

# FINKRX: Establishing Best Practices for Korean Financial NLP

Guijin Son<sup>1\*</sup> Hyunwoo Ko<sup>1</sup> Hanearl Jung<sup>1</sup> Chami Hwang<sup>2†</sup>

OneLineAI<sup>1</sup> KRX<sup>2</sup>  
spthsrbwls123@yonsei.ac.kr hcharm2ing@krx.co.kr

## Abstract

In this work, we present the first open leaderboard for evaluating Korean large language models focused on finance. Operated for about eight weeks, the leaderboard evaluated 1,119 submissions on a closed benchmark covering five MCQA categories: finance and accounting, stock price prediction, domestic company analysis, financial markets, and financial agent tasks and one open-ended qa task. Building on insights from these evaluations, we release an open instruction dataset of 80k instances and summarize widely used training strategies observed among top-performing models. Finally, we introduce FINKRX, a fully open and transparent LLM built using these best practices. We hope our contributions help advance the development of better and safer financial LLMs for Korean and other languages.<sup>1</sup>

## 1 Introduction

Large language models (LLMs) hold significant potential for financial applications (Son et al., 2023; Chen et al., 2024b,a). However, performance issues in this domain can lead to monetary losses, making it imperative to develop reliable evaluation systems prior to deployment. Unfortunately, the inherently closed nature of the financial industry limits the sharing of models (Wu et al., 2023) and dataset (Mahfouz et al., 2024), slowing the development of relevant techniques and often resulting in duplicated efforts across companies and teams.

Existing tools to evaluate LLM performance in the Korean financial domain include KRX-Bench (Son et al., 2024a), which specifically assesses knowledge of Korean listed companies, and KMMLU (Son et al., 2024d), a broader benchmark spanning 45 categories that incorporates a subset of finance and economics. However, these

benchmarks fall short of reflecting the broad potential for LLM applications in the financial sector. To address this gap, we compile a comprehensive finance benchmark consisting of approximately 5.5k multiple-choice questions derived from online exams, LLM-generated questions, and hand-crafted instances. This benchmark covers five key topics: financial markets, finance and accounting, domestic company analysis, financial agents (Hu et al., 2024), and stock price prediction (Soun et al., 2022). Recognizing that multiple-choice questions may not fully represent real-world prompts (Kim et al., 2024), we also include an open-ended QA set featuring 100 challenging prompts.

To encourage the adoption of the benchmark and foster an open research culture, we take a further step by launching an open leaderboard for financial LLMs. It was operated for two months, comprising two stages: a preliminary round and a main round. Over the course of the competition, more than 1,000 models were submitted, with over 600 models remaining publicly accessible to date<sup>2</sup>, thereby laying the groundwork for future research. In addition, we compile the submitted models along with their system cards to document effective tuning strategies. Furthermore, we collect over 200,000 instances from competing teams, filter them, and release a high-quality instruction dataset consisting of 80,000 samples. Finally, after regenerating responses for each instance using Deepseek-R1 (Guo et al., 2025) and training on these trajectories, we release FINKRX, the first reasoning model for the Korean financial domain.

## 2 Motivation and Related Works

The financial industry has witnessed rapid expansion in the adoption of artificial intelligence, with particular emphasis on generative AI technologies driving innovations in enhanced customer

<sup>†</sup>Corresponding author.

<sup>1</sup><https://krxbench.koscom.co.kr/home/main>

<sup>2</sup>As of 2025.02.28

Category	Examples
Financial Markets <b>642 total</b>	다음 중 대한민국 주식시장 매매 제도에 대한 기술로 알맞은 것은 무엇인가? <i>Which of the following descriptions is correct regarding the trading system of the Korean stock market?</i> A. Opening time is 10:00 AM. B. The daily price limit for the KOSPI market is $\pm 15\%$ of the previous day's closing price. [...]
Finance and Accounting <b>1,450 total</b>	다음 중 화폐의 시간가치에 관한 설명으로 옳지 않은 것은 무엇인가? <i>Which of the following statements about the value of money is incorrect?</i> A. In monthly compounding, the monthly interest rate is calculated by dividing the annual [...] B. Given the same initial investment and conditions, compound interest yields higher [...]
Domestic Company Analysis <b>2,039 total</b>	엑세스바이오의 COVID-19 진단 제품의 매출 기여와 미국 시장 판매에 대해서 올바른 것은? <i>What is correct regarding the sales contribution of Access Bio's COVID-19 diagnostic products and their sales in the U.S. market?</i> A. Access Bio's COVID diagnostic products were developed for general health screening [...] B. Access Bio's COVID diagnostic products have demonstrated effectiveness through [...]
Financial Agent <b>46 total</b>	데이터프레임의 '종가' 열의 평균 값을 계산하는 코드를 고르시오. <i>Choose the code that calculates the average value of the 'Closing Price' column in the DataFrame.</i> A. <code>df['Close Price'].mean()</code> B. <code>df['Total Traded Quantity'].median()</code> [...]
Stock Price Prediction <b>1,472 total</b>	주식 A에 대한 분석 결과표를 바탕으로 향후 A의 주가가 상승/하락할지 예측하시오. <i>Based on the analysis report of stock A, predict whether the future price of A will rise or fall.</i>
Open-Ended FinQA <b>100 total</b>	위반행위로 얻은 이익이란 무엇이고 그 범위는 어떻게 정의되는가? <i>What are the profits gained from breach of contract, and how is their scope defined?</i>

Table 1: **Overview of the benchmark used for evaluation.** Each example demonstrates a specific question type for each category. Gray text are English translations provided for better reachability.

service, improved risk management, and overall operational efficiency (McKinsey & Company, 2025). Despite these advancements, Korean financial institutions face significant challenges in harnessing proprietary language models (Jaech et al., 2024; Team et al., 2023). Strict security regulations—such as network separation policies (Financial Services Commission, 2024)—impede their ability to fully leverage these innovations. Moreover, the absence of clear guidelines and robust evaluation frameworks for managing the risks inherent in generative AI—such as hallucinations (Kang and Liu, 2023), biases (Zhou et al., 2024), and information leakage (Liu et al., 2024)—further complicates the integration. In response, Son et al. (2024a) introduced KRX-Bench, the first publicly available benchmark designed to assess the knowledge of LLMs in Korean companies. However, KRX-Bench remains limited in scope and has yet to achieve widespread adoption among Korean financial institutions.

In this work, drawing inspiration from financial benchmarks in various languages (Xie et al., 2024a; Nie et al., 2024; Koncel-Kedziorski et al., 2023), we extend KRX-Bench to develop a more comprehensive benchmark for Korean financial language models by incorporating five additional categories. Moreover, our work distinguishes itself by operating an open leaderboard with a total prize

pool of approximately \$42,000, which has attracted submissions of around 1,000 models, creating the groundwork for future works in financial NLP.

### 3 Leaderboard Construction

In this work, we introduce an open leaderboard for Korean financial LLMs and share lessons learned from its two-month operation, comprised of a preliminary round (October 14–November 7, 2024) and a final round (November 13–December 6, 2024). In total, 1,119 models were submitted—478 in the preliminary round and 641 in the final round, establishing a foundation for open research in Korean financial LLMs, with over 600 models remaining publicly accessible. The following sections describe the benchmark construction process (Section 3.1), present operational details (Section 3.2) and summarize key statistics (Section 3.3).

#### 3.1 Benchmark Details

The benchmark used for evaluation consisted of five categories in the preliminary round: finance and accounting, stock price prediction, domestic company analysis, financial markets, and financial agent tasks. For the final round, only three categories were used: finance and accounting, financial markets, and open-ended finance QA. Table 1 details the examples of each category.

**Finance and Accounting** For this category, we compile four-option MCQA questions, primarily sourced from university exams. In the preliminary round, these questions were presented with four options, while in the final round, the answer set was expanded to eight options. The augmentation uses two methods: (1) grouping questions based on embeddings to mix similar items, and (2) applying rule-based augmentations (Wang et al., 2024; Zhao et al., 2024), such as replacing an answer option with "none of the above" (thereby making it the correct answer) or shuffling the order of options. A manual human check is done post-augmentation to ensure correctness.

**Financial Markets** For this category, we employ an approach similar to the Finance and Accounting category. However, the source questions are collected from exams that assess understanding of the Korean financial system and related laws.

**Stock Price Prediction** This category is inspired by Soun et al. (2022). We randomly sample fixed-length stock price data (OHLCV: Open, High, Low, Close, Volume) from Korean stock markets, using only post-2024 data to mitigate potential contamination. A set of technical indicators is computed and presented in a Markdown table format (e.g., adj-close for adjusted closing price; inc-5, inc-10, inc-15, inc-20, inc-25, and inc-30 for percentage changes over the past 5, 10, 15, 20, 25, and 30 trading days). Models are tasked with a binary classification—predicting whether the price will increase or decrease—and are expected to detect basic signals of momentum (Jegadeesh and Titman, 1993) or mean reversion (Poterba and Summers, 1988) in the time-series data.

**Domestic Company Analysis** For this section, we directly employ KRX-Bench (Son et al., 2024b), an automatically generated benchmark constructed using GPT-4o (Hurst et al., 2024) leveraging annual filings from Korean companies. It consists of 4-option MCQA questions designed to assess knowledge on topics such as Product Offerings, Financial Policy, and Business Strategy.

**Financial Agents** This subset evaluates the capability to function as an automated financial agent by executing code-based tasks on real financial data. Similar to Hu et al. (2024), the model is provided with a CSV file and an instruction to extract specific information and perform corresponding coding operations. The model is presented with

multiple output options, including perturbed variants, and is prompted to select the correct one.

**Open-Ended FinQA** Given that all subsets employ multiple-choice or binary-choice formats, we were concerned that these evaluation methods may not fully capture the diversity of prompts encountered in real-world applications. Drawing inspiration from open-ended evaluations such as MT-Bench (Zheng et al., 2023), we curated a set of 100 challenging prompts from three sources: the legal reasoning subset of KRX-Bench (Son et al., 2024a), advanced math questions from HRM8K (Ko et al., 2025), and graduate-level financial engineering and econometrics exam questions. A gold standard answer was generated using o1-Pro (Jaech et al., 2024), and GPT-4o was utilized as an LLM-as-a-Judge to determine whether competing models produced responses superior to this standard. Figure 6 illustrates the prompts employed in the LLM-as-a-Judge evaluation.

### 3.2 Operation Details

The leaderboard was active for eight weeks, from October 14, 2024, to November 7, 2024, on a dedicated, self-hosted website. The competition was structured in two rounds: a preliminary round and a final round. In the preliminary round, participants uploaded their models publicly on Hugging Face and submitted the corresponding model links. The top 30 teams advanced to the final round, where each team was allowed up to three submissions. Models were evaluated on a server equipped with 2 A6000 Ada GPUs, with capacity scaling up to 8 GPUs depending on the number of submissions. The benchmark dataset was kept private, with only one sample released from each subset.

To ensure consistency and fairness, participants were restricted in the choice of base models to prevent incompatibility issues with the inference engine and to avoid giving larger companies with more training resources an unfair advantage. Allowed models include Qwen (1.5B and 7B) (Yang et al., 2024), Mistral (7B) (Jiang, 2024), GLM-4 (9B) (GLM et al., 2024), Llama 3/3.1 (8B) (Grattafiori et al., 2024), Amber (Liu et al., 2023), Phi 3.5 (mini) (Abdin et al., 2024), and Gemma2 (2B and 9B) (Team et al., 2024). Both base and instruct models were allowed. Teams that advanced to the main rounds were provided \$2500 of AWS credit to help model training.

Participants were required to disclose their

datasets and confirm that they did not include any copyrighted material, as a condition for qualifying for the prize money. For evaluation, we adopt a zero-shot chain-of-thought (CoT) format. Initially, each model is prompted to generate a CoT reasoning. We then concatenate the original prompt with the generated CoT and append “### Answer:” to prompt the model to produce its final answer. In this step, a logit processor is employed to ensure that the model selects from the provided options, thereby preventing evaluation errors due to format mismatches. To prevent spamming, each team is allowed to submit one model per day.

### 3.3 Statistics

**Submission** During the preliminary rounds, 233 accounts signed up, with 71 making at least one submission. A total of 478 models were submitted—averaging nearly seven submissions per active team—and November 5 was the busiest day with 45 entries. Moreover, the largest single-day influx of new registrations occurred on October 14, when 83 accounts joined, highlighting strong early interest. For details on the overall trend, see Figure 5. In the main rounds, a high submission rate was maintained throughout the entire period, with 30 teams contributing a total of 641 submissions.




Figure 1: **Distribution of participants.** The shades of blue bars indicate corporate participants.

**Participants** A total of 71 teams submitted at least one model during the leaderboard period (each team may consist of up to four members). Among these teams, 52.5% were corporate participants, with the remaining teams representing universities and student groups. The corporate participants were further categorized into Tech, Finance, and Other sectors, as presented in Figure 1. These results demonstrate that the competition successfully attracted a diverse range of participants—not only students, but also a substantial number of companies, including publicly listed tech companies,

securities firms, and banks.

## 4 Analysis

### 4.1 Analysis on Data Collection

As our rules prohibit the use of licensed materials for training LLMs, participants focused on collecting license-free financial content—potentially useful for constructing corpora to train Korean LLMs. Table 2 lists the 11 most-used domains, with a strong focus on government (go.kr) and non-profit organizations (or.kr). After collecting the raw corpora from these sources, participants mostly employed either GPT-4o (Hurst et al., 2024) or Qwen2.5-72B-Instruct (Yang et al., 2024) to convert the data into MCQA (Bi et al., 2024) or Instruction-Response formats, with some employing an LLM-as-a-Judge (Zheng et al., 2023; Son et al., 2024c) for validation.

Link	Name
<a href="http://krx.co.kr">krx.co.kr</a>	Korea Exchange
<a href="http://krxverse.co.kr">krxverse.co.kr</a>	KRXverse
<a href="http://fsc.go.kr">fsc.go.kr</a>	Financial Services Commission
<a href="http://bok.or.kr">bok.or.kr</a>	Bank of Korea
<a href="http://law.go.kr">law.go.kr</a>	Korean Law Information Service
<a href="http://kasb.or.kr">kasb.or.kr</a>	Korea Accounting Standards Board
<a href="http://mss.go.kr">mss.go.kr</a>	Ministry of SMEs and Startups
<a href="http://ftc.go.kr">ftc.go.kr</a>	Fair Trade Commission
<a href="http://kifrs.com">kifrs.com</a>	K-IFRS
<a href="http://kiep.go.kr">kiep.go.kr</a>	Korea Institute for International Economic Policy
<a href="http://kocw.net">kocw.net</a>	Korea OpenCourseWare

Table 2: **Data collection sources.**

To ensure reusability, we collect about 200,000 data samples from HuggingFace (released by competing teams) and applied quality filters: the Min-Hash algorithm to remove near-duplicates, a regex filter to exclude time-bound queries (e.g., “What will Kakao’s 2024 sales be?”), and a rule-based filter to remove incomplete or overly short questions. This process yielded a final set of 86,007 instances. For further details see Appendix A.



Models	F&A	Stock	Company	Market	Agent	Average
AnonymousLLMer/krx-qwen2.5-v1106	<b>0.51</b>	<b>0.56</b>	0.94	<b>0.49</b>	<b>0.83</b>	<b>0.67</b>
AnonymousLLMer/krx-qwen2.5-v1105	0.44	<b>0.56</b>	0.92	0.39	0.81	0.62
KR-X-AI/krx-qwen2-7b-instruct-v4_m	0.4	0.55	0.92	0.41	0.77	0.61
2point5p/krx-qwen2.5-7b-it-prompt-v2	0.5	0.55	0.95	0.46	0.57	0.61
TwoSubPlace/krx-qwen2-7b-it-baseline-v6	0.4	0.52	0.90	0.44	0.79	0.61
KR-X-AI/krx-qwen2-7b-instruct-v3	0.4	0.49	0.9	0.44	0.72	0.59
SejongKRX/Sejong-Qwen-v1	0.41	0.45	0.93	0.42	0.66	0.57
2point5p/krx-qwen2.5-7b-it-X-Two	0.44	0.5	<b>0.96</b>	0.41	0.53	0.57
lsw0570168/krx-q25-7b-it-v8	0.41	0.55	0.85	0.43	0.62	0.57
SejongKRX/Sejong-Qwen-v7	0.35	0.45	0.95	0.44	0.6	0.56

Table 3: **Performance of Top-10 models from the preliminary rounds.** The highest performance of each subset is highlighted in **bold** and the second best is underlined.

Models	F&A	Market	Open-Ended	Average
overfit-brothers/hello_world06	0.65	<b>0.83</b>	0.01	0.50
AnonymousLLMer/krx-qwen2.5-v1206-1	0.63	0.65	0.04	0.44
shibainu24/qwen2.5-7B-inst-release-1516wk	0.56	0.67	0.04	0.43
Q-PING/krx_1205_test_model_3	0.58	0.64	0.02	0.42
Hi-Q/krx_1206_test_model_2	0.60	0.61	0.02	0.41
FINKRX (Ours)	<b>0.78</b>	0.66	<b>0.18</b>	<b>0.54</b>

Table 4: **Performance of top 5 models from the main rounds and FINKRX.** FINKRX shows the best average performance with notable improvements in Financial & Accounting and Open-Ended FinQA. The highest performance of each subset is highlighted in **bold** and the second best is underlined.

## 4.2 Analysis on Top-Performing Models

Table 3 presents the performance of the top 10 models from the preliminary rounds, and Figure 2 displays the corresponding score trends. The largest improvement was observed in Domestic Company Analysis, where scores rose from 0.51 to 0.94. However, Financial & Accounting and Financial Markets experienced relatively modest gains. We attribute this to the relatively simple methods used by most teams during the preliminary rounds. All top 10 teams primarily employed supervised fine-tuning (SFT) for model training. Interestingly, team Americano incorporated a brief continual pre-training phase (Xie et al., 2024b) before SFT on 3.7GB of text; however, the performance impact of this additional step remains inconclusive in a small-scale setting. Notably, all top-performing models were based on Qwen2.5-7B-Instruct.

Teams advancing to the main rounds employed multi-step, more complex training methods. For example, Shinbainu used a curriculum-based SFT approach that began with training on easier samples and then proceeded to a second round of SFT on more challenging prompts generated via the Evolve Instruct method (Xu et al., 2023; Luo et al., 2023a,b). The final model was subsequently refined using DPO (Rafailov et al., 2023), lever-




Figure 2: **Preliminary round performance trends.**

aging preference data from the stage-two SFT model—which generated two responses that were then evaluated by an LLM-as-a-Judge (Zheng et al., 2023). Similar strategies were observed among other teams; for instance, Hi-Q and Overfit Brothers implemented KTO (Ethayarajh et al., 2024) and DPO, respectively.

Interestingly, team Hi-Q adopts continual pre-training and demonstrated its effectiveness on a private finance benchmark, as shown in Figure 3. Notably, CPT+SFT scores an average of




Figure 3: **Evaluation results reported but Hi-Q.** Performance of each methodology is represented by boxed numbers, and green numbers indicate the improvement over CPT.

2.7 points higher than plain SFT, indicating that a well-structured continued pre-training approach can benefit LLMs in Korean finance. However, further research is required to establish what marks a good continual pre-training. Details of the benchmarks used by Hi-Q are provided in Appendix C.

## 5 FINKRX: Open LLM for Korean Finance

To aggregate the open resources collected during the competition, we train our own LLM, FINKRX. In line with recent trends toward reasoning LLMs (Jaech et al., 2024; Guo et al., 2025), FINKRX is designed to generate a two-step response: a first step enclosed within `<think>` and `</think>` tags, where the model performs self-correcting reasoning, and a second step enclosed within `<solution>` and `</solution>` tags, which provides the final summary of the reasoning process. It should be noted that this effort is not intended to achieve state-of-the-art language model; rather, it serves to evaluate the quality of the collected resources and provide guidelines for future research.

### 5.1 Details in Training FINKRX

Recent studies have shown that supervised fine-tuning (SFT) is effective enough in training reasoning language models (Muennighoff et al., 2025; Ye et al., 2025; Wen et al., 2025; Sun et al., 2025). Moreover, during the competition, submissions that have combined SFT with preference optimization techniques such as DPO or KTO have successfully adapted models for the Korean financial domain. Accordingly, we adopt a two-stage training approach: SFT followed by DPO. The SFT dataset comprises prompts paired with responses generated by Deepseek-R1, split evenly between English and Korean. For Korean prompts, the so-

lutions are translated into Korean while retaining the reasoning process in English. The dataset is drawn from three sources: (1) English Prompt-R1 responses collected online (Zhao et al., 2025), (2) Korean Prompt-R1 responses collected online (Son et al., 2025), and (3) 86k prompts from Section 3.1, for which we generated responses using R1. We employed GPT-4o to filter correct samples, retrying up to six attempts for incorrect samples, resulting in approximately 400k instances. Post-SFT, the model struggled with everyday prompts, tended to overthink (Kumar et al., 2025), and occasionally displayed formatting issues by treating some queries as if they were MCQA tasks. We attribute these issues to the data distribution, which heavily emphasized academic multiple-choice questions paired with extended reasoning. To address these behaviors, we conducted a final DPO stage, where chosen samples are generated from R1, and rejected samples are drawn from the SFT model.

### 5.2 Performance Analysis

The performance of FINKRX is reported in Table 4. FINKRX demonstrates strong results in the Finance & Accounting category, which includes a diverse range of accounting and econometrics tasks that benefit from robust mathematical and logical reasoning. These capabilities also yield strong performance on open-ended FinQA tasks, where multi-step logical deductions are necessary. In contrast, FINKRX shows weaker performance in the Market category. The Market category relies more on factual and domain-specific knowledge; they do not benefit as strongly from FINKRX’s reasoning-oriented approach. These findings are consistent with those of Ha (2025), who observed that while reasoning-focused models excel at challenging mathematical questions, their performance decline in knowledge-intensive domains as training progresses.

## 6 Conclusion

In this work, we present the largest Korean finance benchmark covering five categories: finance and accounting, stock price prediction, domestic company analysis, financial markets, and financial agent tasks. To encourage adoption, we launched a leaderboard that attracted hundreds of participants from academia and industry, resulting in around 600 publicly available models. We distilled successful strategies from these submissions into an

80k-instruction dataset, which we used to train and release FINKRX, a publicly available reasoning model for Korean finance.

## References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. [arXiv preprint arXiv:2404.14219](#).
- AI HUB. 2025a. Financial and law machine reading comprehension dataset. <https://www.aihub.or.kr/aihubdata/data/view.do?currMenu=115&topMenu=100&dataSetSn=71610>. Accessed: 22 March 2025.
- AI HUB. 2025b. Numerical computation and machine reading comprehension dataset. <https://www.aihub.or.kr/aihubdata/data/view.do?currMenu=115&topMenu=100&dataSetSn=71568>. Accessed: 22 March 2025.
- Axolotl AI. 2025. Axolotl: Scalable fine-tuning framework for llms. <https://axolotl-ai-cloud.github.io/axolotl/>. Github.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. 2024. Deepseek llm: Scaling open-source language models with longtermism. [arXiv preprint arXiv:2401.02954](#).
- Yuemin Chen, Feifan Wu, Jingwei Wang, Hao Qian, Ziqi Liu, Zhiqiang Zhang, Jun Zhou, and Meng Wang. 2024a. [Knowledge-augmented financial market analysis and report generation](#). In [Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track](#), pages 1207–1217, Miami, Florida, US. Association for Computational Linguistics.
- Zhiyu Zoey Chen, Jing Ma, Xinlu Zhang, Nan Hao, An Yan, Armineh Nourbakhsh, Xianjun Yang, Julian McAuley, Linda Petzold, and William Yang Wang. 2024b. A survey on large language models for critical societal domains: Finance, healthcare, and law. [arXiv preprint arXiv:2405.01769](#).
- Hyung Won Chung, Noah Constant, Xavier Garcia, Adam Roberts, Yi Tay, Sharan Narang, and Orhan Firat. 2023. Unimax: Fairer and more effective language sampling for large-scale multilingual pretraining. [arXiv preprint arXiv:2304.09151](#).
- Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. [arXiv preprint arXiv:2402.01306](#).
- Financial Services Commission. 2024. Press release - financial regulatory innovation meeting held. <https://www.fsc.go.kr/no010101/83594>.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. [arXiv preprint arXiv:2406.12793](#).
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. [arXiv preprint arXiv:2407.21783](#).
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. [arXiv preprint arXiv:2501.12948](#).
- Huy Hoang Ha. 2025. [Pensez: Less data, better reasoning – rethinking french llm](#). [Preprint](#), arXiv:2503.13661.
- Pin-Lun Hsu, Yun Dai, Vignesh Kothapalli, Qingquan Song, Shao Tang, Siyu Zhu, Steven Shimizu, Shivam Sahni, Haowen Ning, and Yanning Chen. 2024. Liger kernel: Efficient triton kernels for llm training. [arXiv preprint arXiv:2410.10989](#).
- Xueyu Hu, Ziyu Zhao, Shuang Wei, Ziwei Chai, Qianli Ma, Guoyin Wang, Xuwu Wang, Jing Su, Jingjing Xu, Ming Zhu, et al. 2024. Infiagent-dabench: Evaluating agents on data analysis tasks. [arXiv preprint arXiv:2401.05507](#).
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. [arXiv preprint arXiv:2410.21276](#).
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. [arXiv preprint arXiv:2412.16720](#).
- Narasimhan Jegadeesh and Sheridan Titman. 1993. [Returns to buying winners and selling losers: Implications for stock market efficiency](#). [The Journal of Finance](#), 48(1):65–91.
- Fengqing Jiang. 2024. Identifying and mitigating vulnerabilities in llm-integrated applications. Master’s thesis, University of Washington.
- Haoqiang Kang and Xiao-Yang Liu. 2023. Deficiency of large language models in finance: An empirical examination of hallucination. [arXiv preprint arXiv:2311.15548](#).
- Seungone Kim, Juyoung Suk, Ji Yong Cho, Shayne Longpre, Chaeun Kim, Dongkeun Yoon, Guijin Son, Yejin Cho, Sheikh Shafayat, Jinheon Baek, et al. 2024. The biggen bench: A principled benchmark for fine-grained evaluation of language models with language models. [arXiv preprint arXiv:2406.05761](#).

- Hyunwoo Ko, Guijin Son, and Dasol Choi. 2025. Understand, solve and translate: Bridging the multilingual mathematical reasoning gap. [arXiv preprint arXiv:2501.02448](#).
- Rik Koncel-Kedziorski, Michael Krumdick, Viet Lai, Varshini Reddy, Charles Lovering, and Chris Tanner. 2023. Bizbench: A quantitative reasoning benchmark for business and finance. [arXiv preprint arXiv:2311.06602](#).
- Abhinav Kumar, Jaechul Roh, Ali Naseh, Marzena Karpinska, Mohit Iyyer, Amir Houmansadr, and Eugene Bagdasarian. 2025. Overthinking: Slow-down attacks on reasoning llms. [arXiv preprint arXiv:2502.02542](#).
- Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, et al. 2023. Llm360: Towards fully transparent open-source llms. [arXiv preprint arXiv:2312.06550](#).
- Zhiwei Liu, Xin Zhang, Kailai Yang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024. Fmdl-lama: Financial misinformation detection based on large language models. [arXiv preprint arXiv:2409.16452](#).
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023a. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. [arXiv preprint arXiv:2308.09583](#).
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023b. Wizardcoder: Empowering code large language models with evol-instruct. [arXiv preprint arXiv:2306.08568](#).
- Mahmoud Mahfouz, Ethan Callanan, Mathieu Sibue, Antony Papadimitriou, Zhiqiang Ma, Xiaomo Liu, and Xiaodan Zhu. 2024. [The state of the art of large language models on chartered financial analyst exams](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1068–1082, Miami, Florida, US. Association for Computational Linguistics.
- McKinsey & Company. 2025. The state of ai. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. s1: Simple test-time scaling. [arXiv preprint arXiv:2501.19393](#).
- Ying Nie, Binwei Yan, Tianyu Guo, Hao Liu, Haoyu Wang, Wei He, Binfan Zheng, Weihao Wang, Qiang Li, Weijian Sun, et al. 2024. Cfinbench: A comprehensive chinese financial benchmark for large language models. [arXiv preprint arXiv:2407.02301](#).
- James M. Poterba and Lawrence H. Summers. 1988. [Mean reversion in stock prices: Evidence and implications](#). *Journal of Financial Economics*, 22(1):27–59.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE.
- Guijin Son, Jiwoo Hong, Hyunwoo Ko, and James Thorne. 2025. Linguistic generalizability of test-time scaling in mathematical reasoning. [arXiv preprint arXiv:2502.17407](#).
- Guijin Son, Hyunjun Jeon, Chami Hwang, and Hanearl Jung. 2024a. [Krx bench: Automating financial benchmark creation via large language models](#). In *Proceedings of the Joint Workshop of the 7th Financial Technology and Natural Language Processing, the 5th Knowledge Discovery from Unstructured Data in Financial Services, and the 4th Workshop on Economics and Natural Language Processing@ LREC-COLING 2024*, pages 10–20.
- Guijin Son, Hyunjun Jeon, Chami Hwang, and Hanearl Jung. 2024b. [KRX bench: Automating financial benchmark creation via large language models](#). In *Proceedings of the Joint Workshop of the 7th Financial Technology and Natural Language Processing, the 5th Knowledge Discovery from Unstructured Data in Financial Services, and the 4th Workshop on Economics and Natural Language Processing*, pages 10–20, Torino, Italia. Association for Computational Linguistics.
- Guijin Son, Hanearl Jung, Moonjeong Hahm, Keonju Na, and Sol Jin. 2023. Beyond classification: Financial reasoning in state-of-the-art language models. [arXiv preprint arXiv:2305.01505](#).
- Guijin Son, Hyunwoo Ko, Hoyoung Lee, Yewon Kim, and Seunghyeok Hong. 2024c. Llm-as-a-judge & reward model: What they can and cannot do. [arXiv preprint arXiv:2409.11239](#).
- Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2024d. [Kmmlu: Measuring massive multitask language understanding in korean](#). [arXiv preprint arXiv:2402.11548](#).
- Yejun Soun, Jaemin Yoo, Minyong Cho, Jihyeon Jeon, and U Kang. 2022. Accurate stock movement prediction with self-supervised learning from sparse noisy tweets. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 1691–1700. IEEE.



- Lin Sun, Guangxiang Zhao, Xiaoqi Jian, Yuhang Wu, Weihong Lin, Yongfu Zhu, Linglin Zhang, Jinzhu Wu, Junfeng Ran, Sai-er Hu, et al. 2025. Tinyr1-32b-preview: Boosting accuracy with branch-merge distillation. [arXiv preprint arXiv:2503.04872](#).
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. [arXiv preprint arXiv:2312.11805](#).
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. [arXiv preprint arXiv:2408.00118](#).
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. [arXiv preprint arXiv:2212.10560](#).
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhramil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. In [The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track](#).
- Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu Tang, Xiaowei Lv, et al. 2025. Light-r1: Curriculum sft, dpo and rl for long cot from scratch and beyond. [arXiv preprint arXiv:2503.10460](#).
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhakaran Kambadur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. [arXiv preprint arXiv:2303.17564](#).
- Qianqian Xie, Weiguang Han, Zhengyu Chen, Ruoyu Xiang, Xiao Zhang, Yueru He, Mengxi Xiao, Dong Li, Yongfu Dai, Duanyu Feng, et al. 2024a. Finben: A holistic financial benchmark for large language models. [Advances in Neural Information Processing Systems](#), 37:95716–95743.
- Yong Xie, Karan Aggarwal, and Aitzaz Ahmad. 2024b. Efficient continual pre-training for building domain specific large language models. In [Findings of the Association for Computational Linguistics ACL 2024](#), pages 10184–10201.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhao Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. [arXiv preprint arXiv:2304.12244](#).
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2.5 technical report. [arXiv preprint arXiv:2412.15115](#).
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. 2025. Limo: Less is more for reasoning. [arXiv preprint arXiv:2502.03387](#).
- Han Zhao, Haotian Wang, Yiping Peng, Sitong Zhao, Xiaoyu Tian, Shuaiting Chen, Yunjie Ji, and Xiang-gang Li. 2025. Am deepseek r1 distilled 1.4m.
- Qihao Zhao, Yangyu Huang, Tengchao Lv, Lei Cui, Qinzhen Sun, Shaoguang Mao, Xin Zhang, Ying Xin, Qiufeng Yin, Scarlett Li, et al. 2024. Mmlu-cf: A contamination-free multi-task language understanding benchmark. [arXiv preprint arXiv:2412.15194](#).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. [Advances in Neural Information Processing Systems](#), 36:46595–46623.
- Yuhang Zhou, Yuchen Ni, Yunhui Gan, Zhangyue Yin, Xiang Liu, Jian Zhang, Sen Liu, Xipeng Qiu, Guangnan Ye, and Hongfeng Chai. 2024. Are llms rational investors? a study on detecting and reducing the financial bias in llms. [arXiv preprint arXiv:2402.12713](#).

## A Further details on FINKRX-INSTRUCT

Here we report the average length of questions and responses using the Unimax tokenizer proposed by Chung et al. (2023).




Figure 4: Statistics of prompt and response length in FINKRX-INSTRUCT.

## B Training details for FINKRX

Axolotl (Axolotl AI, 2025) is used for the SFT and DPO training in Section 5.1. We train Qwen2.5-Math-7B-Instruct with DeepSpeed-Zero1 (Rajbhandari et al., 2020) on 8 H100 80GB GPUs for 25 hours. Hsu et al. (2024) is used for optimization. Table 5 and 6 are configurations used for SFT and DPO respectively.

Category	Section 5.1
Sequence Length	16,384
Learning Rate	$4 \times 10^{-5}$
Global Batch (Effective)	256
Learning Rate Scheduler	Cosine Decay
Warmup Ratio	0.05
Training Epochs	2

Table 5: SFT configuration details for Section 5.1.

Category	Section 5.1
Sequence Length	16,384
Learning Rate	$5 \times 10^{-6}$
Global Batch (Effective)	64
Learning Rate Scheduler	Cosine Decay
Warmup Ratio	0.05
Training Epochs	1

Table 6: DPO configuration details for Section 5.1.

## C Further details on private evaluation tools used by team Hi-Q

In Figure 3, we share private evaluation results conducted by Hi-Q. The evaluation is done on a private benchmark consisting of financial questions collected from sources such as AiHub<sup>1</sup> and KMMLU (Son et al., 2024d), to assess the model’s financial knowledge and capability. In particular, the private benchmark comprises the following subsets:

- **Accounting:** A private question set on Korean accounting.
- **Financial Accounting Generated:** Synthetically generated using GPT-4 on sample instances, following a Wang et al. (2022)-like approach (also applied to the Financial Market Generation subset).
- **KMMLU-accounting:** The accounting subset of the KMMLU dataset.
- **AiHUB-NC-MRC:** A dataset provided by AiHUB focusing on numerical computation and machine reading comprehension (AI HUB, 2025b).

<sup>1</sup><https://www.aihub.or.kr/>

- **AiHUB-FL-MRC**): A dataset provided by AiHUB focusing on financial and law machine reading comprehension (AI HUB, 2025a).

The benchmark evaluation results of methodologies attempted by Hi-Q are presented in Table 7.

Category	Subset	CPT	SFT	CPT+SFT	CPT+SFT+DPO
Financial Accounting	Accounting	32.0	39.0	41.0	<b>43.0</b>
	Financial_Accounting_Generated	55.0	71.0	70.0	<b>73.0</b>
	KMMLU_Accounting	37.0	42.0	41.0	<b>44.0</b>
	AiHUB-NC-MRC_calculation	55.0	57.0	60.0	<b>61.0</b>
	AiHUB-NC-MRC_boundary_extraction	85.0	91.0	<b>95.0</b>	<b>95.0</b>
	AiHUB-NC-MRC_multilateral_comparison	50.0	49.0	<b>59.0</b>	56.0
Financial Markets	AiHUB-FL-MRC_mcqa	52.0	<b>67.0</b>	66.0	62.0
	AiHUB-FL-MRC_process	80.0	84.0	84.0	<b>89.0</b>
	AiHUB-FL-MRC_answer_boundary	83.0	90.0	<b>94.0</b>	93.0
	Financial_Market_Generated	52.0	60.0	67.0	<b>69.0</b>
Avg.		58.1	65.0	67.7	<b>68.5</b>

Table 7: The internal benchmark results of Hi-Q. The **bold** font indicates that the highest score of each section.

## D Additional resources

In this section, we present additional resources that were excluded from the main text due to space constraints:

1. Figure 5: Model submission trends during preliminary rounds from Section 3.3.
2. Figure 6: Sample prompt used for LLM-as-a-Judge to evaluate the Open-Ended FinQA subset from Section 3.1 .




Figure 5: Model submission trends during the preliminary rounds.

## E Additional Results

In Tables 8 and 9, we present the performance of baseline models on the benchmarks used for the preliminary and main rounds, respectively. The tables include results for Qwen (1.5B and 7B)(Yang et al., 2024), Mistral (7B)(Jiang, 2024), GLM-4 (9B)(GLM et al., 2024), Llama 3/3.1 (8B)(Grattafiori et al., 2024), Amber (Liu et al., 2023), Phi 3.5 (mini)(Abdin et al., 2024), and Gemma2 (2B and 9B)(Team et al., 2024).

```
[System]
Please act as an impartial judge and evaluate the quality of the responses provided by two AI
assistants to the user question displayed below. You should choose the assistant that follows
the user's instructions and answers the user's question better. Your evaluation should consider
factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of
their responses. Begin your evaluation by comparing the two responses and provide a short
explanation. Avoid any position biases and ensure that the order in which the responses were
presented does not influence your decision. Do not allow the length of the responses to
influence your evaluation. Do not favor certain names of the assistants. Be as objective as
possible. After providing your explanation, output your final verdict by strictly following this
format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a
tie.

[User Question]
{question}

[The Start of Assistant A's Answer]
{answer_a}
[The End of Assistant A's Answer]

[The Start of Assistant B's Answer]
{answer_b}
[The End of Assistant B's Answer]
```

Figure 6: Prompt used for LLM-as-a-Judge to evaluate the Open-Ended FinQA subset



Models	Company	F&A	Stock	Agent	Market	Average
Qwen2-7B-Instruct	0.51	0.27	0.54	0.62	0.26	0.44
gemma-2-9b	0.31	0.25	0.54	0.43	0.27	0.36
Llama-3.1-8B	0.40	0.22	0.56	0.38	0.22	0.36
gemma-2-9b-it	0.39	0.28	0.55	0.32	0.25	0.36
Qwen2-7B	0.29	0.21	0.55	0.45	0.25	0.35
Llama-3.2-3B	0.43	0.23	0.55	0.32	0.20	0.35
Qwen2.5-7B	0.33	0.24	0.54	0.40	0.24	0.35
Meta-Llama-3-8B	0.38	0.23	0.56	0.30	0.23	0.34
Qwen2.5-1.5B-Instruct	0.27	0.26	0.54	0.30	0.21	0.34
Qwen2.5-3B	0.37	0.22	0.54	0.28	0.22	0.33
Qwen2.5-7B-Instruct	0.32	0.28	0.51	0.34	0.22	0.33
Mistral-7B-Instruct-v0.3	0.37	0.24	0.54	0.30	0.21	0.33
Llama-3.2-3B-Instruct	0.27	0.23	0.50	0.40	0.20	0.33
Qwen2.5-3B-Instruct	0.30	0.25	0.54	0.28	0.22	0.32
Llama-3.1-8B-Instruct	0.28	0.25	0.51	0.32	0.24	0.32
Qwen2.5-1.5B	0.30	0.25	0.56	0.26	0.23	0.31
Qwen2-1.5B-Instruct	0.26	0.22	0.53	0.33	0.24	0.31
gemma-2-2b	0.25	0.23	0.55	0.32	0.19	0.31
Llama-3.2-1B	0.27	0.26	0.55	0.23	0.18	0.30
Mistral-7B-Instruct-v0.2	0.28	0.21	0.54	0.26	0.23	0.30
gemma-2-2b-it	0.32	0.24	0.49	0.28	0.18	0.30
Qwen2-1.5B	0.25	0.20	0.54	0.30	0.22	0.30
Llama-3.2-1B-Instruct	0.28	0.23	0.52	0.21	0.19	0.29
AmberChat	0.26	0.23	0.53	0.23	0.21	0.29
Amber	0.25	0.23	0.54	0.23	0.21	0.29
Meta-Llama-3-8B-Instruct	0.28	0.24	0.53	0.21	0.21	0.29
Mistral-7B-Instruct-v0.1	0.22	0.21	0.55	0.28	0.22	0.29
Mistral-7B-v0.3	0.27	0.20	0.51	0.23	0.21	0.28
Phi-3.5-mini-instruct	0.25	0.25	0.54	0.17	0.18	0.28
Mistral-7B-v0.1	0.29	0.20	0.53	0.17	0.21	0.28

Table 8: Performance of base models in preliminary round

<b>Models</b>	<b>F&amp;A</b>	<b>Market</b>	<b>Open-Ended</b>	<b>Average</b>
gemma-2-9b-it	0.43	0.64	0.00	0.36
Qwen2.5-7B-Instruct	0.50	0.56	0.00	0.35
Qwen2-7B-Instruct	0.45	0.53	0.00	0.33
Qwen2.5-3B-Instruct	0.40	0.52	0.00	0.31
Qwen2.5-7B	0.37	0.41	0.00	0.28
Meta-Llama-3-8B-Instruct	0.37	0.43	0.00	0.27
Qwen2-7B	0.37	0.40	0.00	0.26
Llama-3.1-8B-Instruct	0.32	0.45	0.00	0.26
Phi-3.5-mini-instruct	0.38	0.36	0.00	0.25
gemma-2-9b	0.32	0.41	0.00	0.24
Qwen2.5-1.5B-Instruct	0.34	0.35	0.00	0.23
Qwen2.5-3B	0.30	0.36	0.00	0.22
Mistral-7B-Instruct-v0.3	0.32	0.30	0.00	0.21
Llama-3.1-8B	0.24	0.36	0.00	0.20
Llama-3.2-3B-Instruct	0.26	0.34	0.00	0.20
Qwen2-1.5B-Instruct	0.20	0.28	0.00	0.19
Qwen2.5-1.5B	0.27	0.28	0.00	0.18
Meta-Llama-3-8B	0.25	0.30	0.00	0.18
gemma-2-2b-it	0.22	0.32	0.00	0.18
gemma-2-2b	0.30	0.23	0.00	0.18
Mistral-7B-Instruct-v0.1	0.30	0.22	0.00	0.17
Mistral-7B-v0.1	0.24	0.26	0.00	0.17
Llama-3.2-3B	0.24	0.25	0.00	0.16
Llama-3.2-1B	0.24	0.25	0.00	0.16
AmberChat	0.24	0.24	0.00	0.16
Qwen2-1.5B	0.22	0.25	0.00	0.16
Mistral-7B-Instruct-v0.2	0.22	0.24	0.00	0.15
Mistral-7B-v0.3	0.21	0.22	0.00	0.14
Amber	0.24	0.19	0.00	0.14
Llama-3.2-1B-Instruct	0.22	0.20	0.00	0.14

Table 9: **Performance of base models in main round**