



**HAL**  
open science

# FinRL-DeepSeek: LLM-Infused Risk-Sensitive Reinforcement Learning for Trading Agents

Mostapha Benhenda

► **To cite this version:**

Mostapha Benhenda. FinRL-DeepSeek: LLM-Infused Risk-Sensitive Reinforcement Learning for Trading Agents. 2025. hal-04934770

**HAL Id: hal-04934770**

**<https://hal.science/hal-04934770v1>**

Preprint submitted on 10 Feb 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# FinRL-DeepSeek: LLM-Infused Risk-Sensitive Reinforcement Learning for Trading Agents

Mostapha Benhenda

mostaphabhenda@gmail.com

## Abstract

This paper presents a novel risk-sensitive trading agent combining reinforcement learning and large language models (LLMs). We extend the Conditional Value-at-Risk Proximal Policy Optimization (CPPO) algorithm, by adding risk assessment and trading recommendation signals generated by a LLM from financial news. Our approach is backtested on the Nasdaq-100 index benchmark, using financial news data from the FNSPID dataset and the DeepSeek V3, Qwen 2.5 and Llama 3.3 language models. The code, data, and trading agents are available at: [https://github.com/benstaf/FinRL\\_DeepSeek](https://github.com/benstaf/FinRL_DeepSeek)

## 1. Introduction

Automated trading agents increasingly leverage reinforcement learning (RL), but often overlook two aspects:

1. Integration of alternative data sources, such as financial news.
2. Risk management.

This paper introduces a novel hybrid RL-LLM trading agent, incorporating financial news insights at both the action and risk levels. The main contribution of the paper is the introduction of an LLM-based risk assessment score derived from news, in addition to LLM-based trading recommendations. This demonstrates the potential of LLMs for feature extraction from news beyond standard sentiment analysis, achieved through carefully designed prompts.

## 2. Related Work

There is a growing body of work on LLM-augmented RL agents, such as (Unnikrishnan, 2024; Anonymous, 2025). Compared with them, the way LLMs influence the RL agent is simpler here.

In FinGPT (Yang et al., 2023), RL is applied to the training of LLM on stock prices. This approach is more complex

than ours, which has the benefit of leveraging established LLM APIs.

Other hybrid RL-LLM approaches include FinCon (Yu et al., 2024), which makes a complex distillation of information through a synthesized multi-agent collaboration mechanism. In this paper, we only use simple prompts.

These RL-LLM approaches also differ from "pure LLM" agents like (Lopez-Lira and Tang, 2024) and (Ding et al., 2024) that only use LLM recommendations.

## 3. Data and LLM Prompts

We use the FNSPID dataset (Dong et al., 2024), which contains 15.7 million time-aligned financial news records spanning 1999–2023. To reduce LLM API costs, we randomly select a representative news article per stock per day, reducing the dataset to 2 million records. Using LLMs (DeepSeek V3 (DeepSeek-AI et al., 2024), Qwen 2.5 72B (Qwen et al., 2025) Llama 3.3 70B (Grattafiori et al., 2024)), we extract stock recommendations and risk assessment using the following prompts :

**Stock Recommendation Prompt (from (Dong et al., 2024)):**

"You are a financial expert with stock recommendation experience. Based on a specific stock, score for range from 1 to 5, where 1 is negative, 2 is somewhat negative, 3 is neutral, 4 is somewhat positive, 5 is positive"

**Risk Assessment Prompt (new):**

"You are a financial expert specializing in risk assessment for stock recommendations. Based on a specific stock, provide a risk score from 1 to 5, where: 1 indicates very low risk, 2 indicates low risk, 3 indicates moderate risk (default if the news lacks any clear indication of risk), 4 indicates high risk, and 5 indicates very high risk."

## 4. Trading Algorithms

### 4.1. Reinforcement Learning agents on price-only data

#### 4.1.1. PROXIMAL POLICY OPTIMIZATION (PPO)

OpenAI (Schulman et al., 2017) introduced Proximal Policy Optimization (PPO) reinforcement learning algorithm, and (Lele et al., 2020; Liu et al., 2022) applied it to stock trading

The PPO objective ensures stable policy updates by clipping the probability ratio:

$$L_{\text{PPO}}(\theta) = \mathbb{E} [\min (r_t(\theta) \cdot A_t, \text{clip} (r_t(\theta), 1 - \epsilon, 1 + \epsilon) \cdot A_t)]$$

where:

- $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$  is the probability ratio between the new policy  $\pi_\theta$  and the old policy  $\pi_{\theta_{\text{old}}}$ .
- $A_t$  is the advantage estimate at time  $t$ .
- $\epsilon$  is the clipping parameter that restricts large updates to the policy.

#### 4.1.2. CONDITIONAL VALUE AT RISK-PROXIMAL POLICY OPTIMIZATION (CVAR-PPO)

Conditional Value at Risk-Proximal Policy Optimization (CVaR-PPO) (Ying et al., 2022) extends PPO with a risk constraint to penalize high-loss trajectories:

$$L_{\text{CVaR-PPO}}(\theta, \eta, \lambda) = L_{\text{PPO}}(\theta) + \lambda \left( \frac{1}{1 - \alpha} \mathbb{E}[(\eta - D(\pi_\theta))^+] - \eta + \beta \right) \quad (1)$$

where:

- $L_{\text{PPO}}(\theta)$  is the PPO objective.
- $D(\pi_\theta)$  is the trajectory return.
- $\eta$  is the CVaR threshold.
- $(\eta - D(\pi_\theta))^+ = \max(0, \eta - D(\pi_\theta))$  represents the CVaR loss beyond the threshold.
- $\lambda$  is the Lagrange multiplier enforcing the constraint.
- $\alpha$  is the CVaR confidence level (e.g., 0.05 for the worst 5%).
- $\beta$  is an auxiliary penalty parameter.

To our knowledge, we are the first to apply CVaR-PPO to stock trading.

### 4.2. LLM-infused PPO

In LLM-infused PPO, we compute stock-specific recommendation scores ( $S_f$ ) from the FNSPID dataset using an LLM. These scores influence trading actions. The influenced action is:

$$a_t^{\text{mod}} = S_f \cdot a_t$$

where:

- $S_f > 1$ : Amplifies actions under positive recommendation.
- $S_f < 1$ : Dampens actions under negative recommendation.
- $S_f = 1$ : Leaves actions unchanged.

This perturbation  $S_f$  is chosen close to 1 to keep stability of the algorithm. For example:

- $S_f = 1.1$  if stock score is 5 and  $a_t > 0$  or score is 1 and  $a_t < 0$
- $S_f = 1.05$  if stock score is 4 and  $a_t > 0$  or score is 2 and  $a_t < 0$
- $S_f = 0.95$  if stock score is 4 and  $a_t < 0$  or score is 2 and  $a_t > 0$
- $S_f = 0.9$  if stock score is 5 and  $a_t < 0$  or score is 1 and  $a_t > 0$

### 4.3. LLM-infused CVaR-PPO(CPPO)

Financial news from the FNSPID dataset is processed with an LLM to generate risk scores ( $R_f^i$ ) for each stock  $i$  and day.

The perturbation  $R_f^i$  is also chosen close to 1 to keep stability of the algorithm. For example:

- $R_f^i = 1.1$  for very high risk score 5
- $R_f^i = 1.05$  for high risk score 4
- $R_f^i = 1$  for moderate risk score 3
- $R_f^i = 0.95$  for low risk score 2
- $R_f^i = 0.9$  for very low risk score 1

The aggregate risk score  $R_f$  is defined as:

$$R_f = \sum_i w_i R_f^i$$

where  $w_i$  is the financial weight of stock  $i$  in the portfolio (we have:  $\sum_i w_i = 1$ ).

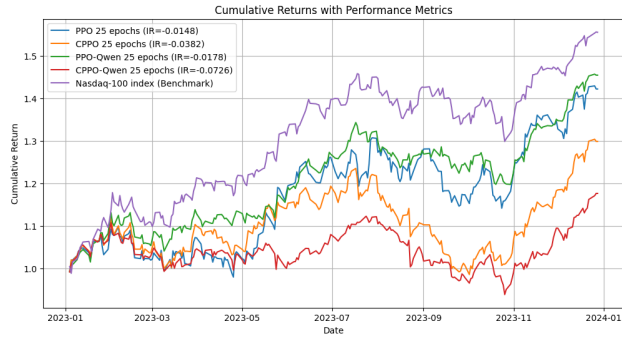


Figure 1. Backtesting after 500k training steps, 3 years training history, 1 year trading

These scores adjust trajectory returns in Cvar PPO to reflect market risk. The adjusted return is defined as:

$$D_{R_f}(\pi_\theta) = R_f \cdot D(\pi_\theta)$$

At the end, financial news influence trading actions in two ways: through  $S_f$  and  $R_f$ .

## 5. Results

### 5.1. Early Stopping: 400-500k training steps

In these tests, we use a 10% infusion of LLM, which means that the original PPO and CPPO up to 10%

#### 5.1.1. FIGURE 1:

**Training Period:** 2019–2022

**Training Iterations:** 500k steps (25 epochs, 20k steps each)

**Trading Period:** 2023

These preliminary results in figure 1 with Qwen 2.5 show that integrating LLM stock recommendations consistently improves PPO in terms of cumulative returns. However, it's not sufficient to beat Nasdaq 100 benchmark so far.

#### 5.1.2. FIGURE 2:

**Training Period:** 2013–2018

**Training Iterations:** 400k steps (20 epochs, 20k steps each)

**Trading Period:** 2019-2023

Results show significant improvement when taking longer training period (6 years vs. 3 years) for PPO and CPPO. However, PPO remains too volatile.

DeepSeek V3 is slightly better than Llama 3.3 all other things being equal.

The use of LLM always worsens performance in this test.

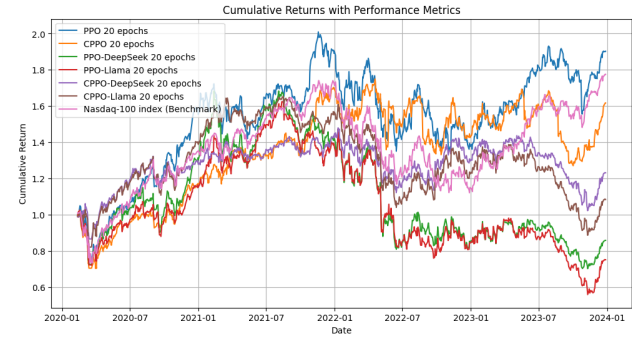


Figure 2. Backtesting after 400k training steps, 6 years training history, 3 years trading

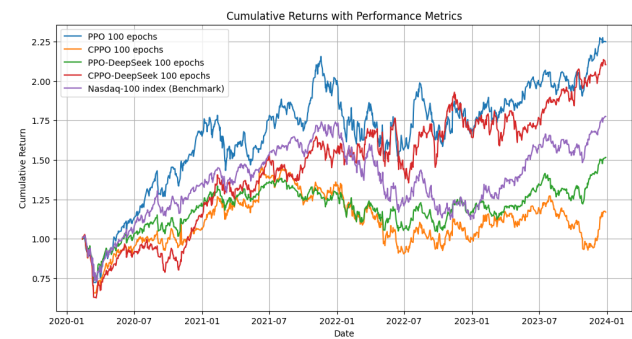


Figure 3. After training for 2 Million steps (100 epochs, 20k steps each)

### 5.2. After training for 2 Million steps

#### 5.2.1. FIGURE 3:

Model	Information Ratio	CVaR	Rachev Ratio
PPO (100 epochs)	0.0100	-0.0394	1.0637
CPPO (100 epochs)	-0.0148	-0.0439	1.0404
PPO-DeepSeek (100 epochs)	-0.0093	-0.0338	0.9890
CPPO-DeepSeek (100 epochs)	0.0078	-0.0437	0.9818

Table 1. Performance metrics for different models over 100 epochs.

#### 5.2.2. FIGURE 4:

In these two runs (RL agents are stochastic), PPO and CPPO-DeepSeek outperform other methods, including the Nasdaq 100 benchmark.

PPO seems to perform better in Bull markets, and CPPO-DeepSeek in Bear markets, the transition being at the end of 2021 before the war in Ukraine and the crisis that followed.

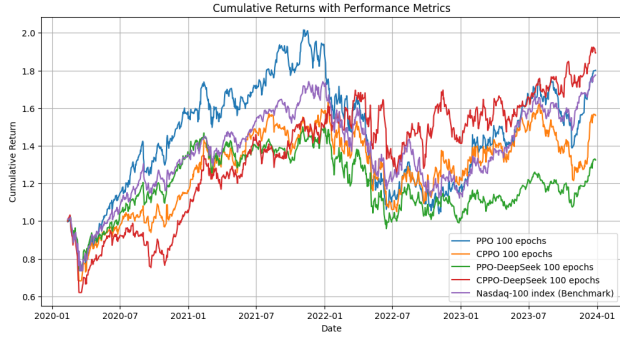


Figure 4. After training for 2 Million steps (100 epochs, 20k steps each) - Second run

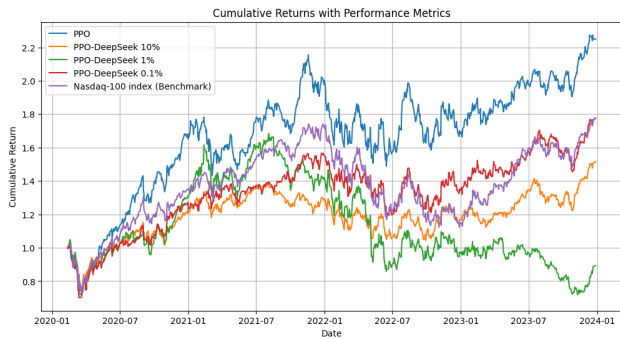


Figure 5. Impact of LLM infusion on PPO

### 5.3. Impact of LLM infusion strength

In the following tests, we tweak LLM infusion strength parameters from 10% down to 0.1% (i.e. LLM perturbation parameters vary from 0.9 – 1.1 to 0.999 – 1.001)

#### 5.3.1. FIGURE 5:

Model	Information Ratio	CVaR	Rachev Ratio
PPO	0.0100	-0.0394	1.0637
PPO-DeepSeek 10%	-0.0093	-0.0338	0.9890
PPO-DeepSeek 1%	-0.0252	-0.0459	1.0394
PPO-DeepSeek 0.1%	-0.0011	-0.0375	0.9536

Table 2. Performance metrics for PPO and PPO-DeepSeek variants.

For PPO-DeepSeek, stronger LLM infusion mostly degrades performance, even for tiny (0.1%) perturbations.

#### 5.3.2. FIGURE 6:

For CPPO-DeepSeek, stronger LLM infusion improves performance

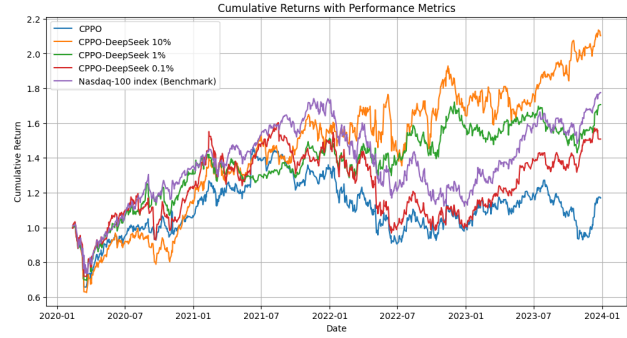


Figure 6. Impact of LLM infusion on CPPO

Model	Information Ratio	CVaR	Rachev Ratio
CPPO	-0.0148	-0.0439	1.0404
CPPO-DeepSeek 10%	0.0078	-0.0437	0.9818
CPPO-DeepSeek 1%	-0.0032	-0.0365	0.9573
CPPO-DeepSeek 0.1%	-0.0060	-0.0441	0.9789

Table 3. Performance metrics for CPPO and CPPO-DeepSeek variants.

## 6. Conclusion

We propose LLM-infused RL agents for algorithmic trading that integrates stock trading recommendations and risk assessments from news. Future work includes the following directions:

- **Optimizing RAM usage:** Longer training requires more RAM. For instance, 500k training steps used 16 GB (+Swap memory), whereas 2 million steps required 128 GB (+Swap memory). Optimizing memory efficiency is crucial for scalability.
- **Reducing timescale:** Since markets react quickly and abruptly to events, reducing the timescale of decision-making could improve trading performance.
- **Improving news signal quality:** Enhancing the quality of news signals from the FNSPID dataset is an important direction to ensure better market performance.

## References

- Anonymous. Advancing algorithmic trading with large language models: A reinforcement learning approach for stock market optimization. <https://openreview.net/forum?id=w7BGq6ozOL>, 2025.
- DeepSeek-AI, A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. URL <https://arxiv.org/abs/2412.19437>.

- 
- Qianggang Ding, Haochen Shi, and Bang Liu. Tradexpert: Revolutionizing trading with mixture of expert llms. *arXiv preprint arXiv:2411.00782*, 2024. URL <https://arxiv.org/abs/2411.00782>.
- Zihan Dong, Xinyu Fan, and Zhiyuan Peng. Fnspid: A comprehensive financial news dataset in time series. *arXiv preprint arXiv:2402.06698*, 2024. URL <https://arxiv.org/abs/2402.06698>.
- A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Shreyas Lele, Kavitha Gangar, Harshal Daftary, and Dewashish Dharkar. Stock market trading agent using on-policy reinforcement learning algorithms. <https://ssrn.com/abstract=3582014>, 2020.
- Xiao-Yang Liu, Hongyang Yang, Qian Chen, Rui Zhang, Liuqing Yang, Bowen Xiao, and Christina Dan Wang. Finrl: A deep reinforcement learning library for automated stock trading in quantitative finance. *arXiv preprint arXiv:2011.09607*, 2022. URL <https://arxiv.org/abs/2011.09607>.
- Alejandro Lopez-Lira and Yuehua Tang. Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*, 2024. URL <https://arxiv.org/abs/2304.07619>.
- A. Qwen, A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, et al. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2025. URL <https://arxiv.org/abs/2412.15115>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. URL <https://arxiv.org/abs/1707.06347>.
- A. Unnikrishnan. Financial news-driven llm reinforcement learning for portfolio management. *arXiv preprint arXiv:2411.11059*, 2024. URL <https://arxiv.org/abs/2411.11059>.
- Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. Fingpt: Open-source financial large language models. *arXiv preprint arXiv:2306.06031*, 2023. URL <https://arxiv.org/abs/2306.06031>.
- ChengYang Ying, Xinning Zhou, Hang Su, Dong Yan, Ning Chen, and Jun Zhu. Towards safe reinforcement learning via constraining conditional value-at-risk. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 3673–3680. International Joint Conferences on Artificial Intelligence Organization, 2022. doi: 10.24963/ijcai.2022/510.
- Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yupeng Cao, Zhi Chen, et al. Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. *arXiv preprint arXiv:2407.06567*, 2024. URL <https://arxiv.org/abs/2407.06567>.