

FINZERO: LAUNCHING MULTI-MODAL FINANCIAL TIME SERIES FORECASTING WITH LARGE REASONING MODEL

Yanlong Wang^{1,2,*}, Jian Xu¹, Fei Ma³, Hongkang Zhang¹, Hang Yu⁴, Tiantian Gao¹, Yu Wang^{2,6}
Haochen You⁵, Shao-Lun Huang^{1†}, Danny Dongning Sun^{2†}, Xiao-Ping Zhang^{1†}

¹ Tsinghua University ² Pengcheng Laboratory ³ Guangming Laboratory ⁴ Ant Group
⁵ Columbia University ⁶ Southern University of Science and Technology

ABSTRACT

Financial time series forecasting is both highly significant and challenging. Previous approaches typically standardized time series data before feeding it into forecasting models, which inherently leads to a loss of important information during the encoding process. Moreover, past time series models generally require fixed numbers of variables or look-back window lengths, which further limits the scalability of time series forecasting. In addition, interpretability and uncertainty in forecasting remain underexplored, as these factors directly impact the reliability and practical value of predictions. To address these issues, we first constructed a diverse financial image-text dataset (FVLDB) and developed the Uncertainty-adjusted Group Relative Policy Optimization (UARPO) method to enable the model not only output predictions but also analyze the uncertainty of those predictions. We then launched FinZero, a multi-modal pre-trained model finetuned by UARPO to perform reasoning, prediction, and analytical understanding on the FVLDB financial time series. Extensive experiments validate that FinZero exhibits strong adaptability and scalability. After fine-tuning with UARPO, FinZero achieves an approximate 13.48% improvement in prediction accuracy over GPT-4o in the high-confidence group, demonstrating the effectiveness of reinforcement learning fine-tuning of multi-modal large model for financial time series forecasting tasks.

Index Terms— Financial Time Series, Reinforced Fine-tuning, Uncertainty Quantification, Reasoning

1. INTRODUCTION

The field of time series forecasting has garnered increasing attention [1, 2, 3], as time-series data is widely present in various real-world industries (e.g., transportation, weather, power, finance, etc.). Extracting future trends from historical information holds significant practical value. Among these, financial time series exhibit more distinctive characteristics

as they are influenced by more complex factors [4, 5]. For instance, the asset price movements are shaped by a broad range of external macro- and micro-level influences [6, 7], as well as the interplay between buyers and sellers in determining transaction prices [8, 9, 10]. This implies that, in such a game-theoretic environment, any discernible patterns or identifiable features (e.g., the pronounced periodicity seen in transportation or power time series) tend to diminish once traders recognize and exploit them for profit. This "adaptive" nature of markets leads to the inability of historical patterns to fully replicate in the future. Predicting such time series is undoubtedly highly challenging. However, even marginal improvements in forecasting performance can yield substantial impacts, particularly in high-frequency trading scenarios [11, 12, 13].

To improve time-series forecasting performance, including financial time series such as exchange rate prediction, specialized models have been designed [14, 15, 16]. However, several challenges remain unresolved. **First**, most current time series models require standardization to transform the data into a numerical range that the model can process, such as the normalization techniques used in RevIN [17]. This inevitably leads to the loss of partial information from the original values. **Second**, patching-based processing is commonly adopted, but this may not fully align with the size and location of critical features [18]. **Third**, time series models usually operate with several fixed configurations, such as look-back window size, variable types and quantities, and data frequency [19, 16], which significantly limit the model generalization. Although some large foundation time series models [20, 21, 22, 23] have partially addressed these issues, the advanced reasoning capabilities of large models have not been fully leveraged in time series applications. **Besides**, the interpretability of reasoning and uncertainty quantification in forecasting results remain critical yet understudied challenges.

To address the aforementioned challenges, we have abandoned traditional model architectures that process raw time series values and instead transformed the original time series into image compositions. Leveraging reinforcement learning fine-tuning development [24, 25, 26, 27, 28, 29, 30], we en-

* Work done during an internship at Ant Group.

† Corresponding author.




Fig. 1. The overall pipeline of the FinZero model, which is fine-tuned via the UARPO method on the FVLDB.

hance the visual reasoning capabilities of multimodal large model (MLM). Our focus is on financial time-series trend prediction and reasoning tasks. To better support this, we constructed the FVLDB dataset, comprising over 10,000 financial time series image-text pairs. To ensure dataset diversity, we performed stratified sampling across multiple dimensions, including asset types, prediction task categories, historical sequence lengths and frequencies, time series indicator varieties, and image styles. To tackle the inherent uncertainty and non-stationarity in financial time series forecasting, we propose the Uncertainty-Adjusted Relative Policy Optimization (UARPO) method. UARPO evaluates both intra-group relative advantage (IGRA, performance within a group) and cross-group relative advantage (CGRA, performance between groups over a recent window). Additionally, it adjusts advantage levels based on prediction uncertainty (Uncertainty-Adjusted Relative Advantage, UARA).

In this work, we propose the FinZero model, as illustrated in Figure 1, which fine-tunes 3B-parameter multimodal large model via the UARPO method in the FVLDB dataset, which enables MLM to explicitly account for prediction uncertainty. Comparative experiments with GPT-4 show a 13.48% improvement in prediction accuracy in the high-confidence group, validating the effectiveness of RL-based cross-modal fine-tuning for financial time-series forecasting and reasoning. By providing confidence score and reasoning traces, FinZero helps users better understand model predictions and their rationale, ultimately supporting more informed financial decision-making, making it particularly valuable for real-world financial applications where risk assessment is paramount.

2. METHODS

2.1. Uncertainty Adjusted Related Policy Optimization

The GRPO [31] is employed to fine-tune the DeepSeek-R1 [32]. As an improvement over the PPO [33], GRPO eliminates the need for an additional model as a policy model (as required by methods like PPO) and leverages Group Relative

Advantage sampled from multiple outputs within a group, thereby avoiding the necessity for extra value function approximation. GRPO primarily focuses on the relative advantages among multiple outputs within each sample group, while other methods like REINFORCE++ [34] utilize discounted cumulative rewards to construct advantage variations that reflect the training process, which helps improve training stability. Additionally, how to reflect the uncertainty in model inference results holds significant importance, as it aids decision-making by assessing the confidence level of reasoning outcomes.

Based on the above, we propose the UARPO algorithm, which introduces two key enhancements: (1) Under the same prediction target, a multi-dimensional advantage function combining In-Group Relative Advantage (IGRA) within samples and Cross-Group Relative Advantage (CGRA) across groups; (2) Construction of an uncertainty function (UA) based on the model’s inference confidence scores, ultimately forming Uncertainty-Adjusted Relative Advantage (UARA). The optimization objective can be expressed as Equation (1):

$$J_{\text{UARPO}}(\theta) = E \left[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q), \tau \in \mathcal{T} \right] \\ \frac{1}{G} \sum_{i=1}^G \frac{1}{|O_i|} \sum_{t=1}^{|O_i|} \left[\min \left\{ \frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})} \left(\hat{A}_{i,t}^I + \hat{A}_t^{S\tau} \right) \hat{U}_{i,t}, \right. \right. \\ \left. \left. \text{clip} \left(\frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}, 1 - \varepsilon, 1 + \varepsilon \right) \left(\hat{A}_{i,t}^I + \hat{A}_t^{S\tau} \right) \hat{U}_{i,t} \right\} \right. \\ \left. - \beta D_{\text{KL}}[\pi_{\theta} || \pi_{\text{ref}}] \right], \quad (1)$$

$$\hat{A}_{i,t}^I \triangleq \tilde{r}_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}, \quad (2)$$

$$\hat{A}_t^{S\tau} \triangleq \tilde{s}_t^{\tau} = \frac{s_t^{\tau} - \text{mean}(\mathbf{s}_{t-1,t}^{\tau})}{\text{std}(\mathbf{s}_{t-1,t}^{\tau})}, \quad (3)$$

where π_{θ} and $\pi_{\theta_{\text{old}}}$ are the current and old policy models, q and o are questions and outputs sampled from the question dataset and the old policy $\pi_{\theta_{\text{old}}}$, respectively. ε is

Algorithm 1 Iterative UARPO

```
1: Input: Initial policy model  $\pi_{\theta_{\text{init}}}$ ; reward model  $r_{\phi}$ ; task prompts  $\mathcal{D}$ ; hyperparameters  $\epsilon, \beta, \mu$ ; stack length  $L$ 
2: Initialize: policy model  $\pi_{\theta} \leftarrow \pi_{\theta_{\text{init}}}$ ; target special stack  $\mathbf{s}_L^{\tau}$ 
3: for iteration = 1 to  $I$  do
4:   Update reference model  $\pi_{\text{ref}} \leftarrow \pi_{\theta}$ 
5:   Initialize stack  $S[0..L-1]$ 
6:   for step = 1 to  $M$  do
7:     Sample batch  $\mathcal{D}_b \subset \mathcal{D}$ 
8:     Update the old policy  $\pi_{\theta_{\text{old}}} \leftarrow \pi_{\theta}$ 
9:     Sample  $G$  outputs  $\{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)$  for each question  $q \in \mathcal{D}_b$ .
10:    Compute rewards  $\{r_i\}_{i=1}^G$  and confidence scores  $\{u_i\}_{i=1}^G$  for each output  $O_i$  by running  $r_{\phi}$ .
11:    Compute current step average reward  $\frac{1}{G} \sum_{i=1}^G r_L^{\tau}$  for current target  $\tau$ .
12:    Compute  $\hat{A}_{i,t}^I$  for the  $t$ -th token of  $o_i$  through group relative advantage estimation.
13:    if  $\text{step} > L$  then
14:      Compute  $\hat{A}_{i,t}^I$  for the  $t$ -th token through latest  $L$  step relative advantage estimation for target  $\tau$ 
15:    Gather two part relative advantage and multiply with corresponding confidence score
16:    for UARPO iteration = 1, ...,  $\mu$  do
17:      Update the policy model  $\pi_{\theta}$  by maximizing the UARPO objective.
18:      Update  $r_{\phi}$  through continuous training using a replay mechanism.
19:      Update  $r_{\phi}$  with replay mechanism
20: Output:  $\pi_{\theta}$ 
```

a clipping-related hyper-parameter introduced in PPO for stabilizing training. $\hat{A}_{i,t}^I$ represents the in-group relative advantage as in GRPO, where $\mathbf{r} = [r_0, r_1, \dots, r_i, \dots, r_G]$, $\hat{A}_t^{S_{\tau}}$ represents the cross-group relative advantage where $\mathbf{s}_{t-l,t}^{\tau} = [s_{t-l}, s_{t-l+1}, \dots, s_t | \tau]$ and $s_t^{\tau} = \frac{1}{G} \sum_{i=1}^G r_{i,t}^{\tau}$, which indicate the advantage of the current group's overall performance relative to the average performance over multiple steps in a recent window period under the same prediction objective. $\mathbf{s}_{t-l,t}^{\tau}$ is a group consisting of multiple steps with window length l . $\hat{U}_{i,t} \triangleq \alpha \cdot \frac{\text{score} - \text{const}}{100}$ is the uncertainty adjustment function, and α denotes an adjustable coefficient. The algorithmic iterative process can be described as Algorithm 1.

2.2. Rewards and Uncertainty

- **Accuracy Reward** Prediction accuracy is commonly used to evaluate the performance of reinforcement learning models and construct loss functions. Specifically, it measures the consistency between the model's predictions and the ground-truth outcomes (rise/fall) of each sample.
- **Completion Length Reward** Previous works have found that text length expansion occurs in large model RL reasoning, which is helpful for improving reasoning time and enabling complex reasoning. Therefore, we provide this type of reward. Specifically, when the text reasoning length is no more than 200 tokens, a gradually increasing reward is offered.
- **Format Reward** Format reward help the model learn the target output format during fine-tuning. FinZero is required to follow the output format as shown in the system prompt of Figure 2.

- **Confidence Score** Prior works [35, 36] have explored the feasibility and methods for large models to learn task uncertainty. Given the high uncertainty inherent in financial decision-making, where uncertainty analysis is critical for model development and real-world application, we integrate model reasoning uncertainty into reinforcement learning fine-tuning. During each image-text reasoning process, the model outputs a confidence score based on the input information and its reasoning. This score quantifies the model's uncertainty about its reasoning result for the given task, enabling it to learn problem difficulty and uncertainty through training.

3. EXPERIMENTS

3.1. FVLDB Dataset

To validate our idea, we specifically construct a financial time-series image-text dataset (FVLDB as Figure 2) with over 10000+ samples. The images in FVLDB contain a wealth of financial assets, along with corresponding text descriptions and questions. To enhance data diversity, FVLDB includes index data from global stock markets, as well as data on cryptocurrency assets such as Bitcoin. The time-series length, sampling frequency, type, and number of features of the assets in each image are variable, and the image styles are also diverse. This flexibility enables the model to process diverse data types.

3.2. Setup

The FinZero utilize the Qwen2.5-VL-3B model as the backbone and fine-tuned it directly on the FVLDB dataset with the UARPO algorithm. For baselines, we select the original Qwen2.5-VL-3B model, the Qwen2.5-VL-7B model, and the larger-scale GPT-4o model. Additionally, we also fine-tuned Qwen2.5-VL-3B with GRPO, and also constructed a Naive Model, which extends the trend of the past period of time. The Adam optimizer was adopted with a learning rate of $1e-6$, and the fine-tuning process ran for two epochs. All experiments were conducted on a server equipped with two 80G Nvidia A100 GPUs.

3.3. Results

As shown in Figure 3, the model's rewards continuously increase during the UARPO fine-tuning process: the format reward and completion length reward rise rapidly in the early stage of training and then stabilize, while the accuracy reward also increases steadily with training; meanwhile, the loss value decreases consistently. The prediction performance of the fine-tuned model on the test set is presented in Table 1. After UARPO fine-tuning, FinZero exhibits more competitive prediction performance compared to baseline models, whether in price prediction tasks or volatility prediction

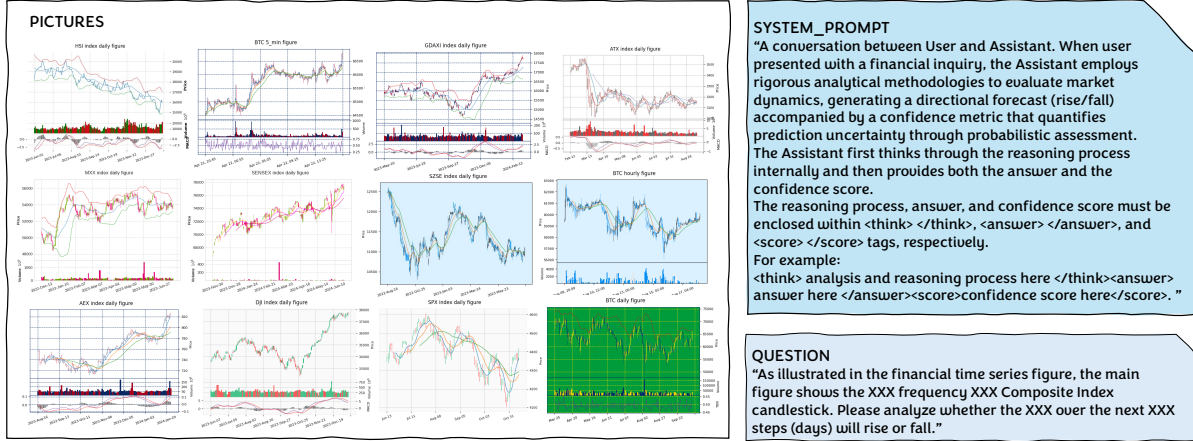


Fig. 2. Overview of Image-Text Pairs for the FVLDB Dataset.

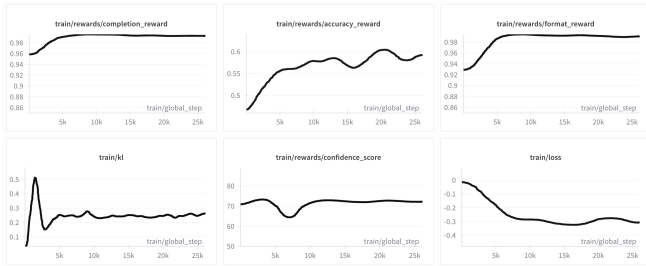


Fig. 3. Overview of the FinZero Training.

Table 1. Main Results of Model Accuracy Comparison.

	Volatility ACC (%)				Price ACC (%)			
Model	5	21	63	Avg	5	21	63	Avg
Naive	48.54	46.23	48.46	47.75	50.00	52.04	50.00	50.68
Qwen2.5-VL-3B	46.67	45.69	50.51	47.62	54.20	51.64	52.54	52.79
Qwen2.5-VL-7B	50.49	43.64	51.16	48.43	<u>55.55</u>	51.91	51.14	53.53
GRPO	53.68	<u>54.86</u>	52.15	<u>53.56</u>	53.24	<u>53.63</u>	<u>53.76</u>	<u>53.54</u>
GPT-4o	<u>54.28</u>	48.26	53.38	51.97	56.16	51.22	51.14	52.84
FinZero	56.31	65.74	<u>52.93</u>	58.33	54.52	56.31	65.88	58.90

Table 2. Model Prediction Accuracy Across Confidence Score Groups.

	Low (%)	Middle (%)	High (%)
Qwen2.5-VL-3B	51.2	51.7	49.3
Qwen2.5-VL-7B	47.38	47.81	54.36
GRPO	53.85	53.19	54.61
GPT-4o	49.85	49.42	54.75
FinZero	54.48	56.67	62.13




Fig. 4. Finetuning Accuracy Trend Comparison.

tasks. While FinZero with 3B parameter size surpasses larger parameter models such as GPT-4o. Additionally, when test set samples are divided into three equal groups based on the model’s uncertainty scores sorted from highest to lowest as in Table 2, it shows that for the FinZero, the prediction accuracy of samples with high confidence scores is further improved - the prediction accuracy of the highest-score group is increased by approximately 13.5% relative to that of GPT-4o.

Furthermore, compared to Qwen2.5-VL-3B fine-tuned by GRPO, FinZero achieves better average prediction performance. Meanwhile, grouping based on confidence scores exhibits a more pronounced positive correlation with prediction accuracy. Besides, we illustrate the accuracy changes of the two models during the fine-tuning process, as shown in Figure 4.

4. CONCLUSIONS

This study introduces FinZero, a model for multimodal financial time-series reasoning. The FinZero achieve competitive forecasting accuracy rivaling larger models like Qwen-7B and GPT-4o. Moreover, FinZero provides uncertainty scores that reliably indicate prediction confidence - higher scores correlate with greater accuracy. This work demonstrates the potential of cross-modal reinforcement learning to advance financial reasoning, offering both a robust method and a practical tool for reliable predictions.

5. REFERENCES

- [1] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long, “Autoforformer: Decomposition transformers with auto-correlation for long-term series forecasting,” *Advances in neural information processing systems*, vol. 34, pp. 22419–22430, 2021.
- [2] Donghao Luo and Xue Wang, “Moderntcn: A modern pure convolution structure for general time series analysis,” in *The twelfth international conference on learning representations*, 2024, pp. 1–43.
- [3] Hongbin Huang, Minghua Chen, and Xiao Qiao, “Generative learning for financial time series with irregular and scale-invariant patterns,” in *The Twelfth International Conference on Learning Representations*.
- [4] Tao Wu, Xiangyun Gao, Sufang An, and Siyao Liu, “Time-varying pattern causality inference in global stock markets,” *International Review of Financial Analysis*, vol. 77, pp. 101806, 2021.
- [5] Guido Baltussen, Sjoerd van Bakkum, and Zhi Da, “Indexing and stock market serial dependence around the world,” *Journal of Financial Economics*, vol. 132, no. 1, pp. 26–48, 2019.
- [6] Andrew W Lo and A Craig MacKinlay, “Stock market prices do not follow random walks: Evidence from a simple specification test,” *The review of financial studies*, vol. 1, no. 1, pp. 41–66, 1988.
- [7] Brian H Boyer, Tomomi Kumagai, and Kathy Yuan, “How do crises spread? evidence from accessible and inaccessible stock indices,” *The journal of finance*, vol. 61, no. 2, pp. 957–1003, 2006.
- [8] Arun Narayanasamy, Humnath Panta, and Rohit Agarwal, “Relations among bitcoin futures, bitcoin spot, investor attention, and sentiment,” *Journal of Risk and Financial Management*, vol. 16, no. 11, pp. 474, 2023.
- [9] Burton G Malkiel, “The efficient market hypothesis and its critics,” *Journal of economic perspectives*, vol. 17, no. 1, pp. 59–82, 2003.
- [10] Turan G Bali and Hao Zhou, “Risk, uncertainty, and expected returns,” *Journal of Financial and Quantitative Analysis*, vol. 51, no. 3, pp. 707–735, 2016.
- [11] Jonathan Brogaard and Abalfazl Zareei, “Machine learning and the stock market,” *Journal of Financial and Quantitative Analysis*, vol. 58, no. 4, pp. 1431–1472, 2023.
- [12] Markus Leippold, Qian Wang, and Wenyu Zhou, “Machine learning in the chinese stock market,” *Journal of Financial Economics*, vol. 145, no. 2, pp. 64–82, 2022.
- [13] Shihao Gu, Bryan Kelly, and Dacheng Xiu, “Empirical asset pricing via machine learning,” *The Review of Financial Studies*, vol. 33, no. 5, pp. 2223–2273, 2020.
- [14] Seunghan Lee, Taeyoung Park, and Kibok Lee, “Soft contrastive learning for time series,” in *The Twelfth International Conference on Learning Representations*.
- [15] Defu Cao, Furong Jia, Serkan O Arik, Tomas Pfister, Yixiang Zheng, Wen Ye, and Yan Liu, “TEMPO: Prompt-based generative pre-trained transformer for time series forecasting,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [16] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long, “Timesnet: Temporal 2d-variation modeling for general time series analysis,” *arXiv preprint arXiv:2210.02186*, 2022.
- [17] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo, “Reversible instance normalization for accurate time-series forecasting against distribution shift,” in *International Conference on Learning Representations*, 2022.
- [18] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam, “A time series is worth 64 words: Long-term forecasting with transformers,” 2023.
- [19] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long, “itransformer: Inverted transformers are effective for time series forecasting,” *arXiv preprint arXiv:2310.06625*, 2023.
- [20] Luke Nicholas Darlow, Qiwen Deng, Ahmed Hassan, Martin Asenov, Rajkarn Singh, Artjom Joosen, Adam Barker, and Amos Storkey, “DAM: Towards a foundation model for forecasting,” in *The Twelfth International Conference on Learning Representations*.
- [21] Mingtian Tan, Mike Merrill, Vinayak Gupta, Tim Althoff, and Tom Hartvigsen, “Are language models actually useful for time series forecasting?,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 60162–60191, 2024.
- [22] Tian Zhou, Peisong Niu, Liang Sun, Rong Jin, et al., “One fits all: Power general time series analysis by pretrained lm,” *Advances in neural information processing systems*, vol. 36, pp. 43322–43355, 2023.
- [23] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y. Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen, “Time-LLM: Time series forecasting by reprogramming large language models,” in *The Twelfth International Conference on Learning Representations*.
- [24] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al., “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” *arXiv preprint arXiv:2501.12948*, 2025.
- [25] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al., “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [26] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al., “Constitutional ai: Harmlessness from ai feedback,” *arXiv preprint arXiv:2212.08073*, 2022.
- [27] Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo, “Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning,” *arXiv preprint arXiv:2502.14768*, 2025.
- [28] Liang Chen, Lei Li, Haozhe Zhao, and Yifan Song, “Vinci. r1-v: Reinforcing super generalization ability in vision-language models with less than 3,” .
- [29] Yingzhe Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, et al., “Lmm-r1: Empowering 3b llms with strong reasoning abilities through two-stage rule-based rl,” *arXiv preprint arXiv:2503.07536*, 2025.
- [30] Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfang Lu, Daocheng Fu, Tiancheng Han, Botian Shi, Wenhai Wang, Junjun He, et al., “Mm-eureka: Exploring the frontiers of multimodal reasoning with rule-based reinforcement learning,” *arXiv preprint arXiv:2503.07365*, 2025.
- [31] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo, “Deepseekmath: Pushing the limits of mathematical reasoning in open language models,” 2024.
- [32] DeepSeek-AI, D. Guo, D. Yang, H. Zhang, and et al., “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” 2025.
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov, “Proximal policy optimization algorithms,” 2017.
- [34] Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen, “Reinforce++: An efficient rlhf algorithm with robustness to both prompt and reward models,” 2025.
- [35] Shudong Liu, Zhaocong Li, Xuebo Liu, Runzhe Zhan, Derek F. Wong, Lidia S. Chao, and Min Zhang, “Can llms learn uncertainty on their own? expressing uncertainty effectively in a self-training manner,” in *EMNLP*, 2024.
- [36] Zhiqiu Xia, Jinxuan Xu, Yuqian Zhang, and Hang Liu, “A survey of uncertainty estimation methods on large language models,” 2025.