

# RPM-MCTS: Knowledge-Retrieval as Process Reward Model with Monte Carlo Tree Search for Code Generation

Yuanyuan Lin<sup>1,3</sup>, Xiangyu Ouyang<sup>2</sup>, Teng Zhang<sup>1\*</sup>, Kaixin Sui<sup>3</sup>

<sup>1</sup> School of Computer Science and Technology, Huazhong University of Science and Technology, China

<sup>2</sup> College of Computer Science and Technology, Xi'an Jiaotong University, China

<sup>3</sup> ByteDance Seed, China

{linyuan, tengzhang}@hust.edu.cn, oyx2019@stu.xjtu.edu.cn, suikaixin@163.com

## Abstract

Tree search-based methods have made significant progress in enhancing the code generation capabilities of large language models. However, due to the difficulty in effectively evaluating intermediate algorithmic steps and the inability to locate and timely correct erroneous steps, these methods often generate incorrect code and incur increased computational costs. To tackle these problems, we propose RPM-MCTS, an effective method that utilizes Knowledge-Retrieval as Process Reward Model based on Monte Carlo Tree Search to evaluate intermediate algorithmic steps. By utilizing knowledge base retrieval, RPM-MCTS avoids the complex training of process reward models. During the expansion phase, similarity filtering is employed to remove redundant nodes, ensuring diversity in reasoning paths. Furthermore, our method utilizes sandbox execution feedback to locate erroneous algorithmic steps during generation, enabling timely and targeted corrections. Extensive experiments on four public code generation benchmarks demonstrate that RPM-MCTS outperforms current state-of-the-art methods while achieving an approximately 15% reduction in token consumption. Furthermore, full fine-tuning of the base model using the data constructed by RPM-MCTS significantly enhances its code capabilities.

## Introduction

Code generation aims at understanding problem descriptions in natural language and generating the corresponding code snippets. In recent years, large language models (LLMs) have demonstrated remarkable performance in code generation tasks (Zhang et al. 2024b). For code generation, the early methods involve dividing code planning and synthesis into two phases using chain-of-thought or tree structures (Wei et al. 2022; Jiang et al. 2024; Zelikman et al. 2023). Wang et al. (2024) have demonstrated that providing LLMs with a *correct* solution can significantly enhance model performance, even when these solutions consist of incomplete plans, i.e., for solutions, correctness is preferred over completeness, and the key to enhancing the code generation capability of LLMs lies in generating correct plans.

Programming languages possess their own inherent logical structures and tightly interconnected knowledge, which

makes it essential not to overlook long-range dependencies within the code. Previous work has shown that rotary position embedding does not always lead to attention weights decaying with relative distance (Barbero et al. 2024). Concurrently, through exploring the attention distribution between tokens, we have experimentally demonstrated that selecting algorithmic steps as the basic units is a superior choice. Therefore, our objective focuses on how to accurately generate intermediate algorithmic steps.

However, a limitation of previous methods lies in the lack of an evaluation and correction mechanism for intermediate algorithmic steps, which fails to guarantee the correctness of these steps (Lu et al. 2025; Li et al. 2025a). One way to tackle this issue is to use a value function or reward model to verify reasoning traces for correctness, which then serves as a learning signal for self-training (Lightman et al. 2024; Wang et al. 2023). However, training a reliable reward model to verify every step in a reasoning trace generally depends on dense human-generated annotations per reasoning step (Lightman et al. 2024), which does not scale well.

Unlike other reasoning tasks, code generation benefits from the homogeneity of algorithmic workflows across different problem categories. This allows us to leverage historical experience from a knowledge base containing numerous correct algorithmic steps to evaluate the process reward of expansion steps. Additionally, code generation typically benefits from detailed feedback provided by compilers. Consequently, in this paper, we propose RPM-MCTS, which optimizes the Monte Carlo Tree Search (MCTS) algorithm using external information feedback. Our method utilizes the knowledge base for intermediate algorithmic step-level evaluation and employs sandbox feedback for result-level assessment of complete code. Specifically, the root node of the Monte Carlo tree represents the coding problem, while all other nodes represent individual algorithmic steps. During each iteration, multiple distinct potential next steps are generated based on the current reasoning path. Node selection is guided by historical experience from the knowledge base, enabling faster discovery of high-value search paths. In the simulation phase, complete code is generated and evaluated using sandbox and model feedback to update node values. Notably, during simulation, we localize erroneous steps within the full algorithmic workflow and incorporate newly generated correct steps into the tree, thereby reducing token

\*Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

consumption. After multiple iterations, the highest-scoring path from root to leaf is selected, ultimately yielding a complete solution alongside its corresponding code. The contributions are summarized as follows:

- We propose RPM-MCTS, which leverages knowledge base retrieval scores to evaluate intermediate algorithmic steps, steering LLMs to explore high-value reasoning paths more effectively.
- We leverage sandbox feedback during the simulation phase to evaluate code generated from reasoning steps, localize errors, and truncate simulations, thereby reducing computational costs.
- We conduct extensive experiments and show that RPM-MCTS is superior to state-of-the-art methods. Moreover, we verify that base models fine-tuned with data generated by RPM-MCTS enjoy greater code capabilities.

## Related Work

**Monte Carlo Tree Search.** As the extension of Chain-of-Thought (CoT) (Wei et al. 2022), Tree-of-Thought (ToT) (Yao et al. 2023) enhances the reasoning and planning capabilities of LLMs by exploring different thought paths within a tree structure. Subsequently, Monte Carlo Tree Search has served as a search algorithm to more effectively guide LLMs in exploring intermediate sub-steps (Zhao, Lee, and Hsu 2023; Hao et al. 2023; Zhou et al. 2024; Ding et al. 2023). ReST-MCTS\* (Zhang et al. 2024a) combines process reward guidance with Monte Carlo Tree Search to collect high-quality reasoning trajectories and step-by-step values for training strategy and reward models. SRA-MCTS (Xu et al. 2025) further extends this to the field of code generation, using Monte Carlo Tree Search to generate intermediate reasoning steps and conducting iterative self-evaluation to synthesize training data for supervised fine-tuning. However, relying solely on model self-evaluation introduces biases and hallucinations, and small-scale LLMs exhibit limited instruction-following capabilities. RethinkMCTS (Li et al. 2025b) is another prior work that also uses execution feedback but employs a patching strategy. If this patch fails, the search may proceed on an incorrect path, making it less suitable for generating high-quality SFT data.

**Process Evaluation.** In heuristic search, a robust reasoning process needs to have self-evaluation capabilities, and the evaluation results are further used to guide the search. Early work mainly focused on outcome-level evaluation (Cobbe et al. 2021), that is, evaluating the complete solution after the reasoning is completed. Outcome-level evaluation is simple to implement but often requires more detailed assessment. Step-level evaluation (Lightman et al. 2024; Wang et al. 2023; Gao et al. 2024) emphasizes the assessment of individual reasoning steps. In tree search algorithms, process evaluation is widely used to guide search trajectories. Logic-RL (Xie et al. 2025) optimizes path selection by implementing state scoring in beam search. Furthermore, step-level evaluation has proven its effectiveness in both error correction and the summarization of reasoning steps. Zheng et al. (2024) developed a method capable of

accurately locating inaccuracies in specific reasoning steps, thereby providing more precise and actionable feedback for comprehensive evaluation.

## Method

In this section, we elaborate on the proposed modified MCTS that incorporates the knowledge base as a process reward model. The methodology comprises three key components: knowledge base construction, RPM-MCTS, and code generation. First, knowledge base retrieval scores circumvent random selection during node expansion. Then, in the expansion phase, nodes are filtered based on similarity metrics to eliminate redundant candidates. Finally, during the simulation phase, the algorithm performs error reflection and retains nodes with verified correct reasoning. These collective strategies enable faster exploration of higher-quality algorithmic steps.

### Knowledge Base Construction

In this section, we introduce the construction of a retrievable global knowledge base designed to mitigate hallucination during the planning process. Due to the homogeneity of algorithms within the same category, where fundamental principles and methods are relatively similar, we utilize a knowledge base containing numerous correct algorithms across diverse categories. This serves as the evaluation model for intermediate algorithmic steps in RPM-MCTS, eliminating the need to train a separate process reward model.

We use the training set data from APPS (Hendrycks et al. 2021) and CodeContests (Li et al. 2022), which contain coding problems paired with their correctly implemented solutions. We utilize the Claude Sonnet 3.7 to generate the correct algorithmic steps corresponding to the correct code and decompose them step by step. We sequentially concatenate the problems by rolling them out according to the algorithmic steps. Specifically, for problem  $p_i$  with  $n_i$  algorithmic steps and  $a_i^{(j)}$  corresponding to the  $j$ -th step, we have

$$\mathcal{K}_i = \{\text{concat}(p_i, a_i^{(1)}, \dots, a_i^{(j)}), j = 1, 2, \dots, n_i\}, \quad (1)$$

and  $\mathcal{K} = \bigcup_{i=1}^n \mathcal{K}_i$  is the knowledge base with cardinality  $n$ .

To enhance retrieval efficiency and improve retrieval precision by distinguishing between problems with similar descriptions but different algorithmic solutions, we organize the knowledge base into 14 distinct algorithm categories and store them as vector database by using the BGE (Xiao et al. 2024) embedding model.

### RPM-MCTS

We propose an enhanced MCTS method, named RPM-MCTS. In this method, the root node represents the problem, while all other nodes represent an algorithmic step. Specifically, the method comprises four distinct phases: Selection, Expansion, Evaluation and Reflection, and Backpropagation, as shown in Figure 1. These phases are performed on a search tree composed of tree nodes and are iterated multiple times, with each iteration generating a concrete algorithmic step.

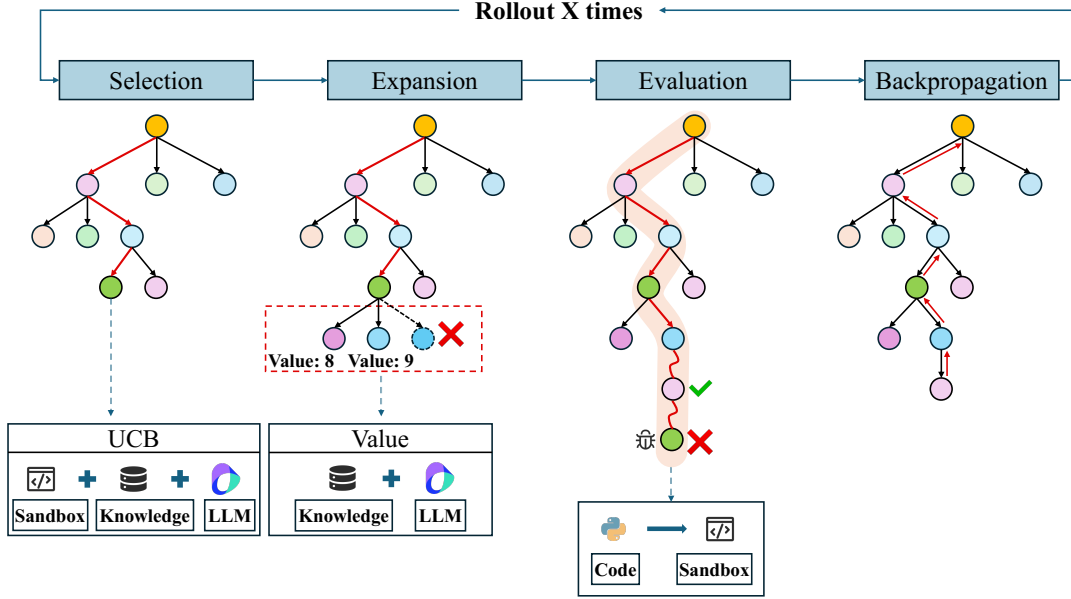


Figure 1: Overview of RPM-MCTS. (a) Selection: Select a leaf node according to Eqn. (2). (b) Expansion: After selecting a node, expand multiple child nodes, and use knowledge base retrieval scores and LLM evaluation to select nodes for simulation. The node color represents similarity magnitude. (c) Evaluation: Generate complete reasoning steps for the selected node, generate code strictly in accordance with these reasoning steps, and use a sandbox for information feedback. (d) Backpropagation: Propagate the reward scores backward. The yellow root node represents the problem, and the remaining nodes represent each reasoning step.

**Selection.** In the selection phase, a leaf node is selected from the current tree for further expansion according to the selection score, which is defined as a weighted combination of the Upper Confidence Bound (UCB) (Silver et al. 2017) and the knowledge base retrieval score:

$$\text{SelectionScore}(s, a) = \text{UCB}(s, a) + \alpha K(s, a), \quad (2)$$

where  $(s, a)$  denotes a state-action pair with  $s$  containing the description of the problem and previously generated algorithmic steps and  $a$  representing the new step at the current node. The parameter  $\alpha$  is for balancing the two terms.

UCB is a classical multi-armed bandit algorithm and well performed in addressing the exploration-exploitation trade-off. UCB selects actions by computing an upper confidence estimate of each action’s potential reward:

$$\text{UCB}(s, a) = Q(s, a) + \beta \sqrt{\frac{\log N(s)}{1 + N(s, a)}}, \quad (3)$$

where  $Q(s, a)$  represents the empirical mean cumulative reward after taking action  $a$  from state  $s$ ,  $N(s)$  is the number of times state  $s$  has been explored in the current context, and  $N(s, a)$  is the number of times action  $a$  has been taken in state  $s$ . The parameter  $\beta$  is for trading off the exploitation (the former term) and exploration (the latter term).

The knowledge base retrieval score  $K(s, a)$  is obtained by retrieving the concatenated  $(s, a)$  pair from the knowledge base. Specifically, let  $f$  denote the embedding model that maps  $(s, a)$  to a vector with the same dimension as the

knowledge base. Given the preceding reasoning path, the knowledge base retrieval score for the current node is calculated as follows:

$$K(s, a) = \max \left( 0, \max_{k \in \mathcal{K}} \frac{f((s, a)) \cdot k}{\|f((s, a))\| \cdot \|k\|} \right). \quad (4)$$

The knowledge base similarity score  $K(s, a)$  enables acquisition of step-wise assessments prior to the evaluation phase. In other words, when newly generated nodes remain unexplored, we prioritize leveraging historically validated solutions through knowledge base retrieval scores to identify higher-value nodes.

Starting from the root node, we recursively select the child node with the maximum SelectionScore value at each branching point. Selection ties are resolved stochastically. Each iteration advances to the highest-scoring child node until reaching a leaf node.

**Expansion.** Upon selecting a leaf node during the selection phase, the expansion phase aims to generate remaining child nodes, thereby expanding the search scope of the entire tree. Since the attention weights between tokens do not always decay with relative distance (Barbero et al. 2024), we conduct an in-depth study on the attention mechanism between tokens in LLMs during code generation to reveal the influencing factors among tokens. As shown in Figure 2, it can be observed that certain tokens have a profound impact on subsequent code generation. It can thus be inferred that these key tokens can summarize and interpret the information of previously generated tokens, and have higher ref-

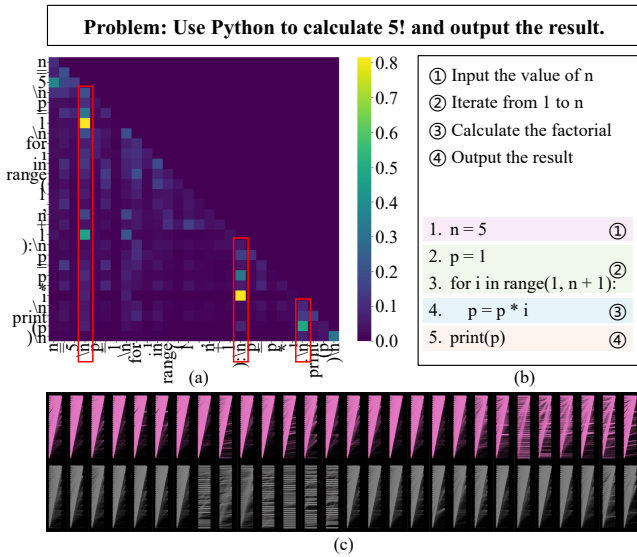


Figure 2: (a) Token-level attention heatmap for code corresponding to the programming problem. (b) Algorithmic steps and corresponding code for the programming problem. (c) Attention sink phenomenon.

erence value for subsequent token generation. Meanwhile, relevant studies (Barbero et al. 2025; Xiao et al. 2023) have shown that modern LLMs exhibit the phenomenon of “attention sink”. Specifically, numerous attention heads allocate a disproportionate share of weights, such as exceeding 30% or even 80%, to the beginning-of-sequence token (bos), despite its primary function as a sequence delimiter with minimal semantic content. Therefore, to facilitate our examination of inter-token dependencies in code generation tasks, we selectively visualize token attention mechanisms at designated layers. Figure 2(c) shows that attention not only sinks to (bos) but also peaks at algorithmic step boundaries, justifying that algorithmic step blocks are more effective basic processing units in code generation tasks. Therefore, we select algorithmic steps as the basic units for expansion.

To ensure diversity in generated steps during the expansion phase, we implement a sampling decoding strategy that sequentially generates each child node. Specifically, to prevent repetitive generation by the LLM, we iteratively provide all previously generated steps as context when producing each new step. The input for the LLM is

$$\text{concat}(s, a_1, \dots, a_i, g), i = 1, 2, \dots, b \quad (5)$$

where  $g$  represents the reflection in the simulation phase, and  $b$  denotes the maximum number of branches each node can expand.

After expanding  $b$  nodes, we employ cosine similarity for filtering to reduce computational costs by avoiding simulations on redundant nodes. Specifically, we map the reasoning steps of the  $b$  nodes to vectors using the embedding model  $\mathcal{E}$  and calculate the cosine similarities between these  $b$  nodes. When the similarity exceeds a predetermined threshold, the node is identified as redundant and filtered out. This

method effectively reduces the search space and enhances algorithmic efficiency while maintaining diversity.

**Evaluation and Reflection.** During the evaluation phase, simulation and evaluation are performed for the selected leaf nodes. We provide the LLM with the algorithmic steps  $s$  already generated for the node and its ancestor nodes, enabling the LLM to strictly follow the generated steps and continue simulating to complete all remaining steps. We search for the thoughts and evaluate with the code generated following the thoughts.

The generated code undergoes sandbox evaluation using public test cases. However, since public test cases only cover a subset of possible scenarios, the code may fail on unseen cases, such as boundary conditions or performance issues. We therefore employ the LLM to analyze the complete algorithmic steps based on sandbox feedback.

We assess the steps generated during the expansion phase through two components, which are the pass rate on public test cases and LLM evaluation. The final evaluation score is obtained by weighted summation of these two scores. The formula is as follows:

$$Q(s, a) = \gamma \cdot r_{\text{exec}} + (1 - \gamma) \cdot r_{\text{LLM}} \quad (6)$$

where  $r_{\text{exec}}$  denotes the pass rate on public test cases,  $r_{\text{LLM}}$  represents the score from LLM evaluation based on the sandbox feedback results and complete steps provided to the LLM, and  $\gamma$  indicates the weight controlling these two parts of the scores.

For code that fails to public test cases, we isolate erroneous algorithmic steps by decomposing the code into blocks and sequentially debugging each block via LLM analysis with public test inputs (Zhong, Wang, and Shang 2024). We retain all correct steps generated during the simulation phase, truncated before the first erroneous step. These validated steps are then incorporated into the MCTS tree as expanded nodes.

The entire RPM-MCTS process is terminated when the solution passes all public test cases and achieves a high LLM evaluation score. Otherwise, node updates and reflection are performed, and the RPM-MCTS process proceeds until the maximum iteration count is reached.

**Backpropagation.** The objective of backpropagation is to update the reward values of nodes upon completion of state value evaluation. We propagate reward values backward from leaf nodes to the root node, updating the state estimates of all nodes along the path. For newly generated nodes during the expansion phase, they collectively update their parent node. As the number of simulations increases, these value estimates become increasingly accurate. This process repeats until the preset maximum simulation count is reached, ultimately resulting in a search tree that records the state value and visit count for each node.

## Generate Code

Termination of the RPM-MCTS process occurs under two conditions: 1) If all public test cases are passed and LLM analysis confirms robustness to unseen edge cases before reaching maximum iterations, the code generated during the

simulation phase is retained. 2) When maximum iterations are reached without meeting termination criteria, the leaf node with the highest state value is selected, its ancestral path is traced, and the LLM is instructed to generate code by rigorously adhering to the algorithmic steps assembled from this path.

## Experiments

### Experimental Settings

**Datasets.** For the construction of the knowledge base, we use the train set splits of APPS (Hendrycks et al. 2021) and CodeContests (Li et al. 2022) as data sources. After validation and filtering, we obtained 11,038 samples with a total of 82,923 steps. For benchmarking, we used the test set splits of APPS and CodeContests, as well as HumanEval+ (Liu et al. 2023) and MBPP+ (Liu et al. 2023). The APPS dataset contains three difficulty levels: introductory, interview, and competition. We selected 150 validated samples from each difficulty level. The CodeContests dataset consists of competitive programming problems collected from contest websites such as Codeforces. Additionally, HumanEval (Chen et al. 2021) and MBPP (Austin et al. 2021) are widely recognized benchmarks in the code generation domain, while HumanEval+ and MBPP+ introduce a larger number of test cases to enable more accurate evaluations. We utilized Claude Sonnet 3.7 to convert all datasets into a unified format, which primarily includes the problem statement, public test cases, private test cases, and standard solution. To facilitate sandbox execution, we transformed datasets with standard input-output problems into function definitions with docstrings. For datasets without public test cases, we selected the first two private test cases as public test cases.

**Baselines.** We selected the following methods as baselines for comparison. Base LLM refers to directly prompting the LLM to output solution code using the problem statement and public test cases as input. LDB (Zhong, Wang, and Shang 2024) leverages the LLM to track intermediate variables during code execution to iteratively improve the code. ToT (Yao et al. 2023) performs a search of thought steps using DFS or BFS before generating the final code. SRA-MCTS (Xu et al. 2025) combines LLM with MCTS to explore intermediate reasoning steps. The complete steps obtained by SRA-MCTS are used as input to prompt the LLM to directly infer and output the solution code for evaluation.

**Implementation Details.** We use two large-parameter backbone models, Qwen3-235B-A22B (Yang et al. 2025) and Claude Sonnet 3.7, alongside a smaller-parameter model, Qwen3-8B. In the code generation domain, pass@k (Chen et al. 2024) is a widely used metric, and we adopted pass@1 as the evaluation metric. The rollout, i.e., maximum number of iterations, was set to 5 for all methods. The branching factor  $b$  for tree-based methods was set to 3. The exploration constant  $\beta$  for UCB was set to 0.5. In RPM-MCTS, the weight of the knowledge base retrieval score  $\alpha$  was set to 0.5, and the similarity filtering threshold was set to 0.85.

### Main Results

Our method achieves the most significant improvements across different backbone models and datasets. As shown in Table 1, Qwen3-8B achieves an average improvement of 11.90%, Qwen3-235B-A22B achieves an average improvement of 7.71%, and Claude Sonnet 3.7 achieves an average improvement of 10.86%. Since the base Qwen3-8B performs worse than the other two larger base LLMs across all datasets, especially on simpler datasets, the Qwen3-8B shows the most significant improvement when using RPM-MCTS. On the two more challenging datasets, APPS-competition and CodeContests, Qwen3-8B achieves an average improvement of 13.3%, Qwen3-235B-A22B achieves an average improvement of 14.67%, and Claude Sonnet 3.7 achieves an average improvement of 18.34%. This is because Qwen3-8B has weaker evaluation scoring capabilities, while larger LLMs have relatively stronger evaluation capabilities, resulting in greater gains. This demonstrates that the more difficult the task, the more accurate evaluation of intermediate algorithm steps is required.

LDB achieves greater improvements on simpler datasets compared to more challenging ones. We found that this is because, for more difficult problems, through multiple rounds of execution feedback, LLMs often only modify code conditions to pass public test cases rather than thinking about modifying the actual logic of the code. SRA-MCTS shows performance improvements on more challenging datasets but declines on simpler ones. The reason is that for simple problems, LLM evaluation scores are always perfect or near-perfect, prematurely ending the search for steps, resulting in incomplete or lower-quality reasoning steps.

Comparing the results across three different difficulty levels in the APPS dataset, it can be observed that for the two larger LLMs, as the difficulty increases, our method brings more significant performance improvements. The higher the difficulty of the problem, the more guidance the LLM needs to avoid getting lost in complex reasoning chains. This demonstrates the effectiveness of our method in evaluating intermediate steps, helping LLMs enhance their evaluation capabilities and further unlocking the vast potential code knowledge and reasoning abilities inherent in LLMs.

For fair comparison, even without using knowledge base retrieval scores as rewards, our method outperforms other baselines. Experimental results show that overall, especially on the two most challenging datasets, incorporating the knowledge base further stabilizes and improves performance. The reason is that LLM evaluation of intermediate steps in complex problems is unreliable, and random exploration struggles to find the correct solution path. Therefore, leveraging the knowledge base to use the reasoning patterns of historically similar problems as guidance helps direct the search. This demonstrates the effectiveness of using knowledge base retrieval scores as rewards for intermediate process evaluation.

On a few simpler datasets, performance slightly improves when knowledge base retrieval scores are not used. We analyze that this is because, in simple tasks, LLMs can already accurately evaluate the quality of generated paths. In this case, introducing knowledge base rewards, while aiming to

Method	APPS-Intro.	APPS-Interv.	APPS-Comp.	CodeContests	HumanEval+	MBPP+	Average
<b>Qwen3-8B</b>							
Base LLM	56.7	35.3	29.3	10.7	75.6	72.2	52.1
LDB	64.0 (+7.3)	42.0 (+6.7)	28.0 (-1.3)	11.3 (+0.7)	78.1 (+2.4)	70.1 (-2.1)	53.5 (+1.4)
ToT	69.3 (+12.7)	54.0 (+18.7)	41.3 (+12.0)	17.3 (+6.7)	82.3 (+6.7)	70.4 (-1.9)	59.0 (+6.9)
SRA-MCTS	67.3 (+10.7)	42.7 (+7.3)	29.3 (+0.0)	16.0 (+5.3)	73.8 (-1.8)	65.9 (-6.4)	52.8 (+0.7)
Ours w/o KB	76.7 (+20.0)	56.7 (+21.3)	40.7 (+11.3)	22.3 (+11.6)	82.3 (+6.7)	<b>78.3 (+6.1)</b>	63.5 (+11.4)
Ours	<b>77.3 (+20.7)</b>	<b>60.0 (+24.7)</b>	<b>43.6 (+14.3)</b>	<b>23.0 (+12.3)</b>	<b>83.5 (+7.9)</b>	76.2 (+4.0)	<b>64.0 (+11.9)</b>
<b>Qwen3-235B-A22B</b>							
Base LLM	78.0	54.7	42.7	24.0	86.0	78.8	64.6
LDB	78.7 (+0.7)	61.3 (+6.7)	47.3 (+4.7)	25.3 (+1.3)	86.0 (+0.0)	78.8 (+0.0)	66.4 (+1.8)
ToT	84.7 (+6.7)	62.7 (+8.0)	57.3 (+14.7)	27.3 (+3.3)	85.4 (-0.6)	75.4 (-3.4)	67.7 (+3.1)
SRA-MCTS	76.0 (-2.0)	52.7 (-2.0)	44.0 (+1.3)	24.7 (+0.7)	85.4 (-0.6)	70.9 (-7.9)	61.7 (-3.0)
Ours w/o KB	<b>88.0 (+10.0)</b>	<b>72.0 (+17.3)</b>	52.0 (+9.3)	34.7 (+10.7)	86.6 (+0.6)	79.9 (+1.1)	71.3 (+6.7)
Ours	86.7 (+8.7)	67.3 (+12.7)	<b>59.3 (+16.7)</b>	<b>36.7 (+12.7)</b>	<b>87.8 (+1.8)</b>	<b>81.2 (+2.4)</b>	<b>72.3 (+7.7)</b>
<b>Claude Sonnet 3.7</b>							
Base LLM	78.7	56.0	59.3	31.3	82.9	77.8	67.3
LDB	82.0 (+3.3)	64.7 (+8.7)	73.3 (+14.0)	33.3 (+2.0)	88.4 (+5.5)	77.0 (-0.8)	71.5 (+4.2)
ToT	84.0 (+5.3)	68.0 (+12.0)	66.0 (+6.7)	39.3 (+8.0)	86.0 (+3.1)	74.6 (-3.2)	70.8 (+3.6)
SRA-MCTS	83.3 (+4.7)	63.3 (+7.3)	62.0 (+2.7)	36.0 (+4.7)	81.1 (-1.8)	74.3 (-3.4)	68.4 (+1.1)
Ours w/o KB	<b>92.0 (+13.3)</b>	73.3 (+17.3)	78.0 (+18.7)	42.7 (+11.3)	86.6 (+3.7)	79.1 (+1.3)	76.2 (+8.9)
Ours	<b>92.0 (+13.3)</b>	<b>74.0 (+18.0)</b>	<b>81.3 (+22.0)</b>	<b>46.0 (+14.7)</b>	<b>89.0 (+6.1)</b>	<b>81.0 (+3.2)</b>	<b>78.1 (+10.9)</b>

Table 1: Performance comparison of all methods across different backbone models on code generation benchmarks. Values in parentheses indicate the improvement over the base LLM.

provide additional prior information, may retrieve historical cases that are textually similar but logically different in their solutions, introducing noise into MCTS node selection. In contrast, for complex tasks, LLM evaluation capabilities for intermediate steps are limited, the search space is vast, and solutions are sparse. The structured priors provided by the knowledge base effectively guide the search direction, significantly improving success rates. This phenomenon indicates that the effectiveness of knowledge base rewards depends on the balance between task difficulty and LLM evaluation confidence.

### Ablation Study

We conduct ablation experiments using Qwen3-235B-A22B as the backbone model to evaluate performance, and the results are shown in Figure 3.

**w/o KB** indicates that only LLM evaluation is used in selection, without knowledge base retrieval. Compared to the complete method, the overall performance slightly decreased, with an average drop of 1.05%. The decline was most significant on the two more challenging datasets, with an average drop of 4.67%. This indicates that large models still face challenges with complex problems. By introducing a knowledge base to compare the generated reasoning steps with the correct reasoning steps of similar problems in the knowledge base, the self-assessment capability for complex problems can be improved.

**w/o ER** means that the execution rewards of public test cases in the sandbox are not used during the simulation phase. This resulted in the largest overall performance drop, highlighting that the core of RPM-MCTS reflection lies in




Figure 3: Ablation study results on different benchmarks.

the detailed feedback provided by the code execution environment. In fact, previous research (Huang et al. 2023) has already pointed out that without external feedback, LLMs lack the ability to self-correct their reasoning processes.

**w/o SF** refers to the removal of similarity filtering, i.e., not discarding similar child nodes during the expansion phase. The results show that filtering out repeated intermediate algorithmic steps based on similarity allows resources to be better allocated to exploring new steps, thereby improving performance while reducing computational costs.

**w/o LDB** denotes not using LDB to locate erroneous steps in our method. The average performance drop was minimal, indicating that removing LDB has little impact on our method. With execution feedback, LLMs are already capable of accurately locating errors. However, in a few cases,

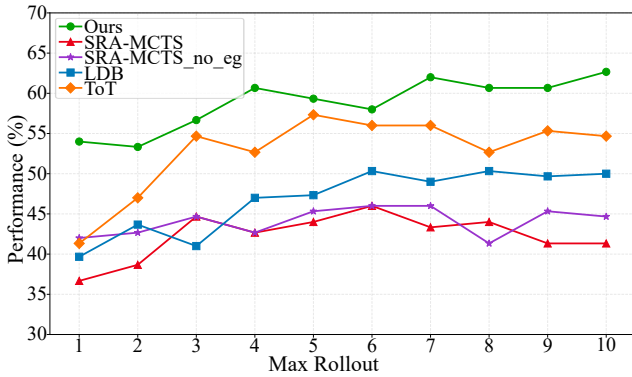


Figure 4: Performance comparison across different maximum rollout values.

LDB still helps in pinpointing erroneous steps.

### Performance vs. Rollout

We explore the results of different values of the hyperparameter rollout on Qwen3-235B-A22B, as shown in Figure 4. Since SRA-MCTS is prone to premature termination due to self-overestimation by the model, we set its end gate value to exceed the maximum possible score, allowing it to reach the maximum number of iterations whenever possible. We denote this variant as SRA-MCTS\_no\_eg. The results show that in the early stages, all methods exhibit significant performance improvements as the rollout increases, after which the performance gradually stabilizes. Notably, RPM-MCTS exhibits better performance even with a rollout of 1. This is because it enjoys two advantages in its first rollout: proactive guidance via its Knowledge Base during the selection phase, and wrong step truncation with rethink-based regeneration during the simulation phase. This allows it to perform at least one round of verification and reflection and generate complete code. Moreover, for simpler problems, RPM-MCTS can often arrive at the correct answer with only a single simulation, whereas traditional tree search methods tend to require multiple unnecessary expansions even for straightforward tasks.

### Token Efficiency Analysis

Figure 5 shows the average token usage of different methods across all benchmark datasets. Our method reduces token consumption by approximately 15% compared to the previous MCTS method on both Qwen3-235B-A22B and Claude Sonnet 3.7. This improvement is attributed to: 1) The knowledge base retrieval scoring prioritizes more correct nodes, avoiding exploration of invalid branches. 2) Similarity filtering eliminates duplicate intermediate reasoning steps, enabling dynamic pruning of the Monte Carlo tree and reducing redundant path generation. 3) The simulation phase leverages sandbox feedback to pinpoint erroneous steps, while retaining the verified correct ones. Overall, RPM-MCTS achieves enhanced search efficiency and generation quality through knowledge base guidance and execution feedback.

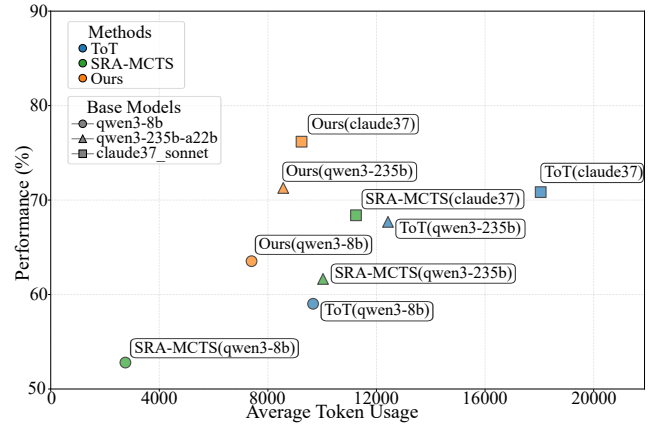


Figure 5: Comparison of token consumption and performance across different methods and models.

Benchmark	Base	Ours	$\Delta$
SWE-Bench (Jimenez et al.)	37.6	38.5	+0.9
MBPP+ (Liu et al.)	75.4	76.7	+1.3
LiveCodeBench (Jain et al.)	46.2	50.5	+4.3
Aider (Aider-AI)	17.3	22.2	+4.9
McEval (Chai et al.)	57.5	61.2	+3.7

Table 2: Supervised fine-tuning results on Doubao-1.5-pro-32K model using RPM-MCTS synthesized data.

### Reasoning Steps for Data Distillation

Our method is training-free, yet the algorithmic synthesized reasoning steps it produces can also be used for supervised fine-tuning. Based on Doubao-1.5-pro-32K model, we utilize RPM-MCTS for data distillation and construct a dataset of 2.4k code generation samples with reasoning steps. This distilled data is then combined with a foundational dataset of 170k samples to perform full fine-tuning on Doubao-1.5-pro-32K. Benchmark results in Table 2 demonstrate that the training data generated using RPM-MCTS significantly enhances the code capabilities of the base model.

### Conclusion

In this paper, we propose RPM-MCTS, which leverages a knowledge base and external sandbox feedback to directly obtain accurate reward values without additional training of a process reward model. During the search process, errors are identified and promptly corrected. Results demonstrate that RPM-MCTS outperforms state-of-the-art methods under more constrained search budgets. Additionally, we construct training data using RPM-MCTS and perform full fine-tuning on base model, which significantly enhances code capabilities of the base model.

A limitation of RPM-MCTS is that code solvable in a single line may be divided into multiple lines due to the step-by-step approach, without impacting correctness. In the future, during the evaluation phase of MCTS, we can dynamically adjust the weights of external rewards from the knowledge base and sandbox, based on the uncertainty of LLMs.

## Acknowledgments

This work was supported by the National Science and Technology Major Project (2022ZD0114803), the Natural Science Foundation of Wuhan (2023010201020229), the Fundamental Research Funds for the Central Universities (NO.NJ2023032), and the Major Program (JD) of Hubei Province (2023BAA024).

## References

- Aider-AI. 2024. Aider’s Polyglot Benchmark. <https://aider.chat/2024/12/21/polyglot.html>. Accessed: 2025-08-01.
- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Barbero, F.; Vitvitskiy, A.; Perivolaropoulos, C.; Pascanu, R.; and Veličković, P. 2024. Round and round we go! what makes rotary positional encodings useful? *arXiv preprint arXiv:2410.06205*.
- Barbero, F.; Álvaro Arroyo; Gu, X.; Perivolaropoulos, C.; Bronstein, M.; Veličković, P.; and Pascanu, R. 2025. Why do LLMs attend to the first token? *arXiv preprint arXiv:2504.02732*.
- Chai, L.; Liu, S.; Yang, J.; Yin, Y.; Jin, K.; Liu, J.; Sun, T.; Zhang, G.; Ren, C.; Guo, H.; et al. 2024. McEval: Massively Multilingual Code Evaluation. *arXiv preprint arXiv:2406.07436*.
- Chen, L.; Guo, Q.; Jia, H.; Zeng, Z.; Wang, X.; Xu, Y.; Wu, J.; Wang, Y.; Gao, Q.; Wang, J.; et al. 2024. A survey on evaluating large language models in code generation tasks. *arXiv preprint arXiv:2408.16498*.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; Pinto, H. P. D. O.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- 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, B.; Cai, Z.; Xu, R.; Wang, P.; Zheng, C.; Lin, R.; Lu, K.; Liu, D.; Zhou, C.; Xiao, W.; et al. 2024. Llm critics help catch bugs in mathematics: Towards a better mathematical verifier with natural language feedback. *arXiv preprint arXiv:2406.14024*.
- Hao, S.; Gu, Y.; Ma, H.; Hong, J. J.; Wang, Z.; Wang, D. Z.; and Hu, Z. 2023. Reasoning with language model is planning with world model. *arXiv preprint arXiv:2305.14992*.
- Hendrycks, D.; Basart, S.; Kadavath, S.; Mazeika, M.; Arora, A.; Guo, E.; Burns, C.; Puranik, S.; He, H.; Song, D.; et al. 2021. Measuring coding challenge competence with apps. *arXiv preprint arXiv:2105.09938*.
- Huang, J.; Chen, X.; Mishra, S.; Zheng, H. S.; Yu, A. W.; Song, X.; and Zhou, D. 2023. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*.
- Jain, N.; Han, K.; Gu, A.; Li, W.-D.; Yan, F.; Zhang, T.; Wang, S.; Solar-Lezama, A.; Sen, K.; and Stoica, I. 2024. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*.
- Jiang, X.; Dong, Y.; Wang, L.; Fang, Z.; Shang, Q.; Li, G.; Jin, Z.; and Jiao, W. 2024. Self-planning code generation with large language models. *ACM Transactions on Software Engineering and Methodology*, 33(7): 1–30.
- Jimenez, C. E.; Yang, J.; Wettig, A.; Yao, S.; Pei, K.; Press, O.; and Narasimhan, K. R. 2024. SWE-bench: Can Language Models Resolve Real-world Github Issues? In *The 12th International Conference on Learning Representations (ICLR’24)*.
- Li, J.; Li, G.; Li, Y.; and Jin, Z. 2025a. Structured chain-of-thought prompting for code generation. *ACM Transactions on Software Engineering and Methodology*, 34(2): 1–23.
- Li, Q.; Xia, W.; Dai, X.; Du, K.; Liu, W.; Wang, Y.; Tang, R.; Yu, Y.; and Zhang, W. 2025b. Rethinkmcts: Refining erroneous thoughts in monte carlo tree search for code generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP’25)*, 8103–8121.
- Li, Y.; Choi, D.; Chung, J.; Kushman, N.; Schrittwieser, J.; Leblond, R.; Eccles, T.; Keeling, J.; Gimeno, F.; Dal Lago, A.; et al. 2022. Competition-level code generation with alphacode. *Science*, 378(6624): 1092–1097.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2024. Let’s verify step by step. In *The 12th International Conference on Learning Representations (ICLR’24)*.
- Liu, J.; Xia, C. S.; Wang, Y.; and Zhang, L. 2023. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems 36 (NeurIPS’23)*, 21558–21572.
- Lu, T.; Li, Y.; Wang, L.; Lin, B.; Tang, J.; Lv, Q.; Xu, W.; Zheng, H.-T.; Li, Y.; Su, X.; et al. 2025. LSR-MCTS: Alleviating Long Range Dependency in Code Generation. *arXiv preprint arXiv:2504.07433*.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.
- Wang, E.; Cassano, F.; Wu, C.; Bai, Y.; Song, W.; Nath, V.; Han, Z.; Hendryx, S.; Yue, S.; and Zhang, H. 2024. Planning in natural language improves llm search for code generation. *arXiv preprint arXiv:2409.03733*.
- Wang, P.; Li, L.; Shao, Z.; Xu, R.; Dai, D.; Li, Y.; Chen, D.; Wu, Y.; and Sui, Z. 2023. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv preprint arXiv:2312.08935*.

- 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. In *Advances in neural information processing systems 35 (NeurIPS'22)*, 24824–24837.
- Xiao, G.; Tian, Y.; Chen, B.; Han, S.; and Lewis, M. 2023. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Xiao, S.; Liu, Z.; Zhang, P.; Muennighoff, N.; Lian, D.; and Nie, J.-Y. 2024. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'24)*, 641–649.
- Xie, T.; Gao, Z.; Ren, Q.; Luo, H.; Hong, Y.; Dai, B.; Zhou, J.; Qiu, K.; Wu, Z.; and Luo, C. 2025. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2502.14768*.
- Xu, B.; Lin, Y.; Li, Y.; and Gao, Y. 2025. SRA-MCTS: Self-driven Reasoning Augmentation with Monte Carlo Tree Search for Code Generation. In *Proceedings of the 34th International Joint Conference on Artificial Intelligence (IJCAI'25)*, 8678–8686.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; and Narasimhan, K. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems 36 (NeurIPS'23)*, 11809–11822.
- Zelikman, E.; Huang, Q.; Poesia, G.; Goodman, N.; and Haber, N. 2023. Parsel: Algorithmic Reasoning with Language Models by Composing Decompositions. In *Advances in Neural Information Processing Systems 36 (NeurIPS'23)*, 31466–31523.
- Zhang, D.; Zhoubian, S.; Hu, Z.; Yue, Y.; Dong, Y.; and Tang, J. 2024a. ReST-MCTS\*: LLM Self-Training via Process Reward Guided Tree Search. In *Advances in Neural Information Processing Systems 37 (NeurIPS'24)*, 64735–64772.
- Zhang, Z.; Chen, C.; Liu, B.; Liao, C.; Gong, Z.; Yu, H.; Li, J.; and Wang, R. 2024b. Unifying the Perspectives of NLP and Software Engineering: A Survey on Language Models for Code. *Transactions on Machine Learning Research*.
- Zhao, Z.; Lee, W. S.; and Hsu, D. 2023. Large language models as commonsense knowledge for large-scale task planning. In *Advances in Neural Information Processing Systems 36 (NeurIPS'23)*, 31967–31987.
- Zheng, K.; Decugis, J.; Gehring, J.; Cohen, T.; Negrevergne, B.; and Synnaeve, G. 2024. What Makes Large Language Models Reason in (Multi-Turn) Code Generation? *arXiv preprint arXiv:2410.08105*.
- Zhong, L.; Wang, Z.; and Shang, J. 2024. Debug like a Human: A Large Language Model Debugger via Verifying Runtime Execution Step by Step. In *Findings of the Association for Computational Linguistics (ACL'24)*, 851–870.
- Zhou, A.; Yan, K.; Shlapentokh-Rothman, M.; Wang, H.; and Wang, Y.-X. 2024. Language Agent Tree Search Unifies Reasoning, Acting, and Planning in Language Models. In *Proceedings of the 41st International Conference on Machine Learning (ICML'24)*, 62138–62160.