Research and Analysis on Chain of Thought (CoT) Reasoning and Interpretability in Large Language Models

Pengyu Liao

College of Information and Engineering, Hefei University of Econmics, Hefei, China
lllzz6662222@gmail.com

Abstract. As an important form of Large Language Models (LLM), Chain-of-Thought (CoT) has made breakthroughs in logical consistency, interpretability, and task accuracy by guiding the model to generate answers by stepwise reasoning. This paper systematically reviews the research progress of CoT reasoning in the last five years (2021-2025) and analyse it from three main dimensions, First, the evolution and limitations of chain-of-thinking reasoning: from zero-sample and few-sample CoT to structured reasoning methods such as Self-Consistency and Tree-of-Thought (ToT), although the depth of reasoning of the model has been improved effectively, there are still problems such as unstable reasoning and insufficient generalisation. Second, the verifiability and consistency of the reasoning chain: researchers have gradually paid attention to how to verify the correctness of the reasoning process and ensure the logical closure between the reasoning steps and the conclusions. Thirdly, the development of tree and multi-path reasoning; through multi-path exploration, path weighting and aggregation mechanism, the reasoning robustness and diversity of the model under complex tasks are improved. This paper concludes that the current CoT still has challenges in terms of stability, realism and cross-domain transferability. Future research should focus on three aspects: establishing a verifiable reasoning framework, developing multi-level structured thought modelling, and constructing a unified interpretability evaluation system, in order to promote the large language model to achieve more reliable and transparent intelligent reasoning.

Keywords: Large Language Models, Chain of Thought, Interpretability, Multi-path reasoning, Trusted AI.

1 Introduction

The application of AI technology in the fields of natural language processing, knowledge reasoning and human-computer interaction has become more and more widespread, among which Large Language Models (LLMs), with their powerful language generation and reasoning capabilities, have driven AI development from the stage of perceptual intelligence to that of cognitive intelligence. However, as the scale of the models increases, their internal decision-making process be-comes more and

© The Author(s) 2026
K. Subramanian (ed.), *Proceedings of the International Workshop on Advances in Deep Learning for Image Analysis and Computer Vision (IWADIC 2025)*, Advances in Computer Science Research 128,
https://doi.org/10.2991/978-94-6239-648-7_56

more complex and inexplicable, and researchers begin to pay attention to how to make the models show traceable reasoning process while generating results.

Against this background, Wei et al, proposed for the first time the Chain-of-Thought (CoT) cueing strategy, which enables language models to output explicit intermediate reasoning steps before answering questions [1]. This innovation not only significantly improves the accuracy of tasks such as arithmetic reasoning and logic quizzing, but also makes the decision-making process of the model more transparent. Subsequently, Wang et al. improved the stability of reasoning through self-consistency strategy [2], and Kojima et al, proposed a "zero-sample CoT" method to prove that the model can

demonstrate reasoning ability even in the unsupervised situation [3], and Gao et al.'s procedural assisted language model PAL verifies the correctness of the results by executing the reasoning procedure, which is a good example for the reasoning process. Gao et al. lay the foundation for verifiable inference by executing the inference procedure to verify the correctness of the results [4].

At the same time, the research in the direction of structured reasoning is also developing. Yao et al. proposed the Tree-of-Thoughts (ToT) model, which realizes multi-path parallel reasoning through tree search [5]; Long et al. further proposed the Graph-of-Thoughts (GoT) model, which introduces graph-structured reasoning to achieve global information sharing [6]. These research have promoted CoT from linear logic to structured exploration, and become an important direction of Al interpretability research in recent years.

In summary, the proposal of chain-of-thinking reasoning is not only a key factor for model performance improvement, but also an important way to achieve trustworthy Al. The paper will systematically review the core research results on CoT reasoning between 2021 and 2025, and start from three main lines: (1) the evolution and limitations of chain-of-thinking reasoning; (2) the verifiability and consistency of interpretation of reasoning chain; and (3) the development trend of tree and multi-path reasoning. Finally, on this basis, this paper puts forward the current problems and challenges in the field, and discusses and looks forward to the future research directions.

2 The evolution and limitations of chain-of-thinking reasoning

Chain of Thought (CoT) was originally proposed by Wei et al. to guide models to improve the correctness of complex tasks by generating intermediate reasoning steps [1], The core idea is to simulate the human thinking process and decompose problems into interpretable logical chains. Experiments show that in arithmetic reasoning tasks such as GSM8K. the accuracy of the model using CoT hints is improved by more than 20% compared with traditional direct answers. However, this method is highly dependent on the model size, and small models are often unable to generate reasonable reasoning chains, reflecting the obvious limitation of "size dependence".

To alleviate this problem. Wang et al. proposed a Self-Consistency strategy, which improves the robustness of the inference results by generating multiple inference paths and using a voting mechanism to determine the final answer [2]. This study shows that diversified inference paths can reduce the randomness error to a certain extent. However, with the increase of the number of paths, the computational complexity increases significantly, which leads to the decrease of reasoning efficiency.

Kojima et al. proposed the "zero-sample CoT" method, which stimulates the inference potential of the model through the simple cue 'Let's think step by step' [3]. The study shows that the large model itself has an implicit reasoning ability, and the prompt can be used as a triggering mechanism. Although this method significantly broadens the scope of CoT, its inference results are highly dependent on the cue expression, which shows the problem of "cue sensitivity".

From a comprehensive point of view, the development of chain-of-thought reasoning has gone through three stages: initial explicit logic expression, statistical consistency enhancement and zero-sample expansion. Its main limitations are: (i) the truthfulness of reasoning is difficult to be quantified; (ii) the stability of cues is low; and (iii) there is a lack of validation mechanism for the reasoning results. These problems have pushed researchers to further explore how to improve the verifiability and consistency of CoT.

3 Verifiability and interpretative consistency of the chain of reasoning

In order to solve the problem of "pseudo-interpretation" in CoT reasoning, Gao et al. proposed a programme-assisted language mode. PAL, which transforms linguistic reasoning into an executable pro-gramme. The model first generates pseudo-code, and then verifies the correctness of the reasoning through the results of the programme execution, so that the interpretation and the results form a closed loop [4]. The method significantly improves the accuracy in mathematical and logical reasoning tasks and provides a viable paradigm for subsequent verification-based reasoning,

Lyu et al. further proposed the Faithful CoT framework, which emphasises the causal consistency between the chain of reasoning and the final answer. They designed a 'fidelity indicator' to measure the actual contribution of intermediate reasoning steps to the outcome, effectively distinguishing between" reasonable explanations" and "true reasoning"[7].

Zheng et al. proposed a reflection mechanism, Reflexion [7], which allows the model to self-correct after generating a chain of reasoning. Through the introduction of metacognitive feedback, the model can identify logical loopholes and generate corrected answers [8], while Burns et al. proposed the Explain- Verify framework [8], which combines "explanation" and "verification" to form a closed loop of

explanation generation and verification, making the reasoning process reviewable and traceable. The explanation-verify framework was proposed by Burns et al. in [8], which combines "explanation" and "verification" to form a closed loop of explanation generation and verification, so that the reasoning process can be reviewed and traced [9].

In general, the research on verifiable reasoning has shifted CoT from "result-oriented" to "process-oriented", emphasising the consistency between the reasoning chain and the decision logic of the model. The future development direction may include combining formal logic, symbolic reasoning and statistical learning to build an automatically verifiable reasoning system.

4 Tree reasoning and multi-path exploration

With the increase of task complexity, a single linear reasoning chain can hardly meet the demand of multi-stage decision-making. Yao et al. proposed the Tree-of-Thoughts (ToT) framework, which realizes multipath parallel reasoning by introducing a search tree structure [6]. The model generates different inference branches at each node and uses heuristic search to select the optimal path [5]. The method significantly improves the performance of multi-step inference and strategy planning tasks.

Long et al, proposed the Graph-of-Thoughts (GoT) model [9], which further extends the tree structure into a graph structure, allowing information sharing and global optimisation across paths [6]. GoT utilizes the graph neural network mechanism to realize the interactions between different reasoning nodes, and demonstrates superior scalability in complex tasks such as reasoning games and theorem proving.

Besta et al. proposed the Reasoning-as-Plannig framework by tormalising thoughtful reasoning as a pianning search problem [10]. This approach achieves structure-aware optimisation of the reasoning process by dynamically evaluating the depth of reasoning and the reasonableness of paths. This method combines the ideas of reinforcement learning and search algorithms to make the reasoning path more adaptive and efficient.

Although tree- and graph-structured reasoning greatly improves the inference ability of models, it also brings challenges such as high computational cost and complex optimisation. Future research can explore the lightweight structured reasoning framework and automated node evaluation mechanism to achieve the balance between performance and efficiency.

5 Discussion

In recent years, with the rapid development of artificial intelligence and natural language processing technology, Large Language Models (LLMs) have shown excellent ability in understanding and generating natural language. However, traditional LLMs are often better at pattern matching and surface semantic generation, and are difficult to complete complex reasoning tasks. The proposed Chain-of-Thought (CoT) mechanism provides an explicit reasoning path for models to improve logical consistency and interpretability by progressively decomposing problems and displaying intermediate reasoning processes. Since Wei et al. proposed CoT in 2022, the method has been widely used in mathematical reasoning, logical judgement, code generation and other task scenarios [1]. At the same time, researchers have gradually realized that CoT is not only a cue engineering method, but also a cognitive heuristic mechanism, whose research significance lies in revealing the potential correlation between the internal reasoning structure of the model and the human thinking process. In this context, scholars at home and abroad have carried out systematic research on structure optimisation, self-consistency validation, interpretability enhancement, etc.

Wang et al. proposed Self-Consistency, which significantly improves inference stability by voting on multiple inference paths [2]; Kojima et al. explored the potential of zero- sample inference, which enables the model to be used in the absence of explicit examples, and the model to be used in the absence of explicit examples. models to migrate reasoning power without explicit examples [3]. In addition, Gao et al. and Yao et al. proposed program-assisted and tree-based reasoning structures, respectively, to provide a framework basis for structured thinking in models [4, 5].

Although chain-of-consciousness reasoning has made significant progress in recent years, there are still several challenges. First, there is a tension between interpretability and truthfulness. The current chain of reasoning tends to stay at the linguistic level of reasonableness, and lacks the mapping to the real decision-making mechanism of the model. Second, although verifiable reasoning can enhance consistency, its adaptability in open-domain tasks is limited. and procedural verification can hardly cover all semantic scenarios.

In addition, structured reasoning methods (e.g., ToT and GoT) bring the problems of increasing computational cost and expanding search complexity. The lack of unified evaluation standard in different re-searches makes it difficult to compare the quality of model inference cross-sectionally. Future research can be promoted in three directions: (1) constructing a unified evaluation index system for reasoning chain; (2) exploring a scalable verification framework; and (3) enhancing model self-reflection and human-machine collaborative reasoning mechanism.

In summary, the proposal of chain-of-thought reasoning has opened up a new path for the study of logical reasoning and interpretability in large language models. This paper reviews its development history, core methodology and future trends. It is foreseeable that with the continuous advancement of model scale and algorithm optimisation, CoT will become an important bridge between cognitive modelling of Al and collaborative decision-making between humans and machines. Future research should continue to deepen the validation mechanism, cross-modal fusion and safety and controllability, in order to promote the formation of a more credible and transparent intelligent reasoning system.

6 Conclusion

This paper systematically reviews the research progress of chain-of-thinking (CoT) reasoning in large language models in the past five years, and provides an in-depth analysis from the evolution of reasoning mechanism, verifiability research to structured reasoning framework. The results show that CoT significantly improves the logical consistency and interpretability of the model, and provides a foundation for building trustworthy Al.

Looking ahead, chain of thought reasoning still faces several key challenges. First, how to improve the generation efficiency while maintaining the accuracy of reasoning is still one of the bottlenecks in the practical application of the model. Second, the problem of verifiability and reproducibility of the reasoning process needs to be solved. Currently, most of the CoT generation results rely on the randomness of the model, and there is a lack of systematic quality assessment and calibration mechanism. In addition, how to make the model achieve adaptive reasoning and cross knowledge domain migration in multi-task and multi-modal environments is also an important research direction in the future.Researchers can try to introduce hybrid reasoning paradigms (e.g, combining symbolic reasoning and neural reasoning), or develop intelligent body frameworks with the ability of 'reflection-verification- rethinking' closed loop. The ultimate goal is to make large language models not only 'think' but also' prove themselves ' in reasoning.

References

1. Wei, J., Wang, X., Schuurmans, D., et al.: Chain-of-thought prompting elicits reasoning in large language models//Advances in Neural Information Processing Systems. (2022)
2. Wang, X., Wei, J., Schuurmans, D., et al.: Self-consistency improves chain-of-thought reasoning in language models. arXiv preprint arXiv:2203.11171, (2022)
3. Kojima, T., Gu, S. S., Reid, M., et al. Large language models are zero-shot reasoners//NeurIPS. (2023)
4. Gao, L., Schuurmans, D., Chen, X., et al.: PAL: Program-aided language models//International Conference on Machine Learning. (2023)
5. Yao, S., Zhao, J., Cao, S., et al.: Tree of thoughts: Deliberate problem solving with large language models//ICML. (2023)
6. Long, S., Liu, X., Zhang, Z., et al.: Graph of thoughts: Integrating structured reasoning in large language models. arXiv preprint arXiv:2402.01302, (2024)
7. Lyu, Q., Li, X., Liu, Z., et al.: Faithful chain-of-thought reasoning//ACL. (2023)
8. Zheng, L., Chen, Z., Zhao, Y., et al.: Reflexion: Language agents with verifiable self-reflection. arXiv preprint arXiv:2404.09952, (2024)
9. Burns, C., et al.: Explain-Verify: A framework for faithful reasoning in language models//ICLR. (2024)
10. Besta, M., et al.: Reasoning as planning: Structure-aware chain-of-thought in large language models//NeurIPS. (2025)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.