

Towards Better Correctness and Efficiency in Code Generation

Yunlong Feng, Yang Xu, Xiao Xu, Binyuan Hui*, Junyang Lin*

Qwen Team, Alibaba Group
{yunlong.fyl,binyuan.hby,junyang.ljy}@alibaba-inc.com

Abstract

While code large language models have demonstrated remarkable progress in code generation, the generated code often exhibits poor runtime efficiency, limiting its practical application in performance-sensitive scenarios. To address this limitation, we propose an efficiency-oriented reinforcement learning framework guided by a novel performance reward. Based on this framework, we take a deeper dive into the code efficiency problem, identifying then proposing methods to overcome key bottlenecks: (1) Dynamic exploration overcomes the static data constraints of offline fine-tuning, enabling the discovery of more efficient code implementations. (2) The error-insensitive reinforcement learning method and high-contrast efficiency signals are crucial for mitigating systematic errors and achieving effective optimization. (3) Online exploration is most effective when starting from a high-correctness baseline, as this allows for efficiency improvements without sacrificing accuracy. With these discoveries, we finally propose a two-stage tuning method, which achieves high and balanced performance across correctness and efficiency. The results of experiments show the effectiveness of the method, which improves code correctness by 10.18% and runtime efficiency by 7.75% on a 7B model, achieving performance comparable to much larger model.

Extended version — <https://arxiv.org/pdf/2508.20124>

1 Introduction

With the rapid development of large language models (LLMs) (Brown et al. 2020; Achiam et al. 2023; Touvron et al. 2023; Dubey et al. 2024; Jiang et al. 2023; Bai et al. 2023; Yang et al. 2024a; Anthropic 2024; OpenAI 2024), code large language models have improved their capabilities in various code-related tasks. However, recent studies have indicated that these large models typically focus on the correctness of the generated code while neglecting its runtime efficiency (Waghjale et al. 2024; Huang et al. 2024b; Du et al. 2024; Liu et al. 2024). As a result, the generated code can be 3 to 13 times slower than human-written versions, limiting the practical application of model-generated code in real-world software development workflows (Huang et al. 2024b).

To this end, recent work has advanced along several key fronts, underpinned by the development of dedicated benchmarks like EvalPerf (Liu et al. 2024), Mercury (Du et al. 2024), and EffiBench (Huang et al. 2024b). Building on these evaluation frameworks, researchers have designed structured optimization workflows, such as the iterative, feedback-driven process in SOAP (Huang et al. 2024a) and the two-stage approach of LLM4EFFI (Ye et al. 2025). In parallel, a major research thrust focuses on directly augmenting model capabilities through fine-tuning and reinforcement learning (RL). This includes methods like Effi-Learner (Huang et al. 2024c) and PIE (Shypula et al. 2024), which leverage curated code samples, as well as approaches like Mercury’s (Du et al. 2024) validation of Direct Preference Optimization (DPO) and ACECode’s (Yang et al. 2024b) use of rule-based PPO to further improve code efficiency during generation.

To better understand the underlying challenges, we conduct a series of explorations and identify the following key insights: (1) Offline fine-tuning methods are bottlenecked by static datasets, as their performance is capped by the implementations present in the training data, limiting the ability to discover novel, more efficient algorithmic solutions. (2) While online methods like Reinforcement Learning (RL) offer a path to dynamic exploration, they introduce instabilities, as the inherently noisy and error-prone efficiency signal can prevent the model from learning effectively. This presents the challenge of designing an error-insensitive RL framework with high-contrast efficiency signals. (3) Our findings indicate that the optimization’s starting point is critical, and that a model with high initial correctness serves as a superior foundation for any subsequent efficiency optimization.

Based on these findings, we design a practical two-stage training strategy, which is structured to first establish a robust foundation of code correctness before optimizing for runtime efficiency. In the first phase, we focus on establishing foundational correctness by fine-tuning a base LLM, which can be done with normal offline methods. This creates a high-accuracy model to serve as a solid starting point. In the second phase, we build upon this foundation to enhance efficiency. We employ the error-insensitive online reinforcement learning algorithm that uses high-contrast efficiency signals to improve runtime performance while maintaining the high level of correctness achieved in the first stage.

Our main contributions are as follows:

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

- We first analyze the limitations of offline methods, revealing that they are bottlenecked by initial data quality and struggle to enhance efficiency without sacrificing the model’s correctness.
- We then investigate techniques for the online RL phase. We find that selecting an error-insensitive algorithm (RLOO) is crucial for preserving correctness, and we introduce a stabilized reward signal and high-contrast inputs to ensure a robust and effective optimization process.
- Based on these insights, we propose a two-stage tuning strategy, which achieves high and balanced performance across correctness and efficiency. The results show the effectiveness of this method, which improves code correctness by 10.18% and runtime efficiency by 7.75% on a 7B model, achieving performance comparable to much larger model.

2 Efficiency-oriented Reinforcement Learning

It’s hard to evaluate the runtime efficiency of code. Manually creating a reference version is labor-intensive and costly, and even then, the notion of “fast” code is inherently relative, as identifying a definitively optimal implementation is often intractable. This means even a carefully hand-tuned version cannot be guaranteed to be optimal (Yang et al. 2024b). Consequently, relying on a single, manually-optimized version as a performance baseline is inadequate.

To facilitate continuous model improvement, we therefore adopt a group comparison methodology for performance assessment. As our performance metric, we use CPU instruction counts rather than wall-clock time to mitigate sources of non-determinism, such as operating system scheduling (Liu et al. 2024). Besides, we introduce inputs with high complexity differentiation to overcome the insensitivity of trivial inputs to algorithmic efficiency differences during execution. However, the distribution of these raw counts is typically highly skewed and exhibits large variance, which can impede stable model training. To address this, we apply a logarithmic transformation to the counts to compress their value range and stabilize the variance (Enter 2008).

2.1 Reward Function

Our objective is to ensure that correct and faster code results in higher scores, with improvement on efficiency consistently contributing to the reward. The reward should be benchmarked against the efficiency of the code generated by the current model for the same query.

To formalize this, we first define our parameters.

- $C = -\log(N)$: The performance metric for the code being evaluated, where N is the number of CPU instructions. Due to the negative sign, a higher value of C indicates faster execution.
- $G = [C_1, C_2, \dots, C_m]$: An array of performance metrics for all correct codes generated for the same query.
- $\max(G) / \min(G)$: The maximum / minimum value in G , representing the most efficient (fastest) / least efficient (slowest) code among all codes generated for the query.

Based on these parameters, we designed the reward function \mathcal{R} . If the code passes all tests, it receives a score calculated by $\mathcal{F}(C, G)$. This function maps the performance metric of the code to a reward where the slowest passing code gets a score of 1.0 and the fastest gets a score up to 2.0. Otherwise, it receives a penalty based on the error type.

This idea is formulated as follows:

$$S_{min} = 1.0$$

$$S_{max} = \min(2.0, S_{min} + \max(G) - \min(G))$$

$$\mathcal{F}(C, G) = S_{min} + (C - \min(G)) \cdot \frac{S_{max} - S_{min}}{\max(G) - \min(G)}$$

The final reward $\mathcal{R}(C, G)$ is assigned based on the execution outcome, which connects the performance score with the error classification:

- **Pass All Tests:** The code receives the calculated performance score, $\mathcal{R} = \mathcal{F}(C, G)$.
- **Test Error:** A score of 0.0 is assigned if any test fails.
- **No Entity Error:** A penalty of -0.5 is applied if the code cannot locate the corresponding test function.
- **Format Error:** A penalty of -1.5 is assigned if the response does not follow the specified format.

This is summarized by the complete reward function:

$$\mathcal{R}(C, G) = \begin{cases} \mathcal{F}(C, G) & \text{if Pass All Tests,} \\ 0.0 & \text{for a Test Error,} \\ -0.5 & \text{for a No Entity Error,} \\ -1.5 & \text{for a Format Error.} \end{cases}$$

3 Identify the Bottlenecks in Code Efficiency

This chapter aims to conduct a series of preliminary experiments to deeply analyze the inherent characteristics and challenges of existing mainstream methods in code optimization tasks. By understanding their strengths and limitations, we aim to achieve a better balance between correctness and efficiency in code generation.

3.1 Experiments

To contextualize our analysis, we compare several common fine-tuning methods across two standard benchmarks. These methods can be broadly categorized into two groups: offline methods, which learn from a static, precollected dataset, and online methods, which dynamically generate data during the training process.

Offline Methods

- **SFT:** A standard approach where the model is fine-tuned exclusively on the fastest correct code solution available for each problem in the training set.
- **DPO (Rafailov et al. 2023):** For each problem, we construct correctness pairs (correct vs. incorrect code) and efficiency pairs (faster vs. slower correct code).

Online Methods

- **GRPO** (Shao et al. 2024): A method that optimizes the policy based on the relative performance ranking within a group of generated solutions.
- **RLOO** (Ahmadian et al. 2024): A reinforcement learning technique that uses a leave-one-out strategy to estimate the reward baseline.

Datasets To comprehensively evaluate the model’s code optimization capabilities, we utilize several datasets for training and benchmarking. The Table 1 below summarizes the datasets used for training and evaluation.

Dataset	Source	Metric	Num (#)	Pass@1
Train Dataset				
EffiBench	LeetCode	-	790	62.99
Effi-Learner	Synthetic	-	2,978	81.39
Benchmark				
Mercury	LeetCode	Beyond	256	-
Evalperf	Evalplus	DPS _{norm}	109	-

Table 1: The subset of the training dataset. Notably, the queries from Effi-Learner are simpler, resembling the difficulty of Evalplus problems, while EffiBench’s queries are more akin to Mercury’s, since they all originate from LeetCode. The passrate is based on the Qwen-Coder-7B-Instruct.

The metrics DPS_{norm} and Beyond are used to represent the model’s performance in terms of code efficiency. A higher value for these metrics indicates better efficiency. For a more detailed explanation, please refer to the appendix of Extended version.

3.2 Results of Preliminary Experiment

Method	EvalPerf		Mercury	
	DPS _{norm}	Pass@1	Beyond	Pass@1
Untuned	77.60	77.50	74.80	77.20
SFT	77.85	79.19	66.36	83.20
DPO	87.61	81.89	72.23	85.47
GRPO	83.96	81.39	79.15	87.50
RLOO	81.97	82.46	76.00	89.92

Table 2: The results of different methods on Qwen-2.5-Coder-Instruct-7B model.

We evaluated the different methods on the Qwen-2.5-Coder-Instruct-7B model, with the results presented in Table 2. Below, we analyze the performance of each method to identify the challenges in the scene and guide future research.

3.3 Offline Methods Bottlenecked by Initial Data

Offline methods, such as Supervised Fine-Tuning (SFT) and Direct Preference Optimization (DPO), are common strategies for improving a model’s performance. These approaches typically rely on generating a pool of candidate solutions

from a base model and then selectively using them for fine-tuning. However, their ultimate success is fundamentally constrained by the quality of this initial sample pool. If the base model struggles to generate solutions that are both correct and efficient for certain problems, the offline tuning process is inherently limited, as it lacks high-quality data to learn from. In the following analysis, we demonstrate how this sampling bottleneck creates a crucial trade-off between correctness and efficiency, especially when tackling complex coding tasks.

Limited by Inefficient Samples from Challenging Problems Our experiments first explored fine-tuning using only the fastest correct code generated by the base model. This method boosted correctness but at the cost of significantly reduced runtime efficiency on the more challenging Mercury benchmark.

A plausible explanation arises from the data in Table 1 and Table 2. The base model’s correctness is much lower on the difficult LeetCode problems in EffiBench (62.99%) than on the simpler synthetic problems in Effi-Learner (81.39%). This indicates that for complex problems, the few correct solutions the model manages to generate are often naive, brute-force implementations. While functionally correct, they are computationally inefficient. Consequently, when we use these difficult samples to create the fine-tuning data, we inadvertently select for these “correct-but-slow” solutions, teaching the model to replicate inefficient code.

DPO Hits an Efficiency Ceiling As the Table 2 and Table 3 shows, DPO makes more effective use of negative samples than SFT, leading to significant improvements in code correctness without compromising efficiency. However, on more challenging problems, such as those in the Mercury benchmark, DPO faces a similar limitation to SFT. Both methods are constrained by the base model’s initial capabilities, making it difficult to sample solutions that are both correct and highly efficient. This inherent limitation hinders DPO’s ability to further enhance code efficiency for complex tasks.

Proportion	EvalPerf		Mercury	
	DPS _{norm}	Pass@1	Beyond	Pass@1
0.0	84.60	87.20	71.45	87.11
0.5	87.61	81.89	72.23	85.47
1.0	84.80	74.70	72.21	77.30

Table 3: Results of DPO with different settings on Qwen-2.5-Coder-Instruct-7B. The Proportion represents the proportion of efficiency pairs in the training data.

3.4 Data Composition Dictates the Trade-off Between Correctness and Efficiency

While considering execution efficiency, code correctness must also be taken into account. Consequently, we seek to understand how the proportion of “efficiency pairs” within the training data influences the model’s performance.

As the Figure 1 shows, there is a core trade-off in model performance between code correctness (Y-axis) and runtime

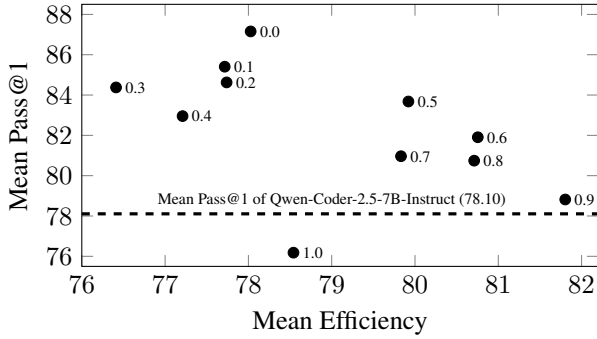


Figure 1: Impact of the proportion of efficiency pairs in training data on code correctness and efficiency in DPO training.

efficiency (X-axis), which can be tuned by varying the proportion of “efficiency pairs.”

- As the proportion of efficiency pairs increases from 0.0 to 0.9, runtime efficiency steadily improves while code correctness declines, demonstrating that this ratio can be used to steer the model’s focus.
- Notably, at a 1.0 proportion (using only efficiency pairs), performance deteriorates on both metrics, with correctness falling below the baseline. This indicates that solely pursuing efficiency at the complete expense of correctness leads to an overall degradation in performance.

From these results, we conclude that a significant emphasis on efficiency ($p < 0.9$) can be applied without substantially degrading model correctness, as this principle fails only under extreme conditions, for example, when efficiency is exclusively pursued at $p = 1.0$. Consequently, we conjecture that within an online reinforcement learning, the dynamic reward mechanism can adeptly manage this balance, thereby improving efficiency without compromising model correctness.

3.5 Online RL is Constrained by Initialization, Reward, and Algorithm

High Complexity	EvalPerf		Mercury	
	DPS _{norm}	Pass@1	Beyond	Pass@1
From Qwen-2.5-Coder-7B				
Y	90.81	78.06	80.59	88.75
N	81.94	80.85	76.90	90.70
From Qwen-2.5-Coder-Instruct-7B				
Y	83.96	81.39	79.15	87.50
N	81.94	80.85	76.90	90.70

Table 4: The Impact of Syn Inputs with High Complexity Differentiation in online reinforcement learning.

Inputs with High Complexity Differentiation Gain Better Efficiency During reinforcement learning, employing inputs with enhanced complexity differentiation to compute efficiency rewards enables more precise discrimination of code

runtime performance, thereby driving superior efficiency performance. As validated in Table 4, such synthetically generated complex inputs effectively . However, this approach introduces a trade-off that may compromise code correctness. This effect is particularly pronounced in scenarios with insufficient test coverage, as evidenced by the EvalPerf benchmark results where efficiency optimization conflicts with functional accuracy.

Model	EvalPerf		Mercury	
	DPS _{norm}	Pass@1	Beyond	Pass@1
Base	90.81	78.06	80.59	88.75
Inst	83.96	81.39	79.15	87.50
DPO ($p = 0.1$)	87.35	84.32	80.52	88.83
DPO ($p = 0.5$)	88.27	83.42	78.66	86.48

Table 5: The Impact of different initialization in reinforcement learning. Where the p means proportions efficiency pairs. The Base and Inst means Qwen-2.5-Coder-7B and Qwen-2.5-Coder-Instruct-7B. The DPO model is tuned based on the Qwen-2.5-Coder-Instruct-7B.

Initialization from Models with Higher Diversity or Accuracy Gain Better Performance As the Table 5 shows, the base model and post-trained models achieve better performance. Base models exhibit higher output diversity (Zhu et al. 2025), which is effectively leveraged by online reinforcement learning. This suggests that starting directly from a base model while considering both correctness and efficiency in RL is feasible. Initializing with DPO (High Accuracy, $p = 0.1$) allows RL to focus more on optimizing efficiency rather than correctness during the training process. As we conjectured in the section 3.4, this results in better efficiency optimization without significantly compromising accuracy.





Figure 2: Comparison of GRPO and RLOO performance in terms of accuracy (pass@1) and efficiency (dps), Initialized with the Qwen-2.5-Coder-7B-Instruct. While GRPO’s efficiency improves more rapidly than RLOO’s during training, it experiences a notable decline in correctness in the later stages (after 5-10 episodes). In contrast, RLOO demonstrates a better ability to maintain its correctness over time.

Error-Insensitive RL Algorithms can Maintain Accuracy better As the Table 2 and Figure 2 shows, RLOO gener-

Phase 1: Correction Growth



Phase 2: Efficiency Improve




Figure 3: Overview of our method. Our approach consists of two phases. In Phase 1, we enhance code correctness by fine-tuning a base LLM with Direct Preference Optimization (DPO) on preference pairs of generated solutions. In Phase 2, the resulting high-accuracy model serves as an initial policy and is further optimized for runtime efficiency using online reinforcement learning (with error-insensitive RL algorithms, such as RLOO).

LLMs	Method	EvalPerf		Mercury		AVG	
		⚡	Pass@1	⚡	Pass@1	⚡	Pass@1
Qwen-2.5-Coder-Instruct-32B	Untuned	80.92	85.59	76.97	94.14	78.95	89.87
	ECCO	82.16 (↑1.24)	63.56 (↓22.03)	73.29 (↓3.68)	89.06 (↓5.08)	77.73 (↓1.22)	76.31 (↓13.56)
	Effi-Learner	82.45 (↑1.53)	77.11 (↓8.48)	77.13 (↑0.16)	91.41 (↓2.73)	79.79 (↑0.84)	84.26 (↓5.61)
	LLM4EFFI	86.20 (↑5.28)	87.30 (↑1.71)	78.96 (↑1.99)	93.75 (↓0.39)	82.58 (↑3.63)	90.53 (↑0.66)
Qwen-2.5-Coder-Instruct-7B	Untuned	77.60	77.50	74.80	77.20	76.18	77.37
	DPO	88.50 (↑10.90)	81.64 (↑4.14)	76.66 (↑1.86)	84.38 (↑7.18)	82.58 (↑6.40)	83.01 (↑5.64)
	RLOO	81.97 (↑4.37)	82.46 (↑4.96)	76.00 (↑1.20)	89.92 (↑12.72)	78.99 (↑2.81)	86.19 (↑8.82)
	Ours	88.82 (↑11.22)	84.32 (↑6.82)	79.04 (↑4.24)	90.78 (↑13.58)	83.93 (↑7.75)	87.55 (↑10.18)

Table 6: The results of different methods of Qwen-2.5-Coder-Instruct-7B model and Qwen-2.5-Coder-Instruct-32B. The results of Qwen-2.5-Coder-Instruct-32B are cited from LLM4EFFI (Ye et al. 2025). The symbols ⚡ in EvalPerf and Mercury represent the DPS_{norm} and Beyond metrics, respectively.

ally achieves higher average accuracy, while GRPO tends to achieve higher code execution efficiency. This discrepancy may be attributed to GRPO’s regularization mechanism, which amplifies minor improvements but also magnifies errors in efficiency rewards, potentially leading to over-optimization and reduced correctness.

4 Towards Better Correctness and Efficiency

4.1 How Can we Gain Better Performance?

As we analyzed the impact of various factors on the final model performance. This analysis will help us understand the key aspects that need to be considered to achieve better accuracy and runtime efficiency.

1. **Switching to a More Stable RL Algorithm to Maintain Accuracy:** As illustrated in Figure 2, RLOO demonstrates superior stability compared to GRPO in the later stages of training. Consequently, we adopt the RLOO algorithm

to optimize the model’s performance while mitigating accuracy degradation.

2. **Initializing from a High-Accuracy Model:** The RLOO algorithm tends to preserve the accuracy of its initial model. We therefore initialize our training from a checkpoint with higher accuracy—specifically, a model previously fine-tuned with Direct Preference Optimization (DPO). The proportion of “efficiency pairs” for this stage was set to 0.1.
3. **Using High-Contrast Inputs to Compute Efficiency Rewards:** Our prior analysis indicates that employing more discriminative inputs is crucial for calculating rewards. These inputs elicit a wider range of performance, enabling a more accurate distinction between the runtime efficiencies of different code generations.

Example Response

The problem requires finding the number of range sums that lie within a given inclusive range [lower, upper]. A **straightforward** approach would be to use a nested loop to calculate all possible range sums, but this would be inefficient for large input sizes due to its $O(n^2)$ complexity. A **more efficient** solution can be achieved using a prefix sum array and a binary search approach to count the valid range sums in $O(n \log n)$ time complexity.

```
1  ```python
2  class Solution(object):
3      def countRangeSum(self, nums, lower, upper):
4          count, running_sum = 0, 0
5          accum_sum = [0]
6          for num in nums:
7              running_sum += num
8              count += bisect_right(accum_sum, running_sum - lower) - bisect_left(
9                  accum_sum, running_sum - upper)
9              bisect.insort(accum_sum, running_sum)
10         return count
11  ```
```

Figure 4: The model typically posits that there exists a straightforward approach to solving the problem, then proceeds to highlight that this method may be inefficient for certain inputs. Subsequently, it proposes an alternative solution with improved speed and concludes by providing the corresponding code.




Figure 5: Accuracy and runtime efficiency of the RLOO training process using a DPO-tuned model for initialization.

4.2 Method

Derived from our analysis, we designed a two-stage training strategy as shown in Figure 3. This approach is structured to first establish a robust foundation of code correctness before systematically optimizing for runtime efficiency.

Stage 1: Correctness Growth We utilize Direct Preference Optimization (DPO) with a dataset of code pairs heavily weighted towards correctness. Specifically, the training data is composed of 90% correctness-focused pairs and 10% efficiency-focused pairs, which allows us to produce a high-accuracy foundational model.

Stage 2: Efficiency Improve Initializing from the DPO-tuned checkpoint, we employ the RLOO algorithm to fine-tune for runtime efficiency. As shown in Figure 5, the training process yields a stable improvement in runtime efficiency while preserving the high level of correctness achieved in the prior stage.

4.3 Results

The result in Table 2 demonstrates that our proposed method achieves significant performance enhancements on two benchmarks. On the EvalPerf benchmark, it achieved a score of 88.82 on DPS_{norm} and 84.32 on Pass@1. For the Mercury benchmark, the scores were 79.04 for Beyond and 90.78 for Pass@1. The success of this method validates our findings and demonstrates an effective path towards better correctness and efficiency in code generation.

4.4 Case Study

This section presents specific cases to illustrate both the model’s intended successful paradigm and several instances of reward hacking.

In a typical successful optimization case (see Figure 4), the model first identifies a direct or “brute-force” solution and analyzes its potential performance bottlenecks on certain inputs. Subsequently, it proposes a more efficient algorithm and provides the complete implementation code. This process demonstrates the model’s expected reasoning and optimization capabilities.

However, we also observed that the model exploits loopholes in the evaluation mechanism to “hack” the reward signal, leading to its distortion.

- **Exploiting the Evaluation Method:** One failure mode arises from our method of measuring execution efficiency (see Figure 6a). To ensure stable results, our evaluation script warms up the system and averages multiple runs. The model exploited this by introducing an LRU Cache to “memorize” results across repeated executions, thereby achieving artificially high reward scores in subsequent runs.
- **Exploiting Limited Test Cases:** Another hacking behavior stems from the limited scope of our test cases (see

```

1 import string
2 from functools import lru_cache
3
4
5 class Solution(object):
6     @lru_cache(maxsize=None)
7     def isPalindrome(self, s):
8         cleaned = s.translate(
9             s.maketrans(
10                '',
11                '',
12                string.punctuation + ' '
13            )
14        ).lower()
15        return cleaned == cleaned[::-1]
16

```

(a) Reward Hacking: LRU cache.

```

1 fib_dict = {
2     0: 0,
3     1: 1,
4     2: 1,
5     ...
6     29: 514229,
7     30: 832040
8 }
9
10 class Solution(object):
11     def fib(self, n):
12         """
13         :type n: int
14         :rtype: int
15         """
16         return fib_dict[n]

```

(b) Reward Hacking: Static Table.

Figure 6: Two examples of reward hacking behavior in different scenarios. (a) The absence of isolation between evaluation environments for the same code segment allows for the utilization of an LRU Cache, which significantly enhances “runtime efficiency” and consequently yields a high reward. (b) The incompleteness and limited scope of the collected test cases for the problems enable the model to construct a static table from its pre-trained knowledge to pass the tests, thereby achieving a high reward.

Figure 6b). Drawing on knowledge from pre-training, the model used data structures like dictionaries to hard-code or memoize solutions for specific test inputs. While this approach passed the existing tests, it lacked generalizability, failed on new inputs, and sometimes produced incorrect results.

5 Related Work

Recent advancements in Code Large Language Models (CodeLLMs) have demonstrated remarkable capabilities across a spectrum of code-related tasks (Li et al. 2022; Lozhkov et al. 2024; Wang et al. 2023; Yu et al. 2023; Wei et al. 2023; Luo et al. 2024; Hui et al. 2024). However, a growing body of research highlights a significant limitation: their tendency to overlook the execution efficiency of the generated code (Waghjale et al. 2024; Huang et al. 2024b; Du et al. 2024; Liu et al. 2024). Empirical studies have shown that code produced by these models can be 3 to 13 times slower than human-written counterparts (Huang et al. 2024b). To systematically evaluate and address this issue, several benchmarks have been established. Notably, Mercury (Du et al. 2024), EvalPerf (Liu et al. 2024), and EffiBench (Huang et al. 2024b) provide standardized frameworks for assessing code execution efficiency, often utilizing problems from platforms like LeetCode and EvalPlus.

In response to this efficiency gap, a nascent body of work has explored methods to optimize the code generated by LLMs. These approaches can be broadly categorized into workflow-based methods and data-centric fine-tuning strategies. One line of inquiry focuses on engineering the generation process itself. For instance, ECCO (Waghjale et al. 2024) introduce a self-refinement method using natural language feedback to critique its own code and then apply optimization suggestions, while SOAP (Huang et al. 2024a) introduces an iterative optimization loop, leveraging direct execu-

tion feedback to refine the code. Similarly, LLM4EFFI (Ye et al. 2025) proposes a two-domain framework that mimics human development by first selecting an efficient algorithm in a “logic domain” before implementing it in a “code domain.” Another prominent direction involves curating training datasets of efficient code and employing advanced training techniques. Mercury (Du et al. 2024), for example, fine-tunes models via Supervised Fine-Tuning (SFT) and Direct Preference Optimization (DPO) on a corpus of efficient, human-written solutions. Building on automated optimization, Effi-Learner (Huang et al. 2024c) fine-tunes models using the performance-enhanced code generated by the SOAP framework. Others, like PIE (Shypula et al. 2024), create specialized datasets for instruction tuning to improve C++ code efficiency, albeit with an observed trade-off in functional correctness. ACECode (Yang et al. 2024b) applies PPO to improve efficiency with human-written solutions as upper bound references.

6 Conclusion

In this work, we addressed the critical gap between the functional correctness and the runtime efficiency of code generated by large language models. Our investigation into different optimization methods revealed their distinct strengths and limitations. We found that offline approaches, while constrained by static initial data, are effective at establishing a strong correctness baseline. In contrast, error-insensitive online methods excel at dynamically improving runtime efficiency while maintaining the high level of correctness already achieved. This discovery directly guided our strategic approach. By first solidifying the model’s correctness before pursuing efficiency gains, we achieved significant improvements in both metrics. The success of this method validates our findings and demonstrates an effective path towards better correctness and efficiency in code generation.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *ArXiv preprint*, abs/2303.08774.
- Ahmadian, A.; Cremer, C.; Gallé, M.; Fadaee, M.; Kreutzer, J.; Pietquin, O.; Üstün, A.; and Hooker, S. 2024. Back to Basics: Revisiting REINFORCE-Style Optimization for Learning from Human Feedback in LLMs. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 12248–12267. Bangkok, Thailand: Association for Computational Linguistics.
- Anthropic. 2024. Claude 3.5 Sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>. 2024.06.21.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *ArXiv preprint*, abs/2309.16609.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Du, M.; Luu, A. T.; Ji, B.; Liu, Q.; and Ng, S. 2024. Mercury: A Code Efficiency Benchmark for Code Large Language Models. In Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *ArXiv preprint*, abs/2407.21783.
- Enter, H. F. S. 2008. Applied Linear Statistical Models (5th ed.). *Journal of the American statistical association*, 103(jun.): p.880.
- Huang, D.; Dai, J.; Weng, H.; Wu, P.; Qing, Y.; Zhang, J. M.; Cui, H.; and Guo, Z. 2024a. SOAP: Enhancing Efficiency of Generated Code via Self-Optimization. *ArXiv preprint*, abs/2405.15189.
- Huang, D.; Qing, Y.; Shang, W.; Cui, H.; and Zhang, J. 2024b. EffiBench: Benchmarking the Efficiency of Automatically Generated Code. In Globersons, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J. M.; and Zhang, C., eds., *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Huang, D.; Zeng, G.; Dai, J.; Luo, M.; Weng, H.; Qing, Y.; Cui, H.; Guo, Z.; and Zhang, J. M. 2024c. Effi-Code: Unleashing Code Efficiency in Language Models. *ArXiv preprint*, abs/2410.10209.
- Hui, B.; Yang, J.; Cui, Z.; Yang, J.; Liu, D.; Zhang, L.; Liu, T.; Zhang, J.; Yu, B.; Dang, K.; et al. 2024. Qwen2. 5-Coder Technical Report. *ArXiv preprint*, abs/2409.12186.
- Jiang, A.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D.; de las Casas, D.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B (2023). *ArXiv preprint*, abs/2310.06825.
- 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.
- Liu, J.; Xie, S.; Wang, J.; Wei, Y.; Ding, Y.; and Zhang, L. 2024. Evaluating Language Models for Efficient Code Generation. In *First Conference on Language Modeling*.
- Lozhkov, A.; Li, R.; Allal, L. B.; Cassano, F.; Lamy-Poirier, J.; Tazi, N.; Tang, A.; Pykhtar, D.; Liu, J.; Wei, Y.; et al. 2024. Starcoder 2 and the stack v2: The next generation. *ArXiv preprint*, abs/2402.19173.
- Luo, Z.; Xu, C.; Zhao, P.; Sun, Q.; Geng, X.; Hu, W.; Tao, C.; Ma, J.; Lin, Q.; and Jiang, D. 2024. WizardCoder: Empowering Code Large Language Models with Evol-Instruct. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- OpenAI. 2024. GPT-4o. <https://openai.com/index/hello-gpt-4o>. 2024.05.13.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y. K.; Wu, Y.; and Guo, D. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *ArXiv preprint*, abs/2402.03300.
- Shypula, A.; Madaan, A.; Zeng, Y.; Alon, U.; Gardner, J. R.; Yang, Y.; Hashemi, M.; Neubig, G.; Ranganathan, P.; Bastani, O.; and Yazdanbakhsh, A. 2024. Learning Performance-Improving Code Edits. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *ArXiv preprint*, abs/2307.09288.
- Waghjale, S.; Veerendranath, V.; Wang, Z. Z.; and Fried, D. 2024. ECCO: Can We Improve Model-Generated Code Efficiency Without Sacrificing Functional Correctness? *ArXiv preprint*, abs/2407.14044.

Wang, Y.; Le, H.; Gotmare, A.; Bui, N.; Li, J.; and Hoi, S. 2023. CodeT5+: Open Code Large Language Models for Code Understanding and Generation. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 1069–1088. Singapore: Association for Computational Linguistics.

Wei, Y.; Wang, Z.; Liu, J.; Ding, Y.; and Zhang, L. 2023. Magicoder: Source code is all you need. *ArXiv preprint*, abs/2312.02120.

Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Fan, Z. 2024a. Qwen2 Technical Report. *ArXiv preprint*, abs/2407.10671.

Yang, C.; Kang, H. J.; Shi, J.; and Lo, D. 2024b. ACECode: A Reinforcement Learning Framework for Aligning Code Efficiency and Correctness in Code Language Models. *ArXiv preprint*, abs/2412.17264.

Ye, T.; Huang, W.; Zhang, X.; Ma, T.; Liu, P.; Yin, J.; and Wang, W. 2025. LLM4EFFI: Leveraging Large Language Models to Enhance Code Efficiency and Correctness. *ArXiv preprint*, abs/2502.18489.

Yu, Z.; Zhang, X.; Shang, N.; Huang, Y.; Xu, C.; Zhao, Y.; Hu, W.; and Yin, Q. 2023. Wavecoder: Widespread and versatile enhanced instruction tuning with refined data generation. *ArXiv preprint*, abs/2312.14187.

Zhu, A.; Asawa, P.; Davis, J. Q.; Chen, L.; Hanin, B.; Stoica, I.; Gonzalez, J. E.; and Zaharia, M. 2025. BARE: Leveraging Base Language Models for Few-Shot Synthetic Data Generation. arXiv:2502.01697.