

ReCode: Updating Code API Knowledge with Reinforcement Learning

Haoze Wu¹, Yunzhi Yao¹, Wenhao Yu², Ningyu Zhang^{1 3*}

¹Zhejiang University

²Tencent AI, Seattle Lab

³State Key Lab. for Novel Software Technology, Nanjing University, P.R. China

wuhz1020@gmail.com, yyzodd@zju.edu.cn, wenhaoyu97@gmail.com, zhangningyu@zju.edu.cn

Abstract

Large Language Models (LLMs) exhibit remarkable code generation capabilities but falter when adapting to frequent updates in external library APIs. This critical limitation, stemming from reliance on outdated API knowledge from their training data, even with access to current documentation, impedes reliable code generation in dynamic environments. To tackle this issue, we propose ReCode (rule-based **R**einforcement learning for **C**ode **U**ppdate), a novel framework that mimics human programmer adaptation to API changes. Specifically, we construct a dataset of approximately 2,000 data entries to train the LLMs to perform version migration based on updated information. Then, we introduce a modified string similarity metric for code evaluation as the reward for reinforcement learning. Our experiments demonstrate that ReCode substantially boosts LLMs' code generation performance in dynamic API scenarios, especially on the unseen CodeUpdateArena task. Crucially, compared to supervised fine-tuning, ReCode has less impact on LLMs' general code generation abilities. We apply ReCode on various LLMs and reinforcement learning algorithms (GRPO and DAPO), all achieving consistent improvements. Notably, after training, Qwen2.5-Coder-7B outperforms that of the 32B parameter code instruction-tuned model and the reasoning model with the same architecture.

Code — <https://github.com/zjunlp/ReCode>

1 Introduction

Large Language Models (LLMs) have recently demonstrated remarkable code generation abilities (Chen et al. 2021b; Zan et al. 2023; Rozière et al. 2024; Guo et al. 2024; Hui et al. 2024; Team et al. 2024; OpenAI et al. 2024; Jiang et al. 2024; Team et al. 2025a; Yang, Li et al. 2025). This capacity enables LLMs to solve data science-related tasks by calling external libraries in the generated code (Lai et al. 2022; Wang et al. 2025c; Hong et al. 2025). However, the APIs of external libraries are updated very frequently, while the model still has outdated information in their parameters (Wu et al. 2024; Isiah et al. 2024; Wang et al. 2025a; Yao et al. 2023; Liu et al. 2025b). As shown in Figure 1-top, the

*Corresponding author.

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

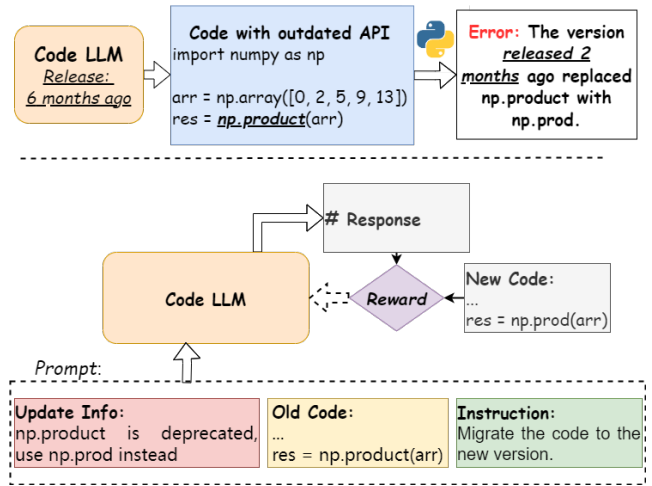


Figure 1: **Top:** LLMs cannot be aware of API updates that occur after their release date, which may lead to code errors. **Bottom:** Simply incorporating update information into the prompt cannot effectively alleviate the issue of outdated APIs. ReCode enhances their ability to migrate code to new versions through rule-based RFT.

LLM released before would generate code containing outdated APIs, leading to task failure in the up-to-date environment. This poses a challenge for the application of LLMs in scenarios that require high-quality code generation, such as AI4Research (Lu et al. 2024; Starace et al. 2025; Team et al. 2025c; Liu et al. 2025a), Software Engineering (Zhang et al. 2023; Jimenez et al. 2024; Deng et al. 2025), and Human-Computer Interaction (Zhang et al. 2024; Li et al. 2025).

The root cause of these challenges lies in the fact that LLMs are trained on static datasets, making it difficult for them to adapt to dynamically evolving scenarios such as API updates. While supervised fine-tuning (SFT) can partially alleviate this issue (Liu et al. 2025b), the high frequency of updates may result in exorbitant costs and catastrophic forgetting of previously learned knowledge (Biderman et al. 2024; Luo et al. 2025). An alternative approach involves embedding updated information directly into the prompt as a dynamic knowledge supplement (Liu et al. 2025b; Wang et al. 2025a). This method can be further enhanced by inte-

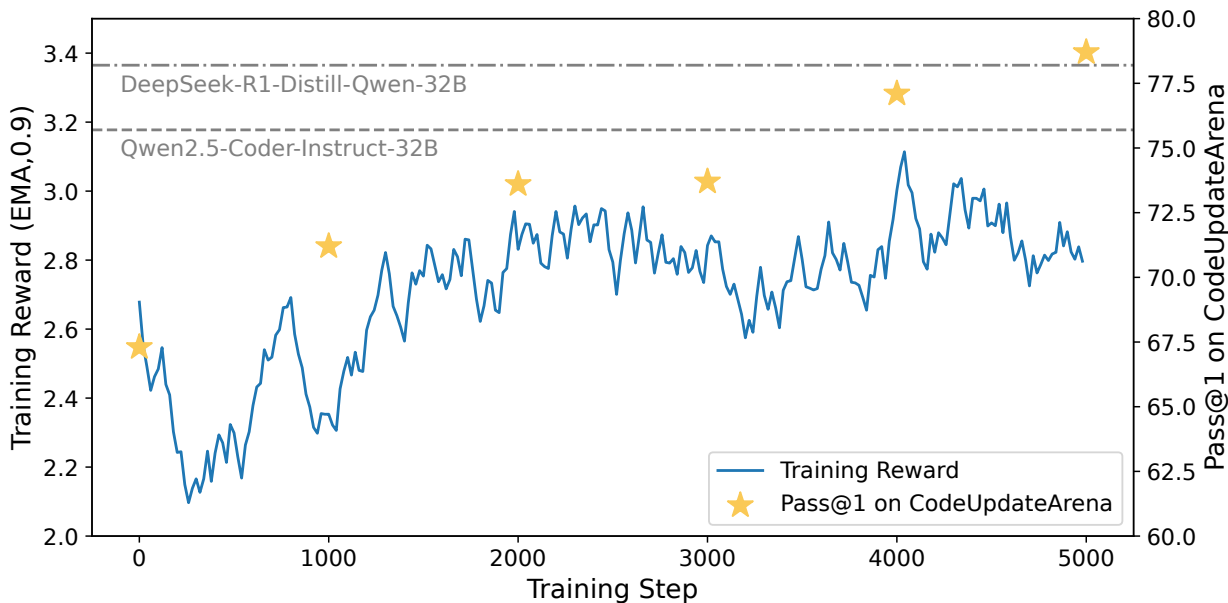


Figure 2: **Training Reward and Test Pass Rate during RL Fine-Tuning.** It demonstrates that Qwen2.5-Coder-7B-Instruct can enhance performance on the unseen CodeUpdateArena, even surpassing 32B code model and reasoning model after training. The two dashed lines in the figure represent the Pass@1 of the corresponding models on CodeUpdateArena.

grating it with retrieval-augmented generation (RAG), offering greater potential in handling dynamic scenarios (Lewis et al. 2021; Gao et al. 2024; Gupta, Ranjan, and Singh 2024). Furthermore, inputting updated knowledge through prompts avoids the risk of erasing other unrelated knowledge. Nevertheless, significant room for improvement remains in version-related code evaluation benchmarks, particularly for open-source code models (Liu et al. 2025b). This may be attributed to the conflict between inherent (parameters) and external (prompt) knowledge, where LLMs tend to prioritize their internal knowledge. In summary, while prompting is better suited for this scenario than SFT, it is necessary to address the challenge of knowledge conflicts.

Reinforcement Fine-Tuning (RFT) has been shown to enhance the model’s ability to integrate retrieval in RAG systems, boosting response accuracy (Jin et al. 2025). Inspired by this, we design ReCode, a rule-based Reinforcement learning for Code Update approach (Kaelbling, Littman, and Moore 1996; Ghasemi, Moosavi, and Ebrahimi 2025) to enhance the model’s code migration capabilities when encountering conflicting information. To take data science as an example, programmers first learn a specific version of the library, such as NumPy. Then, after being informed of the API updates, they can map the old version of the code in their minds to the new version. Our goal is to enable LLMs to use new APIs to complete tasks based on the updated information provided in the prompt. Similar to the Version-Aware Code Migration task in Versicode (Wu et al. 2024), we train the model to migrate code from an old version to a new version based on the update information. As shown in

Figure 1-bottom, we put the updated document in the prompt and fine-tune the model to better understand the prompt by using reinforcement learning based on the improved string similarity reward for code evaluation.

We evaluate the model’s performance on CodeUpdateArena (Liu et al. 2025b), a more challenging task, where the model must solve practical tasks using API update information. Our experimental results indicate that ReCode significantly enhances LLMs’ code generation capabilities in dynamic API scenarios. As shown in Figure 2, **the trained Qwen2.5-Coder-7B model outperforms Qwen2.5-Coder-32B and achieves a higher Pass@1 score than DeepSeek-R1-Distill-Qwen-32B.** Additionally, ReCode’s impact on the model’s general code generation abilities is less pronounced than that of supervised fine-tuning (SFT). In summary, our contributions are: (1) ReCode is the first to explore the application of rule-based RFT in dynamic API scenarios. Particularly, we combine the use of prompts with RFT; (2) We construct a training dataset comprising approximately 2,000 data entries, specifically designed to train models in performing version migration based on updated information; (3) Through extensive experiments and analyses, we highlight the potential of ReCode in code generation and knowledge update scenarios.

2 API Knowledge Update

Problem formula. During the pre-training phase, the LLM accumulates code knowledge within its parameter θ via an autoregressive approach (Yenduri et al. 2023). However, θ remains static post-pre-training and thus fails to incorporate

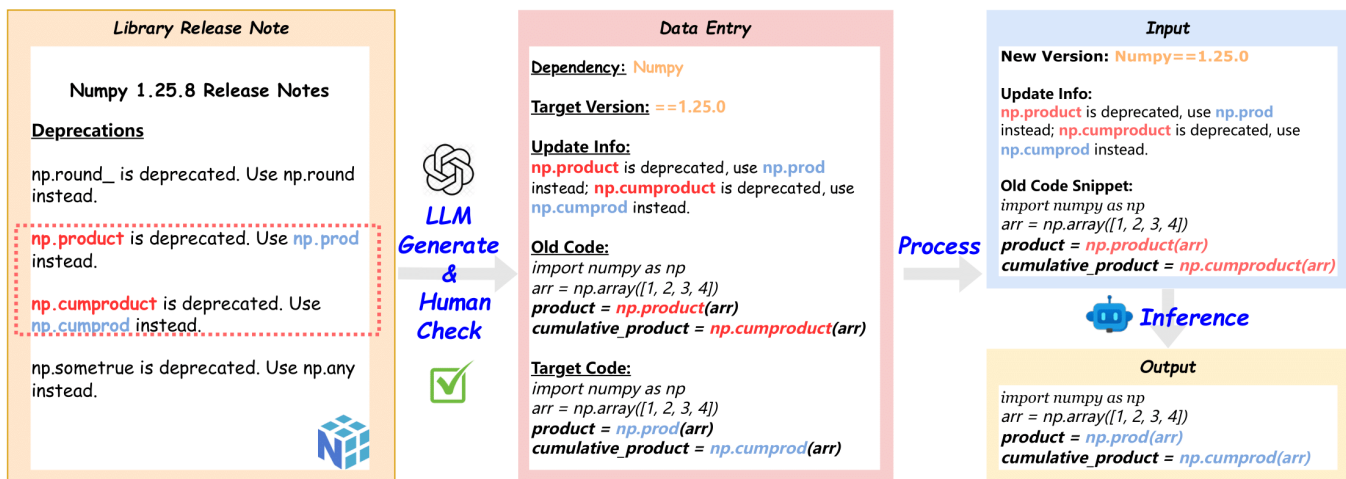


Figure 3: The pipeline of data collection and training task with a running example.

subsequent API updates. Consequently, when tasked with question q , the model may produce code containing outdated APIs: $code\ with\ outdated\ API \leftarrow P_{\theta}(\cdot|q)$.

One basic strategy is to embed updated API information within the prompt c_{update} . The model then generates code based on both the question q and the embedded update details: $code \leftarrow P_{\theta}(\cdot|q, c)$. Nevertheless, this method introduces potential conflicts between the model’s internal knowledge (housed in θ) and external information (provided in the prompt c). Such conflicts can lead to the model overlooking the updated API details in c , even when additional updates are supplied. Unlike directly writing the updated knowledge into the model’s parameter θ , we employ reinforcement learning to fine-tune $\theta \rightarrow \theta'$, enabling the model to more effectively leverage the updated API information presented in the prompt. This refined approach allows the model to generate code that better aligns with current API standards: $update\ code \leftarrow P_{\theta'}(\cdot|q, c)$.

CodeUpdateArena. CodeUpdateArena (Liu et al. 2025b) is designed to evaluate the ability of LLMs to handle API updates. The dataset comprises 670 program synthesis tasks, covering updates to 54 functions across seven different Python packages. Unlike other datasets, CodeUpdateArena is a synthetic dataset generated with LLM assistance. It prevents overlap with the model’s training data and allows for the assessment of LLMs’ adaptability to completely new updates. In addition, each entry in the dataset includes at least three test cases that can be directly executed for verification.

3 ReCode

3.1 Overview

As we described in Section 1, providing API update information to LLMs via prompts is the most promising solution, as it is more aligned with the behavior of human programmers and can be seamlessly integrated with RAG. However, current LLMs usually miss the information provided in the instructions. Our goal is to enhance the model’s ability to follow the updated information provided in the prompt.

Training. For training, we train models to perform version migration based on updated information. As shown in Figure 3-right, given a data entry $e_i = [d_i, v_i, u_i, c_i^{(old)}, c_i^{(target)}]$, which corresponds to [*Dependency*, *Target Version*, *Update Info*, *Old Code*, and *Target Code*], the input of training task is $x_i = [d_i, v_i, u_i, c_i^{(old)}]$ and the output is $y = c_i^{(target)}$. The actual prompts with templates we used are provided in the Appendix. We provide a detailed discussion of the training dataset construction in Section 3.2. Moreover, we detail the rewards utilized in the training process in Section 3.3.

Testing. During testing, each piece of data in CodeUpdateArena (Liu et al. 2025b) includes: Dependency and API update doc u_i , a real-world question q_i , and a function signature to be implemented s_i . The test task requires the model to generate the complete function code for s_i based on the above three pieces of information. Since the dataset provides test cases, we use Pass@ k as the evaluation metric. The prompts with templates are provided in the Appendix.

3.2 Training Dataset Construction

Since no existing dataset provides input-output pairs, we construct our own training dataset. As shown in Figure 3-left, the data construction process is as follows:

1. We access the release notes of major data science libraries (e.g., Numpy, Pandas, PyTorch, matplotlib) to identify paragraphs that detail specific API updates.
2. We leverage GPT-4 to generate two code snippets with equivalent functionality: one using the old API and the other utilizing the updated API.
3. Human experts review these code snippets to ensure they incorporate the updated API correctly. More details can be seen in Appendix.
4. Only those code snippets that pass the expert review are included in the dataset.

The final dataset comprises approximately 2K entries, with detailed statistics provided in Appendix. It is worth not-

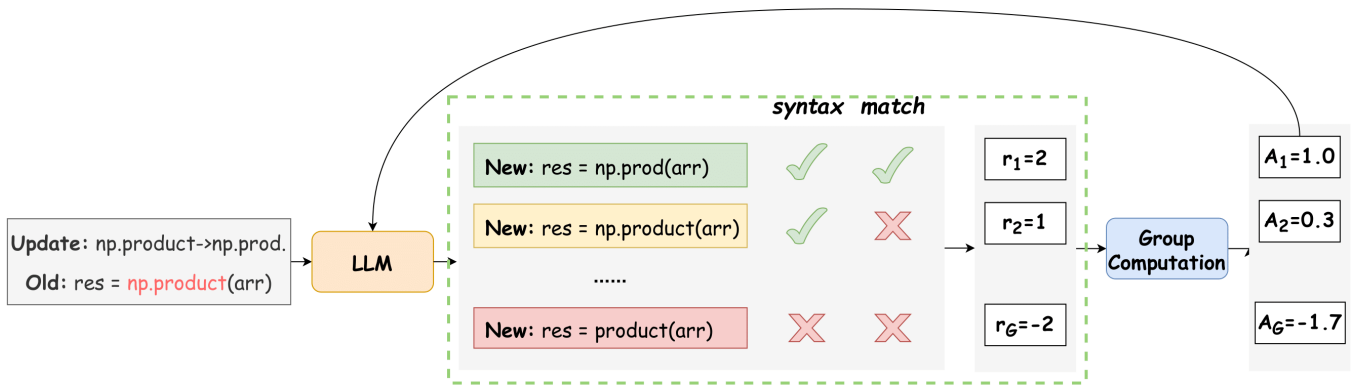


Figure 4: **Correctness Reward and Training Pipeline (taking GRPO as an example)**. The dashed box shows the correctness reward of our design, which includes two parts: syntax checking and string matching. It is worth mentioning that ReCode can be adapted to any reinforcement learning algorithm and is not limited to GRPO.

ing that the API updates in our dataset encompass a diverse range of changes, including but not limited to API renaming, parameter addition, and functionality modification. A detailed discussion is included in Appendix.

Isolation of training and testing data. In addition to the different tasks of training and testing, there are also fundamental differences in the sources of their API updates. The data we collected all come from real API updates; whereas CodeUpdateArena (Liu et al. 2025b) is a completely synthetic dataset generated by LLMs. We believe that this fundamentally ensures the isolation between datasets.

3.3 Reward Design

Similar to the existing works (DeepSeek-AI 2025; Xie et al. 2025), the reward in ReCode consists of two parts: format and correctness.

Format Reward. We hope that the output of our model meets the format:

`<think>...</think><answer>...</answer>`.

That is, to output the thinking process within the `<think>` tag, and to output the target code within the `<answer>` tag. The format reward is defined as follows:

$$R_{format}(x) = \begin{cases} +1, & \text{if } x \text{ meets the format} \\ -1, & \text{else} \end{cases} \quad (1)$$

where x is the output of the model.

Correctness Reward. Code, unlike math (Hendrycks et al. 2021; Liu et al. 2024a; Chernyshev et al. 2025), does not have a verifiable, unique standard answer. When improving coding skills, code quality can be verified by the pass rate of test cases (Liu and Zhang 2025). However, we argue that the pass rate of test cases is not a suitable reward metric for code migration task. Our training objective is to migrate correctly generated code to a new version, with the focus on “migration” rather than the inherent correctness of the code itself. Following previous works (Wu et al. 2024; Wang et al. 2025a), we use string matching metrics to evaluate the code’s cross-version migration capability:

- **Edit Similarity (ES):** ES assesses the similarity between predicted completions and target codes by analyzing the edit operations needed to transform one into the other.

- **Exact Match (EM):** EM calculates the rate at which predicted completions exactly match the target codes after normalizing return values.

Versicode (Wu et al. 2024) introduced the Critical Diff Check (CDC) metric by adding rules based on EM. However, since it only considered API name updates, CDC is not applicable to diverse API updates. Drawing on Logic-RL’s (Xie et al. 2025) reward mechanisms, we incorporated Abstract Syntax Tree (AST)-based code syntax checking into the string matching metrics, as illustrated in Figure 4. Taking EM as an example:

$$EM^*(x) = \begin{cases} +2.0, & \text{if } x \text{ match the target} \\ -1.5, & \text{elif } x \text{ is syntactically valid} \\ -2.0, & \text{else} \end{cases} \quad (2)$$

Similarly, ES^* is:

$$ES^*(x) = \begin{cases} -2.0, & \text{if } x \text{ isn't syntactically valid} \\ ES(x) \times 3.5 - 1.5, & \text{else} \end{cases} \quad (3)$$

The reward range when the syntax is correct is also $[-1.5, 2.0]$. Our reward is applicable to any RL algorithm that uses policy gradient updates. In our experiments, we select GRPO (Shao et al. 2024) and its modified version, DAPO (Yu et al. 2025), as the training algorithms.

4 Experiment

4.1 Experimental Setup

Model. The community commonly opts for rule-based reinforcement learning training of models in the Qwen family (Yu et al. 2025; Liu and Zhang 2025; Zeng et al. 2025). In our experiments, we use two code models to evaluate our method. This is to verify that our method can bring improvements not only to the Qwen model. The models we used are Qwen-2.5-Coder-7B-Instruct (Hui et al. 2024) and DeepSeek-v1.5-Coder-7B-Instruct (Guo et al. 2024). Both models demonstrate superior performance among models with fewer than 10B parameters. Notably, instruction-tuned models are found to retain R1 characteristics (Xie et al. 2025). Additionally, in the experimental

Model	Method	CodeUpdateArena		HumanEval+(Δ)
		Pass@ 1(Δ)	Pass@ 5(Δ)	
Qwen2.5-Coder-32B-Instruct DeepSeek-R1-Distill-Qwen-32B	Untrained	75.7 (+0.0)	84.3 (+0.0)	-
	Untrained	78.2 (+0.0)	86.1 (+0.0)	-
DS-v1.5-Coder-7B-Instruct	Untrained	59.1 (+0.0)	72.5 (+0.0)	71.3 (+0.0)
	SFT	53.4 (-5.7)	67.3 (-5.3)	64.0 (-7.3)
	ReCode GRPO	63.6 (+4.5)	77.3 (+4.7)	67.7 (-3.6)
	ReCode DAPO	63.6 (+4.5)	78.2 (+5.6)	68.9 (-2.4)
Qwen2.5-Coder-7B-Instruct	Untrained	67.3 (+0.0)	74.0 (+0.0)	84.1 (+0.0)
	SFT	69.4 (+2.1)	78.2 (+4.1)	70.2 (-11.7)
	ReCode GRPO	74.6 (+7.4)	82.1 (+8.0)	82.3 (-1.8)
	ReCode DAPO	78.7 (+11.3)	84.3 (+10.2)	81.7 (-2.4)

Table 1: The performance results using the GRPO and DAPO algorithms on CodeUpdateArena and HumanEval+.

section, we employ Qwen2.5-Coder-Instruct-32B (Hui et al. 2024) and DeepSeek-R1-Distill-Qwen-32B (DeepSeek-AI 2025) as baseline models. The performance metrics for these two models are obtained via the API of SiliconFlow¹. A detailed discussion on baseline selection is left in Appendix.

Training. In our experiment, we utilize the DoRA (Liu et al. 2024b) to update the model due to the computational limits. The hyperparameters are configured as $r = 64, \alpha = 64$. For the RL component, both GRPO and DAPO are set with $G = 8$, while in GRPO, the parameter β is assigned a value of 0.001. The training procedure spans 5000 steps, targeting the training tasks with a batch size set at 8 and a learning rate of 5×10^{-5} . Specifically, the initial 150 training steps incorporate a learning rate warm-up schedule, which is subsequently followed by a cosine schedule.

Test Dataset and Metric. We select the CodeUpdateArena benchmark (Liu et al. 2025b) as our test task and filter inaccurate cases from this public dataset². A detailed account of these errors is provided in Appendix. We evaluate our corrected data using the **Pass@ k** metric, with k values of 1 and 5 in our experiments. All tests are conducted using the inputs with updated information as described in Section 3.1. This in itself represents the setting of RAG. To examine how additional training affects the model’s general code capabilities, we employed HumanEval+ (Liu et al. 2023).

4.2 Main Result

Table 1 presents the results obtained from CodeUpdateArena after training.

ReCode enhances the model’s pass rate within the arena. We conduct an ablation study on the reward components in Section 4.3 and adopt ES* as the training reward. As shown in Table 1, both GRPO and DAPO yield consistent gains on the arena leaderboard. Most notably, the Qwen2.5-Coder-7B-Instruct model ultimately surpasses both the 32B-parameter instruction-tuned model and the distilled reasoning model of the same architecture on Pass@1. Its Pass@5

is slightly lower than that of DeepSeek-R1-Distill-Qwen-32B, a finding that echoes prior work (Yue et al. 2025): the ceiling of reasoning improvements delivered by RL is capped by the base model’s intrinsic capacity. Crucially, the benefits of ReCode are not confined to Qwen models: DeepSeek-Coder also enjoys a substantial uplift (Pass@1 +4.5, Pass@5 +5.6). Considering that DeepSeek-Coder is inherently weaker than Qwen2.5-Coder, the observed disparity in performance gains remains entirely reasonable.

SFT exhibits limited generalization capabilities when transitioning from the code migration task to real-world code generation tasks. As Table 1 shows, SFT’s performance gains lag behind those of ReCode (Qwen), and in some cases even degrade the pre-trained model’s efficacy (DS), consistent with observations by Liu et al. 2025b. This limitation stems from SFT’s tendency to minimize loss by memorizing prompt–answer pairs while ignoring update documentation, leading to poor robustness under task discrepancy. In contrast, RL encourages the model to earn high rewards only by correctly interpreting documentation and generating valid code, thereby fostering genuine understanding and enabling consistent generalization across dynamic, decoupled training–testing settings.

ReCode has less impact on the general capabilities of LLMs than SFT. Training the model on the updated API data may degrade its general-purpose coding proficiency. To quantify this, we evaluate the model on HumanEval+ (Liu et al. 2023), a benchmark specifically designed to assess general code-generation abilities, after training. As shown in Table 1, comparing pre-training and post-training results reveals that ReCode exerts a far milder influence on general coding capabilities than SFT. Consequently, we maintain that ReCode remains the most viable and promising solution for dynamic API scenarios.

4.3 Reward Design Ablation

In Section 3.3, we outline the design space for the correctness reward, which encompasses EM, ES, EM*, and ES* metrics. Using Qwen2.5-Coder in conjunction with GRPO (Shao et al. 2024), we conduct a comparative analysis to evaluate how these different reward mechanisms in-

¹<https://www.siliconflow.cn/>

²<https://github.com/leo-liuzy/CodeUpdateArena/blob/main/data/arena-ungrouped.json>

	format	+EM	+ES	+EM*	+ES*
Pass@1(Δ)	-2.3	-1.2	+5.4	+1.1	+7.4
Pass@5(Δ)	-3.0	-3.2	+5.2	+2.0	+8.0

Table 2: Ablation study on reward design using Qwen2.5-Coder and GRPO, comparing the impact of format reward, exact match (EM), edit similarity (ES), and their syntax-checked variants (EM*, ES*) on final performance metrics (Pass@1 and Pass@5).

fluence the final performance.

Table 2 illustrates the variations in the final metrics when training with different rewards. Below are three key findings:

Using only the format reward does not lead to improvement. Without correctness rewards, the policy model cannot receive feedback on whether the code migration is correct. Moreover, instruction-tuned models can follow instructions well without additional training; further training using only format rewards would be redundant. This underscores the necessity of introducing correctness rewards in the API update scenario.

The inclusion of AST-based syntax checking is necessary. It is evident that regardless of whether using the strict matching metric EM or the loose similarity measure ES, the inclusion of syntax checking enhances performance. We speculate that pure string matching may result in the model’s understanding of the task degrading. Additionally, Versicode (Wu et al. 2024) also observed that when the target code length increases, the correlation coefficient between EM and the pass rate of test cases decreases.

Similarity measurement is better than strict matching. We find that although ES and EM ultimately share the same goal of aligning the output with the target code, ES and ES* achieve superior results. This may be related to the calculation of the inter-group advantage in naive GRPO. When none of the outputs in a group strictly match the target code, each output receives the same reward, leading to a zero advantage for each output. DAPO (Yu et al. 2025) also highlighted this issue. In contrast, ES offers more flexible values. Even if none of the outputs in a group strictly match the target code, different outputs can still receive varying rewards. To verify this, we calculate the proportion of data with inter-group reward variance being 0 before training started. To better highlight the differences, we discard the data with an initial reward of 2. As shown in Figure 5-left, ES* has a lower proportion than EM*. This means the model can get reward feedback on more data it couldn’t handle before (reward < 2) when using GRPO.

5 Analysis

If not specifically emphasized otherwise, the analyses are based on Qwen2.5-Coder-7B-Instruct w/ DAPO ReCode.

5.1 Training Dynamics

Reward and Pass Rate. As depicted in Figure 2, during the initial few hundred training steps, the reward undergoes a decline rather than an increase. This may appear

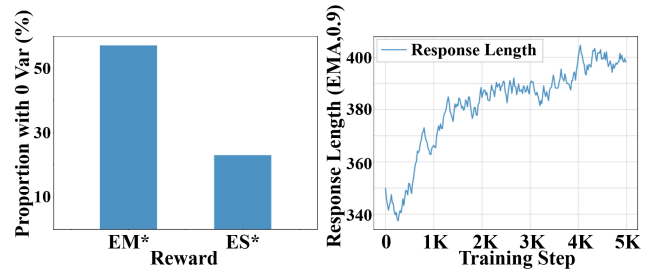


Figure 5: **Left:** the proportion of data with inter-group reward variance being 0. **Right:** The changes in response length during the training process.

counterintuitive, but our analysis reveals that the primary driver of this early-stage reduction is the decrease in format reward. Instruction-tuned models inherently possess robust instruction-following abilities, enabling them to readily comply with reasoning formats without undergoing training. Nevertheless, this very strength in adhering to instructions might concurrently restrict the exploration space of reinforcement learning algorithms, potentially entrenching specific modes of thinking. We posit that the initial reward drop reflects the model’s reluctant shift as it commences exploration. It discards its pre-existing instruction-following capabilities and established thought processes to embark on a new learning phase. The subsequent reward increase is attributed to the model having sufficiently expanded its exploration space, which in turn enables the gradual enhancement of both reward and performance metrics. However, in our experimental observations, even with this initial reward decline, instruction-tuned models demonstrated a faster learning pace compared to their base model counterparts.

Response Length. As shown in Figure 5-right, the model’s output length tends to increase during training. However, we do not believe there is a direct relationship between the model’s reasoning ability and response length. In our training tasks, the length of the model’s Chain of Thought (CoT) output is inherently limited (around 300), and even after training, the model’s output length only increases to approximately 400. This may be due to the limited thinking required for the migration task, with the model not engaging in excessive thinking throughout the process.

5.2 Case Study: The Impact of RL?

A natural question arises: Why does ReCode result in improvements? To explore this question, we divide the error cases before training into two categories. One category includes cases that are corrected after training, and the other includes those that remain wrong.

ReCode can help LLMs overcome their laziness in acquiring external knowledge. We observe that the pre-trained model exhibits ”laziness” regarding the updates in the prompt. Specifically, it either ignores the mentioned updated APIs or overlooks the new API parameters. In the Appendix, we show an erroneous case where the model disregards the updated API *math.sin* and misinterprets the angle unit ”gradian,” producing incorrect code.

Hypothesis. When generating code, an LLM draws from two information sources: internal knowledge (parameters) and provided prompts. Similar to humans, the model tends to rely more on its internal knowledge.

This propensity limits improvements when using prompts as the sole information source in dynamic scenarios. This is also one of the challenges faced by RAG systems. Longpre et al. 2022 found that when encountering conflicting information between context and internal knowledge, QA models tend to rely on internal knowledge. Xu et al. 2024 emphasized that knowledge obsolescence can lead to a conflict between parameters and prompts. Fortunately, ReCode corrects these issues, as demonstrated in the Appendix, where the model accurately uses the updated API information. ReCode enables the LLM to overcome its tendency to be “lazy” in utilizing external knowledge.

The capacity to address problems is fundamentally constrained by the pre-trained model. Most uncorrected cases are unrelated to API updates. In the Appendix, we show a case where the model is supposed to develop a pendulum simulation program. It ignores the physical quantities in the parameters and fails to describe the physical laws as required in the prompt. Unsurprisingly, even after training, the model still couldn’t write this function correctly under the condition of Pass@1. Our method doesn’t focus on enhancing physical reasoning abilities. Thus, Qwen2.5-Coder-7B-Instruct, which originally lacked this ability, still couldn’t complete the task after training.

5.3 Evaluation Under Complex Scenarios

Since our training and testing tasks only involve a single API update at a time, we explore two more complex scenarios:

Multiple updates to a single API. `torch.gels()` was deprecated in version 1.2 in favor of `torch.lstsq()`. Later, `torch.lstsq()` was also deprecated in version 1.9, replaced by `torch.linalg.lstsq()`. To update code from `torch.gels` to `torch.linalg.lstsq`, one could do it in two steps following the version sequence. However, as shown in Appendix, with update-aware prompting, the model can directly migrate from `torch.gels` to `torch.linalg.lstsq` without going through `torch.lstsq`, which is more align with human programmers.

Multiple API Updates. Can the model handle updates to multiple APIs at the same time? In the Appendix, we provide an example where the model correctly generated code when facing multiple API updates. It also doesn’t show laziness in using the knowledge from the prompt. To demonstrate the improvement achieved by ReCode, we construct 20 similar questions involving multiple API updates to test models’ performance. As shown in Table 3, ReCode enhances the performance of Qwen2.5-Coder, bringing a 7B instruct model closer to the performance of a 32B reasoning model.

6 Related Work

Benchmarking Code Generation. Most current code completion benchmarks (Chen et al. 2021a; Austin et al. 2021;

Model	Pass@1	Pass@5
DeepSeek-R1-Distill-Qwen-32B	65	80
Qwen2.5-Coder-7B-Instruct +ReCode	35 60	55 75

Table 3: The results tested on 20 questions involving multiple API updates.

Cassano et al. 2023; Nijkamp et al. 2023; Zheng et al. 2024; Zhuo et al. 2025) fail to test models’ cross-version code migration skills. To address this issue, several version-related code completion benchmarks have been proposed recently. Versicode (Wu et al. 2024) includes over 9,000 code samples from websites but only covers API name changes and lacks executable test cases. In contrast, GitChameleon (Isilah et al. 2024) collects over 100 real-world update samples to evaluate whether models can generate correct code based on library versions. CodeUpdateArena (Liu et al. 2025b) is a dataset of synthetic data with 670 samples generated by LLMs that includes executable test cases.

Reinforcement Fine-Tuning. Reinforcement Learning (RL) remains crucial for LLM post-training, with RLHF being key to GPT’s evolution into ChatGPT (Ouyang et al. 2022). With the emergence of slow-thinking reasoning models (OpenAI 2024; DeepSeek-AI 2025; Team et al. 2025b), the community has recognized the importance of rule-based RL. Particularly, the “Aha moments” observed in mathematical tasks are quite fascinating. Recently, the community has found that RFT is also effective in other domain tasks, such as video understanding (Feng et al. 2025a; Chen et al. 2025), code generation (Liu and Zhang 2025; Ma et al. 2025), tool using (Li, Zou, and Liu 2025; Qian et al. 2025; Wang et al. 2025b), machine translation (Feng et al. 2025b), and others (Jin et al. 2025).

7 Discussion and Conclusion

To the best of our knowledge, ReCode is the first to adopt the rule-based RFT post-training method for dynamic API scenarios. By training models on the code migration task, we’ve addressed the issue of AI programmers struggling to migrate code based on updates without test data. ReCode helps models make correct choices when encountering knowledge conflicts. AI programmers are tireless and efficient but lack dynamic adaptability to new knowledge, a reasoning ability once thought unique to humans. How to endow AI with this ability remains a significant challenge.

ReCode represents an exploratory step in code generation. Through ReCode, Qwen2.5-Coder-7B-Instruct has achieved cross-scale performance, surpassing the strong reasoning model with 32B parameters in the Pass@1 metric. In the future, we aim to explore this paradigm’s application in other knowledge update scenarios and further enhance reward design and training algorithms.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62576307, No. NSFCU23B2055, No. NSFCU19B2027), the Fundamental Research Funds for the Central Universities (226-2023-00138), Ningbo Natural Science Foundation (2024J020), Yongjiang Talent Introduction Programme (2021A-156-G), Tencent AILab Rhino-Bird Focused Research Program (RBFR2024003), Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

References

- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; and Sutton, C. 2021. Program Synthesis with Large Language Models. *arXiv:2108.07732*.
- Biderman, D.; Portes, J.; Ortiz, J. J. G.; Paul, M.; et al. 2024. LoRA Learns Less and Forgets Less. *arXiv:2405.09673*.
- Cassano, F.; Gouwar, J.; Nguyen, D.; Nguyen, S.; Phipps-Costin, L.; Pinckney, D.; Yee, M.-H.; Zi, Y.; Anderson, C. J.; Feldman, M. Q.; Guha, A.; Greenberg, M.; and Jangda, A. 2023. MultiPL-E: A Scalable and Polyglot Approach to Benchmarking Neural Code Generation. *IEEE Transactions on Software Engineering*, 49(7): 3675–3691.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; de Oliveira Pinto, H. P.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; Ray, A.; Puri, R.; Krueger, G.; Petrov, M.; Khlaaf, H.; Sastry, G.; Mishkin, P.; Chan, B.; Gray, S.; Ryder, N.; Pavlov, M.; Power, A.; Kaiser, L.; Bavarian, M.; Winter, C.; Tillet, P.; Such, F. P.; Cummings, D.; Plappert, M.; Chantzis, F.; Barnes, E.; Herbert-Voss, A.; Guss, W. H.; Nichol, A.; Paino, A.; Tezak, N.; Tang, J.; Babuschkin, I.; Balaji, S.; Jain, S.; Saunders, W.; Hesse, C.; Carr, A. N.; Leike, J.; Achiam, J.; Misra, V.; Morikawa, E.; Radford, A.; Knight, M.; Brundage, M.; Murati, M.; Mayer, K.; Welinder, P.; McGrew, B.; Amodei, D.; McCandlish, S.; Sutskever, I.; and Zaremba, W. 2021a. Evaluating Large Language Models Trained on Code.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; et al. 2021b. Evaluating Large Language Models Trained on Code. *arXiv:2107.03374*.
- Chen, Y.; Huang, W.; Shi, B.; Hu, Q.; Ye, H.; et al. 2025. Scaling RL to Long Videos. *arXiv:2507.07966*.
- Chernyshev, K.; Polshkov, V.; Artemova, E.; Myasnikov, A.; Stepanov, V.; Miasnikov, A.; and Tilga, S. 2025. U-MATH: A University-Level Benchmark for Evaluating Mathematical Skills in LLMs.
- DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv:2501.12948*.
- Deng, X.; Da, J.; Pan, E.; He, Y. Y.; Ide, C.; Garg, K.; et al. 2025. SWE-Bench Pro: Can AI Agents Solve Long-Horizon Software Engineering Tasks? *arXiv:2509.16941*.
- Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wang, B.; and Yue, X. 2025a. Video-R1: Reinforcing Video Reasoning in MLLMs. *arXiv:2503.21776*.
- Feng, Z.; Cao, S.; Ren, J.; Su, J.; Chen, R.; Zhang, Y.; Xu, Z.; Hu, Y.; Wu, J.; and Liu, Z. 2025b. MT-R1-Zero: Advancing LLM-based Machine Translation via R1-Zero-like Reinforcement Learning. *arXiv:2504.10160*.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Wang, M.; and Wang, H. 2024. Retrieval-Augmented Generation for Large Language Models: A Survey. *arXiv:2312.10997*.
- Ghasemi, M.; Moosavi, A. H.; and Ebrahimi, D. 2025. A Comprehensive Survey of Reinforcement Learning: From Algorithms to Practical Challenges. *arXiv:2411.18892*.
- Guo, D.; Zhu, Q.; Yang, D.; Xie, Z.; Dong, K.; Zhang, W.; Chen, G.; Bi, X.; Wu, Y.; Li, Y.; Luo, F.; Xiong, Y.; and Liang, W. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming – The Rise of Code Intelligence.
- Gupta, S.; Ranjan, R.; and Singh, S. N. 2024. A Comprehensive Survey of Retrieval-Augmented Generation (RAG): Evolution, Current Landscape and Future Directions. *arXiv:2410.12837*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. *arXiv:2103.03874*.
- Hong, S.; Lin, Y.; Liu, B.; Liu, B.; Wu, B.; et al. 2025. Data Interpreter: An LLM Agent For Data Science.
- Hui, B.; Yang, J.; Cui, Z.; Yang, J.; Liu, D.; Zhang, L.; et al. 2024. Qwen2.5-Coder Technical Report. *arXiv preprint arXiv:2409.12186*.
- Islah, N.; Gehring, J.; Misra, D.; Muller, E.; Rish, I.; Zhuo, T. Y.; and Caccia, M. 2024. GitChameleon: Unmasking the Version-Switching Capabilities of Code Generation Models. *arXiv:2411.05830*.
- Jiang, J.; Wang, F.; Shen, J.; Kim, S.; and Kim, S. 2024. A Survey on Large Language Models for Code Generation. *arXiv:2406.00515*.
- Jimenez, C. E.; Yang, J.; Wettig, A.; Yao, S.; Pei, K.; Press, O.; and Narasimhan, K. 2024. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? *arXiv:2310.06770*.
- Jin, B.; Zeng, H.; Yue, Z.; Yoon, J.; Arik, S.; Wang, D.; Zamani, H.; and Han, J. 2025. Search-R1: Training LLMs to Reason and Leverage Search Engines with Reinforcement Learning. *arXiv:2503.09516*.
- Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. 1996. Reinforcement Learning: A Survey. *arXiv:cs/9605103*.
- Lai, Y.; Li, C.; Wang, Y.; Zhang, T.; Zhong, R.; Zettlemoyer, L.; tau Yih, S. W.; Fried, D.; Wang, S.; and Yu, T. 2022. DS-1000: A Natural and Reliable Benchmark for Data Science Code Generation. *arXiv:2211.11501*.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; et al. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv:2005.11401*.
- Li, J.; Zhao, W.; Zhao, J.; Zeng, W.; Wu, H.; et al. 2025. The Tool Decathlon: Benchmarking Language Agents for Diverse, Realistic, and Long-Horizon Task Execution. *arXiv:2510.25726*.
- Li, X.; Zou, H.; and Liu, P. 2025. ToRL: Scaling Tool-Integrated RL. *arXiv:2503.23383*.
- Liu, H.; Zheng, Z.; Qiao, Y.; Duan, H.; Fei, Z.; Zhou, F.; Zhang, W.; Zhang, S.; Lin, D.; and Chen, K. 2024a. MathBench: Evaluating the Theory and Application Proficiency of LLMs with a Hierarchical Mathematics Benchmark. *arXiv:2405.12209*.
- Liu, J.; Li, Y.; Zhang, C.; Li, J.; Chen, A.; Ji, K.; et al. 2025a. WebExplorer: Explore and Evolve for Training Long-Horizon Web Agents. *arXiv:2509.06501*.
- Liu, J.; Xia, C. S.; Wang, Y.; and ZHANG, L. 2023. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Liu, J.; and Zhang, L. 2025. Code-R1: Reproducing R1 for Code with Reliable Rewards.

- Liu, S.-Y.; Wang, C.-Y.; Yin, H.; Molchanov, P.; Wang, Y.-C. F.; Cheng, K.-T.; and Chen, M.-H. 2024b. DoRA: Weight-Decomposed Low-Rank Adaptation. *arXiv:2402.09353*.
- Liu, Z. L.; Pandit, S.; Ye, X.; Choi, E.; and Durrett, G. 2025b. CodeUpdateArena: Benchmarking Knowledge Editing on API Updates. *arXiv:2407.06249*.
- Longpre, S.; Perisetla, K.; Chen, A.; Ramesh, N.; DuBois, C.; and Singh, S. 2022. Entity-Based Knowledge Conflicts in Question Answering. *arXiv:2109.05052*.
- Lu, C.; Lu, C.; Lange, R. T.; Foerster, J.; Clune, J.; and Ha, D. 2024. The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. *arXiv preprint arXiv:2408.06292*.
- Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; and Zhang, Y. 2025. An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-tuning. *arXiv:2308.08747*.
- Ma, P.; Zhuang, X.; Xu, C.; Jiang, X.; Chen, R.; and Guo, J. 2025. SQL-RL: Training Natural Language to SQL Reasoning Model By Reinforcement Learning. *arXiv:2504.08600*.
- Nijkamp, E.; Pang, B.; Hayashi, H.; Tu, L.; Wang, H.; Zhou, Y.; Savarese, S.; and Xiong, C. 2023. CodeGen: An Open Large Language Model for Code with Multi-Turn Program Synthesis. *arXiv:2203.13474*.
- OpenAI. 2024. Introducing OpenAI o1. <https://openai.com/o1/>.
- OpenAI; Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; et al. 2024. GPT-4 Technical Report. *arXiv:2303.08774*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C. L.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. *arXiv:2203.02155*.
- Qian, C.; Acikgoz, E. C.; He, Q.; Wang, H.; Chen, X.; Hakkani-Tür, D.; Tur, G.; and Ji, H. 2025. ToolRL: Reward is All Tool Learning Needs. *arXiv:2504.13958*.
- Rozière, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; et al. 2024. Code Llama: Open Foundation Models for Code. *arXiv:2308.12950*.
- 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:2402.03300*.
- Starace, G.; Jaffe, O.; Sherburn, D.; Aung, J.; Chan, J. S.; Maksin, L.; Dias, R.; Mays, E.; Kinsella, B.; Thompson, W.; Heidecke, J.; Glaese, A.; and Patwardhan, T. 2025. PaperBench: Evaluating AI's Ability to Replicate AI Research. *arXiv:2504.01848*.
- Team, C.; Zhao, H.; Hui, J.; Howland, J.; Nguyen, N.; Zuo, S.; et al. 2024. CodeGemma: Open Code Models Based on Gemma. *arXiv:2406.11409*.
- Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; et al. 2025a. Gemini: A Family of Highly Capable Multimodal Models. *arXiv:2312.11805*.
- Team, K.; Du, A.; Gao, B.; Xing, B.; Jiang, C.; Chen, C.; Li, C.; Xiao, C.; et al. 2025b. Kimi k1.5: Scaling Reinforcement Learning with LLMs. *arXiv:2501.12599*.
- Team, T. D.; Li, B.; Zhang, B.; Zhang, D.; Huang, F.; et al. 2025c. Tongyi DeepResearch Technical Report. *arXiv:2510.24701*.
- Wang, C.; Huang, K.; Zhang, J.; Feng, Y.; Zhang, L.; Liu, Y.; and Peng, X. 2025a. LLMs Meet Library Evolution: Evaluating Deprecated API Usage in LLM-based Code Completion. *arXiv:2406.09834*.
- Wang, H.; Qian, C.; Zhong, W.; Chen, X.; Qiu, J.; Huang, S.; Jin, B.; Wang, M.; Wong, K.-F.; and Ji, H. 2025b. OTC: Optimal Tool Calls via Reinforcement Learning. *arXiv:2504.14870*.
- Wang, Z.; Danek, B.; Yang, Z.; Chen, Z.; and Sun, J. 2025c. Can Large Language Models Replace Data Scientists in Biomedical Research? *arXiv:2410.21591*.
- Wu, T.; Wu, W.; Wang, X.; Xu, K.; Ma, S.; Jiang, B.; Yang, P.; Xing, Z.; Li, Y.-F.; and Haffari, G. 2024. VersiCode: Towards Version-controllable Code Generation. *arXiv:2406.07411*.
- Xie, T.; Gao, Z.; Ren, Q.; Luo, H.; Hong, Y.; Dai, B.; Zhou, J.; Qiu, K.; Wu, Z.; and Luo, C. 2025. Logic-RL: Unleashing LLM Reasoning with Rule-Based Reinforcement Learning. *arXiv:2502.14768*.
- Xu, R.; Qi, Z.; Guo, Z.; Wang, C.; Wang, H.; Zhang, Y.; and Xu, W. 2024. Knowledge Conflicts for LLMs: A Survey. *arXiv:2403.08319*.
- Yang, A.; Li, A.; et al. 2025. Qwen3 Technical Report. *arXiv preprint arXiv:2505.09388*.
- Yao, Y.; Wang, P.; Tian, B.; Cheng, S.; Li, Z.; Deng, S.; Chen, H.; and Zhang, N. 2023. Editing large language models: Problems, methods, and opportunities. *arXiv preprint arXiv:2305.13172*.
- Yenduri, G.; M, R.; G, C. S.; Y, S.; Srivastava, G.; et al. 2023. Generative Pre-trained Transformer: A Comprehensive Review on Enabling Technologies, Potential Applications, Emerging Challenges, and Future Directions. *arXiv:2305.10435*.
- Yu, Q.; Zhang, Z.; Zhu, R.; Yuan, Y.; Zuo, X.; Yue, Y.; Fan, T.; Liu, G.; Liu, L.; Liu, X.; Lin, H.; Lin, Z.; Ma, B.; Sheng, G.; Tong, Y.; Zhang, C.; Zhang, M.; Zhang, W.; Zhu, H.; Zhu, J.; Chen, J.; Chen, J.; Wang, C.; Yu, H.; Dai, W.; Song, Y.; Wei, X.; Zhou, H.; Liu, J.; Ma, W.-Y.; Zhang, Y.-Q.; Yan, L.; Qiao, M.; Wu, Y.; and Wang, M. 2025. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. *arXiv:2503.14476*.
- Yue, Y.; Chen, Z.; Lu, R.; Zhao, A.; Wang, Z.; Yue, Y.; Song, S.; and Huang, G. 2025. Does Reinforcement Learning Really Incentivize Reasoning Capacity in LLMs Beyond the Base Model? *arXiv:2504.13837*.
- Zan, D.; Chen, B.; Zhang, F.; Lu, D.; Wu, B.; Guan, B.; Wang, Y.; and Lou, J.-G. 2023. Large Language Models Meet NL2Code: A Survey. *arXiv:2212.09420*.
- Zeng, W.; Huang, Y.; Liu, Q.; Liu, W.; He, K.; Ma, Z.; and He, J. 2025. SimpleRL-Zoo: Investigating and Taming Zero Reinforcement Learning for Open Base Models in the Wild. *arXiv:2503.18892*.
- Zhang, Q.; Fang, C.; Xie, Y.; Zhang, Y.; Yang, Y.; Sun, W.; Yu, S.; and Chen, Z. 2023. A Survey on Large Language Models for Software Engineering. *arXiv preprint arXiv:2312.15223*.
- Zhang, Z.; Chen, C.; Liu, B.; Liao, C.; Gong, Z.; Yu, H.; Li, J.; and Wang, R. 2024. Unifying the Perspectives of NLP and Software Engineering: A Survey on Language Models for Code. *Transactions on Machine Learning Research*.
- Zheng, Q.; Xia, X.; Zou, X.; Dong, Y.; Wang, S.; Xue, Y.; Wang, Z.; Shen, L.; Wang, A.; Li, Y.; Su, T.; Yang, Z.; and Tang, J. 2024. CodeGeeX: A Pre-Trained Model for Code Generation with Multilingual Benchmarking on HumanEval-X. *arXiv:2303.17568*.
- Zhuo, T. Y.; Vu, M. C.; Chim, J.; Hu, H.; Yu, W.; Widayarsi, R.; et al. 2025. BigCodeBench: Benchmarking Code Generation with Diverse Function Calls and Complex Instructions. *arXiv:2406.15877*.