

Thinking Before Running! Efficient Code Generation with Thorough Exploration and Optimal Refinement

Xiaoqing Zhang^{1,2*} Yuhan Liu^{1†} Flood Sung²
Xiuying Chen³ Shuo Shang⁴ Rui Yan^{1,5,6†}

¹Gaoling School of Artificial Intelligence, Renmin University of China ²Moonshot AI

³Mohamed bin Zayed University of Artificial Intelligence

⁴University of Electronic Science and Technology of China ⁵School of Artificial Intelligence, Wuhan University

⁶Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MoE

{xiaoqingz, yuhan.liu, ruiyan}@ruc.edu.cn floodsung@moonshot.cn
xy-chen@pku.edu.cn jedi.shang@gmail.com

Abstract

Code generation is crucial in software engineering for automating the coding process efficiently. While test-time computation methods show promise, they suffer from high latency due to multiple computation rounds. To overcome this, we introduce **ThinkCoder**, a framework that combines thorough exploration with optimal refinement. The exploration phase diversifies the solution space by searching for potential solutions, followed by a refinement phase that enhances precision. This approach allows us to select the best solution through careful consideration before taking action, avoiding excessive trial and error. To further minimize test-time computation overhead, we introduce preference-driven optimization with Reinforced Self-Training (ReST), which uses exploration trajectories from ThinkCoder to guide LLM’s evolution. This approach enhances LLM’s exploration efficiency via preference learning, cutting costs while maintaining accuracy. ThinkCoder boosts the performance with a single LLM, excelling on benchmarks like HumanEval and MBPP. Compared to SOTA models, it improves Pass@1 by 3.0% over MapCoder with just 6.4% of the computation cost. Against AgentCoder, ThinkCoder achieves a 0.5% higher Pass@1 after 2 rounds, outperforming AgentCoder’s 5 rounds. Additionally, ReST with success trajectories enhances efficiency, allowing models like LLaMA2-7B to achieve competitive results using only 20% of the computational resources. These results highlight the framework’s effectiveness and scalability. ¹

1 Introduction

Recent research advances indicate that large language models (LLMs) have demonstrated remark-

^{*}This work was done during the internship at Moonshot AI.

[†]Corresponding authors.

¹The code is available at <https://github.com/xiaoqzhwu/ThinkCoder>

able capabilities in various programming-related domains, such as code generation (Zheng et al., 2023; Chaudhary, 2023; Dong et al., 2024; Li et al., 2024a), code refinement (Chen et al., 2023; Guo et al., 2024; Zheng et al., 2024; Ridnik et al., 2024; Liu et al., 2024a; Zhang et al., 2024), and software testing (Li et al., 2023; Jalil et al., 2023; Wang et al., 2024; Jones et al., 2024). By utilizing extensive training data, LLMs can understand complex programming tasks, generate syntactically accurate code, and even enhance the quality of their solutions through iterative improvement.

In the context of code generation, researchers have made significant progress in improving prompts to encourage LLMs to generate higher-quality answers. This includes two main directions: optimizing the path for answer exploration and incorporating rich reflection information. In the first direction of path optimization, for example, Self-Planning (Jiang et al., 2024) reduces problem complexity in code generation by introducing a planning phase, which refines the solution generation process. Similarly, CodeCoT (Huang et al., 2023a) improves the exploration by employing a Chain-of-Thought(CoT) prompt, enabling models to reason through each step systematically. In the second direction, the integration of reflection, test cases, and feedback is progressively being used to steer and enhance the generation process. For instance, RethinkMCTS (Li et al., 2024b) integrates verbal feedback into its Monte Carlo Tree Search framework to correct errors and improve code generation. Similarly, AgentCoder (Huang et al., 2023b) boosts efficiency and accuracy by using interactions between multiple agents to provide valuable input.

Despite the significant potential of using CoT for planning code improvements, the thought process incurs a substantial token overhead. To enhance efficiency and minimize trial-and-error costs, it is crucial to clarify the optimization direction before implementing refinement. Existing methods, such

as MapCoder, which leverages four LLM-based agents, consume a considerable amount of computational resources.

Based on these insights, we propose ThinkCoder, a novel framework for efficient code generation that combines thorough exploration and optimal refinement. First, we simplified the multi-agent framework into a single Exploration Agent and a CodeVerifier. The Exploration Agent is responsible for in-depth analysis and generating various codes, while exploring diverse test cases to build a robust testing pool. The CodeVerifier, using the testing pool, independently determines the optimization direction without relying on LLMs. By continuously improving the testing pool throughout the refinement process, ThinkCoder ensures accurate verification of generated code in each round, providing efficient guidance for subsequent optimizations.

Second, to tackle the additional computational load resulting from extensive exploration, we incorporate a preference-driven optimization phase using Reinforced Self-Training (ReST) in ThinkCoder. ReST leverages successful exploration trajectories to train LLM, enabling it to generate optimal solutions more efficiently during exploration. This approach significantly reduces the computational cost of test-time computations, enabling models like LLaMA2-7B to achieve notable improvements in both performance and efficiency.

Our contributions are as follows:

- We introduce ThinkCoder, an efficiency framework for code generation that integrates an LLM-based exploration agent with a non-LLM-based CodeVerifier. We utilize a self-evolving testing pool, significantly enhancing the speed of code verification and improvement.
- We introduce preference-driven optimization with ReST, using exploration trajectories from ThinkCoder to enhance LLMs’ generation and verification capabilities, further reducing computational overhead during testing.
- We have demonstrated the effectiveness of ThinkCoder on multiple benchmark datasets and across different scales of LLMs, achieving state-of-the-art performance in code generation while optimizing computational efficiency.

2 Related Work

Test-time Compute. Test-time computing methods use self-play and additional verifiers to enhance performance, despite the computational overhead.

For example, training a Verifier with Preference Reward Model data and combining strategies like Best-of-N Weighted, Beam Search, and Lookahead Search can help (Snell et al., 2024). ToolLLM (Qin et al., 2023) evaluates tool usage by comparing success rates and solution quality to ChatGPT-ReACT, employing ToolEval as a Verifier for optimal API calls using depth-first search. However, API-based Verifiers are limited by training data and struggle with generalization. In contrast, our approach employs a task-based testing tool as a Verifier, offering a simpler and more effective solution that is closely aligned with the specific task requirements.

Multi-Agent Code Generation. Deriving multiple agents by mimicking the different stages of human programming can effectively improve both the quality and efficiency of code generation, as agents based on large language models possess strong reasoning capabilities (Liu et al., 2024b,c; Zhang et al., 2025). For instance, AgentCoder (Huang et al., 2023b) introduces programmer agents, test design agents, and test execution agents during the code generation process. Similarly, MapCoder (Islam et al., 2024) introduces retrieval agents, planning agents, coding agents, and debugging agents, iterating on code generation quality through more frequent agent communication and reflection. However, MapCoder incurs higher computational overhead and relies solely on sample I/O for code execution, limiting test case coverage and robustness. In contrast, our framework, ThinkCoder, addresses these limitations by self-verification efficiently without excessive computational costs.

Instruction Tuning for Code. High-quality code generation models often require instruction fine-tuning based on large-scale pre-trained language models. The fine-tuning data generally comes from real-world collections or is artificially synthesized. For example, Code Alpaca (Chaudhary, 2023) combines the SELF-INSTRUCT strategy with code fine-tuning, CODELLAMA (Roziere et al., 2023) generates questions, solutions, and test cases through prompts, and OctoPack (Muennighoff et al., 2023) directly collects data from Git. LLMs fine-tuned with instructions possess stronger code generation capabilities. However, existing methods often lack comprehensive guidelines for data collection, volume management, and fine-tuning strategies. To address this gap, we present a detailed fine-tuning strategy for Llama2-7B using instruction-tuning data collected via CodeQwen (Bai et al., 2023), pro-

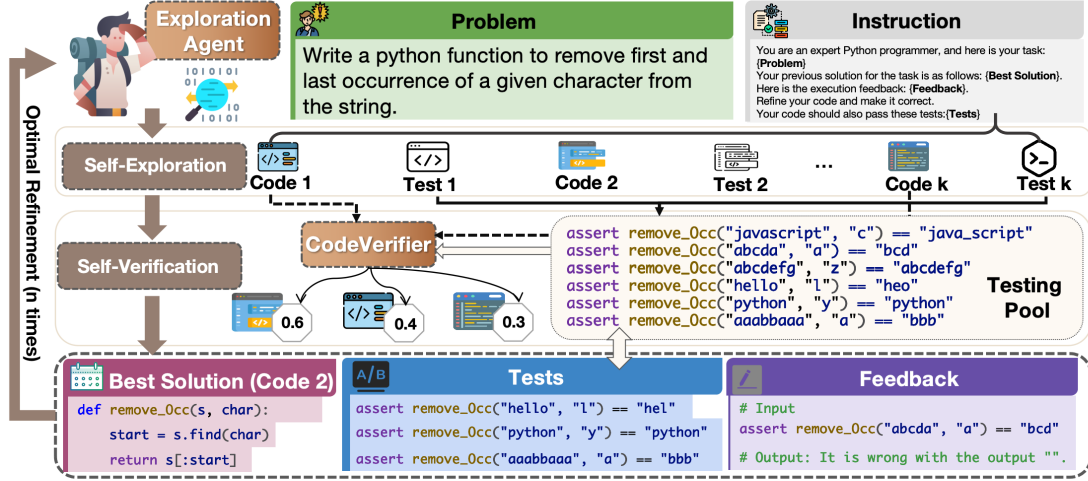


Figure 1: The end-to-end process of ThinkCoder involves k thorough exploration steps followed by n optimal refinement cycles. The Exploration Agent generates k codes and tests simultaneously during self-exploration, storing results in a Testing Pool. The optimal refinement phase includes self-verification that selects the optimal solution with the CodeVerifier and aggregate reflection for the instruction update. The optimal refinement will be repeated recursively for n cycles, ultimately leading to the final solution.

viding a transparent and reproducible methodology that supports the advancement of instruction-tuned code generation models.

3 ThinkCoder

3.1 Overall Architecture

Figure 1 illustrates the structure of ThinkCoder. Given an instruction $I = [p, \text{prompt}]$ for problem p , where prompt is the task description for code generation, the Exploration Agent produces k candidate solutions. Each solution is paired with m test cases, and the temperature t of the LLM’s generation is adjusted to promote diversity. These test cases are then aggregated into a Testing Pool TP . We define the exploration results as $E = [G, TP]$, where $G = \{g_i \mid g_i = \text{LLM}(p, t) \text{ for } i \in \{1, 2, \dots, k\}\}$, and $TP \in \mathbb{T}^{k \times m}$, where \mathbb{T} represents each test. The CodeVerifier sequentially executes each candidate solution in G using the test cases from TP to compute pass rates. The code g_s with the highest pass rate is selected for further refinement. Here we define the pass rate of g_i as $r_{g_i} = \frac{\text{number of tests passed by } g_i}{|TP|}$ and $g_s = \arg \max_i r_{g_i}$. Feedback f from failed test case f_t , along with the current code g_s , generated test cases $tests$, and problem p , is used to update the instructions $I = [p, g_s, f_t, f, tests, \text{prompt}]$ for the next exploration. This iterative process allows the Exploration Agent to generate improved solutions and test cases over n iterations. To avoid unnecessary computation on well-solved problems, if r_{g_i} is

greater than θ , the refinement process is terminated, and the code g_s is returned as the final solution. For more details of ThinkCoder please refer to the algorithm 1 in Appendix A.1. Below, we detail the implementation of Thorough Exploration, Optimal Refinement, Exploration Agent&CodeVerifier, and Testing Pool.

3.2 Thorough Exploration

During the thorough exploration phase, ThinkCoder conducts extensive thinking before running code and performing corrections. The self-exploration process generates k answers simultaneously to expand the search area by increasing the LLM’s temperature, which flattens the probability distribution and enhances diversity. The first iteration of self-exploration aims to generate diverse code and comprehensive test cases, including regular tests, boundary tests, and performance tests. While the subsequent iterations, focus on the exploration of code correction for specific errors. The only LLM-based agent conducts self-exploration with prompt design detailed in Appendix A.4.

3.3 Optimal Refinement

To solve the problem efficiently, it is essential to first identify the optimal solution through careful consideration before making any code corrections. Therefore, optimal refinement involves selecting the best solution via self-verification and gathering the necessary information for deeper exploration




Figure 2: Trajectory collection with ThinkCoder and its application in ReST training for LLMs. We collect success trajectories offline based on the verification with ground truth tests, ensuring the solution aligns with human preferences. For reflection, we use LLM-generated tests, ensuring LLMs specifically address their own mistakes.

and improvement in the next exploration. During self-verification, code quality is assessed by the pass rate, calculated by the CodeVerifier running candidate code against all test cases. The code with the highest pass rate is chosen for refinement. We then randomly select an unresolved error to target in the current round of improvements, ensuring all errors are eventually addressed. The refinement instruction is generated based on the problem, the current best solution, feedback, and test cases. After each refinement, the solution’s pass rate is evaluated. If it exceeds previous solutions, it becomes the new best. Through multiple iterations, the best solution is continuously corrected and optimized, resulting in a gradual increase in the pass rate within the testing pool. For more details, please refer to the Appendix A.4.

3.4 Exploration Agent& CodeVerifier

ThinkAgent consists of a single exploration agent and a CodeVerifier. The exploration agent uses LLMs to search for diversity code and test cases, while the CodeVerifier runs the provided code snippets and test cases in a local environment to generate feedback. To improve the efficiency of the CodeVerifier, we leverage multiprocessing to validate the code in parallel.

3.5 Testing Pool

The Testing Pool includes numerous test cases generated by the exploration agent, categorized into regular, boundary, and performance tests. Our goal is to create a robust pool to assess complex situations beyond simple cases. A higher pass rate indicates stronger code robustness, so we ensure test case diversity and accuracy. Several key operations contribute to achieving this. Abstract Syntax Tree (AST) (Backus et al., 1960) is applied to identify and remove duplicate test cases in the testing pool, eliminating unnecessary time costs associated with redundant tests. Additionally, after each itera-

tion, the testing pool only selects test cases that can improve the code’s pass rate, thus enhancing the diversity and accuracy of the pool. Finally, the testing pool will select test cases with distinct errors and use feedback to update the code’s improvement goals, driving continuous optimization.

4 Preference-driven Optimization

In ThinkCoder, the generated code evolves as the pass rate on the testing pool improves. Each increase in the pass rate signifies the development of better codes. This progression demonstrates that the exploration agent is gradually producing higher-quality outputs. By collecting the trajectory of these improvements, we aim to fine-tune the model, enabling the LLM to generate superior code and test cases more effectively. This approach helps to refine the exploration process, clarify its direction, and reduce the time required for exploration. Figure 2 describes the process of trajectory collection and fine-tuning of the LLM. We will explain this process from the perspectives of Data Collection, the Model Training and Inference.

4.1 Data Collection

The process of data collection for LLM fine-tuning involves two key steps. First, an exploration agent, denoted as \mathcal{M}_0 , generates diverse low-probability data $D_{\mathcal{M}_0}$ by sampling at various temperatures. Second, this data is further refined using reflection to produce an enhanced dataset $R_{\mathcal{M}_0}$. In our data generation task, offline data collection starts with a training set $D = \{x_i\}_{i=1}^N$, where N is the number of instances in the dataset. We generate high-quality preference data $D_{\mathcal{M}_0} \cup R_{\mathcal{M}_0}$ for LLM fine-tuning through the temperature-based and reflection-based generation process.

Temperature-based Generation. During temperature-based generation, we obtain $D_{\mathcal{M}_0}$ to enhance the diversity of responses. For a given problem x , the language model \mathcal{M}_0 is used

Model	n	MBPP↑	MBPP-ET↑	HumanEval↑	HumanEval-ET↑	Avg↑
LLama2-7B-Chat	0	27.8	22.9	12.9	11.2	18.7
	1	35.3(27.0%)	29.0(26.6%)	15.9(23.3%)	11.6(3.6%)	23.0(23.0%)
	2	37.5(34.9%)	29.5(28.8%)	17.1(32.6%)	12.2(8.9%)	24.1(28.9%)
	5	38.1(37.1%)	29.8(30.1%)	21.3(65.1%)	12.9(15.2%)	25.5(35.4%)
	0	73.1	67.8	79.5	71.0	72.9
CodeQwen1.5-7B-Chat	1	80.2(9.7%)	74.7(10.2%)	90.9(14.3%)	80.5(13.4%)	81.6(11.9%)
	2	84.4(15.5%)	75.0(10.6%)	92.1(15.8%)	80.5(13.4%)	83.0(13.9%)
	5	86.0(17.6%)	75.0(10.6%)	93.3(17.4%)	80.5(13.4%)	83.7(14.8%)
	0	70.0	67.4	81.6	73.4	73.1
	1	77.2(10.3%)	73.5(9.1%)	87.2(6.9%)	79.9(8.9%)	79.5(8.8%)
Kimi	2	78.7(12.4%)	74.5(10.5%)	87.2(6.9%)	79.3(8.4%)	79.9(9.3%)
	5	79.5(13.6%)	74.8(11.0%)	86.6(6.1%)	79.3(8.4%)	80.1(9.6%)
	0	81.2	74.4	85.7	81.0	80.6
GPT-4-Turbo	1	89.5(5.1%)	78.2(5.1%)	90.2(5.3%)	84.8(4.7%)	85.7(6.3%)
	2	89.5(5.1%)	79.9(7.4%)	90.8(6.0%)	84.8(4.7%)	86.3(7.1%)
	5	89.5(5.1%)	80.0(7.5%)	91.5(6.8%)	84.8(4.7%)	86.5(7.3%)
	0	84.9	78.2	92.5	85.8	85.4
	1	88.5(4.2%)	79.9(2.2%)	96.3(4.1%)	87.8(2.3%)	88.1(3.2%)
GPT-4o	2	89.6(5.5%)	79.2(1.3%)	97.0(4.9%)	89.0(3.7%)	88.7(3.9%)
	5	90.0(6.0%)	79.2(1.3%)	97.0(4.9%)	89.6(4.4%)	89.0(4.2%)

Table 1: ThinkCoder performance of various LLMs. $n = 0$ indicates results without the framework. $n > 0$ indicates results after n iterations of ThinkCoder, with each exploration generating $k = 5$ solutions for optimal refinement at temperature $t = 0.5$. Percentages in parentheses show improvement over the base LLMs.

to explore k times, generating $\{y^j\}_{j=1}^k$, where $y^j \sim \mathcal{M}_0(y|x)$. During exploration, a relatively large temperature value t is set for \mathcal{M}_0 . The generated outputs are then evaluated using the CodeVerifier, which computes $E(x, y^j) = (\text{pass_rate}, \text{feedback})$. For any sample where $\text{pass_rate} == 1$, the corresponding instance (x, y^j) is added to the training dataset D_{M_0} .

Reflection-based Generation. The primary purpose of generating R_{M_0} is to enhance the ability of error correction. For a given sample x and the feedback $E(x, y^j) = (\text{pass_rate}, \text{feedback})$ from the CodeVerifier, for samples with $\text{pass_rate} = 0$, we add the feedback to x and use the language model \mathcal{M}_0 to perform temperature-based exploration on the input $x' = (x, \text{feedback})$ and get the output y'^j . Then we also use the CodeVerifier to obtain the right correction answer with $E(x', y'^j) = (\text{pass_rate}', \text{feedback}')$. Finally, we add the instance (x', y'^j) with the $\text{pass_rate}' == 1$ to the reflection-based trajectories R_{M_0} .

4.2 Model Training and Inference

Generating high-probability text does not necessarily align well with human preferences, which may lead to biases in various tasks (Becker et al., 2024). Reinforced Self-Training (ReST) mitigates this is-

sue effectively by collecting offline data aligned with human preferences for online training (Gulcehre et al., 2023). We train the base model \mathcal{M} with the success trajectory dataset $D_{M_0} \cup R_{M_0}$:

$$\mathcal{L}_{MLE}(\theta_{\mathcal{M}}) = -\mathbb{E}_{(x,y) \sim D_{M_0} \cup R_{M_0}} \log p_{\theta_{\mathcal{M}}}(y | x)$$

In inference, we compare the performance with baselines at two configurations. The first uses a fine-tuned base model, feeding input x directly to generate text at temperature $t = 0$. This setup tests if ThinkCoder data improves output quality. The second configuration uses the fine-tuned base model as the exploration agent in ThinkCoder. It generates and refines solutions with k explorations and n refinements, using different temperature settings t . This aims to assess the framework’s ability to improve solutions through successful trajectories and the ThinkCoder framework.

5 Experiments

5.1 Dataset

We utilized the MBPP (Austin et al., 2021) and HumanEval (Chen et al., 2021) datasets to evaluate the effectiveness of our approach. To ensure a comprehensive assessment, we tested their enhanced versions—MBPP-ET and HumanEval-ET (Dong et al., 2023)—which include 80 and 35 times more

	Model	MBPP \uparrow	MBPP-ET \uparrow	HumanEval \uparrow	HumanEval-ET \uparrow	Avg \uparrow
GPT-3.5-Turbo	MapCoder	78.3	54.4	80.5	70.1	70.8
	AgentCoder	89.9	89.1	79.9	77.4	84.1
	ThinkCoder(Ours)	90.9	78.8	89.0	81.0	84.9
GPT-4-Turbo	MapCoder	83.1	57.7	93.9	82.9	79.4
	AgentCoder	91.4	91.4	89.6	76.2	87.2
	ThinkCoder(Ours)	90.5	80.5	93.9	86.0	87.7

Table 2: Pass@1 for MapCoder, AgentCoder, and ThinkCoder on MBPP, MBPP-ET, HumanEval, and HumanEval-ET benchmarks with GPT-3.5-Turbo and GPT-4-Turbo as the backbone.

	Model	MBPP \uparrow	MBPP-ET \uparrow	HumanEval \uparrow	HumanEval-ET \uparrow	CodeContests \uparrow	Avg \uparrow
GPT-4o	LDB	82.4	65.4	98.2	81.7	29.3	71.4
	LPW	84.8	65.8	98.2	84.8	34.7	73.7
	ThinkCoder(Ours)	90.0	79.2	97.0	89.6	40.1	79.2

Table 3: Pass@1 for LDB, LPW, and ThinkCoder on MBPP, MBPP-ET, HumanEval, HumanEval-ET, and CodeContests benchmarks with GPT-4o as the backbone.

test cases than the original datasets, respectively. Regarding the test data, the original problem set sizes for MBPP and HumanEval are 257 and 164, while their extended versions contain 378 and 164 problems, respectively. For training, we started with 374 MBPP training examples and 90 validation examples, combining them as the initial dataset. Using temperature-based and reflection-based generation methods, we expanded this set to over 2,000 examples for ReST training. We also evaluate our framework on two major benchmarks designed for harder problems: LiveCodeBench and CodeContests. LiveCodeBench consists of the 167 most recent samples collected between October 2024 and January 2025, while CodeContests includes 165 problems.

5.2 Baselines

In terms of LLM base model selection for the application of ThinkCoder, we considered a range of advanced instruction-tuned models, including the general-purpose models LLama2-7B-Chat, Kimi², GPT-4-Turbo(gpt-4-1106-preview) and GPT-4o(gpt-4o-2024-05-13), as well as the code-specific model CodeQwen1.5-7B-Chat. We conducted a comprehensive evaluation of ThinkCoder’s Pass@1 performance on MBPP and HumanEval across LLMs of different parameter sizes. For the test-time optimization process, we used the base model LLama2-7B to evaluate performance changes on the MBPP dataset. We compared results obtained before and after fine-tuning, using either the original training data or the success trajectories and incorporating or excluding

²<https://platform.moonshot.cn>

the ThinkCoder framework. For the composite model, we compared it with the current SOTA models, MapCoder (Islam et al., 2024), AgentCoder (Huang et al., 2023b), LDB (Zhong et al., 2024), and LPW (Lei et al., 2024).

Model	n	LiveCodeBench \uparrow	CodeContests \uparrow
GPT-4-Turbo	0	33.4	16.1
	1	43.7(30.8%)	26.1(62.1%)
	2	46.1(38.0%)	27.3(69.6%)
	5	44.9(34.4%)	28.5(77.0%)
	0	22.5	21.0
GPT-4o	1	35.3(56.9%)	34.1(62.4%)
	2	38.3(70.2%)	38.9(85.2%)
	5	41.3(83.6%)	40.1(91.0%)

Table 4: ThinkCoder performance on LiveCodeBench and CodeContests. n represents the self-refinement budget, with the self-exploration budget fixed at $k = 5$.

5.3 Implementation details

We set different parameter combinations for different experimental settings. To validate the improvement of the base LLMs, we set the exploration budget $k = 5$, the temperature received by the exploration agent $t = 0.5$, the number of optimal refinements $n = 5$, and the budget control threshold $\theta = 0.8$. The exploration agent generates $m = 3$ test cases per code. To compare ThinkCoder with SOTA results, we set the hyperparameters k , t , n , θ , and m to 20, 1.0, 2, 1.0, and 3, respectively. In the data collection phase, the exploration agent \mathcal{M}_0 is specified as CodeQwen1.5-7B-Chat because it is free and performs well. During ReST training, we set the learning rate to 2×10^{-5} and train for 2 epochs. We employ Low-Rank Adaptation




Figure 3: (a) The Pass@1 metric of the baseline models under different temperature t . (b) The Pass@k metric of the baseline models, where k represents the exploration budget. (c) The variation of the Pass@1 metric for each baseline model under the ThinkCoder framework as the optimal refinement budget n increases.

(LoRA) (Hu et al., 2021) to train LLMs.

6 Evaluation

Table 1 illustrates the performance changes of various base LLMs on the MBPP and HumanEval datasets as the Optimal Refinement iteration n increases. Table 2 and Table 3 compare the performance of ThinkCoder and SOTA models, using GPT-4-Turbo and GPT-4o as the base LLMs. Table 4 shows great potential on harder benchmarks. Table 5 compares ThinkCoder’s computation load with MapCoder at similar settings. Table 6 presents the performance changes of Llama2-7B before and after fine-tuning with trajectory data collected by ThinkCoder. Table 7 illustrates the ablation study of self-exploration and self-refinement. Figure 3 shows the impact of changes in the hyperparameter temperature t , exploration budget k , and refinement iterations n on the Pass@1 and Pass@k metric. Figure 4 depicts the trends in the accuracy of generated test cases and codes as the optimal refinement iterations in the ThinkCoder increase.

6.1 Performance on Code Generation

Table 1 shows that ThinkCoder achieves significant improvements in scenarios with fewer standard test cases, with performance on MBPP and HumanEval improving by 4%-70% after 5 iterations of refinement. In contrast, gains in more extensive scenarios, like MBPP-ET and HumanEval-ET, are more modest, ranging from 1%-35%. ThinkCoder’s generated test cases effectively verify code, reducing the need for manual tests, but its performance saturates in complex scenarios, making it better suited for simpler tasks. Table 2 further demonstrates the potential of ThinkCoder. When the number of exploration iterations k is extended to 20, ThinkCoder surpasses AgentCoder’s performance, which needs 5 refinement rounds, in just 2 iterations.

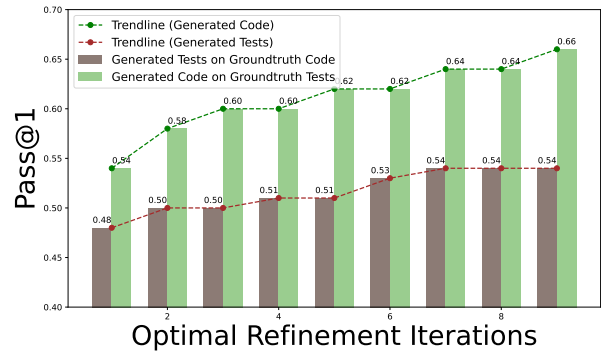


Figure 4: The variation in Pass@1 performance of the exploration code and test cases on the MBPP dataset under ThinkCoder, as evaluated by CodeQwen1.5-7B-Chat as the base LLM.

6.2 Ablation Study of Self-exploration and Self-refinement

Figures 3(b) and 3(c) highlight the performance gains from both self-exploration and self-refinement. The gain from self-exploration is greater than that from self-refinement. As illustrated in Table 7, self-refinement significantly enhances efficiency when paired with self-exploration. However, increasing the exploration budget excessively (e.g., $k = 20$) leads to inefficiency. A moderate increase in the exploration budget accelerates the optimization of the self-refinement process.

6.3 Performance with ReST Training

From the results presented in Table 6, it is clear that ReST fine-tuning on Llama2-7B using successful trajectories collected by ThinkCoder significantly enhances the model’s core capabilities. Firstly, Llama2-7B(ReST) demonstrates superior performance compared to both Llama2-7B(SFT) and Llama2-7B at iteration 0, highlighting that ReST training improves the exploration agent’s

Models	HumanEval				Mbpp				CodeContest			
	n	k	Pass@1↑	Tokens (k)↓	n	k	Pass@1↑	Tokens (k)↓	n	k	Pass@1↑	Tokens (k)↓
MapCoder	5	5	93.9	21.3	3	3	83.1	5.58	3	3	28.5	18.3
ThinkCoder(Ours)	20	1	93.9	0.26	5	1	89.5	0.24	5	5	28.5	2.94

Table 5: The token usage for agent responses and the required refinement iterations in ThinkCoder for HumanEval and MBPP, compared to MapCoder (as ThinkCoder and MapCoder share similar configurations), with GPT-4-Turbo as the base LLM. The tiktoken package is used to calculate the response token usage.

Models	Optimal Refinement Iterations					
	0	1	2	3	4	5
LLama2-7B	14.7	22.7	24.9	27.7	32.7	36.9
LLama2-7B(SFT)	27.8	35.3	37.5	37.6	38.0	38.1
LLama2-7B(ReST)	32.3	46.2	44.8	46.1	45.7	46.1

Table 6: Performance comparison of the LLama2-7B model under different settings: baseline without fine-tuning, supervised fine-tuning (SFT) with original train dataset, and Reinforced Self-Training (ReST) with success trajectories from ThinkCoder on the MBPP dataset.

ability to identify better solutions. Secondly, with a temperature setting of 0.5, the model achieves optimal performance with only a single round of refinement, underscoring the efficiency of ReST fine-tuning in reducing computational overhead.

Self-exploration Budget k	Self-refinement Budget n					
	1	2	3	5	10	20
1	91.7	90.0	91.7	93.3	93.3	95.0
2	93.3	95.0	95.0	96.7	96.7	96.7
5	93.3	93.3	95.0	95.0	96.7	96.7
10	96.7	96.7	96.7	96.7	96.7	96.7
20	96.7	96.7	96.7	96.7	96.7	96.7

Table 7: The performance trends on the randomly 60 Mbpp samples on GPT-4 with varying compositions of self-exploration and self-refinement budgets.

6.4 Evolution of Generated Code and Tests

After each round of optimal refinement, we evaluated the generated solutions and testing pools. Using the ground truth test set, we analyzed the performance changes of the solutions in each round. Meanwhile, the ground truth code was used to assess the changes in the testing pools. As shown in Figure 4, the accuracy of both the solutions and testing pools improved simultaneously. This demonstrates that ThinkCoder is capable of generating continuously optimized test cases while also refining the quality of the solutions.

6.5 Hyperparameters Study

We conducted hyperparameter selection experiments on the MBPP dataset. As shown in Figure 3(a), models exhibit optimal performance at different temperature settings. LLama2-7B-Chat, CodeQwen1.5-7B-Chat, and Kimi improve in Pass@1 as temperature increases, while GPT-4-Turbo and GPT-4o perform better at lower temperatures. Figure 3(b) indicates that Pass@k performance of base LLMs improves with increasing exploration budget k but eventually stabilizes. Figure 3(c) shows that Pass@1 performance improves with more refinement iterations n , but exploration gains are more significant. For details on hyperparameters θ and m , see Appendix A.2 and A.3.

6.6 Cost Analysis

In Table 5, we compare the computational overhead between MapCoder and ThinkCoder. Compared to MapCoder, which uses four LLM-based agents executed serially, ThinkCoder employs only one LLM-based agent and one non-LLM-based CodeVerifier, resulting in lower time overhead. Therefore, we use token overload as a measure of cost consumption. It reveals that ThinkCoder demonstrates more efficient resource utilization compared to MapCoder with a similar pass@1.

7 Conclusion

In this paper, we introduce ThinkCoder, a framework designed to improve the code generation efficiency of LLMs. The proposed approach incorporates LLM’s exploration with a non-LLM-based CodeVerifier that greatly reduces the token usage. To further minimize computational costs during exploration, we capture successful trajectories through ThinkCoder to fine-tune the LLM. This process enhances code generation performance while reducing refinement iterations for optimal solutions. We evaluate our framework across LLMs of varying sizes and demonstrate its superior performance over current composite models.

Limitations

In our current approach to managing the testing pool, we primarily focus on selective expansion but lack effective strategies for cleaning noise from the test set. Moreover, test set diversity relies heavily on LLMs' generation capabilities. Our work highlights the importance of tests and the gap between generated and annotated cases. We plan to enhance the testing pool in future work. Furthermore, while ReST training has demonstrated excellent results in enhancing the code generation capabilities of foundational LLMs, validation in ThinkCoder reveals rapid convergence. This indicates that fine-tuning teaches the LLM to find good answers, but the improvement in exploring the overall solution space remains limited. To address this, we look forward to introducing on-policy training methods in future work, combining data collection with fine-tuning to continuously evolve the model during training. This will generate more diverse answers and enhance the LLM's online exploration capabilities.

Acknowledgments

This work is also supported by the Public Computing Cloud, Renmin University of China and by fund for building worldclass universities (disciplines) of Renmin University of China.

References

- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- John W Backus, Friedrich L Bauer, Julien Green, Charles Katz, John McCarthy, Alan J Perlis, Heinz Rutishauser, Klaus Samelson, Bernard Vauquois, Joseph Henry Wegstein, et al. 1960. Report on the algorithmic language algol 60. *Communications of the ACM*, 3(5):299–311.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jonas Becker, Jan Philip Wahle, Bela Gipp, and Terry Ruas. 2024. Text generation: A systematic literature review of tasks, evaluation, and challenges. *arXiv preprint arXiv:2405.15604*.
- Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. *GitHub repository*.
- Angelica Chen, Jérémy Scheurer, Tomasz Korbak, Jon Ander Campos, Jun Shern Chan, Samuel R Bowman, Kyunghyun Cho, and Ethan Perez. 2023. Improving code generation by training with natural language feedback. *arXiv preprint arXiv:2303.16749*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Yihong Dong, Jiazheng Ding, Xue Jiang, Ge Li, Zhuo Li, and Zhi Jin. 2023. Codescore: Evaluating code generation by learning code execution. *arXiv preprint arXiv:2301.09043*.
- Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. 2024. Self-collaboration code generation via chatgpt. *ACM Transactions on Software Engineering and Methodology*, 33(7):1–38.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. 2023. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*.
- Qi Guo, Junming Cao, Xiaofei Xie, Shangqing Liu, Xiaohong Li, Bihuan Chen, and Xin Peng. 2024. Exploring the potential of chatgpt in automated code refinement: An empirical study. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pages 1–13.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Dong Huang, Qingwen Bu, and Heming Cui. 2023a. Codecot and beyond: Learning to program and test like a developer. *arXiv preprint arXiv:2308.08784*.
- Dong Huang, Qingwen Bu, Jie M Zhang, Michael Luck, and Heming Cui. 2023b. Agentcoder: Multi-agent-based code generation with iterative testing and optimisation. *arXiv preprint arXiv:2312.13010*.
- Md Ashraf Islam, Mohammed Eunus Ali, and Md Rizwan Parvez. 2024. Mapcoder: Multi-agent code generation for competitive problem solving. *arXiv preprint arXiv:2405.11403*.
- Sajed Jalil, Suzzana Rafi, Thomas D LaToza, Kevin Moran, and Wing Lam. 2023. Chatgpt and software testing education: Promises & perils. In *2023 IEEE international conference on software testing, verification and validation workshops (ICSTW)*, pages 4130–4137. IEEE.

- Xue Jiang, Yihong Dong, Lecheng Wang, Zheng Fang, Qiwei Shang, Ge Li, Zhi Jin, and Wenpin Jiao. 2024. Self-planning code generation with large language models. *ACM Transactions on Software Engineering and Methodology*, 33(7):1–30.
- Bryan F Jones, HH Sthamer X Yang, and DE Eyres. 2024. The automatic generation of software test data sets using adaptive search techniques. *WIT Transactions on Information and Communication Technologies*, 14.
- Chao Lei, Yanchuan Chang, Nir Lipovetzky, and Krista A Ehinger. 2024. Planning-driven programming: A large language model programming workflow. *arXiv preprint arXiv:2411.14503*.
- Jia Li, Ge Li, Yongmin Li, and Zhi Jin. 2024a. Structured chain-of-thought prompting for code generation. *ACM Transactions on Software Engineering and Methodology*.
- Qingyao Li, Wei Xia, Kounianhua Du, Xinyi Dai, Ruiming Tang, Yasheng Wang, Yong Yu, and Weinan Zhang. 2024b. Rethinkmcts: Refining erroneous thoughts in monte carlo tree search for code generation. *arXiv preprint arXiv:2409.09584*.
- Tsz-On Li, Wenxi Zong, Yibo Wang, Haoye Tian, Ying Wang, Shing-Chi Cheung, and Jeff Kramer. 2023. Finding failure-inducing test cases with chatgpt. *arXiv preprint arXiv:2304.11686*.
- Yue Liu, Thanh Le-Cong, Ratnadira Widayarsi, Chakkrit Tantithamthavorn, Li Li, Xuan-Bach D Le, and David Lo. 2024a. Refining chatgpt-generated code: Characterizing and mitigating code quality issues. *ACM Transactions on Software Engineering and Methodology*, 33(5):1–26.
- Yuhan Liu, Xiuying Chen, Xiaoqing Zhang, Xing Gao, Ji Zhang, and Rui Yan. 2024b. From skepticism to acceptance: simulating the attitude dynamics toward fake news. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 7886–7894.
- Yuhan Liu, Zirui Song, Xiaoqing Zhang, Xiuying Chen, and Rui Yan. 2024c. From a tiny slip to a giant leap: An llm-based simulation for fake news evolution. *arXiv preprint arXiv:2410.19064*.
- Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro Von Werra, and Shayne Longpre. 2023. Octopack: Instruction tuning code large language models. *arXiv preprint arXiv:2308.07124*.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. Toollm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.
- Tal Ridnik, Dedy Kredo, and Itamar Friedman. 2024. Code generation with alphacodium: From prompt engineering to flow engineering. *arXiv preprint arXiv:2401.08500*.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*.
- Junjie Wang, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. 2024. Software testing with large language models: Survey, landscape, and vision. *IEEE Transactions on Software Engineering*.
- Huan Zhang, Wei Cheng, Yuhan Wu, and Wei Hu. 2024. A pair programming framework for code generation via multi-plan exploration and feedback-driven refinement. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, pages 1319–1331.
- Juntian Zhang, Yuhan Liu, Wei Liu, Jian Luan, Rui Yan, et al. 2025. Weaving context across images: Improving vision-language models through focus-centric visual chains. *arXiv preprint arXiv:2504.20199*.
- Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhui Chen, and Xiang Yue. 2024. Opencodeinterpreter: Integrating code generation with execution and refinement. *arXiv preprint arXiv:2402.14658*.
- Zibin Zheng, Kaiwen Ning, Yanlin Wang, Jingwen Zhang, Dewu Zheng, Mingxi Ye, and Jiachi Chen. 2023. A survey of large language models for code: Evolution, benchmarking, and future trends. *arXiv preprint arXiv:2311.10372*.
- Li Zhong, Zilong Wang, and Jingbo Shang. 2024. Ldb: A large language model debugger via verifying runtime execution step-by-step. *arXiv preprint arXiv:2402.16906*.

A Appendix

A.1 Algorithm of ThinkCoder

Algorithm 1 shows the pseudo-code of ThinkCoder.

A.2 Exploration Budget Control




Figure 5: The relationship between the computational cost and performance of ThinkCoder at different budget control thresholds, where ‘Pass@1’ indicates the performance on the MBPP dataset, and ‘ θ ’ refers to the budget control threshold that allows the task to execute the next exploration process. ‘Exploration Budget’ represents the ratio of total requests during each iteration.

Figure 5 outlines the number of tasks for the exploration agent per iteration and the overall Pass@1 metric under different budget control thresholds. By increasing the exploration budget control threshold θ , the number of tasks requiring correction in each iteration is significantly reduced, leading to a notable decrease in ThinkCoder’s runtime. However, this approach risks leaving some incorrectly validated tasks uncorrected, which could negatively impact performance. In our experiments, we set $\theta = 0.8$, as this value not only progressively improves Pass@1 but also reduces the number of tasks per iteration by nearly half.

A.3 Exploration for Tests

Metrics	Tests m /Exploration			
	1	3	5	10
Pass@1	80.0	81.7	81.7	86.7
Overload(s)	28.7	86.0	112.3	364.5

Table 8: The variation in the Pass@1 metric and the CodeVerifier’s overload is influenced by changes in the number of test cases m during each exploration.

Algorithm 1 ThinkCoder

- 1: Initialize the testing pool $TP=[]$, set the globally max pass rate $g_p = 0$ and best solution $g_s = \text{None}$
- 2: Set the hyperparameters $t = 0.5, k = 5, n = 5, m = 3, \theta = 0.8$
- 3: Set the instruction prompt for code generation task and instruction for task p as $I = [p, \text{prompt}]$
- 4: **for each** $iter \in n$ **do**
- 5: Initialize the locally max pass rate $l_p = 0$ and best solution $l_s = \text{None}$ for iteration $iter$
- 6: Randomly select m successful test cases from TP as tests
- 7: Select a new type of error feedback for test f_t from TP as f
- 8: **if** f_t and f is not None **then**
- 9: Update the instruction $I = [p, g_s, f_t, f, \text{tests}, \text{prompt}]$
- 10: **end if**
- 11: Get code solutions G' and test cases TP' from exploration agent for the instruction I within k iterations
- 12: Deduplicate testing pool TP' by AST
- 13: **for each** $i \in |G'|$ **do**
- 14: Get the pass rate r_{g_i} for solution g_i from CodeVerifier with the input TP'
- 15: **if** $r_{g_i} > l_p$ **then**
- 16: Set $l_p = r_{g_i}, l_s = g_i$
- 17: **end if**
- 18: **end for**
- 19: **if** $l_p > g_p$ **then**
- 20: Set $g_p = l_p, g_s = l_s$
- 21: Add TP' to TP
- 22: **end if**
- 23: **if** $g_p > \theta$ **then**
- 24: return g_s as the best solution
- 25: Terminate optimal refinement iteration
- 26: **end if**
- 27: **end for**

We randomly selected 60 tasks from the MBPP test data to observe the performance improvement and code execution overhead as the number of test cases m increases during each exploration. Table 8 illustrates this trend. The base model we used is Kimi. We observed that increasing m leads to improved Pass@1 performance and results in longer execution times due to the additional effort required for self-verification. We set m for all experiments to balance time efficiency and performance.

A.4 Prompts for Exploration Agent

Figure 6 illustrates the input received by the exploration agent, including the problem, the best solution, the test cases from the ground truth, the generated test cases from testing pooling, and feedback on test failures.

A.5 Inputs for CodeVerifier

Figure 7 shows the input of the executable Python environment, including the code and all test cases from the testing pool.

Exploration Prompt

Below is an instruction that describes a task. Write a response that appropriately completes the request.

Instruction:

You are an expert Python programmer. Here is your task:

```
def find_adverb_position(text):  
    """Write a function to find the first adverb and their positions in a given sentence.
```

Your previous solution for the task is as follows:

```
def find_adverb_position(text):  
    """ Write a function to find the first adverb and their positions in a given sentence.  
    """  
    import re  
    # Define a regex pattern to match adverbs  
    adverb_pattern =  
r"\b(absolutely|actually|basically|clearly|definitely|easily|eventually|generally|hardly|just|literally|loudly|maybe|obviously|possibly|probably|q  
uickly|really|seriously|suddenly|truly|unfortunately|usually|very|well)\b"  
    # Find all matches of adverbs  
    matches = re.finditer(adverb_pattern, text, re.IGNORECASE)  
    # Get the first match if it exists  
    first_match = next(matches, None)  
    if first_match:  
        return (first_match.start(), first_match.end(), first_match.group())  
    else:  
        return None
```

Here is your generated test case: "assert find_adverb_position(long_text) == (0, 7, 'quickly')", it is wrong with the feedback "name 'long_text' is not defined".

Refine your code and make it correct.

Your code should also pass these tests:

```
def check(candidate):  
    assert candidate("clearly!! we can see the sky")==(0, 7, 'clearly')  
    assert candidate("seriously!! there are many roses")==(0, 9, 'seriously')  
    assert candidate("unfortunately!! sita is going to home")==(0, 13, 'unfortunately')
```

```
def test_check():  
    check(find_adverb_position)
```

```
test_check()
```

Your code for the task should start with a [PYTHON] tag and end with a [/PYTHON] tag. Write another 3 test cases for the code. The test cases should include a variety of tests related to regular testing, boundary testing and stress testing. Test cases are directly executable assert statements, without any comments. And your whole test cases should start with a [TEST] tag and end with a [/TEST] tag.

Figure 6: ThinkCoder exploration agent prompt example.

A.6 Example of Thorough Exploration

Figure 8 illustrates multiple solutions explored by the exploration agent.

A.7 Example of Optimal Refinement

To better understand the process of optimal refinement, Figure 9 illustrates a single refinement iteration. Before refinement, we leveraged the best solution output from the previous iteration, along with failed tests and unmet feedback sampled from the testing pool. Using this reflection information, the solution was updated and revalidated on the test cases in the testing pool. The success rate improved from 0.217 to 0.356, leading us to select the current solution as the best solution.

A.8 Example of Testing Pool

Figure 10 illustrates the test cases in the testing pool generated through multiple explorations by ThinkCoder.

Execution Prompt

```
from typing import *

def find_adverb_position(text):
    """ Write a function to find the first adverb and their positions in a given sentence.
    """
    import re
    # Define a regex pattern to match adverbs
    adverb_pattern =
r'\b(absolutely|actually|basically|clearly|definitely|easily|eventually|generally|hardly|just|literally|loudly|maybe|obviously|possibly|probably|q
uickly|really|seriously|suddenly|truly|unfortunately|usually|very|well)\b'
    # Find all matches of adverbs
    matches = re.finditer(adverb_pattern, text, re.IGNORECASE)
    # Get the first match if it exists
    first_match = next(matches, None)
    if first_match:
        return (first_match.start(), first_match.end(), first_match.group())
    else:
        return None

def check(candidate):
    assert candidate("clearly!! we can see the sky")===(0, 7, 'clearly')
    assert candidate("seriously!! there are many roses")===(0, 9, 'seriously')
    assert candidate("unfortunately!! sita is going to home")===(0, 13, 'unfortunately')

def test_check():
    check(find_adverb_position)

test_check()
```

Figure 7: ThinkCoder CodeVerifier input example.

Exploration Examples

```
# 1
def find_adverb_position(text):\n """\n Write a function to find the first adverb and their positions in a given sentence.\n """\n
import re\n adverbs = re.findall(r'\b(\w+)\b', text)\n if adverbs:\n first_adverb = adverbs[0]\n pos =
text.find(first_adverb)\n return (pos, pos + len(first_adverb), first_adverb)\n else:\n return None

# 2
def find_adverb_position(text):\n """\n Write a function to find the first adverb and their positions in a given sentence.\n """\n
import re\n # Define a regex pattern for adverbs\n adverb_pattern =
re.compile(r'\b(very|well|quickly|slowly|carefully|loudly|softly|quickly|briefly|slowly|clearly|seriously|unfortunately)\b',
re.IGNORECASE)\n # Find the first match\n match = adverb_pattern.search(text)\n if match:\n # Return the start position, end
position, and the matched adverb\n return match.start(), match.end(), match.group()\n else:\n # Return None if no adverb is
found\n return None

# 3
def find_adverb_position(text):\n """\n Write a function to find the first adverb and their positions in a given sentence.\n """\n
import re\n # Regular expression pattern to match adverbs\n adverb_pattern =
r'\b(?:<|w)(very|well|just|already|really|quite|hardly|barely|merely|rarely|nearly|only|often|sometimes|always|never|ever|already|yet|then|t
here|here|where|when|how|why|so|too|either|neither|also|like|until|before|after|since|besides|despite|together|apart|around|between|throu
gh|across|along|over|under|inside|outside|onto|into|throughout|despite|regardless|despite)(?!w)\b'\n match =
re.search(adverb_pattern, text, re.IGNORECASE)\n if match:\n return match.start(), match.end(), match.group()\n else:\n
return None, None, None

# 4
def find_adverb_position(text):\n """\n Write a function to find the first adverb and their positions in a given sentence.\n """\n
import re\n # Regular expression pattern for finding adverbs\n adverb_pattern =
r'\b(very|well|slowly|quickly|happily|sadly|often|sometimes|always|never|hardly|clearly|seriously|unfortunately)\b'\n \n # Search for
the first adverb in the text\n match = re.search(adverb_pattern, text)\n if match:\n # Return the start index, end index, and the
adverb itself\n return match.start(), match.end(), match.group()\n else:\n # Return None if no adverb is found\n return None

# 5
def find_adverb_position(text):\n """\n Write a function to find the first adverb and their positions in a given sentence.\n """\n #
Define a list of common adverbs\n adverbs = ['always', 'already', 'also', 'ever', 'even', 'hardly', 'just', 'little', 'much',
'never', 'only', 'quite', 'really', 'so', 'such', 'too', 'very', 'already', 'also', 'anyhow', 'anywhere', 'barely', 'beneath',
'beside', 'besides', 'beyond', 'earlier', 'elsewhere', 'even', 'ever', 'everywhere', 'far', 'further', 'hardly', 'here', 'just',
'later', 'less', 'little', 'lower', 'many', 'more', 'most', 'nearly', 'never', 'nowhere', 'once', 'only', 'other', 'otherwise',
'over', 'presently', 'quickly', 'quietly', 'quite', 'rather', 'really', 'seemingly', 'seriously', 'several', 'slightly', 'so',
'somehow', 'somewhere', 'still', 'such', 'terribly', 'than', 'that', 'the', 'there', 'these', 'this', 'those', 'through', 'thus',
'today', 'together', 'too', 'toward', 'under', 'unfortunately', 'until', 'up', 'upon', 'very', 'well', 'where', 'whereas',
'wherever', 'whether', 'which', 'whichever', 'while', 'whoever', 'whole', 'wholly', 'why', 'willingly', 'with', 'within',
'without', 'wonderfully', 'would']\n # Iterate through the sentence to find the first adverb\n for i, word in enumerate(text.split()):\n
if word.lower() in adverbs:\n return i, i + len(word), word\n return None, None, None
```

Figure 8: ThinkCoder exploration examples.

Refinement Examples

```
# Problem
\`\`\`Write a function to find the first adverb and their positions in a given sentence.

# Iteration 0
# Best Solution
def find_adverb_position(text):\n \`\`\`\n Write a function to find the first adverb and their positions in a given sentence.\n \`\`\`\n
import re\n # Define a regex pattern to match adverbs\n adverb_pattern =
r'\b(absolutely|actually|basically|definitely|easily|eventually|generally|hardly|just|literally|maybe|obviously|possibly|probably|quickly|really|
suddenly|truly|usually|very|well)\b'\n # Find all matches of adverbs\n matches = re.finditer(adverb_pattern, text, re.IGNORECASE)\n
# Get the first match if it exists\n first_match = next(matches, None)\n if first_match:\n return (first_match.start(),
first_match.end(), first_match.group())\n else:\n return None

# Test from Testing Pool
assert find_adverb_position("loudly!! the crowd cheered") == (0, 0, 'loudly')

# Feedback
It is wrong with the output "None"

# Pass_Rate
[0.217]

# Iteration 1
# Best Solution
def find_adverb_position(text):\n \`\`\`\n Write a function to find the first adverb and their positions in a given sentence.\n \`\`\`\n
import re\n # Define a regex pattern to match adverbs\n adverb_pattern =
r'\b(absolutely|actually|basically|clearly|definitely|easily|eventually|generally|hardly|just|literally|loudly|maybe|obviously|possibly|probably|
quickly|really|seriously|suddenly|truly|unfortunately|usually|very|well)\b'\n # Find all matches of adverbs\n matches =
re.finditer(adverb_pattern, text, re.IGNORECASE)\n # Get the first match if it exists\n first_match = next(matches, None)\n if
first_match:\n return (first_match.start(), first_match.end(), first_match.group())\n else:\n return None

# Test from Testing Pool
assert find_adverb_position("loudly!! the crowd cheered") == (0, 0, 'loudly!')

# Feedback
It is wrong with the output "not all arguments converted during string formatting".

# Pass_Rate
[0.356]
```

Figure 9: ThinkCoder refinement examples.

Testing Pool

```
assert find_adverb_position("A very long sentence with no adverbs") == (None, None, None)
assert find_adverb_position("The cat is extremely lazy.") == (8, 17, 'extremely')
assert find_adverb_position("clearly!! we can see the sky") == (0, 7, 'clearly')
assert find_adverb_position("An adverb might be missing.") == None
assert find_adverb_position(", indeed! it's quite a day") == (2, 7, 'indeed')
assert find_adverb_position("The cat sat quietly.") == (12, 19, 'quietly')
assert find_adverb_position("An adverb like 'very' is not found here.") == (10, 15, 'very')
assert find_adverb_position("The dog very quickly jumped over the fence.") == (8, 14, 'very')
assert find_adverb_position("Never quickly and always") == (0, 4, 'never')
assert find_adverb_position("Never forget this") == (0, 4, 'never')
assert find_adverb_position("always quickly she was running") == (0, 5, 'always')
assert find_adverb_position("The adverb 'far' is at the beginning.") == (4, 7, 'far')
assert find_adverb_position("Rarely does he visit.") == (4, 9, 'Rarely')
assert find_adverb_position("No adverbs here!") == (None, None, None)
assert find_adverb_position("A sentence without an adverb") == (None, None, None)
assert find_adverb_position("seriously!! there are many roses") == (0, 9, 'seriously')
assert find_adverb_position("An adverb 'already' is here, but not at the start") == (3, 10, 'already')
assert find_adverb_position("The cat is always sleeping.") == (8, 14, 'always')
assert find_adverb_position("quickly the fox jumps over the lazy dog") == (0, 7, 'quickly')
assert find_adverb_position(long_text) == (0, 7, 'quickly')
assert find_adverb_position("No adverbs here!") == None
assert find_adverb_position("He quickly ran to the store.") == (3, 9, 'quickly')
assert find_adverb_position("unfortunately!! sita is going to home") == (0, 13, 'unfortunately')
assert find_adverb_position("quickly, she ran away") == (0, 7, 'quickly')
assert find_adverb_position("The cat 123 quickly 456 jumped over 789 the lazy dog.") == (13, 20, 'quickly')
assert find_adverb_position("Rapidly, she walked to the store") == (0, 7, 'rapidly')
assert find_adverb_position("she was always on time") == (8, 13, 'always')
assert find_adverb_position("amazingly, the cat quickly jumped over the lazy dog") == (0, 8, 'amazingly')
```

Figure 10: ThinkCoder testing pool examples.