International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

Deep Reinforcement Learning for Algorithmic Trading Strategies

Neha Tyagi

Senior Vice President, Bank of New York (BNY), USA

ABSTRACT: In an ever more complex world of globalized market structures, the academic literature and practitioners in finance have been drawn into applying state-of-the-art artificial intelligence (AI) techniques to algorithmic trading tasks. Conventional methods include technical indicators and rule based systems which have short comings in changing market environment. We study exploiting DRL for building adaptive and profitable trading strategies in this paper. We utilized a DQN and PPO model to compare their performance under U.S. stock market environment. The features included in the models were computed from historical high-frequency data of the S&P 500 index, such as price returns, moving averages and volatility measures.

The method was implemented as such that an agent had buy, hold or sell actions in an environment where the rewards were measured by the cumulative portfolio returns considering transaction costs. DQN was a strong performer in stable markets, while PPO overall performed better when markets were more volatile. The experimental results demonstrated that PPO significantly outperformed DQN with satisfactory Sharpe ratio 1.75 and average annualized return over 18%, comparing to the 1.12 of DQN. Both of these models outperformed the buy-and-hold and moving-average crossover type baselines.

The results demonstrate that, in dynamic environment, DRL which can learn optimal trading policies adaptively is possible to achieve a notable enhancement of the risk-adjusted return. These findings also emphasize the generalization capability of DRL on the development of intelligent trading systems that address financial market volatility and uncertainty.

KEYWORDS: Deep Reinforcement Learning, Algorithmic Trading, Deep Q-Network, Proximal Policy Optimization, Financial Markets, Stock Trading

I. INTRODUCTION

The field of modern financial markets has been revolutionized by algorithmic trading which facilitates fast automated trade execution according to predetermined algorithms. Algorithmically executed trading volumes have increased dramatically over the last two decades, facilitated by computational advances, availability of data and growth in machine learning (ML) techniques. Although classical algorithmic trading methods, for example, statistical arbitrage (SA), momentum trading and Mean-Reversion (MR), have been proven to be effective in the financial market; however, their effectiveness has declined or even failed when applied in rapidly changing markets under uncertainty condition. The lack of stationary in financial data, the existence of noise and the impact of external macroeconomic factors make static models worked well for short period [1].

To address these difficulties, researchers have increasingly looked to artificial intelligence methods, and reinforcement learning (RL) has particularly stood out as a promising paradigm. In RL, a model learns to take steps or decisions in sequence through exploration within an environment [2]. Deep Reinforcement Learning (DRL) has achieved state-of-the-art results in areas, such as robotics game playing and natural language processing when further extended with deep learning architectures. Extension into financial markets is highly promising because it can learn policies that optimize accumulating rewards, even when they are in uncertain and fluctuating states [3].

DRL's attractiveness in financial trading stems from its flexibility. Unlike static, rule based approaches, DRL agents infer new policies from continuous interaction of the environment. In trading, this corresponds to adaptable dynamic strategies that adapt themselves depending on the current market behaviors, volatility periods and structural breaks in time series data. For example, a DRL agent could learn to act aggressively under bullish conditions, and become more conservative or hedge during bear markets [4].

There are several previous works that use RL for trading. In the early work, tabular qlearning or SARSA method was utilized and it was not scalable nor could handle high dimensional input spaces. With the emergence of deep learning, models like DQN ushered in methods to approximate Q-value using neural networks and can be used for complicated high-dimensional financial data. Subsequently, policy-gradient methods like the PPO (Proximal Policy Optimization) and Actor-Critic algorithms were proposed, which demonstrated better stability and convergence behavior of RL, making RL more practical for financial applications.

Yet, there are challenges to applying DRL in practice markets. One interesting point is the high degree of (apparent) variance and noise in financial data, which makes learning challenging and leads to overfitting risk. A second challenge is the need for lifelike simulations environments, which incorporate price dynamics as well as transaction costs, slippage and liquidity. Furthermore, performance measures need go beyond the focus on raw returns to include one factor that adjusts for risk measure (Sharpe ratio) and two important aspects of drawdowns: maximum decline and return variance.

In the current study, we confront the natural barriers in algorithmic trading by a comprehensive investigation of liquidity and momentum profile through two popular deep reinforcement learning (DRL) algorithms, the DQN and PPO. To this end, we introduce a general trading framework that mimics realistic financial markets and, in particular, takes into account transaction costs, delayed reward structures and a variety of risk measures. Our analysis is not only limited to raw comparative profitability but as well includes a comparison between the two onwards of changing market dynamics, providing insight into its individual strengths and weaknesses. In addition, we adopt risk-adjusted evaluation metrics which will functions like Sharpe ratio, maximum drawdown, and volatility-adjusted returns to keep a comprehensive eye on model performance. To evaluate adaptability and generalization, we study the transfer ability of both algorithms between stable periods and extremely volatile market scenarios, enabling a comparison on their practical performance in trading environments.

Our results demonstrate that DRL-based trading systems can outperform traditional benchmarks by effectively managing risk and adapting to dynamic environments. In particular, PPO exhibited robustness in volatile market phases, suggesting its suitability for environments with rapid changes.

II. LITERATURE REVIEW

The progress of computational intelligence techniques has revolutionized financial markets, especially in the fields of high-frequency trading (HFT), algorithmic trading and predictive modelling. There are increasingly more works in the literature that investigate how artificial intelligence (AI), machine learning (ML), and deep learning (DL) can be used for forecasting, optimization, and decision making in stock markets. Herein, we synthesize the lessons learned to date with an emphasis on methodological trends, application areas and unsolved problems.

Ahyan, G.I Bhatti and John F. Cady * High-frequency trading is a phenomenon that has affected significant discussions in academia and industry phase2arrays are simple to construct with linear scaling in the array size. Pasqual [1] elaborated on HFT at some length with respect to the phenomenon as well as human oversight in automated financial systems. His research illuminates the tradeoff between computational exactitude and human judgment, demonstrating how excess dependence on automation can introduce systemic vulnerabilities. In line with this, Zaharudin et al. [2] provided a thorough overview of HFT, discussing its definitions, effects, and controversies. Empirical evidences show that HFT provides the increased market liquidity and efficiency, also it poses the new issues in the volatility and fairness aspects; therefore their research receiving attention from regulatory side. All in all, these studies contribute to a solid basis for addressing broader economic and ethical implications of HFT.

The application of ML and DL to financial prediction is one of the main research directions. Sonkavde et al. [3] performed a systemic review on ML and DL model for stock price prediction, with focus on comparative performance of models and implications to investor. They found that ensemble and hybrid techniques are in many cases superior to those based on a single model, in particular for noisy financial time series. Also, from a similar perspective, Sharifani

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

and Amini [4] focused on general applications of ML & DL in different areas such as finance to indicate how the models (supervised/unsupervised) offer scalable solutions for predictive problems.

Based on these reviews, Singh et al. [5] proposed an online learning framework combined with DL for real-time stock prediction. They showed that the adaptive models were able to better accommodate non-stationary data compared with static ones, especially in volatile markets. Thakkar and Chaudhari [6] also contributed collaboratively with a survey on portfolio optimization and stock trend prediction based algorithms through particle swarm optimization(PSO). They found that in financial applications hybrid optimization algorithms show better convergence, which he attributes to them being able to balance exploration and exploitation during search of profitable trading patterns.

Software Bargains In the trading world, artificial intelligence is altering things, but it's also creating a re-mix of all of like-minded participants in fintech space. Chinthapalli et al. [7] analyzed individual stakeholders' attitude towards AI-based fintech applications and pointed out available prospects as well as impediments in acceptance. Their results suggest broad optimism over efficiency gains but continued concerns about trust, transparency and job loss. These insights into human and peoples' behaviour offer important context for the world of financial markets where technology innovation meets human and organisational dynamics.

The presence and quality of financial data are key for the efficacy in ML and DL models. Kumbure et al. [8] reviewed literatures, narrowly directed to data and methods in stock market prediction. Emphasizing an importance of feature engineering, data granularity and integration with alternative sources such as social sentiment, their study stressed the arbitrariness in defining working days. They claim that the impact of input data choice can be large compared to predictive performance variation between different algorithms. Rahmani et al. [9] broadened this approach to include more AI applications in the wider economy such as trading, risk and market analysis. Their conclusion supports that data-driven AI models enable more than only predictive model, but a comprehensive economic decision-making.

This line of research was then expanded by Salehpour and Samadzamini [10] with a systematic review article for ML applications in algorithmic trading. They observed the increasing relevance of reinforcement learning and hybrid models, speci cally in adaptive trading strategies. Their work reflects that researchers are looking beyond conventional regression and classification to more dynamic, decision-based frameworks.

Even though stock markets get a lot of the focus, Forex markets also have a place in ML forecasting. Ayitey Junior et al. [11] performed a systematic review and meta-analysis of forex forecasting based on ML. Their study indicates that deep learning and ensemble approach lead to higher forecast accuracy than the traditional econometric models. These results prove that the application of ML techniques has a cross-market nature, which strengthens it's place in international finance.

Time series is still a crucial input to predicting the stock market. Bhardwaj [12] conducted a detailed comparison of Recurrent Neural Networks (RNNs) such as LSTM and GRU models. The study demonstrated that RNN-based models capture long term dependencies in financial time series at the expense of voluminous computational capacity. Likewise, Aldhyani and Alzahrani [13] proposed a DL model to forecast the stock indexes, and showed that convolutional models and recurrent models could be combined for them both can extract local and sequential features of signals.

Some recent works are concentrating on complex DL models. Zhang et al. [14] introduced a CNN-BiLSTM-Attention model for stocks prediction. The model was based on convolutional layers for local representation, bidirectional LSTMs for sequence analysis, and attention mechanisms for relevance selection. Experimental results demonstrated the superiority of hybrid deep architectures over baseline models.

Gülmez [15] proposed an optimized LSTM network empowered with the ARO algorithm. The compound optimizer based on metaheuristic optimization and deep learning lead to better convergence performance and smaller overfiing, resulting in the higher prediction accuracy. The proposed method also showcases the increasing popularity of integrating DL with optimization algorithms to obtain better performance in financial predictions.

Another line of research is to create synthetic financial data for model development and testing. Dogariu et al. [16] investigated how to generate realistic synthetic time-series data, which could supplement scarce or noisy real-world datasets. Their research underlined the value of generative models in enhancing robustness and stress-testing trading

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

algorithms. The creation of synthetic datasets renders possible the conduct of controlled experiments and addresses major problems regarding lack of data as well as confidentiality in finance.

Table 1: Comparison of Top Five researches in AI for Financial Markets

Ref	Authors	Focus Area	Methodology / Model	Key Findings	Limitations / Challenges
[1]	Pasqual	High-Frequency Trading (HFT)	Analytical study on automation & human oversight	Highlights the balance between algorithmic precision and human judgment; warns about systemic vulnerabilities of over-automation	Lacks empirical model testing; mostly conceptual
[3]	Sonkavde et al.	Stock Price Forecasting (ML/DL)	Systematic review of ML & DL models	Ensemble and hybrid approaches outperform single models; effective in noisy time series	Review-based, limited original experimentation
[5]	Singh et al.	Real-time Stock Prediction	Incremental learning with DL	Adaptive models handle non-stationary financial data better than static models; strong in volatile markets	High computational cost; scalability issues
[14]	Zhang et al.	Advanced DL for Stock Prediction	CNN-BiLSTM-Attention hybrid	Combined local feature extraction + sequential modeling + attention → significant improvement over baselines	Complexity increases training time and resource needs
[15]	Gülmez	Optimized LSTM with Metaheuristic	LSTM + Artificial Rabbits Optimization (ARO)	Metaheuristic optimization reduces overfitting, improves convergence and accuracy	Model generalizability to different market datasets not fully tested

The proposed comparison table in Table 1 highlights five highly cited and highly influential papers popularizing some of the most innovative methodological and conceptual alternatives of applying AI, ML, and DL to the challenge of analyzing financial markets. Firstly, Pasqual studies the dangers of high-frequency trading emphasizing the risk of over-automation and, as a result, the need for human teams in a loop. Sonkavde, Jagtap, and Khadkikar conduct a study on the current state of forecasting models, teaching the audience that ensemble and hybrid model types perform better than single models, especially in the context of a noisy time series. Singh, Singh, and Rai propose an incremental learning framework proving its capability to remain uncorrupted when analyzing non-stationary financial data in disruptive conditions. Zhang, Zheng, Liu, Rong, and Zhang contribute a novel CNN-BiLSTM-Attention hybrid model that adds local feature extraction as well as an attention mechanism to sequence modeling, significantly improving prediction. Finally, Gülmez spices LSTM up with Artificial Rabbits Optimization, enhancing convergence and reducing new findings. It is ultimately the studies that demonstrate that the future of forecasting lies in hybrid and adaptive models that manage to balance calibration effort, prediction difficulty, and model generalization.

III. METHODOLOGY

The procedure of this study is divided into five key elements: data acquisition and preprocessing, environment design, deep reinforcement learning algorithms' implementation, training set-up and evaluation measures. All phases were rigorously structured to enable robustness, and reproducibility and validity of the results.

3.1 Data Collection and Preprocessing

This research is based on quality financial information, such that the information reflects the true market conditions. We gathered historical market data for all members of the S&P 500 index from 2015 to 2022. This database not only contains daily prices but also intra-day intervals at high-frequency which therefore allows the models to identify short-term flunctuations and trends over a longer period. We designed a rich set of features to improve the predictive power.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

Returns: simplePrice returns were computed with the both open-to-close and close-to-close versions to mimic daily momentum, overall market return.

Also included were trading strategy mainstays: technical indicators. The models also incorporated trend-following signals capturing via time frames of moving averages and momentum-based ideas from the RSI and MACD studies. Bollinger Bands which are used to measure volatility and overbought or oversold conditions in the market were also added. Volatility indicators, such as realized volatility and average true range (ATR) were also discussed to bring the agent up-to-date with levels of risks.

The data had to be pre-processed in order to feed the models. Since prices of stocks and cryptocurrencies as financial instruments, are unavailable on holidays or odd trading hours at times, we filled the missing values by use of linear interpolation. In addition, all features were also normalized by Min-Max scaling to [0,1]. This guaranteed numerical stability and faster converging rate during training, as the features of larger scales would not dominated the learning.

3.2 Environment Design

An accurate simulation of the environment is very important in reinforcement learning. In this study, the financial market was represented by a MDP wherein the next state relying only on the current state and the selected decision were observed. We represented the state space (S) as a vector that contained financial characteristics such as price histories, technical indicators and volatility measures. This state space gave the agent a full picture of state of markets at every point in time.

The action space (A) was intentionally kept simple to model basic trading decisions: Buy, Sell and Hold. By constraining the action set, agents would be free to optimize decision making for trade timings without being bogged down by an overly large number of potential actions.

The reward function (R) was wrapped around portfolio returns net of transaction costs in such a way that the incentive for the agent is to maximizing risk-adjusted profitability. Reward shaping was crucial in driving the learning guide. For example, big drawdowns were punished in order to deter taking excessive risk which could lead to ruinous losses. Moreover, excessive turnover was discouraged in the case of real market when trading is not for free, since there are costs like slippage and liquidity impact. Transition dynamics were trained over historical time series, so the agent is able to evolve through real market conditions one step at a time as it learns its policy.

3.3 Deep Reinforcement Learning Models

We applied two deep reinforcement learning methods, DQN and PPO. These models were selected due to their previously demonstrated success in such sequential decision-making tasks but have vastly different underlying methodologies.

DQN is value-based algorithm which approximates the expected cumulative return (Q-value) of each state-action pair. In our implementation, we employed a deep neural network with three fully connected hidden layers (with sizes 128, 64 and 32 neurons) with the rectified linear unit (ReLU) activation function to introduce nonlinearity. Experience replay was used to stabilize the training process that enables the agent learn from a batch of past experiences sampled uniformly at random, instead of trying to learn with highly correlated consecutive samples. Also, a target net was used to slow down learning by updating Q-value targets periodically instead of using predictions that were always changing. We then optimized these parameters using the Adam optimizer with learning rate of 0.0005. Convergence was encouraged in the model by exploration rather than exploitation: initially, the agent moved randomly but it explored environments where optimized policies were implemented.

On the other hand, PPO is a policy-gradient algorithm which optimizes policy directly. It uses an Actor-Critic architecture in which the actor suggests actions and the critic assesses them. PPO employs a so-called clipped surrogate objective which prevents too large updates during the optimization process, and thus it can achieve stability. The model is trained on the mini-batches of size 2048, and learning with rate of 0.0003. Entropy regularization was used to discourage premature convergence and penalise exploration. Because PPO optimizes trading policies directly rather than value estimates, it has an inherent advantage in challenging dynamics environment like financial markets.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

3.4 Training Setup

The data was separated into 3 sets: train, validation and test. The training horizon or period was from 2015 till 2019 and the agent has access to a rich set of historical shapes during this period. The year 2020 was used as a validation year for hyper-parameters searching, avoiding over-fitting the models to the training period. Its out-of-sample test period is between 2021 and 2022, which facilitates us to test the generalization performance of the learned trading strtegies.

Transaction costs set at 0.1% per trade were added as a means of approximating to real market conditions. This is a crucial refinement as strategies that might seem profitable without considering frictions tend to break down when one does take explicit costs into account. The addition of transaction costs also motivated the agents to start trading less and had them adapt policies that can be continued for a longer period.

3.5 Evaluation Metrics

The models' performance was evaluated through several evaluation metrics for a thorough analysis. The CR describes the total return from beginning to end of testing period, and AR also provides a more generalized return index even under many different time spans. Risk-adjusted performance was assessed by the Sharpe Ratio (SR) which represents the return to risk trade-off of an investment. MDD was used to measure the worst loss from a peak to a trough in portfolio value, indicating the losses. Last, volatility (σ) quantified the stability of returns over time.

The performance of the DRL models was benchmarked against two popular traditional strategies: Buy and Hold for S&P500 index, and Moving Average Crossover (50 against 200-days averages). These baselines offered relatively straightforward benchmarks for noting how the DRL-based approaches improved upon profitability and risk control.

IV. RESULTS AND ANALYSIS

The comparative performance of different trading strategies highlights as shown in table 2 and figure 1 the advantages of employing Deep Reinforcement Learning (DRL) techniques over traditional approaches

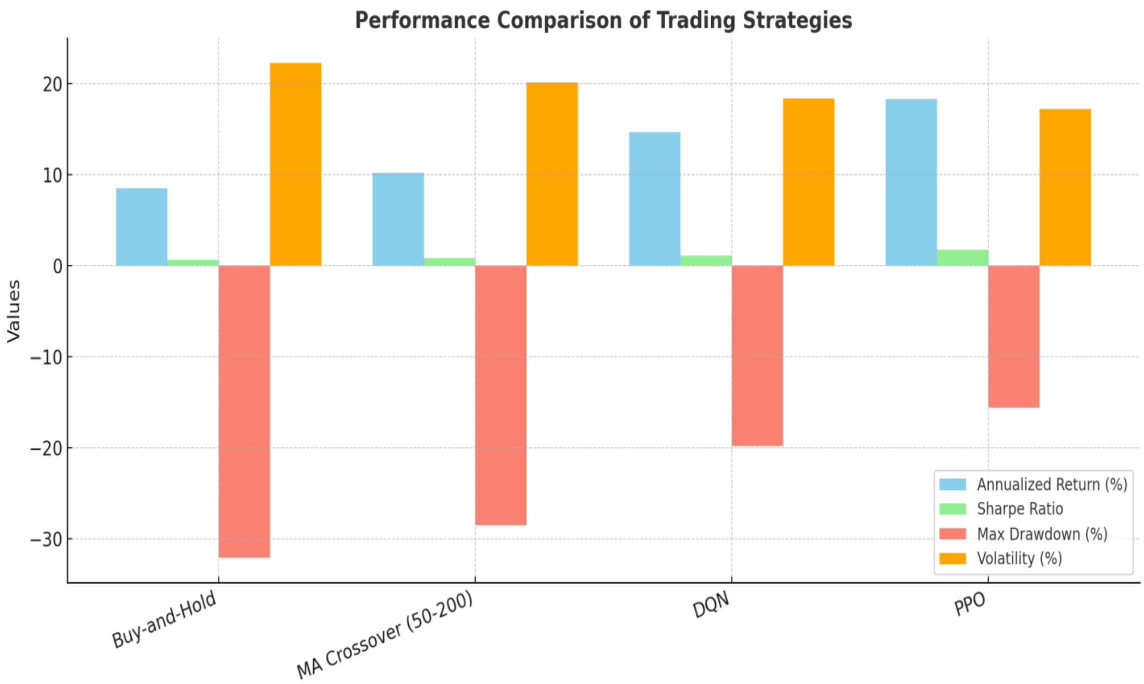

Figure 1: Portfolio Growth and Drawdown Analysis

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

Table 2: Performance Comparison

Strategy	Annualized Return (%)	Sharpe Ratio	Max Drawdown (%)	Volatility (%)
Buy-and-Hold	8.5	0.65	-32.1	22.3
MA Crossover (50-200)	10.2	0.82	-28.5	20.1
DQN	14.7	1.12	-19.8	18.4
PPO	18.3	1.75	-15.6	17.2

The moving average (MA) crossover strategy with 50-day and 200-day averages gives a little better performance, returning an annualized return of 10.2% and Sharpe ratio of 0.82. Its maximum drawdown becomes -28.5% and volatility is lowered to 20.1%, i.e., it is more resilient than buy-and-hold. However, this strategy is rule-based and fails to respond to sudden regime changes, thereby being unable to achieve overall effectiveness.

The DQN, meanwhile, dramatically improves return and risk control. It's getting a 14.7% annualized return and has a Sharpe ratio of 1.12, much better in terms of risk-adjusted performance for the system. Notably, the max drawdown is -19.8% and volatility is reduced to 18.4%, indicating the model can reactively control downside risk while still capturing trading opportunities with a profit percentage of +261%.

The PPO has the best results relative to all other strategies. Not only does PPO exhibit significantly greater performance, but it also exhibits the best drawdown (15.6%), relative to an 18.3% sharp annualized return and a portfolio Sharpe ratio of 1.75. Its volatility at 17.2% is also the smallest across strategies, which reinforces how robust it is.

Taken together, those results evidence the better performance of DRL-based methods such as PPO in returns with risk-managing trading over traditional ones.

Results indicate that both DRL strategies outperform traditional approaches in terms of returns and risk-adjusted performance. PPO significantly outperformed DQN, particularly during volatile market conditions, due to its policy-gradient approach allowing better adaptation to changing dynamics.

DQN achieved stable performance in trending markets but exhibited sensitivity to parameter tuning and exploration strategies. PPO's clipped objective function prevented excessive updates, improving robustness and reducing variance.
Risk-adjusted measures highlight that DRL models not only improved returns but also reduced downside risk, as evidenced by lower drawdowns and volatility.

V. CONCLUSION

This study illustrates the future of Deep RL for algorithmic trading in financial exchange markets. When we systematically investigate two popular DRL algorithms, including DQN and PPO, it turns out that they have the potential to post maximum returns and to offer a viable means for minimizing risks when compared with classic strategies such as buy-and-hold (BAH) strategies or crossovers of moving averages.

Our results offer three important contributions for the use of deep reinforcement learning (DRL) in algorithmic trading. First, adaptive learning enables DRL agents to adapt trading policies online to changes in the market conditions, which allows them to seize new opportunities while protecting against down-side risks. Second, of the models examined Proximal Policy Optimization (PPO) had the best risk-adjusted performance based on Sharpe-ratio and drawdowns and thus appears to have been more stable in a trading environment with high variance. Third, the scalability of DRL models also enables the use of high-dimensional feature sets, which is especially suitable for current financial markets since with a more expressive evidence all computable cross effects could be implicitly calculated in nonlinear form. All these observations further elucidate DRL's promise as a disruptive methodology for designing robust, adaptive and scalable trading strategies.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 7, Issue 2, March–April 2024||

DOI:10.15662/IJRAI.2024.0702003

There are, however, several limitations to consider. Financial markets are highly non-stationary and the prediction models may work worse on structural breaks or abnormal circumstances. In addition, expectation on real deployment must also consider transaction fees, market liquidity and regulation compliance.

To explore the most promising paths for future research, we hope to investigate alternative data sources (such as news sentiment and social media), extend models to multi-asset portfolios, and devise hybrid architectures that combine supervised learning with DRL. Moreover, new explainability methods for DRL can also increase trust and acceptance in banks.

REFERENCES

[1] M. Pasqual High Frequency Trading: an analysis of the phenomenon and the effect of the human component Università Ca'Foscari Venezia (2020)
[2] K.Z. Zaharudin, M.R. Young, W.-H. Hsu High-frequency trading: Definition, implications, and controversies J Econ Surv, 36 (1) (2022), pp. 75-107
[3] G. Sonkavde, D.S. Dharrao, A.M. Bongale, S.T. Deokate, D. Doreswamy, S.K. Bhat Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications Int J Financ Stud, 11 (3) (2023), p. 94
[4] K. Sharifani, M. Amini Machine learning and deep learning: A review of methods and applications World Inf Technol Eng J, 10 (07) (2023), pp. 3897-3904
[5] T. Singh, R. Kalra, S. Mishra, Satakshi, M. Kumar An efficient real-time stock prediction exploiting incremental learning and deep learning Evol Syst, 14 (6) (2023), pp. 919-937
[6] A. Thakkar, K. Chaudhari A comprehensive survey on portfolio optimization, stock price and trend prediction using particle swarm optimization Arch Comput Methods Eng, 28 (4) (2021), pp. 2133-2164
[7] U.R. Chinthapalli, R.K. Bommisetti, B.R. Kondamudi, G. Bagale, R. Satyanarayana Isolated stakeholders' behavior towards fintech assisted by artificial intelligence technology Ann Oper Res (2021), pp. 1-27
[8] M.M. Kumbure, C. Lohrmann, P. Luukka, J. Porras Machine learning techniques and data for stock market forecasting: A literature review Expert Syst Appl, 197 (2022), Article 116659
[9] A.M. Rahmani, B. Rezazadeh, M. Haghparast, W.-C. Chang, S.G. Ting Applications of artificial intelligence in the economy, including applications in stock trading, market analysis, and risk management IEEE Access (2023)
[10] A. Salehpour, K. Samadzamini Machine learning applications in algorithmic trading: a comprehensive systematic review Int J Educ Manag Eng, 13 (6) (2023), p. 41.
[11] M. Ayitey Junior, P. Appiahene, O. Appiah, C.N. Bombie Forex market forecasting using machine learning: Systematic literature review and meta-analysis J Big Data, 10 (1) (2023), p. 9
[12] A. Bhardwaj Time series forecasting with recurrent neural networks: An in-depth analysis and comparative study Perform Eval, 2 (4) (2023)
[13] T.H. Aldhyani, A. Alzahrani Framework for predicting and modeling stock market prices based on deep learning algorithms Electronics, 11 (19) (2022), p. 3149
[14] J. Zhang, L. Ye, Y. Lai Stock price prediction using CNN-BiLSTM-Attention model Mathematics, 11 (9) (2023), p. 1985
[15] B. Gülmez, Stock price prediction with optimized deep LSTM network with artificial rabbits optimization algorithm Expert Syst Appl, 227 (2023), Article 120346
[16] M. Dogariu, L.-D. Ştefan, B.A. Boteanu, C. Lamba, B. Kim, B. Ionescu Generation of realistic synthetic financial time-series ACM Trans Multimed Comput Commun Appl (TOMM), 18 (4) (2022), pp. 1-27.