Option-Driven Sentiment in FinRL: a PPO Approach to Trading

Shenjian Li
Accountancy and Finance
University of Otago
Dunedin, New Zealand
shenjian.li@postgrad.otago.ac.nz

Mingxuan Yu
Accountancy and Finance
University of Otago
Dunedin, New Zealand
yumi7517@student.otago.ac.nz

Freddie Dossor
Accountancy and Finance
University of Otago
Dunedin, New Zealand
dosfr332@student.otago.ac.nz

***Abstract*—This paper addresses Task 1 of the FinRL Contest 2025 by applying a sentiment-enhanced stock trading agent based on Proximal Policy Optimization (PPO) augmented with market sentiment indicators, such as Put-Call Ratio, as key input, which is an option-based market sentiment indicator. The agent was tested on NASDAQ-100 constituents from 2019 to 2023 by using the trained dataset from 2013 to 2018. We found the PPO model outperforms both the baseline index and CVaR-PPO (CPPO) in cumulative return, Sharpe ratio, and Sortino ratio. Notably, the Sharpe ratio improved from 0.69 to 1.08. The results suggest that incorporating market sentiment indicators into PPO can significantly enhance risk-adjusted performance.**

Keywords—Reinforcement Learning, Automated Trading, PPO, Sentiment, Options trading

I. INTRODUCTION

Over the past several decades, equity trading and expected return prediction have attracted considerable attention from academic researchers and industry practitioners. Many trading strategies, including fundamental, momentum, and sentiment-based approaches, have been extensively explored, tested, and validated in real-world financial markets. Recently, with the rapid growth and application of artificial intelligence (AI), trading methods based on reinforcement learning algorithms have become important and popular.

Numerous studies have shown that algorithm trading methods can enhance portfolio performance by maximizing alpha [1][2]. However, prior studies have not focused on identifying and using the signals and indicators that impact alpha in real-world equity trading. Therefore, recognizing signals and indicators with high predictive power is critical, particularly in the increasing uncertainty in the macroeconomic environment and equity markets.

This paper employs a reinforcement learning algorithm to develop an intelligent trading agent following the requirements of FinRL Contest 2025 Task 1 to address the above challenges. Built on the deep learning library FinRL [3][4], we integrate technical indicators, daily share price information, and daily market sentiment indicators, including the VIX index, Turbulence and Put-Call Ratio, into the trading as factors that intelligent agents should consider when making trading strategies. Martin E. Zweig [5] initially proposed the Put-Call Ratio, which has been widely used as a sentiment indicator and tested across various financial markets for its ability to predict short-term price movements. It has been applied to assess structural breaks in investor sentiment surrounding major climate policy events affecting the oil and gas sector [6]. Furthermore, our reward function incorporates total asset growth criteria, encouraging the agent to maximize portfolio performance.

We use PPO as the reinforcement learning model to construct an intelligent agent. We selected PPO due to its well-known stability and ease of implementation [7]. Our results demonstrate that our PPO-based agent outperforms the stocks in the NASDAQ-100 index and the CPPO-trained model in cumulative returns, Sharpe ratio, and Sortino ratio during the trading period. These results suggest that even under conditions of increasing macroeconomic uncertainty, our PPO agent can capture patterns in stock price movements and make trading decisions that are more likely to generate higher alpha.

II. METHODOLOGY

This paper explores utilizing the Put-Call Ratio as an additional metric inside the PPO framework to augment stock trading strategies potentially. Our dataset covers the period from January 1st, 2013, through December 31st, 2023, split into a training set from January 2013 until December 2018 and the trading set from January 2019 through December 2023. The training and trading dataset was downloaded from the organizer-provided Hugging Face website, benstaf/Nasdaq_2013_2023 [8], which included the indicators VIX and Turbulence, which were default in the dataset. We also included these two indicators in our training and trading. The additional indicator, Put-Call Ratio was sourced from the OptionMetrics via Wharton Research Data Services (WRDS) database covering 1st January 2013 to 31st August 2023 collected, ticker, options volume, Call/Put Flag and however the rest of days, due to that is the latest update from WRDS is up to 31st August 2023, so the rest of dataset that needs to be covered for the trading dataset of the Put-Call Ratio were directly collected by using PUT_CALL_VOLUME_RATIO_CUR_DAY function in Bloomberg Terminal from 1st September to 31st December 2023 which covers the rest of trading period.

The Put-Call Ratio is an assessment of put and call option transaction volumes that provides insight into market participants' perspectives and outlook. The higher proportion

979-8-3315-9661-3/25/$31.00 ©2025 IEEE
DOI 10.1109/IDS66066.2025.00020

typically represents bearish sentiment among investors who purchase more put options in expectation of downward movements. In contrast, the lower ratio stands for bullish views as traders favor call options anticipating an ascent. Specifically, the magnitude of measured put trading activity versus call trading activity was derived through the calculation in formula 1.

$$PutCall\ Ratio = \frac{Put\ Volume}{Call\ Volume} \tag{1}$$

By integrating the Put-Call Ratio with traditional technical indicators, such as MACD, Bollinger Bands, RSI, CCI, DX, SMA, VIX, and Turbulence, we trained both PPO and CPPO models, along with the NASDAQ-100 index, to evaluate the model's respective performances. For the missing value, we keep it as N/A rather than imputation; the summary statistics are shown in Table 1.

The PPO model was trained by using a two-layer neural network (512, 512) complied with ReLU activation. And the key hyperparameters included a discount factor of γ=0.995, learning rates of 3e-5 (for policy) and 1e-4 (for value), 100 epochs with 20,000 steps per epoch, and a clipping ratio of 0.7. Reward scaling (1e-4) and a transaction cost of 0.1% per trade were applied to stabilize training and simulate realistic market conditions.

TABLE I. SUMMARY STATISTICS OF ADDED INDICATORS

Statistic	VIX	Turbulence	Put-Call Ratio
count	232,344	232,344	133,156
mean	17.77	600,316.66	1.33
std	7.08	31,559,185.86	10.30
min	9.14	-	-
0.25 (Q1)	13.13	67.98	0.35
0.5 (Median)	15.74	104.74	0.66
0.75 (Q3)	20.65	165.70	1.07
max	82.69	1,660,081,851.38	2,011.13

III. EXPERIMENT AND RESULT

Our findings indicate significant improvements if the Put-Call Ratio is included in training. It achieved superior Sharpe, Sortino, and Information Ratios compared to CPPO and the NASDAQ-100 benchmark in the trading set. The PPO strategy provides notable risk-adjusted returns with a Sharpe Ratio of 1.08. This significantly outperformed the performance of CPPO, 0.89, and the performance of the NASDAQ-100 index benchmark, 0.69. In addition, Sortino ratios are 1.86, 1.41, and 1.09, respectively, with PPO having the highest. These results can be found in Figure 1.

Additionally, the cumulative return analysis from January 2020 through December 2023 demonstrated that the complex PPO model incorporating the fluctuating Put-Call Ratio significantly outperformed other less intricate DeepSeek-based PPO variants employing differing sentiment thresholds, as minute as 10%, 1%, and 0.1%. This proposed the shifting Put-Call Ratio's preferable prognostic aptitude for investment selections, as manifested in the intricate Figure 2.

To ensure consistent comparison across strategies, we recalibrated PPO, CPPO, and the criterion succession utilizing the intersection of internal exchange dates and Yahoo Finance statistics. After reindexing with the standard dates, the multifaceted PPO also displayed the most elevated returns, 256.97%, compared to the less involved CPPO, 215.80%, and the criterion NASDAQ-100, 177.54%, which can be perceived in Figure 3.

These results confirm that PPO augmented with Put-Call Ratio indicators offers a more robust and profitable trading strategy in volatile market environments.

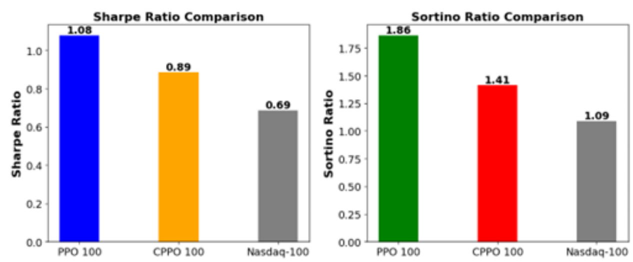

Fig. 1. Sharpe and Sortino ratios for PPO, CPPO and Benchmark.

Fig. 2. Comparison between Put-Call PPO and other variants.

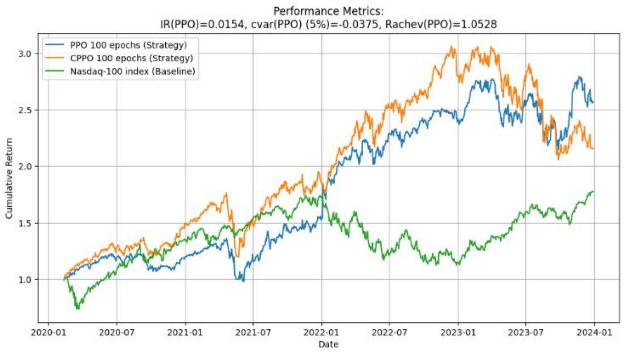

Fig. 3. Comparison between PPO, CPPO and Benchmark.

IV. CONCLUSION AND FUTURE WORK

Overall, our results demonstrate that market sentiment variables, such as the Put-Call Ratio, are essential for enabling intelligent agents to construct high-alpha portfolios. Including an option-based sentimental indicator significantly improves the predictive performance of the PPO trading agent. This finding is consistent with previous studies [9][10]. Moreover, compared to CPPO and the NASDAQ-100 benchmark index, the PPO model exhibits the most potent ability to capture market movements, particularly during complex periods like the COVID-19 crisis. Our findings provide practical insights and contributions to retail investors and fund managers in the investment portfolio construction process.

Future work could explore high-frequency sentimental data from social media platforms such as X (formerly Twitter) and Facebook, to further investigate how the high-frequency sentimental indicators impact the prediction performance of trading agents. Moreover, our current study still has space for comparing additional baseline reinforcement learning models such as Deep Q-Network (DQN) and Advantage Actor-Critic (A2C). Both of these two reinforcement learning algorithms have been discussed that have the potential to improve the performance in tasks of portfolio management [11][12]. Including alternative Reinforcement Learning-based trading agents to evaluate the predictive performance in multi-asset or cross-market settings can be a future research direction.

V. ACKNOWLEDGEMENT

We want to express our sincere gratitude to Dr. Olena Onishchenko for her invaluable instruction and guidance. She encouraged us to participate in the FinRL Contest 2025 and helped us explore potential indicators to identify the key indicator in our model. We also gratefully acknowledge that Otago Business School provides data access and financial support for the conference registration fee.

REFERENCES

[1] V. M. Ngo, H. H. Nguyen, and P. Van Nguyen, "Does reinforcement learning outperform deep learning and traditional portfolio optimization models in frontier and developed financial markets?" *Research in International Business and Finance*, vol. 65, p. 101936, Apr. 2023, doi: https://doi.org/10.1016/j.ribaf.2023.101936.

[2] W. Junfeng, L. Yaoming, T. Wenqing, and C. Yun, "Portfolio management based on a reinforcement learning framework," *Journal of Forecasting*, vol. 43, no. 7, pp. 2792–2808, May 2024, doi: https://doi.org/10.1002/for.3155.

[3] X.-Y. Liu *et al.*, "FinRL: A Deep Reinforcement Learning Library for Automated Stock Trading in Quantitative Finance," *arXiv:2011.09607 [cs, q-fin]*, Nov. 2020, Available: https://arxiv.org/abs/2011.09607

[4] K. Wang, K. Xiao, and Y. X.-Y. Liu, "Parallel Market Environments for FinRL Contests," *arXiv.org*, 2025. https://arxiv.org/abs/2504.02281 (accessed Apr. 11, 2025).

[5] M. E. Zweig, Winning on Wall Street. New York: Warner Books, 1994.

[6] I. Diaz-Rainey, S. A. Gehricke, H. Roberts, and R. Zhang, "Trump vs. Paris: The impact of climate policy on U.S. listed oil and gas firm returns and volatility," *International Review of Financial Analysis*, vol. 76, p. 101746, Jul. 2021, doi: https://doi.org/10.1016/j.irfa.2021.101746.

[7] H. Yang, X.-Y. Liu, S. Zhong, and A. Walid, "Deep reinforcement learning for automated stock trading," *Proceedings of the First ACM International Conference on AI in Finance*, Oct. 2020, doi: https://doi.org/10.1145/3383455.3422540.

[8] M. Benhenda, "FinRL-DeepSeek: LLM-infused risk-sensitive reinforcement learning for trading agents," arXiv preprint arXiv:2502.07393, Feb. 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2502.07393

[9] K. Kim, D. Ryu, and J. Yu, "Is a sentiment-based trading strategy profitable?" *Investment Analysts Journal*, vol. 51, no. 2, pp. 94–107, 2022. doi: 10.1080/10293523.2022.2076373.

[10] P. Houlihan and G. G. Creamer, "Can sentiment analysis and options volume anticipate future returns?" *Computation Economics*, vol. 50, pp. 669–685, 2017.

[11] Z. Gao, Y. Gao, Y. Hu, Z. Jiang, and J. Su, "Application of Deep Q-Network in Portfolio Management," *IEEE Xplore*, May 01, 2020. https://ieeexplore.ieee.org/abstract/document/9101333/ (accessed Nov. 30, 2022).

[12] M. Rezaei and H. Nezamabadi-Pour, "A taxonomy of literature reviews and experimental study of deepreinforcement learning in portfolio management," *Artificial Intelligence Review*, vol. 58, no. 3, Jan. 2025, doi: https://doi.org/10.1007/s10462-024-11066-w.