Enhancing FinRL Trading Agents with Advanced LLM-Processed Financial News: An Improved Approach Using DeepSeek-V3

Satish Chandra
PhD Scholar
Department of Computer Science and Engineering, Anurag University
Hyderabad, India
gullenasatish@gmail.com

Dr. G. Balakrishna
Associate Professor
Department of Computer Science and Engineering, Anurag University
Hyderabad, India
balakrishnacse@anurag.edu.in

***Abstract*—The use of AI techniques to improve stock market related operations has been benefiting traders in the stock market. FinRL is a framework that has been shown to be effective in building RL Trading Agents. Use of LLMs for performance improvement of the Agents is one of the recent initiatives. LLMs are being used to generate signals that could be combined with historical stock prices for better prediction of stock prices movement and profitable trading. This paper endeavours to develop efficiently engineered prompts and generate better signals to improve the performance of RL Trading Agents.**

Keywords*—Financial Reinforcement Learning, FinRL, trading agent, prompt, large language model, DeepSeek, stock market, financial news, stock price trend

I. Introduction

Traditionally, prediction of trends of stock prices had been based on features like Open, High, Low and Close prices. A recent trend has been to consider alternate features like sentiments. However, with the rising awareness of the importance of adding alternate data for stock price trends prediction several researchers have developed models adding alternate data, like sentiments of market participants. One such approach is to build Reinforcement Learning Trading Agents trained using various features. Mostapha Benhenda has developed a framework for building such Reinforcement Trading Agents[1]. The framework leverages the Stock Recommendation Prompt developed by Zihan Dong in [2]. This study performs sentiment analysis by processing financial news using advanced prompting approaches with DeepSeek-V3 to provide prompt based LLM signals for training FinRL trading agents. This is an improved approach for using DeepSeek-V3 to generate sentiment signals.

II. Related Works and Research Question

A. FinRL Trading Agents

Reinforcement Learning enables learning of effective trading strategies. The AI4Finance Foundation is the world's largest AI finance open-source community. The Foundation has hosted FinRL, which is the first open-source framework for financial reinforcement learning.

FinRL has three layers: market environments, agents, and applications (figure 1). For a trading task (on the top), an agent (in the middle) interacts with a market environment (at the bottom), making sequential decisions [7].

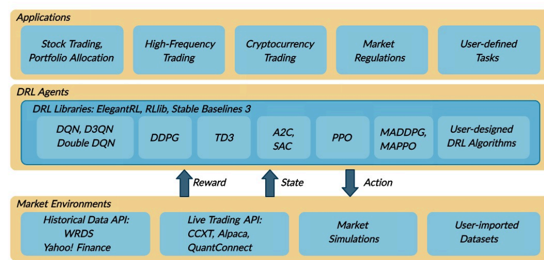

Fig. 1. FinRL Framework

B. Large Language Models for Reinforcement Learning of Trading Agents

One of the recent innovations in the training approach is to include LLM based signals for features such as sentiment and risk [1], [2] for training FinRL based trading agents.

C. Proposed Research Question

The generation of signals from financial text is one of the key challenges identified in [10].

The proposed research question is "Can advanced prompt engineering and novel integration strategies for DeepSeek-V3–processed financial news signals improve the performance of FinRL-based trading agents compared to the baseline FinRL-DeepSeek approach?"

III. Methodology

A. Dataset Selection and Preprocessing

As the project of [1] is the basis of task 1 in FinRL Contest 2025, this paper uses the datasets and framework of [1] and [2], from [8] and [9]. We have used data like 'sentiment_nasdaq_news.csv' , 'trade_data_2019_2023.csv' in generating sentiment signals from DeepSeek-V3.

B. Model Selection

As DeepSeek-V3 outperformed other LLMs [1] we have considered DeepSeek-V3 as the model for generating sentiment based trading signals.

979-8-3315-9661-3/25/$31.00 ©2025 IEEE
DOI 10.1109/IDS66066.2025.00016

C. Prompt Engineering for DeepSeek-V3

{Role-based Prompting has been used in [1] and [2]. We propose to develop and test prompts based on Few-shot prompting, Chain-of-Thought reasoning, Self-Consistency and Counterfactual prompting to generate signals using the DeepSeek-V3. Apart from seeing which prompts perform better we would also like to explore the explainability aspects of the prompts.

Algorithm: Sentiment Signal Scoring }

D. Market Environment

Keyi Wang, et. al. [10] have discussed the importance of high-quality market environments for developing robust trading agents. The signals generated from DeepSeek-V3 need the right market environment. The stable market based environment, detailed in [10] is our environment of choice.

E. Evaluation Metrics

We have used Sharpe ratio and other metrics, as discussed later in this paper.

IV. Experiment and results

To address the sampling bottleneck during training, [10] introduced GPU-based parallel market environments to improve sampling speed. However, for the time being we are not using any GPU based infrastructure and the current focus is on utilizing advanced prompting techniques for better results.

The DeepSeek-V3 model would be fed with various types of prompts to generate signals that would be combined with the historical prices data and then used for the learning of the RL Trading Agents. Data at [8] and [9] would be used to generate the signals. However, as we are facing issues due to the large size of the datasets and have commenced using smaller but relevant datasets.

The code in [6] would be reused to the extent possible by carrying out suitable modifications that would be required due to the introduction of newly designed prompts as well as considerations on account of different/additional datasets, if required.

At a high level, the sentiment scores generated by the DeepSeek-V3 would be combined with the stock data and relevant model training pipelines outlined in [6] would be reused/modified. As we are considering only the DeepSeek-V3, only the code related to DeepSeek would be reused/modified and the code pertaining to other LLMs would be ignored. Also, the stock trading environment python code would be suitably modified. The program 'env_stocktrading_llm_risk.py' gives high importance to LLM signals and is a high risk profile scenario emphasizing the sentimental impact. We propose to modify this program to begin with. We populate the llm_sentiment_col using various types of prompts and adjust trading actions based on sentiment scores (1–5: strong sell to strong buy) to perform a comparative analysis and see which of the prompting techniques is the most suitable for generating efficient results.

Evaluation would be done by using the backtesting code in [6] by carrying out modifications to address only the affected paths in the code.

The key steps in the workflow for the experiment include:

1. Data preparation: using OHLCV data along with the signals generated by the DeepSeek LLM
2. Training
3. Backtesting

V. Conclusion and Future Work

We expect to find performance improvements due to the use of different types of prompts. However, there is an erratic behaviour in the different runs of the code which we are fixing.

For future work, we could consider additional dimensions like explainability analysis, bias detection and fine-tuning DeepSeek through domain adaptation

Acknowledgment

We would like to thank the AI4Finance Foundation for the pioneering initiative of development and maintenance of FinRL framework in addition to more such finance related frameworks/tools. We would like to thank Mostapha for providing the baseline framework. Further, we would like to thank all the authors listed in the references section and beyond whose contributions have led to the use of AI in finance, especially the ones related to stock trading.

References

[1] Mostapha Benhenda, "FinRL-DeepSeek: LLM Infused Risk-Sensitive Reinforcement Learning for Trading Agents," arXiv:2502.07393v1.

[2] Zihan Dong, et al., "FNSPID: A Comprehensive Financial News Dataset in Time Series," arXiv:2402.06698v1.

[3] Xiao-Yang Liu, "FinRL: A Deep Reinforcement Learning Library for Automated Stock Trading in Quantitative Finance," arXiv:2011.09607v2.

[4] John Schulman, et. al., "Proximal Policy Optimization Algorithms," arXiv:1707.06347v2.

[5] Xiao-Yang Liu, et. al., "Practical Deep Reinforcement Learning Approach for Stock Trading," arXiv:1811.07522v3.

[6] https://github.com/benstaf/FinRL_DeepSeek.

[7] https://github.com/AI4Finance-Foundation/FinRL.

[8] https://huggingface.co/datasets/Zihan1004/FNSPID/tree/main/Stock_news.

[9] https://huggingface.co/benstaf .

[10] Keyi Wang, Kairong Xiao and Xiao-Yang Liu Yanglet. Parallel Market Environments for FinRL Contests. arXiv preprint arXiv:2504.02281.

[11] X.-Y.Liu ,Z.Xia,J.Rui,J.Gao,H.Yang,M.Zhu,C.D.Wang,Z.Wang, and J. Guo, "FinRL-meta:market environments and benchmarks for data-driven financial reinforcement learning," in International Conference on Neural InformationProcessingSystems,NY,USA,2022.

[12] Antonin Raffin, Ashley Hill, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, and Noah Dormann. Stable baselines3. https://github.com/DLR-RM/stable-baselines3, 2019.

[13] Arnav Grover. FinRLlama: A Solution to LLM-Engineered Signals Challenge at FinRL Contest 2024. Conference'24, November 2024, NYC, NY, USA.