Frontiers of Information Technology & Electronic Engineering
www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com
ISSN 2095-9184 (print); ISSN 2095-9230 (online)
E-mail: jzus@zju.edu.cn

Correspondence:

# When DeepSeek-R1 meets financial applications: benchmarking, opportunities, and limitations

Shuoling LIU[1,2], Liyuan CHEN[2], Jiangpeng YAN[2,4], Yuhang JIANG[2], Xiaoyu WANG[2], Xiu LI[4], Qiang YANG[‡1,3]

[1] *The Hong Kong University of Science and Technology, Hong Kong 999077, China*
[2] *E Fund Management Co., Ltd., Guangzhou 510000, China*
[3] *Webank AI, Shenzhen 518054, China*
[4] *Tsinghua Shenzhen International Graduate School, Shenzhen 518055, China*
E-mail: {liushuoling chenly, yanjiangpeng, jiangyuhang, wangxiaoyu02}@efunds.com.cn; li.xiu@sz.tsinghua.edu.cn; qyang@cse.ust.hk
Received Apr.10, 2025; Revision accepted Apr. 28, 2025; Crosschecked

Unedited

**Abstract:** The ways in which the recent progresses of reasoning large language models (LLMs), especially the new open-source model DeepSeek-R1, can benefit financial services is an underexplored problem. While LLMs have ignited numerous applications within the financial sector, including financial news analysis and general customer interactions, DeepSeek-R1 further unlocks the advanced reasoning ability with multiple reinforcement learning–integrated training steps for more complex financial queries and provides distilled student models for resource-constrained scenarios. In this paper, we first introduce the technological preliminaries of DeepSeek-R1. Subsequently, we benchmark the performance of DeepSeek-R1 and its distilled students on two public financial question–answer datasets as a starting point for interdisciplinary research on financial artificial intelligence (AI). Then, we discuss the opportunities that DeepSeek-R1 offers to current financial services, its current limitations, and three future research directions. In conclusion, we argue for a proper approach to adopt reasoning LLMs for financial AI.

**Key words:** Large language models; Reasoning; Artificial intelligence; Financial technology; FinTech
https://doi.org/10.1631/FITEE.2500227 **CLC number:**

## 1 Introduction

Fueled by Web-scale datasets, high-performance computing devices, and deep learning algorithms, recent years have witnessed the emergence of large language models (LLMs) (Chang et al., 2024), which are paving the way toward artificial general intelligence (AGI) (Goertzel, 2014). Built on the Transformer (Vaswani et al., 2017) architecture and trained with extensive text data corpora, GPT-3 (Brown et al., 2020), released by OpenAI in 2020, first demonstrated an unprecedented ability to understand, generate, and manipulate human language in a manner that closely mimics human-like intelligence, with a number of follower LLMs including Google's Gemini (Team et al., 2023), Alibaba's Qwen (Bai JZ et al., 2023), and so on. However, early LLMs struggled with hallucination problems (Huang et al., 2025) and usually failed to solve complex reasoning problems. With the notable breakthrough of the "chain-of-thought" (CoT) prompting technique (Wei et al., 2022), OpenAI released the first reasoning LLM called OpenAI-o1 (Jaech et al., 2024) in 2024 by training this model to increase the length of self-generated step-by-step human-like reasoning processes during the post-training process. Despite OpenAI-o1's remarkable

‡ Corresponding author
© Zhejiang University Press 2025

capabilities, which can achieve human-level or superior performance on various tasks (including competitive programming, advanced mathematics, and scientific research), access to OpenAI-o1 is exclusively available to a limited number of subscribers, restricting its widespread utilization and impact across broader communities. In response to this demand, the introduction of DeepSeek-R1 (Guo DY et al., 2025) in January 2025 marked a significant milestone. As the first open-weight reasoning LLM that can achieve a similar performance as OpenAI-o1, DeepSeek-R1 gained global attention by democratizing access to advanced artificial intelligence (AI) technologies, making them nearly cost-free. To take it a step further, DeepSeek also open-sourced a series of distilled student models that can be deployed in resource-constrained conditions.

DeepSeek-R1 and its deployment-friendly distilled student models have fostered swift applications in various fields (Peng et al., 2025; Qian et al., 2025), where financial contexts arguably demand the highest levels of reasoning capabilities and data-processing prowess due to their inherent complexity and risk. In the FinTech sector, while models such as BloombergGPT (Wu et al., 2023) and XuanYuan (Zhang XY and Yang, 2023) have been previously developed specifically for finance, they often fall short when it comes to sophisticated reasoning tasks such as finance numeric analysis, law and regulation logic thinking, and so on. This paper aims to discuss DeepSeek-R1 from the perspective of the FinTech industry. Given its open-source nature and advanced reasoning abilities, DeepSeek-R1 offers a promising foundation upon which more specialized models tailored to the exacting demands of financial service can be constructed.

In this paper, we first introduce the technological preliminaries of DeepSeek-R1, emphasizing its contributions to multistep reinforcement-learning (RL)-based training strategy and the exploration of distilling reasoning ability using student models with fewer parameters. Subsequently, we benchmark the performance of DeepSeek-R1 and its distilled students on two public financial question–answer (QA) datasets as a starting point for interdisciplinary research on financial AI, compared to common baseline non-reasoning LLMs. Then, we discuss how DeepSeek-R1 can benefit current financial services with its advanced reasoning ability. While acknowledging the strides made, we also conclude that there are potential drawbacks to current reasoning LLMs and discuss three future research directions. In conclusion, we argue for properly adopting reasoning LLMs for financial AI research.

## 2 Preliminaries of DeepSeek-R1

In this section, we briefly introduce some preliminaries of reasoning LLMs, which are divided into three parts: the concept of CoTs, the DeepSeek-R1 training process, and DeepSeek-R1 distillation exploration.

### 2.1 Concept of CoTs

Following the emergence of GPT-3, researchers recognized the latent reasoning capabilities (Webb et al., 2023) inside LLMs, despite unreliability in some tasks. Specifically, it was discovered that even simple prompt engineering, such as adding “think step by step” (Kojima et al., 2022), could guide LLMs to generate human-like sequential reasoning texts and improve their performance for reasoning tasks. Fig. 1 shows an example where DeepSeek-V3 (Liu AX et al., 2024) is prompted to raise and answer a financial question with and without CoT. This foundational approach evolved an active research topic, that is, how to generate high-quality CoT of different types, including CoT with self-consistency (Wei et al., 2022), “tree-of-thought” (Yao et al., 2023), “graph-of-thought” (Besta et al., 2024), and so on. According to OpenAI-o1’s technology report (Jaech et al., 2024), OpenAI researchers integrated CoT reasoning deeply into its training process, encouraging their model to make longer contemplation of CoT sequences before giving final answers through RL (Bai YT et al., 2022; Lightman et al., 2023; Zelikman et al., 2024). This advancement allows OpenAI-o1 to tackle more-complex reasoning tasks and provide clearer insights into its problem-solving processes, surpassing previous models such as GPT-4 in multistep reasoning tasks.

Unedited

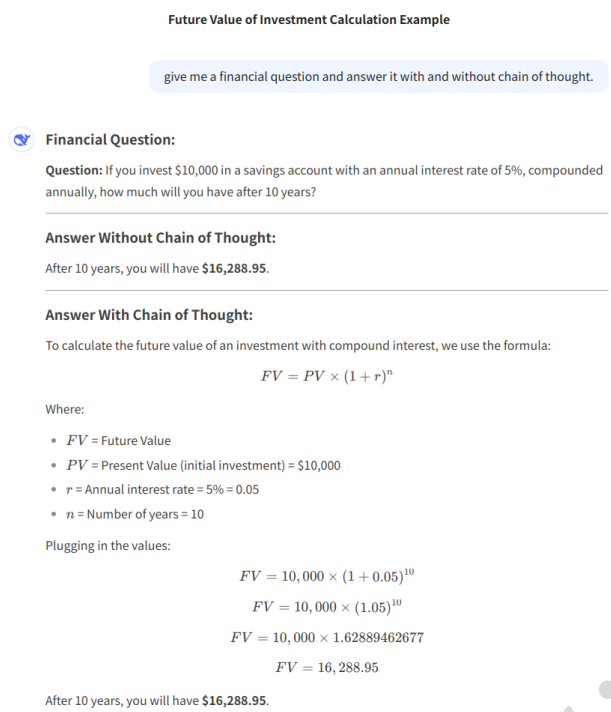


**Fig. 1 Prompting DeepSeek-V3 to raise and answer a financial question with and without chain-of-thought**

## 2.2 DeepSeek-R1 training process

To learn to generate accurate step-by-step reasoning CoTs, DeepSeek-R1 was trained under four carefully designed stages, as shown in Fig. 2a, including the following:

(1) Cold start supervised fine-tuning (SFT): thousands of handcrafted high-quality CoTs were collected to train DeepSeek-V3 noted as $M_0$, a powerful LLM with 671B parameters, and obtain model $M_1$ as a cold start point.

(2) Reasoning-focused RL: the DeepSeek team proposed a large-scale RL training strategy, whereby the GRPO algorithm (Shao et al., 2024) and rule-based reward models are integrated, to guide $M_1 \rightarrow M_2$ on reasoning-focused tasks. As a result, $M_2$ can learn how to perform multistep thinking processes on coding, mathematics, science, and logic reasoning.

(3) Rejection sampling and SFT: the generated high-quality CoTs ($n$=600 k) were sampled and refined from the responses of $M_2$, with additional data on more reasoning tasks. These data together with 200k nonreasoning data, including writing, factual question answering, and role-playing tasks, were used to supervise the $M_2 \rightarrow M_3$ training process.

(4) Final RL alignment: $M_3$ further went through a secondary RL stage to ensure its alignment to human preferences for helpfulness and safety, to reach DeepSeek-R1 $M_4$.

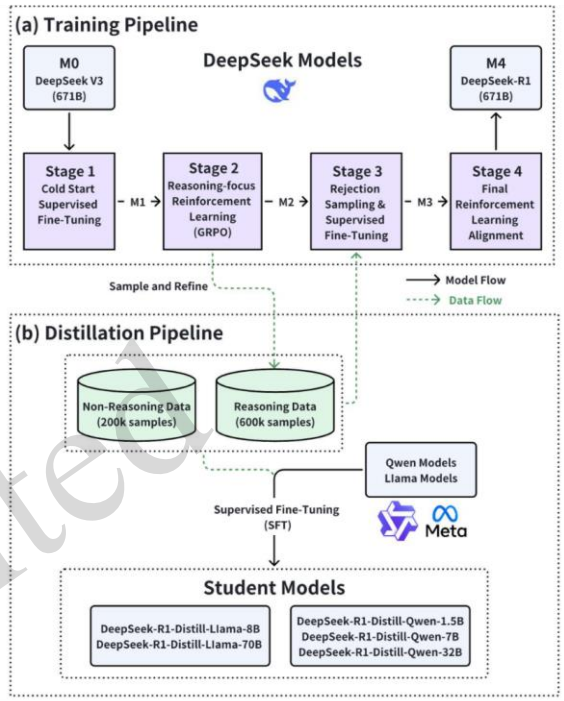


**Fig. 2 (a) Four training stages of DeepSeek-R1, where $M_0$–$M_4$ stand for models in different stages, and group relative policy optimization (GRPO) algorithm (Shao** et al.**, 2024) serves as the core solution for Stage 2, and (b) knowledge distillation pipeline**

Given that the DeepSeek founder has a background as a quant investing researcher, it is reasonable to assume that the DeepSeek team incorporated finance-related reasoning data at Stages 2 and 3. This addition is of two-fold significance. Firstly, financial knowledge, being highly structured (Lusardi and Mitchell, 2014), can enhance the model's overall reasoning capabilities. The well-organized nature of financial concepts and rules provides a solid framework for the model to learn and improve its logical thinking. Secondly, as a result, the model is enabled to perform exceptionally well on financial reasoning-type questions. Therefore, DeepSeek-R1 performs well in our following financial QA experiments. It also can be concluded that high-quality CoTs play an essential role in the success of DeepSeek-R1. Note that DeepSeek teams

Unedited

open-sourced the weight of DeepSeek-R1 $M_4$ without any data used in their training process.

### 2.3 DeepSeek-R1 distillation exploration

According to the training process, the original DeepSeek-R1 shares the same model size as DeepSeek-V3 with 671B parameters, which is not suitable to be adopted in resource-constrained conditions. To equip smaller LLMs with reasoning ability, DeepSeek researchers explored performing SFT on 1.5B-70B Qwen (Bai JZ et al., 2023) and Llama (Touvron et al., 2023) LLMs on the 800k samples in Stage 3, as shown in Fig. 2b, obtaining models including DeepSeek-R1-Distill-Qwen-32B, DeepSeek-R1-Distill-Llama-8B, and so on. The DeepSeek team has shown that distilling knowledge from these carefully sampled data is effective, and smaller LLMs can achieve competitive results on math and coding tasks, compared to DeepSeek-R1.

It is reminded that these student models are actually distilled from $M_2$, as we described above, rather than from $M_4$. DeepSeek's exploration is an early trial to demonstrate the effectiveness of knowledge distillation (Hinton et al., 2015). We believe that ways to obtain better student models from DeepSeek-R1 $M_4$ will be a promising future research direction.

## 3 Benchmarking reasoning LLMs on financial question answering

In this section, we benchmark DeepSeek-R1 and its distilled student models on two public financial QA datasets with common LLM baselines, including DeepSeek-V3 (Liu AX et al., 2024), OpenAI-o1-preview (Jaech et al., 2024), and OpenAI GPT-4 (Achiam et al., 2023), to compare their performance in the finance area.

### 3.1 The dataset and evaluation protocols

(1) Datasets: FinEval (Zhang XY and Yang, 2023) contains a total of 1321 single-choice QA pairs, which are categorized into four primary types: finance, economy, accounting, and certificate. Note that FinEval includes 4661 questions in total, where the answers to 3340 questions are not released. FinanceIQ (Zhang XY and Yang, 2023) includes a total of 7173 single-choice QA pairs from 10 subjects, such as the China Insurance Certification Exam, the US Certified Public Accountant exam, financial mathematics exams, and so on.

(2) Evaluation protocols: all the LLMs are prompted to give final answers with a JSON format CoTs and answers using the prompting template shown in Fig. 3. It is noted that DeepSeek-R1 answers questions with a long reasoning content with "<think>...</think>" tag. We report the accuracy on two datasets, where accuracy= $\frac{\text{number of correct answered questions}}{\text{total number of questions}}$.

**Fig. 3 Two cases of DeepSeek-R1's answers to finance math questions. The original Chinese and translated English contents are shown for readability. DeepSeek-R1's reasoning contents (>1000 characters) are in gray and shortened for space limitation. The same prompting template is used for all LLMs**

### 3.2 Main results

The experimental results are summarized in Table 1. We can find that reasoning LLMs DeepSeek-R1/OpenAI-o1-preview significantly outperform

Unedited

their predecessor general LLMs DeepSeek-V3/GPT-4, indicating that answering financing questions needs advanced reasoning ability. Furthermore, it can be observed that DeepSeek-R1/DeepSeek-V3 perform better than OpenAI-o1-preview/GPT-4 by a significant margin. We assume that this is because there existed data biases in the datasets used, whereby most financial questions were sampled from Chinese finance-related examinations and OpenAI-o1/GPT-4 lacking Chinese financial knowledge. We can also find that DeepSeek-R1-Distill-Qwen-32B has very competitive overall performance on these two datasets, considering that its learning parameters (32B) are much fewer than in DeepSeek-V3/R1 (671B). With 7B learning parameters, which are much fewer than in other models, DeepSeek-R1-Distill-Qwen-7B cannot achieve satisfactory results.

**Table 1 Benchmark different LLMs on two financial QA datasets with accuracy rate in terms of correctly answering questions**

| Models | FinEval | FinanceIQ | FinanceIQ's Financial Math |
|---|---|---|---|
| DeepSeek-V3 | 82.20% | 73.46% | 54.84% |
| DeepSeek-R1 | 91.52% | 86.35% | 82.80% |
| DeepSeek-R1-Distill-Qwen-32B | 82.73% | 77.19% | 56.99% |
| DeepSeek-R1-Distill-Qwen-7B | 51.17% | 45.63% | 18.28% |
| GPT-4 | 70.03% | 63.40% | 37.50% |
| OpenAI-o1-preview | 78.11% | 79.20% | 72.73% |

Note: Best results are in bold.

Unedited

In addition to reporting the overall accuracy on FinEval and FinanceIQ, we specifically highlight the performance of the Financial Mathematics subset of FinanceIQ, as all models achieve their lowest scores in this category compared to other subjects. It can be noted that DeepSeek-R1 shows a relatively more pronounced performance advantage in financial mathematics over other models. This observation further underscores the importance of specialized knowledge and advanced reasoning capabilities in handling complex financial calculations.

We present visualizations of two DeepSeek-R1 responses to financial mathematics questions in Fig. 3. DeepSeek-R1 can calculate precisely and handle basic compound interest calculations in finance. Nevertheless, because of its restricted advanced financial knowledge, it may be confused by irrelevant conditions in certain complex problems.

## 4 Discussion

### 4.1 Opportunities for financial application of DeepSeek-R1

Based on the experimental results, DeepSeek-R1 is able to play the following three roles with its current financial knowledge.

(1) Financial text processor: DeepSeek-R1 is capable of analyzing financial news, company announcements, and other textual contents in a more structured manner through logical reasoning. In our experiments, DeepSeek-R1 achieves an accuracy >80% on questions related to the judgment and interpretation of Chinese economic laws and regulations. Therefore, DeepSeek-R1 can serve as a financial text processor to help financial officers in scenarios, such as financial meeting summary, research report analysis, and compliance review.

(2) Financial numerical analysis aide: DeepSeek-R1 significantly outperforms other models in FinanceIQ's Financial Mathematics category, as shown in Table 1. Accordingly, DeepSeek-R1 can provide users with numerical analysis insights that have a certain level of credibility in various financial services, serving as an analysis assistant. However, it is important to note that while the calculations are deemed credible, they are not infallible and should be used with due caution.

(3) Financial educator**:** considering that the text questions were collected from multiple qualification examinations for various financial professions, and because DeepSeek-R1 can provide not only highly accurate choices but also detailed problem-solving analysis steps, DeepSeek-R1 can act as a talented educator, tutoring beginners in mastering relevant financial knowledge. However, during this process, the involvement of seasoned professionals is still required to identify any errors in DeepSeek-R1's responses.

### 4.2 Limitations of DeepSeek-R1 in financial applications

While DeepSeek-R1 has achieved satisfactory performance, it is essential to recognize that there is still a gap between the examined questions and real-world financial practices. We identify the following four limitations in applying DeepSeek-R1 to the finance domain:

(1) Lack of evaluation on large-scale financial reasoning datasets: our experiments are based on small-scale, relatively simple single-choice question datasets as an early step, and DeepSeek-R1 has not been comprehensively evaluated for financial fields. It is necessary to establish more large-scale datasets with a greater diversity of financial question types and more advanced reasoning application scenarios. Such datasets can be used to not only evaluate the performance of reasoning LLMs but also lay the data foundation for further improving the performance of reasoning LLMs in financial domains.

(2) Reduced but still-existing hallucination: the financial sector is an area under high regulation, and even a minor error can have far-reaching consequences. Although DeepSeek-R1 can reduce the impact of hallucinations by demonstrating step-by-step reasoning process, hallucination risks still persist and are more difficult to detect. There is a need to explore more-advanced strategies to enhance the interpretability of DeepSeek-R1's black-box nature and reduce hallucinations, such as combining neural and symbolic operations (Chervonyi et al., 2025), developing trustworthy LLMs (Liu Y et al., 2023), and so on.

(3) Lack of model collaboration: for elementary financial text analysis tasks, such as corporate entity extraction, strong reasoning capabilities are not always necessary. Generation of excessively long thought chains by DeepSeek-R1 may instead reduce processing efficiency, while DeepSeek-V3, which directly provides answers, may be more cost-effective. Thus, DeepSeek-R1 needs to learn more from humans, who can seamlessly switch between fast (System 1) and slow (System 2) thinking (Hua and Zhang, 2022) and thus strike a balance through model collaboration.

(4) Lack of multimodality ability: beyond text-based data, the financial services domain requires handling multimodal data, including charts, images, audio, and price time-series data. At present, DeepSeek-R1 is limited to the text modality. Expanding its capabilities from text-centric domains to multimodal scenarios remains for exploration.

### 4.3 Future directions for DeepSeek-R1 in financial applications

DeepSeek-R1 has already led the discussion among AI practitioners across various industries. From the perspective of FinTech, here, we put forward three potential future research directions related to DeepSeek-R1:

(1) Federated financial reasoning LLMs: considering the DeepSeek-R1 training process, the scaling law (Kaplan et al., 2020) of high-quality reasoning data still holds. However, building large-scale financial datasets is challenging, considering privacy laws and financial regulations. To address this issue, the paradigm of federated learning (Fan et al., 2023; Kuang et al., 2024) can be adopted for training financial reasoning LLMs, where data privacy can be preserved. An implementation of federated financial LLM training frameworks is illustrated in Fig. 4. Moreover, protecting the high-quality reasoning CoT data generated by financial LLMs is an important research topic for federated financial reasoning LLMs. Exploring how to implicitly inject "copyright-protection watermark" information into the reasoning CoT is a worthy research direction.

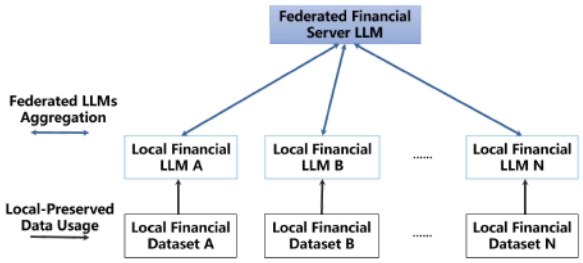


**Fig. 4 An implementation of federated financial LLM training frameworks**

(2) Multimodel collaborative financial AI: currently, DeepSeek-R1 requires extensive CoTs to solve every given problem. It is necessary to consider integrating fast thinking to more efficiently address problems in financial scenarios. Therefore, future financial AI should be constructed through the collaboration of fast-thinking and slow-thinking models, and there have been some early attempts (Sun et al., 2024; Miao et al., 2024). In addition to this

Unedited

paradigm combination, considering that DeepSeek-R1 has also explored distilling easy-to-deploy student models with certain reasoning capabilities, methods for the cooperation between large and small language models to achieve a more computation-friendly and low-carbon financial AI need to be explored.

(3) Multimodality financial agent: building of advanced agent systems for financial services requires handling data from multiple modalities. Key research directions include developing multimodality reasoning large models that can understand and generate different modalities of financial data. Previous studies (Zhang WT et al., 2024; Daiya and Lin, 2021) have demonstrated that integrating multimodality features can further enhance performance compared to relying solely on single-modality text data.

## 5 Conclusions

In conclusion, DeepSeek-R1 is set to inspire in-depth research and diverse applications in financial services. This paper details DeepSeek-R1's technology, benchmarks its performance on two financial QA datasets, and discusses its opportunities and limitations. As an early effort to spark interest in integrating reasoning LLMs into financial AI, our study has the following limitation: it has only been tested on single-choice questions with insufficient model comparisons. We leave the upgradation of this research for our future work and also as open problems to the research community. This work proves DeepSeek-R1's potential in finance, and we advocate proper use of reasoning LLMs and suggest future financial AI research directions.

### Contributors

Shuoling LIU, Liyuan CHEN, and Jiangpeng YAN designed the research. Shuoling LIU, Jiangpeng YAN, Yuhang JIANG, and Xiaoyu WANG processed the data. Shuoling LIU, Jiangpeng YAN, Yuhang JIANG, and Xiaoyu WANG drafted the paper. Liyuan CHEN and Xiu LI helped organize the paper. Xiu LI and Qiang YANG revised and finalized the paper.

### Conflict of interest

All the authors declare that they have no conflict of interest.

### References


Achiam J, Adler S, Agarwal S, et al., 2023. GPT-4 technical report. https://arxiv.org/abs/2303.08774

Bai JZ, Bai S, Chu YF, et al., 2023. Qwen technical report. https://arxiv.org/abs/2309.16609

Bai YT, Jones A, Ndousse K, et al., 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. https://arxiv.org/abs/2204.05862

Besta M, Blach N, Kubicek A, et al., 2024. Graph of thoughts: solving elaborate problems with large language models. Proc 38th AAAI Conf on Artificial Intelligence, p.17682-17690. https://doi.org/10.1609/aaai.v38i16.29720

Brown TB, Mann B, Ryder N, et al., 2020. Language models are few-shot learners. Proc 34th Int Conf on Neural Information Processing Systems, article 159.

Chang YP, Wang X, Wang JD, et al., 2024. A survey on evaluation of large language models. *ACM Trans Intell Syst Technol*, 15(3):39. https://doi.org/10.1145/3641289

Chervonyi Y, Trinh TH, Olšák M, et al., 2025. Gold-medalist performance in solving Olympiad geometry with AlphaGeometry2. https://arxiv.org/abs/2502.03544

Daiya D, Lin C, 2021. Stock movement prediction and portfolio management via multimodal learning with transformer. Proc IEEE Int Conf on Acoustics, Speech and Signal Processing, p.3305-3309. https://doi.org/10.1109/ICASSP39728.2021.9414893

Fan T, Kang Y, Ma GQ, et al., 2023. FATE-LLM: a industrial grade federated learning framework for large language models. https://arxiv.org/abs/2310.10049

Gemini Team, Google, 2023. Gemini: a family of highly capable multimodal models. https://arxiv.org/abs/2312.11805

Goertzel B, 2014. Artificial general intelligence: concept, state of the art, and future prospects. *J Artif Gen Intell*, 5(1):1-48. https://doi.org/10.2478/jagi-2014-0001

Guo DY, Yang DJ, Zhang HW, et al., 2025. DeepSeek-R1: incentivizing reasoning capability in LLMs via reinforcement learning. https://arxiv.org/abs/2501.12948

Guo X, Xia HT, Liu ZW, et al., 2023. FinEval: a Chinese financial domain knowledge evaluation benchmark for large language models. https://arxiv.org/abs/2308.09975

Hinton G, Vinyals O, Dean J, 2015. Distilling the knowledge in a neural network. https://arxiv.org/abs/1503.02531

Hua WY, Zhang YF, 2022. System 1+system 2=better world: neural-symbolic chain of logic reasoning. Proc Findings of the Association for Computational Linguistics, p.601-612. https://doi.org/10.18653/v1/2022.findings-emnlp.42

Huang L, Yu WJ, Ma WT, et al., 2025. A survey on hallucination in large language models: principles, taxonomy, challenges, and open questions. *ACM Trans Inf Syst*, 43(2):42. https://doi.org/10.1145/3703155


Unedited

Jaech A, Kalai A, Lerer A, et al., 2024. OpenAI o1 system card. https://arxiv.org/abs/2412.16720

Kaplan J, McCandlish S, Henighan T, et al., 2020. Scaling laws for neural language models. https://arxiv.org/abs/2001.08361

Kojima T, Gu SS, Reid M, et al., 2022. Large language models are zero-shot reasoners. Proc 36th Int Conf on Neural Information Processing Systems, article 1613.

Kuang WR, Qian BC, Li ZT, et al., 2024. FederatedScope-LLM: a comprehensive package for fine-tuning large language models in federated learning. Proc 30th ACM SIGKDD Conf on Knowledge Discovery and Data Mining, p.5260-5271. https://doi.org/10.1145/3637528.3671573

Lightman H, Kosaraju V, Burda Y, et al., 2024. Let's verify step by step. Proc 12th Int Conf on Learning Representations. p.1-24.

Liu AX, Feng B, Xue B, et al., 2024. DeepSeek-V3 technical report. https://arxiv.org/abs/2412.19437

Liu Y, Yao YS, Ton JF, et al., 2023. Trustworthy LLMs: a survey and guideline for evaluating large language models' alignment. https://arxiv.org/abs/2308.05374

Lusardi A, Mitchell OS, 2014. The economic importance of financial literacy: theory and evidence. *J Econ Lit*, 52(1): 5-44. https://doi.org/10.1257/jel.52.1.5

Miao XP, Oliaro G, Zhang ZH, et al., 2024. SpecInfer: accelerating large language model serving with tree-based speculative inference and verification. Proc 29th ACM Int Conf on Architectural Support for Programming Languages and Operating Systems, p.932-949. https://doi.org/10.1145/3620666.3651335

Peng YF, Malin BA, Rousseau JF, et al., 2025. From GPT to DeepSeek: significant gaps remain in realizing AI in healthcare. *J Biomed Inform*, 163:104791. https://doi.org/10.1016/j.jbi.2025.104791

Qian LF, Zhou WP, Wang Y, et al., 2025. Fino1: on the transferability of reasoning enhanced LLMs to finance. https://arxiv.org/abs/2502.08127

Shao ZH, Wang PY, Zhu QH, et al., 2024. DeepSeekMath: pushing the limits of mathematical reasoning in open language models. https://arxiv.org/abs/2402.03300

Sun GY, Jin MY, Wang ZT, et al., 2024. Visual agents as fast and slow thinkers. https://arxiv.org/abs/2408.08862

Touvron H, Martin L, Stone K, et al., 2023. Llama 2: open foundation and fine-tuned chat models. https://arxiv.org/abs/2307.09288

Vaswani A, Shazeer N, Parmar N, et al., 2017. Attention is all you need. Proc 31st Int Conf on Neural Information Processing Systems, p.6000-6010.

Webb T, Holyoak KJ, Lu HJ, 2023. Emergent analogical reasoning in large language models. *Nat Hum Behav*, 7(9):1526-1541. https://doi.org/10.1038/s41562-023-01659-w

Wei J, Wang XZ, Schuurmans D, et al., 2022. Chain-of-thought prompting elicits reasoning in large language models. Proc 36th Int Conf on Neural Information Processing Systems, article 1800.

Wu SJ, Irsoy O, Lu S, et al., 2023. BloombergGPT: a large language model for finance. https://arxiv.org/abs/2303.17564

Yao SY, Yu D, Zhao J, et al., 2023. Tree of thoughts: deliberate problem solving with large language models. Proc 37th Int Conf on Neural Information Processing Systems, article 517.

Zelikman E, Harik GR, Shao YJ, et al., 2024. Quiet-STaR: language models can teach themselves to think before speaking. Proc First Conf on Language Modeling.

Zhang WT, Zhao LX, Xia HC, et al., 2024. A multimodal foundation agent for financial trading: tool-augmented, diversified, and generalist. Proc 30th ACM SIGKDD Conf on Knowledge Discovery and Data Mining, p.4314-4325. https://doi.org/10.1145/3637528.3671801

Zhang XY, Yang Q, 2023. XuanYuan 2.0: a large Chinese financial chat model with hundreds of billions parameters. Proc 32nd ACM Int Conf on Information and Knowledge Management, p.4435-4439. https://doi.org/10.1145/3583780.3615285

Unedited