

RATT: A Thought Structure for Coherent and Correct LLM Reasoning

Jinghan Zhang¹, Xiting Wang², Weijieying Ren³, Lu Jiang⁴, Dongjie Wang⁵, Kunpeng Liu^{1*}

¹Portland State University

²Renmin University of China

³Pennsylvania State University

⁴Dalian Maritime University

⁵University of Kansas

{jinghanz,kunpeng}@pdx.edu, xitingwang@ruc.edu.cn, wjr5337@psu.edu, jiangl761@dmlu.edu.cn, wangdongjie@ku.edu

Abstract

Large Language Models (LLMs) gain substantial reasoning and decision-making capabilities from thought structures. However, existing methods such as Tree of Thought and Retrieval Augmented Thoughts often fall short in complex tasks due to the limitations of insufficient local retrieval of factual knowledge and inadequate global selection of strategies. These limitations make it challenging for these methods to balance factual accuracy and comprehensive logical optimization effectively. To address these limitations, we introduce the Retrieval Augmented Thought Tree (RATT), a novel thought structure that considers both overall logical soundness and factual correctness at each step of the thinking process. Specifically, at every point of a thought branch, RATT performs planning and lookahead to explore and evaluate multiple potential reasoning steps, and integrate the fact-checking ability of Retrieval-Augmented Generation (RAG) with LLMs' ability to assess overall strategy. Through this combination of factual knowledge and strategic feasibility, the RATT adjusts and integrates the thought tree structure to search for the most promising branches within the search space. This thought structure significantly enhances the model's coherence in logical inference and efficiency in decision-making, and thus increases the limit of the capacity of LLMs to generate reliable inferences and decisions based on thought structures. A broad range of experiments on different types of tasks showcases that the RATT structure significantly outperforms existing methods in factual correctness and logical coherence.

Introduction

Large Language Models (LLMs) have shown impressive reasoning and decision-making capabilities in complex tasks like mathematical reasoning and creative writing by processing and generating thoughts based on token-level predictions (Huang et al. 2022a; Zhang et al. 2023). However, this approach limits their ability to extend higher-level or multi-perspective reasoning (Huang et al. 2022b). Research on LLM reasoning indicates that implementing structured thought processes significantly enhances their performance in inference and decision-making (Yu et al. 2023).

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

These thought structures help LLMs to organize and generate their responses both contextually and hierarchically, as well as logically coherent reason across extended narratives and complex problem spaces.

Researchers have developed various thought structures for LLMs. Compared to the vanilla input-output (IO) (Zhang et al. 2024) prompting, where the model directly responds to the prompt, the basic idea of thought structures like Chain-of-Thought (CoT) (Wei et al. 2022) and Self-consistency with CoT (CoT-SC) (Wang et al. 2022) is to guide the reasoning process by generating a sequential and coherent reasoning framework according to the given task. These methods usually involve generating a series of intermediate thought steps for analyzing and solving the initial question. They significantly improve LLM reasoning performance and the reasoning process transparency.

Despite the advancements in improving LLMs' reasoning capability, current thought structures still encounter a major challenge in balancing local factual accuracy and global strategy planning effectiveness. These limitations could restrict their applicability in complex and dynamic scenarios. For example, in Figure 1, CoT and Tree of Thoughts (ToT) (Yao et al. 2024), which simulate coherent reasoning processes, lack an effective fact-checking mechanism. As a result, errors that occur in the early stage of reasoning can propagate through the entire thought chain and lead the conclusions to deviate far away from reality. Moreover, the global selection mechanism of the optimal branch of ToT is not effective and efficient enough, which could result in generating excessive texts and branches (Ding et al. 2023). While Retrieval Augmented Thoughts (RAT) (Wang et al. 2024) combine CoT and Retrieval-Augmented Generation (RAG) (Lewis et al. 2020) to perform fact checking, this method faces challenges in global strategy optimization, as it lacks systematic planning and lookahead to search for the optimal reasoning path in the search space and hence is limited in complex tasks where a broader perspective is necessary (Zhao et al. 2024; Renze and Guven 2024).

Our Targets. To solve the aforementioned issues and enhance the reliability and accuracy of LLM reasoning, we aim to develop a method that seamlessly unifies local and global optimization. At the **local** level, the method should utilize external knowledge early and continuously to prevent and correct factual errors and thus avoid leading the

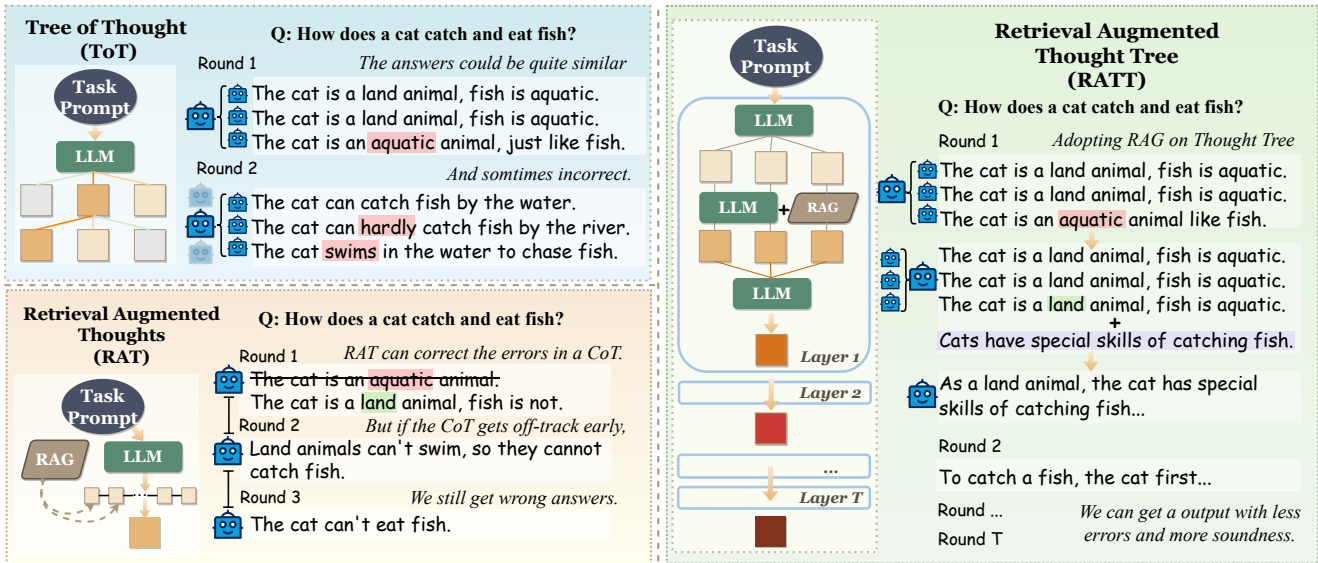


Figure 1: Comparison of LLM thought structures. With optimized tree structure and RAG, our method is capable of reducing factual and logical errors.

model to incorrect search areas. If the factual errors are not identified early, they accumulate during subsequent generations and eventually lead to unreliable conclusions. At the **global** level, the method should structure the reasoning process with planning and lookahead abilities. This allows for the identification of logically coherent and globally optimal reasoning paths even when there are frequent corrections of factual errors. Solely relying on comprehensive correction after the initial reasoning often leads to suboptimal solutions, as post-processing may not correct all factual errors or even introduce new errors due to delayed fact-checking.

Our Approach. To achieve these goals, we develop the Retrieval Augmented Thought Tree (RATT), a novel thought structure that simultaneously ensures both local factual correctness and global comprehensive logical soundness at each step of the reasoning process in one unified framework. In RATT, we perform planning and lookahead to multiple potential reasoning steps at each step of reasoning and integrate the fact-correction capability of RAG along with LLM’s assessment of its overall strategy. Then, RATT integrates factual correctness and logical strategy feasibility to optimize the reasoning process and guide us toward the most promising branches within the search space. This approach greatly improves the logical coherence and inference reliability of LLM’s decision-making and precise reasoning.

In summary, our contribution includes:

1. We introduce a new thought structure called Retrieval Augmented Thought Tree (RATT), which considers and improves comprehensive logical soundness and factual correctness at each step of the reasoning process. This structure significantly enhances the logical coherence and decision-making efficiency of LLMs in complex reasoning tasks.
2. We develop a novel paradigm incorporating RAG into a

tree-structured thought process. This paradigm conducts lookahead and general-to-detail fact-checking analyses at each node of the thought tree.

3. We conduct a series of experiments to validate the effectiveness and robustness of our method (RATT) across different tasks. Our experiments show that our method has clear advantages over existing methods.

Related Works

Thought Structures for LLMs

Thought structures are a series of prompt engineering methods that guide models to generate more specific, accurate, and high-quality content (Zhang et al. 2022; Minaee et al. 2024). Among these methods, Chain of Thought (CoT) (Wei et al. 2022) is a milestone development that guides a reasoning process by generating a series of logically coherent intermediate steps of inference. Based on CoT, Wang et al. (2022) (Wang et al. 2022) introduce the Self-consistency with CoT (CoT-SC) method, which enhances reasoning accuracy and stability by independently generating multiple thought chains and selecting the most reliable answer. Furthermore, the ToT extends CoT by constructing a thought tree with multiple reasoning branches. The thought tree structure is capable of planning, looking ahead, and backtracking, thus providing a broad and global view of the solution space. These developments make multi-step thinking and reasoning similar to human cognitive processes.

Retrieval-Augmented Generation for LLM Reasoning

Retrieval-Augmented Generation (RAG) (Lewis et al. 2020) for LLM reasoning has become a vital approach to enhance the quality of output. The LLM generates a response after retrieving information in an external library, which is

a set of documents or knowledge related to the task, with a query relevant to the task (Lewis et al. 2020; Shuster et al. 2021). This approach helps the LLM to produce responses that are more accurate, contextually relevant, and with fewer hallucinations (Yang et al. 2024; Wu et al. 2024). Following RAG, (Wang et al. 2024) developed the Retrieval-Augmented Thoughts (RAT) approach, which incorporates the RAG retrieval process into a reasoning chain-of-thought. The RAT approach first generates a reasoning chain and retrieves relevant information using LLM-generated queries based on the prompt and each reasoning step. Then RAT corrects and refines the reasoning chain step by step. The primary advantage of RAT is its ability to correct errors in the reasoning process, which is one notable step for enhancing the performance of LLM reasoning. However, as RAT only thinks and refines following one certain and complete path of thought, one potential problem of RAT is to fall into a local suboptimal solution in the search space.

Methodology

In this section, we introduce **Retrieval Augmented Thought Tree (RATT)**, an automated novel thought structure for language models that prioritizes both logical coherence and factual correctness. We aim to enhance the reliability and accuracy of LLM reasoning in two aspects: robust local fact-checking to prevent the accumulation of errors and global planning and lookahead to improve the logical coherence of reasoning paths. As shown in Figure 2, we implement this RATT through several steps, including (1) **Thought Node Generation**, (2) **Retrieval and Selection**, and (3) **RAG Correction and Integration**.

Problem Formulation

To formalize our target and solution, first let us define the elements and functions. Given a task \mathcal{A} and a pre-trained LLM p_θ with parameters θ , we aim to enhance the performance of LLM on task \mathcal{A} with performance metric $\epsilon = \{\epsilon_1, \epsilon_2, \dots, \epsilon_j\}$. The task \mathcal{A} can be any problem that can be solved with a logical sequence of thoughts or decisions. We have an input x which is a language sequence and an output y . Our target is to construct the best thought tree structure $\mathcal{T} = \mathcal{T}^*$ capable of navigating the search space and optimizing the reasoning process with task \mathcal{A} and prompt x :

$$\mathcal{T}^* = \arg \max_{\mathcal{T}} \epsilon(y), \quad \text{where } y = p_\theta(x|\mathcal{T}). \quad (1)$$

Here we denote a thought tree $\mathcal{T} = (\mathcal{N}, \mathcal{E})$ to be a hierarchical tree structure consists of nodes and edges where each nodes represents a thought or decision, and an edge represents a logical flow between two nodes. Here \mathcal{N} is the set of nodes in the tree and $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ is the set of directed edges connecting the nodes.

Thought Node Generation

Our first step is to generate a thought tree that balances both exploration and exploitation adeptly and effectively. Given an prompt x with task information, the LLM p_θ first embeds

the prompt:

$$q_x = p_\theta(\text{embed}(x)). \quad (2)$$

Then the LLM generates initial thoughts utilizing multiple logical strategies $s^{(m)}$ to explore possible solutions in the search space of task \mathcal{A} from a broad range of perspectives and logical approaches. Each initial thought represents a distinct logical viewpoint, and these thoughts collectively cover various possible or potential directions and solutions to tasks \mathcal{A} . Specifically, in the t -th iteration, we define the initial thoughts to be the initial nodes of the tree: $\mathbf{n}_t = \{n_t^{(1)}, n_t^{(2)}, \dots, n_t^{(m)}\}$. In this way, the thought tree takes every step with a broad and diverse range of *logical possibilities*. We denote a reasoning branch \mathcal{B} of the tree as a sequence of connected nodes starting from the root node and extending to any leaf node. Each branch \mathcal{B} represents a feasible solution for task \mathcal{A} , and each node in the branch sequence is connected by an edge \mathcal{E} from its parent node, formulated as:

$$\mathcal{B} = \{n_{root}, \dots, n_{leaf}\}, (n_i \xrightarrow{p_\theta} n_{i+1}) \in \mathcal{E}. \quad (3)$$

After generating initial nodes, the thought tree employs a planning and lookahead strategy, which is essential for dynamically evaluating the potential of branches. For each node n_t^m , the LLM evaluates the quality of the immediate next steps as well as simulates the future states and consequences of the node’s choice. The simulation provides insights into the likely outcomes to the model. This lookahead strategy helps the model to *plan several moves ahead*, refine and adjust generation strategies, and keep aligning the generation direction with the overall goal of the task.

Retrieval and Selection

Given a Library with I candidate documents $\mathcal{R} := \{R_i\}_{i=1}^I$, the LLM embeds each document as:

$$r_i = p_\theta(\text{embed}(R_i)) \in \mathbb{R}^K, \quad (4)$$

where K is the dimension of the embedding and $\text{embed}(\cdot)$ is the embedding process. We then form a query Q_t^m of each n_t^m and the input x :

$$q_t^m = \text{embed}\{Q_t^m\} = \{q_x, q_{n_t^m}\} \in \mathbb{R}^K, \quad \text{where} \quad (5)$$

$$q_x = p_\theta(\text{embed}(x)), \quad \text{and} \quad (6)$$

$$q_{n_t^m} = p_\theta(\text{embed}(n_t^m)). \quad (7)$$

We then retrieve the whole Library to find the top- k most relevant documents. Here we evaluate the relevance of each document R_i to the query Q_t^m by the cosine similarity of the embeddings:

$$\text{sim}(q_t^m, r_i) = \frac{q_t^m \cdot r_i}{\|q_t^m\| \|r_i\|}. \quad (8)$$

These selected documents $\mathcal{R}_{\text{selected}}$ are then adept to enhance the information content of the initial nodes and *support the fact-checking*. Analogizing human learning and reasoning processes that go top to bottom from general understanding to specific details, we adopt a similar strategy

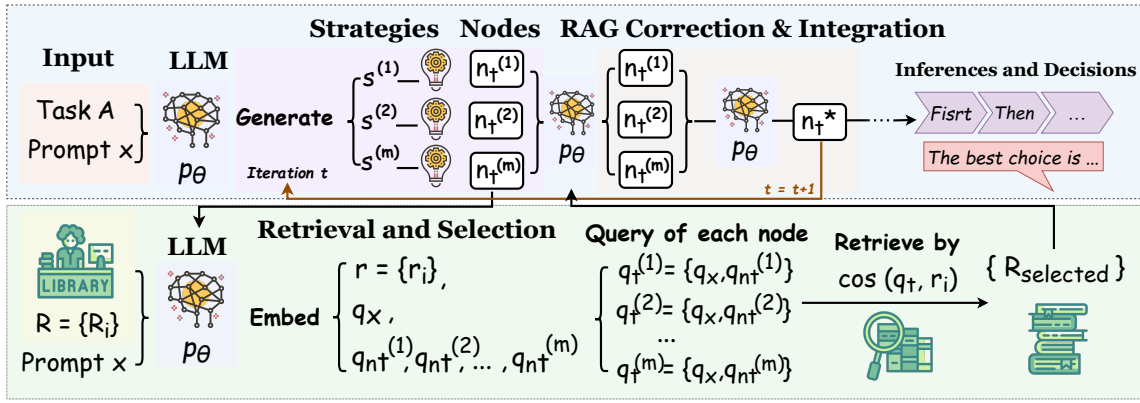


Figure 2: Framework of RATT. Given an input including a task, a prompt, and a set of documents, the LLM generates multiple thought nodes with different strategies and embeds the prompt and documents. Then, the model integrates the nodes into one optimized node. We use the embedding of the optimized node to retrieve and select relative documents. Finally, the LLM corrects possible mistakes and enriches the node with selected documents’ information. This generation and optimization process iterates until reaching the maximum rounds of iteration.

in the retrieval process. We denote all nodes in the same round of generation to be in the same layer L . Specifically, at the top layers $L_1, L_2, \dots, L_{l_1}, l_1 \in [1, l]$, we utilize a *broad and high-level* retrieval strategy to gather basic concepts and background information; At the middle layers $L_{l_1+1}, \dots, L_{l_2}, l_2 \in [l_1 + 1, l]$, we utilize a *targeted* retrieval strategy to acquire deeper information; At the leaf layers $L_{l_2+1}, \dots, L_l, l_2 \in [l_1 + 1, l]$, we utilize a *detailed and specific* retrieval strategy to correct specific errors and enrich the details.

RAG Correction and Integration

We then correct and integrate the initial nodes with relevant documents $\mathcal{R}_{\text{selected}}$. In this process, the LLM p_θ analyzes the documents $\mathcal{R}_{\text{selected}}$ and extracts the most critical and supplementary information relevant to $n_t^1, n_t^2, \dots, n_t^m$. The information includes factual details, background explanations, logical support, and counter-examples. The model p_θ then integrates them into a single, optimized node:

$$n_t^* = p_\theta(n_t^1, n_t^2, \dots, n_t^m, \{\mathcal{R}_{\text{selected}}\}). \quad (9)$$

Here the role of p_θ is to ensure that information integration utilizes the content of documents correctly while keeping the backbones of the reasoning sequence. In this way, we enhance the depth and breadth of the node’s information, and the optimized node n_t^* becomes more *reliable and comprehensive* in subsequent reasoning and decision-making processes. For all nodes generated in this layer, n_t^* represents the most valuable solution that LLM finds in the search space.

After generating the refined node n_t^* , we combine n_t^* with the original prompt x to be the new prompt of task \mathcal{A} , and the model starts a new round of generation. This iterative process continues until the maximum number of iterations T is reached. At the end of the iterations, depending on the specific requirements of task \mathcal{A} , the model output either the node n_T^* or the complete inference of n_T^* as decisions and inference.

Discussions

We further discuss why our method achieves the goals mentioned in the Introduction. First, to address the local challenge of correcting factual errors and avoiding the accumulation of errors, we need to perform fact-checking and corrections in a timely and continuous manner. Our RATT method utilizes RAG technology to help LLMs efficiently and rapidly access information from relevant documents in an external library. The LLM integrates external knowledge to dynamically correct the factual errors caused by its possible limited knowledge, outdated information, or hallucinations within the reasoning context. Second, to address the global challenge of strategy optimization and searching for the optimal solution, we need to comprehensively evaluate the overall generating and searching strategy. Our RATT method integrates factual accuracy and strategic feasibility to perform correction and optimization on each generated thought. In this way, the RATT identifies and navigates the most promising branch within the search space. Finally, to reduce the risk of hallucinations, we need to make corrections during the generation process rather than after the entire reasoning process is completed. This is because the errors that develop and spread through the structure are difficult to entirely correct afterward, as they may have already polluted and magnified the whole generation. Thus, we design an online, incremental generation and correction tree structure. The experimental details in the next section demonstrate the effectiveness of our RATT’s capability in preventing the spread of errors and reducing hallucinations.

Experiments

Experimental Setup

In this section, we propose four particularly challenging or representative tasks for LLM performance evaluation, and demonstrate the effectiveness of our RATT approach. These multi-view tasks provide standard benchmarks to evaluate and compare the performance of our approach and baseline

Algorithm 1: Retrieval Augmented Thought Tree (RATT)

```
1: Input: Task prompt  $x$ , Node number  $m$ , Iteration time  $T$ , LLM  $p_\theta$ 
2:  $n_t \leftarrow ""$   $\triangleright$  Initialize the node as an empty string
3: for  $t = 1$  to  $T$  do
4:   for  $l = 1$  to  $m$  do
5:      $n_t^{(l)} \leftarrow p_\theta(\cdot | x, n_t)$   $\triangleright$  Generate thoughts based on node  $n_t$  and prompt  $x$ 
6:      $\{q_x, q_{n_t}^{(l)}\} \leftarrow \text{embed}(\{n_t^{(l)}, x\})$   $\triangleright$  Embed the node and prompt
7:      $q_t^{(l)} \leftarrow \text{CONCAT}(q_x, q_{n_t}^{(l)})$   $\triangleright$  Concatenate the question and answer into a query vector
8:      $\mathcal{R}_{\text{selected}} \leftarrow \text{RetrieveFromLibrary}(q_t^{(l)})$   $\triangleright$  Retrieve documents  $\mathcal{R}_{\text{selected}}$  from Library
9:      $n_t^{(l)*} \leftarrow p_\theta(\cdot | n_t^{(l)}, \mathcal{R}_{\text{selected}})$   $\triangleright$  Generate a refined node
10:     $n_t \leftarrow \text{CONCAT}(n_t, n_t^{(l)*})$   $\triangleright$  Append the next refined node
11:   end for
12:    $n_t^* \leftarrow p_\theta(\cdot | x, n_t)$   $\triangleright$  Refine and enhance the node
13: end for
14:  $n_T^* \leftarrow p_\theta(\cdot | x, n_t^*)$ 
15: return  $n_T^*$   $\triangleright$  Output  $n_T^*$  as the final generation
```

methods. Then we conduct an analysis of the results and discuss the strengths of our approach shown in the performance enhancement.

Task Description. We test our RATT and baseline methods across four distinct tasks, each evaluating different aspects of the quality of LLM response. The *Code Generation* and *Creative Writing* are two comprehensive tasks, which *Game of 24* focus on testing the logical and numerical reasoning capabilities. We also process a standard *Hallucination Detection* to demonstrate the RATT’s ability to reduce hallucinations.

Baselines Algorithms and Environmental Settings. For baselines, we compare our approach with several prompting methods, including IO, CoT (Wei et al. 2022), CoT-SC (Wang et al. 2022), ToT (Gomez 2023), and RAT (Craft-Jarvis 2024). For the external library and database, we employ the `codeparrot/github-jupyter` (Hugging Face 2024)¹ as the search vector library of task *Code Generation*, and the English Wikipedia² as library of task *Creative Writing* and *Hallucination Detection*. We do not employ any external library for task *Game of 24*, as the goal of this task is to examine the numerical reasoning ability of RATT’s thought tree structure. For the language model, we perform all the experiments on OpenAI API, GPT-3.5 Turbo (OpenAI 2023) model through the OpenAI platform³. For the implements, we perform the tasks on NVIDIA 4090.

¹<https://huggingface.co/datasets/codeparrot/github-jupyter>

²<https://en.wikipedia.org>

³<https://openai.com/index/openai-api/>

Code Generation

Code generation is a task where the model is tasked to understand programming prompts and produce functional and correct code based on those prompts. This task requires precise comprehension and application of programming logic, algorithms, and language syntax of the LLM. We conduct the code generation evaluation in HumanEval. HumanEval is a programming task benchmark specifically for the evaluation of coding capabilities of code generation models. In this task, we first construct a local dataset with CodeParrot and segment long texts into multiple chunks. Then we transform each chunk into embeddings with embedding model `text-embedding-ada-002` (OpenAI 2024)⁴. After introducing tasks from HumanEval, the RATT generates multiple potential answers. Finally, we validate the effectiveness of our RATT approach using the original evaluation script from the HumanEval GitHub repository by *pass@k* method (Liu et al. 2024). This *pass@k* method tests if there is at least one correct answer within k generated responses to assess the efficacy and accuracy of the model. Our experiments on HumanEval have shown the highest performance improvement of 38.05% for *pass@1* and 15.12% for *pass@5* relatively, as summarized in Table 1 and Figure 3.

Method	pass@1	pass@5
IO	50.49%	72.56%
CoT	47.31%	75.88%
ToT	49.36%	77.73%
RAG_1 shot	50.61%	76.22%
RAG_5 shot	45.49%	74.39%
RAT	59.27%	80.49%
RATT	69.70%	83.53%

Table 1: Comparison of code generation performance on HumanEval of different methods.

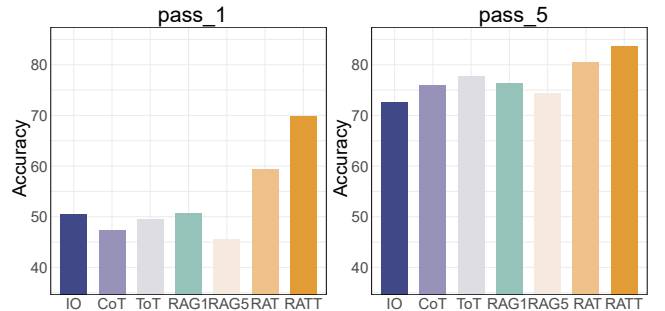


Figure 3: Comparison of code generation performance on HumanEval of different methods.

Creative Writing

Creative writing is a task that challenges LLMs to generate imaginative, coherent, and contextually rich text based on a

⁴<https://openai.com/index/new-and-improved-embedding-model/>

variety of prompts. This task measures the LLMs’ reasoning capabilities to innovate and think logically and coherently while enriching the context information. To further evaluate the RATT’s comprehensive performance on LLMs, we compare it with IO, CoT, ToT and RAT prompting. To assist this process, we utilize the English Wikipedia library by searching this vast and diverse dataset to identify and retrieve the top few web pages that are most relevant to the given writing prompts.

To ensure the consistency and integrity of the evaluation, we employ GPT-4 (OpenAI 2023; Naismith, Mulcaire, and Burstein 2023) to assist in assessing the quality of model outputs as 50% of the reviewers of this task while other reviewers are human annotation experts. Our scoring criteria encompass *Soundness (Sd)*, *Information Relevance (Rel)*, *Content Coherence of Reasoning (Corr)*, and *Clarity of Expression (Expr)*, which are uniformly applied. Each dimension receives a score up to 10, with increments of 0.5. The *Overall (Oa)* score is the average of four metrics.




Figure 4: Comparing methods in Creative Writing.




Figure 5: Comparing methods on GPT models.

Here, we first conduct a comparison of various methods on GPT-3.5 Turbo to assess the performance across different metrics. As shown in Figure 4, our approach demonstrates significant advantages over baseline methods in all dimensions and composite scores. We further study the improvements our method brings to both the GPT-3.5 Turbo and GPT-4o (OpenAI 2024) models. As shown in Figure 5, our approach and RAT significantly outperform the GPT-4o model. The results demonstrate the effectiveness of employing RAG on LLM in text generation and boosting the creative output of LLMs. In Figure 6, we present the detailed prompting and guidelines of the creative writing task.

Hallucination Detection

Following the method introduced in (Ding et al. 2024), we perform the standard hallucination detection. Hallucination

detection refers to the process of identifying and assessing instances where language models generate text that is not grounded in reality (Luo et al. 2024). This task is crucial for evaluating the reliability of LLMs’ text generation, as we target generating content that is accurate and trustworthy (Gao et al. 2023; Rawte, Sheth, and Das 2023). In this task, we utilized the TruthfulQA (Lin, Hilton, and Evans 2021) dataset to measure the truthfulness of outputs. Each “correct answer” in the dataset aims to reflect the truthfulness of a response to a given query. Our model generates answers to these queries, and then we measure the truthfulness of these answers in two ways:

First, we directly calculate the similarity between the generated answers and the correct answers using BLEU and ROUGE scores. Here BLEU and ROUGE scores are the metrics designed to evaluate the quality of text by comparing machine-generated outputs against human references. Then we employ GPT-judge in (SyLinRL 2021) to predict the truthfulness of human-like answers in an end-to-end manner as part of the TruthfulQA project.

Game of 24

The “Game of 24” is a mathematical puzzle task that challenges an LLM’s numerical reasoning capabilities. The game involves using the combination of addition, subtraction, multiplication, and division to manipulate four integers between 1 and 9 to arrive at the final result of 24. The rules are straightforward: players must use all four numbers exactly once, and the order of operations can vary as needed to solve the puzzle. This task is a benchmark for LLM mathematical reasoning (Kim et al. 2023; Ding et al. 2023), as it requires arithmetic operations and strategic planning abilities to explore various combinations and operation sequences for the optimal solution.

Method	Success (%)
IO	7.6
CoT	5.1
CoT-SC	11.2
ToT	14.0
RAT	9.0
RATT-3	25.9
RATT	31.8
Improvement	24.2

Table 2: Success rates on Game of 24.

We test the models on 100 randomly selected games from 4nums⁵. For every test, if the model produces a valid equation that results in 24, we mark it as a success. From Table 2, we can see that the RATT achieves the highest success rate of 31.8%, which significantly outperforms the GPT-3.5 Turbo model by 24.2% of improvement. Here the success of RATT-3 is a success answer of 24 in 3 steps of generation, and RATT in 5. IO, CoT, and CoT-SC can hardly help the model find a solution, while RAT can hardly bring

⁵<https://www.4nums.com/game/difficulties/>

Question: Introduce Jin-Yong's Life.

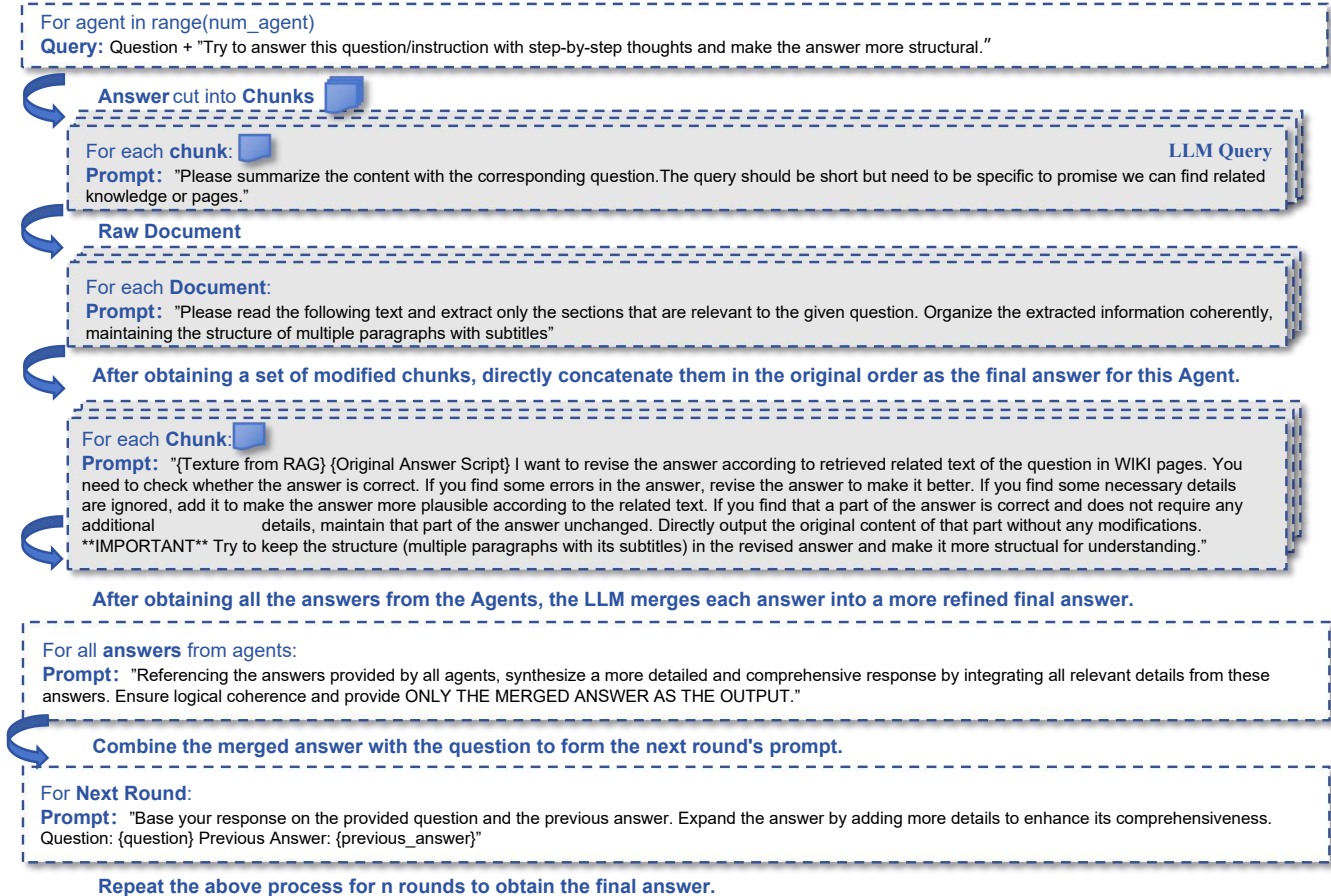


Figure 6: Prompting case of task Creative Writing.




Figure 7: Comparison of hallucination metrics of different methods.

advantages for LLM to win. These results underline our method's reasoning capability to enhance the numerical reasoning ability.

Conclusion

In this paper, we introduce RATT, a novel thought tree structure that integrates both factual correctness and strategic coherence at every step of the reasoning process. This structure significantly enhances LLMs' logical coherence and decision-making efficiency in complex reasoning tasks

by solving the existing problem of thought structures logically and globally. Our extensive experiments across various tasks demonstrate the effectiveness and superiority over existing methods, especially in improving factual correctness and logical coherence. By adopting RATT on LLMs, we extend the boundaries of abilities that LLM performs in handling complex tasks. For future work, we plan to extend this LLM thought structure paradigm across various reasoning and LLM alignment tasks, and to conduct a thorough investigation into the generation strategies that our approach applies to the thought structure and solution space to study further insights into its underlying mechanisms and impacts.

While RATT shows significant advancements and wide adaptability, there are several limitations that require further exploration, including high computational demands and limited scalability with very complex tasks or large external libraries. Second, the effectiveness of the reasoning process and final output heavily relies on the quality of the task prompt and the external library, which can affect the model's performance in scenarios with poorly organized prompts or limited external knowledge. Finally, adopting RATT in tasks with unique requirements may have inherent limitations.

Acknowledgments

The author Xiting Wang was supported by the National Natural Science Foundation of China (NSFC) (NO. 62476279), Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Renmin University of China, and the Fundamental Research Funds for the Central Universities, and the Research Funds of Renmin University of China No. 24XNKJ18.

References

- CraftJarvis. 2024. RAT - Repository. <https://github.com/CraftJarvis/RAT>. Accessed: 2024-05-20.
- Ding, H.; Pang, L.; Wei, Z.; Shen, H.; and Cheng, X. 2024. Retrieve Only When It Needs: Adaptive Retrieval Augmentation for Hallucination Mitigation in Large Language Models. *arXiv preprint arXiv:2402.10612*.
- Ding, R.; Zhang, C.; Wang, L.; Xu, Y.; Ma, M.; Zhang, W.; Qin, S.; Rajmohan, S.; Lin, Q.; and Zhang, D. 2023. Everything of thoughts: Defying the law of penrose triangle for thought generation. *arXiv preprint arXiv:2311.04254*.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; and Wang, H. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Gomez, K. 2023. Tree of Thoughts. <https://github.com/kyegomez/tree-of-thoughts>.
- Huang, W.; Abbeel, P.; Pathak, D.; and Mordatch, I. 2022a. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, 9118–9147. PMLR.
- Huang, W.; Xia, F.; Xiao, T.; Chan, H.; Liang, J.; Florence, P.; Zeng, A.; Tompson, J.; Mordatch, I.; Chebotar, Y.; et al. 2022b. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*.
- Hugging Face. 2024. CodeParrot. <https://huggingface.co/codeparrot>. Accessed: 2024-05-20.
- Kim, S.; Moon, S.; Tabrizi, R.; Lee, N.; Mahoney, M. W.; Keutzer, K.; and Gholami, A. 2023. An LLM compiler for parallel function calling. *arXiv preprint arXiv:2312.04511*.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33: 9459–9474.
- Lin, S.; Hilton, J.; and Evans, O. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Liu, J.; Xia, C. S.; Wang, Y.; and Zhang, L. 2024. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems*, 36.
- Luo, J.; Li, T.; Wu, D.; Jenkin, M.; Liu, S.; and Dudek, G. 2024. Hallucination detection and hallucination mitigation: An investigation. *arXiv preprint arXiv:2401.08358*.
- Minaee, S.; Mikolov, T.; Nikzad, N.; Chenaghlu, M.; Socher, R.; Amatriain, X.; and Gao, J. 2024. Large language models: A survey. *arXiv preprint arXiv:2402.06196*.
- Naismith, B.; Mulcaire, P.; and Burstein, J. 2023. Automated evaluation of written discourse coherence using GPT-4. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, 394–403.
- OpenAI. 2023. GPT-3.5 Turbo Fine-Tuning and API Updates. <https://openai.com/index/gpt-3-5-turbo-fine-tuning-and-api-updates/>. Accessed: 2024-05-20.
- OpenAI. 2023. GPT-4. <https://openai.com/index/gpt-4-research/>.
- OpenAI. 2024. Hello GPT-4. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-05-20.
- OpenAI. 2024. New and Improved Embedding Model. <https://openai.com/index/new-and-improved-embedding-model/>.
- Rawte, V.; Sheth, A.; and Das, A. 2023. A survey of hallucination in large foundation models. *arXiv preprint arXiv:2309.05922*.
- Renze, M.; and Guven, E. 2024. Self-Reflection in LLM Agents: Effects on Problem-Solving Performance. *arXiv preprint arXiv:2405.06682*.
- Shuster, K.; Poff, S.; Chen, M.; Kiela, D.; and Weston, J. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- SyLinRL. 2021. TruthfulQA. <https://github.com/sylinrl/TruthfulQA>.
- Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; and Zhou, D. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Wang, Z.; Liu, A.; Lin, H.; Li, J.; Ma, X.; and Liang, Y. 2024. Rat: Retrieval augmented thoughts elicit context-aware reasoning in long-horizon generation. *arXiv preprint arXiv:2403.05313*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, S.; Xiong, Y.; Cui, Y.; Wu, H.; Chen, C.; Yuan, Y.; Huang, L.; Liu, X.; Kuo, T.-W.; Guan, N.; et al. 2024. Retrieval-Augmented Generation for Natural Language Processing: A Survey. *arXiv preprint arXiv:2407.13193*.
- Yang, K.; Swope, A.; Gu, A.; Chalamala, R.; Song, P.; Yu, S.; Godil, S.; Prenger, R. J.; and Anandkumar, A. 2024. Leandojo: Theorem proving with retrieval-augmented language models. *Advances in Neural Information Processing Systems*, 36.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.

Yu, Z.; He, L.; Wu, Z.; Dai, X.; and Chen, J. 2023. Towards better chain-of-thought prompting strategies: A survey. *arXiv preprint arXiv:2310.04959*.

Zhang, J.; Wang, B.; Li, L.; Nakashima, Y.; and Nagahara, H. 2024. Instruct me more! random prompting for visual in-context learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2597–2606.

Zhang, S.; Chen, Z.; Shen, Y.; Ding, M.; Tenenbaum, J. B.; and Gan, C. 2023. Planning with large language models for code generation. *arXiv preprint arXiv:2303.05510*.

Zhang, Z.; Zhang, A.; Li, M.; and Smola, A. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.

Zhao, P.; Zhang, H.; Yu, Q.; Wang, Z.; Geng, Y.; Fu, F.; Yang, L.; Zhang, W.; and Cui, B. 2024. Retrieval-augmented generation for ai-generated content: A survey. *arXiv preprint arXiv:2402.19473*.