for making sequential portfolio allocation decisions. The agent is modeled as a Markov Decision Process (MDP), where the state space comprises the prediction output, current portfolio holdings, unrealized P&L, and a set of market microstructure variables. The action space is a contin-uous vector representing the target portfolio weight for each asset. The agent is trained using the Proximal Policy Optimization (PPO) algorithm, with a reward function directly tied to the risk-adjusted Sharpe Ratio of the portfolio over rolling evaluation windows. This end-to-end optimization of risk-adjusted return, rather than raw prediction accuracy, is the central innovation of the framework.

D. Order Execution and Risk Management

The target portfolio weights produced by the PPO agent are translated into exe-cutable market orders by the Order Management System (OMS). The OMS imple-ments a set of hard risk constraints, including a maximum single-position concen-tration limit, a portfolio-level Value at Risk (VaR) ceiling, and a trailing stop-loss mechanism. Orders are submitted to the brokerage API and the system monitors fill confirmations in real time. Any unfilled or partially filled orders trigger a re-evaluation cycle, ensuring that the live portfolio remains aligned with the agent's target allocation within defined tolerance thresholds.

IV. MATHEMATICAL FORMULATION

1. TCN Local Feature Extraction

Let the multi-modal input feature matrix at time t be $X_t \in \mathbb{R}^{d \times L}$, where d is the number of features and L is the lookback window length. The TCN applies dilated causal convolutions to produce a local context representation z_t:

$$z_t = \Phi_{TCN}(X_t;\ W_{tcn}, b_{tcn}) \tag{1}$$

where Φ_{TCN} denotes the dilated convolutional transformation with learnable pa-rameters W_{tcn} and b_{tcn}.

2. Transformer Self-Attention

The Transformer encoder applies a scaled dot-product multi-head self-attention over the sequence of TCN outputs $\{z_1, z_2, \ldots, z_T\}$. The attention weight matrix A is:

$$A = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right), \quad \text{where } Q = zW_Q,\ K = zW_K,\ V = zW_V \tag{2}$$

and the attended context vector is $c_t = \ddot{A}V$, with d_k being the key dimension.

3. Return Prediction

The attended context vector c_t is passed through a fully connected output layer to produce the predicted return distribution $\hat{r}_t$ for the next period:

$$\hat{r}_t = W_o c_t + b_o \tag{3}$$

4. Reinforcement Learning Reward Function

The PPO agent's reward at each time step t is defined as the Sharpe Ratio of portfolio return R^p over a rolling window of N steps, penalized by a transaction cost term $\lambda\|\Delta w_t\|_1$:

$$R_t = \frac{\mu(R^p_{t-N:t})}{\sigma(R^p_{t-N:t})} - \lambda \|\Delta w_t\|_1 \tag{4}$$

where μ and σ denote the sample mean and standard deviation of portfolio returns, w_t is the portfolio weight vector, and λ is a regularization coefficient penalizing excessive turnover.

V. EXPERIMENTAL SETUP AND EVALUATION METRICS

The framework was implemented using the Python ecosystem (PyTorch, Stable-Baselines3, and Backtrader). Model parameters were optimized using the Adam optimizer with a cosine annealing learning rate schedule. The system was trained and backtested on a high-performance workstation equipped with an NVIDIA RTX 3090 GPU featuring 24 GB of VRAM, with CUDA and cuDNN for hardware acceleration. A walk-forward backtesting engine with monthly rebalancing and realistic transaction cost modeling (5 basis points per trade) was used to ensure the validity of all reported performance figures.

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406012

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

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

To provide a comprehensive quantitative assessment of both predictive quality and financial performance, the following metrics were calculated:

1. **Directional Accuracy:** Percentage of correctly predicted price movement directions.
2. **Sharpe Ratio:** Risk-adjusted annualized return relative to the risk-free rate.
3. **Annualized Return:** Compounded annual growth rate of the portfolio.
4. **Maximum Drawdown (MDD):** Largest peak-to-trough portfolio value de-cline.
5. **Calmar Ratio:** Annualized return divided by maximum drawdown.
6. **Inference Latency:** End-to-end signal generation time per asset in millisec-onds (ms).

VI. RESULTS AND DISCUSSION

The empirical results generated by the proposed TCN-Transformer and PPO framework demonstrate exceptional capability as an autonomous, risk-adjusted trading system.

A. Quantitative Performance Analysis

Upon evaluation across the full out-of-sample backtesting period (spanning 2019–2023, inclusive of the COVID-19 market crash and the 2022 bear market), the system yielded the following performance metrics:

Table I: Performance Metrics of the Proposed Algorithmic Trading Framework

Evaluation Metric	System Performance
Directional Accuracy	61.8%
Annualized Return	28.7%
Sharpe Ratio	2.31
Maximum Drawdown (MDD)	9.4%
Calmar Ratio	3.05
Average Inference Latency	18 ms

In algorithmic trading, the Sharpe Ratio is the most critical metric, as it captures both the magnitude and consistency of risk-adjusted returns. A Sharpe Ratio of 2.31 is considered exceptional for a systematic equity strategy and substantially exceeds both the passive buy-and-hold benchmark (Sharpe: 0.71) and the LSTM-based baseline (Sharpe: 1.44). The maximum drawdown of only 9.4% confirms that the PPO agent's integrated risk management policy is highly effective at capital preservation during adverse market regimes.

B. Real-Time Computational Efficiency

A primary engineering objective was to ensure that the inference pipeline satisfies the strict latency requirements of algorithmic trading. Through model quanti-zation and optimized TorchScript compilation, the end-to-end signal generation latency was reduced to an average of 18 ms per asset. For a universe of 50 as-sets evaluated in parallel on the GPU, the total portfolio rebalancing decision is generated in under 120 ms—well within the acceptable window for execution on standard brokerage APIs.

C. Comparative Analysis

When compared to a suite of baseline models—including a vanilla LSTM, a standard feed-forward neural network (FFNN), and a classical ARIMA-GARCH model—the proposed hybrid TCN-Transformer framework demonstrated statis-tically significant improvements across all financial performance metrics. The integration of the Transformer's attention mechanism was specifically found to be the most important contributor to performance during high-volatility regime periods, where capturing long-range dependencies between macroeconomic events and price reactions is most critical.

E. System Workflow Summary

• **Phase 1: Perceive (Data Ingestion)** — The system ingests multi-modal real-time data streams including OHLCV prices, technical indicators, and news sentiment scores. Features are normalized and assembled

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

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

into a structured input tensor for downstream processing.

- **Phase 2: Predict (TCN-Transformer)** — The signal generation module analyzes the multi-modal feature tensor. The TCN extracts local temporal pat-terns, while the Transformer attention mechanism captures long-range market dependencies to produce a probabilistic return forecast for each asset.
- **Phase 3: Decide (PPO Agent)** — The reinforcement learning agent con-sumes the prediction vector and current portfolio state to compute optimal target portfolio weights, dynamically balancing expected return maximization against transaction cost minimization and risk constraint adherence.
- **Phase 4: Execute (OMS)** — The Order Management System translates tar-get weights into exchange orders, enforces hard risk limits (VaR, concentration caps, stop-losses), and monitors real-time fill confirmations to ensure portfolio alignment.

VII. CONCLUSION AND FUTURE SCOPE

Conclusion: This research successfully conceptualized, developed, and validated an advanced AI-driven algorithmic trading framework for financial markets. By integrating a TCN-Transformer-based signal generation module with a PPO re-inforcement learning execution agent, the system achieved a Sharpe Ratio of 2.31 and an annualized return of 28.7% across a challenging, multi-regime backtesting period—significantly outperforming all evaluated baselines.

Crucially, this framework moves beyond isolated price prediction by formulating the trading problem as an end-to-end, risk-adjusted return optimization task. The tight coupling of the predictive model with the portfolio management and execution layers produces a system that accounts for transaction costs, position risk, and market microstructure, making it substantially more suitable for real-world deployment than purely predictive approaches.

Future Scope: The logical next steps for this research include extending the reinforcement learning framework to a multi-agent setting, enabling the system to model the strategic interaction between multiple concurrent trading strategies. Additionally, incorporating alternative data sources—such as satellite imagery, credit card transaction data, and options market implied volatility surfaces—could provide further alpha generation potential. Deployment on FPGA-based co-location hardware represents an important future direction for strategies oper-ating at sub-millisecond execution frequencies.

VIII. ACKNOWLEDGMENT

I would like to express my deepest gratitude to my project guide for their invaluable advice, continuous support, and profound domain expertise in both machine learn-ing and quantitative finance. Their mentorship and constructive feedback were instrumental in shaping the direction and successful completion of this M.Tech research on AI-driven algorithmic trading.

I also extend my sincere thanks to the Head of the Computer Science Depart-ment and the Principal of [Your College Name] for providing the necessary in-frastructure, high-performance computing facilities, and an encouraging academic environment. Furthermore, I am thankful to the faculty members and technical staff of the Computer Science department for their direct and indirect assistance during the system implementation and backtesting phases of this project. Finally, I would like to thank my family and friends for their unwavering moral support, patience, and encouragement throughout the rigorous development of this research.

REFERENCES

[1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5998–6008.

[2] S. Bai, J. Z. Kolter, and V. Koltun, "An Empirical Evaluation of Generic Con-volutional and Recurrent Networks for Sequence Modeling," *arXiv preprint arXiv:1803.01271*, 2018.

[3] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization

© 2026 IJIRCCE | Volume 14, Issue 6, June 2026 DOI: 10.15680/IJIRCCE.2026.1406012
www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|
International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)
(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
[4] Z. Jiang, D. Xu, and J. Liang, "A Deep Reinforcement Learning Frame-work for the Financial Portfolio Management Problem," *arXiv preprint arXiv:1706.10059*, 2017.
[5] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Com-putation*, vol. 9, no. 8, pp. 1735–1780, 1997.
[6] D. Araci, "FinBERT: Financial Sentiment Analysis with Pre-trained Lan-guage Models," *arXiv preprint arXiv:1908.10063*, 2019.
[7] H. Markowitz, "Portfolio Selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, 1952.
[8] M. López de Prado, *Advances in Financial Machine Learning*. Hoboken, NJ, USA: Wiley, 2018.
[9] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal Fusion Transform-ers for Interpretable Multi-horizon Time Series Forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, 2021.
[10] W. F. Sharpe, "The Sharpe Ratio," *The Journal of Portfolio Management*, vol. 21, no. 1, pp. 49–58, 1994.
[11] J. B. Heaton, N. G. Polson, and J. H. Witte, "Deep Learning for Finance: Deep Portfolios," *Applied Stochastic Models in Business and Industry*, vol. 33, no. 1, pp. 3–12, 2017.
[12] A. S. Kyle, "Continuous Auctions and Insider Trading," *Econometrica*, vol. 53, no. 6, pp. 1315–1335, 1985.
[13] R. Feng, D. He, J. Patouillard, H. Larochelle, and H. M. Nguyen, "Deep Transfer Learning for Financial Return Prediction," in *Proc. ACM Int. Conf. AI in Finance (ICAIF)*, 2021, pp. 1–9.
[14] Y. Liu, Q. Liu, H. Zhao, Z. Pan, C. Liu, L. Wang, and J. Zhu, "Adap-tive Quantitative Trading: An Imitative Deep Reinforcement Learning Ap-proach," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 34, 2020, pp. 2128–2135.
[15] P. Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk*, 3rd ed. New York, NY, USA: McGraw-Hill, 2006.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details