

# On the Impacts of Contexts on Repository-Level Code Generation

Nam Le Hai<sup>1,2,†</sup>, Dung Manh Nguyen<sup>1</sup>, Nghi D. Q. Bui<sup>1</sup>

<sup>1</sup>FPT Software AI Center, Vietnam

<sup>2</sup>Hanoi University of Science and Technology, Viet Nam

namlh@soict.hust.edu.com, dungnm31@fpt.com, bdqngghi@gmail.com

## Abstract

CodeLLMs are widely used for code generation, yet their ability to handle repository-level dependencies remains underexplored. We introduce **REPOEXEC**, a benchmark for evaluating repository-level code generation, focusing on executability, functional correctness, and dependency utilization. Our study evaluates 18 models, revealing that retaining full dependency context yields the best performance, while smaller context sizes can be misleading. Pretrained LLMs excel in correctness but often reimplement dependencies, while instruction-tuned models better utilize dependencies but sometimes introduce unnecessary complexity. We propose an instruction-tuning dataset that improves dependency handling and introduce a new metric, *Dependency Invocation Rate (DIR)*, to measure context utilization. Experiments show that instruction-tuned models improve DIR by over 10%, and multi-round debugging further enhances both correctness and dependency use. REPOEXEC provides a comprehensive framework to advance CodeLLMs for real-world applications. The dataset and source code are available at <https://github.com/FSoft-AI4Code/RepoExec>.

## 1 Introduction

Code Large Language Models (CodeLLMs) have emerged as powerful tools for assisting with coding tasks (Wang et al., 2021, 2023; Feng et al., 2020; Allal et al., 2023; Li et al., 2023; Lozhkov et al., 2024; Guo et al., 2024; Pinnaparaju et al., 2024; Zheng et al., 2024; Roziere et al., 2023; Nijkamp et al., 2022; Luo et al., 2023; Xu et al., 2022; Bui et al., 2023; Manh et al., 2024; Liu et al., 2024a). While these models excel at generating code from natural language requirements or completing individual lines of code, their application

in real-world, professional software development scenarios presents more complex challenges. A critical aspect of this complexity lies in leveraging relevant *contexts/dependencies* (e.g., function calls, imports, class hierarchy) across an entire software repository, which raises two pivotal questions. First, to what extent are the retrieved contexts accurate and relevant, rather than potential noise in the input? Second, how effectively do LLMs process and incorporate the provided dependencies into their generated code? These inquiries are central to understanding the capabilities and limitations of CodeLLMs in repository-level code generation, where completing a single line of code might require making API calls to functions within the same file (in-file context) or across different files (cross-file context).

Existing repository-level code generation benchmarks, such as RepoBench (Liu et al., 2023b), RepoCoder (Zhang et al., 2023a), CrossCodeEval (Ding et al., 2023), CoCoMIC (Ding et al., 2024), and DevEval (Li et al., 2024), have advanced the evaluation of code generation at a repository level. However, they exhibit key limitations: (1) *Lack of an executable environment*, leading to dependence on match-based metrics that fail to assess functional correctness adequately; (2) *Insufficient control over unit test quality*, reducing the data pipeline scalability and robustness of the evaluation; and (3) *Overemphasis on functional correctness metrics like pass@k*, which is inadequate for comprehensive repository-level code generation evaluation. In real-world development, code often needs to call predefined modules (functions, classes, or variables) in the repository to align with developer intent. While LLMs may generate code that passes tests, it can result in inefficient implementations or unnecessary duplication of predefined functions, leading to issues like technical debt and code smells (Maldonado and Shihab, 2015; Sierra et al., 2019; Santos et al., 2018).

\*A majority of this work was done while the first author was an AI resident at FPT Software AI Center.

†Corresponding author: [namlh@soict.hust.edu.vn](mailto:namlh@soict.hust.edu.vn)

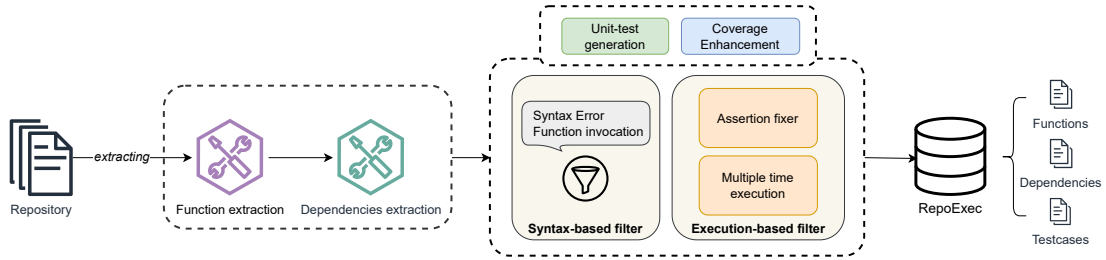


Figure 1: Data Processing Pipeline of REPOEXEC

To address these gaps, we propose REPOEXEC, a novel benchmark designed to ensure both functional correctness and effective dependency utilization. It offers an executable environment for reliable functional evaluation and incorporates a mechanism to generate high-coverage test cases, enhancing the robustness of functional correctness assessment. Additionally, we also propose *Dependency Invocation Rate (DIR)*, a metric that evaluates the extent to which generated code leverages predefined dependencies, offering a more comprehensive measure of model performance beyond merely passing test cases.

Our experiments with REPOEXEC offer several key insights into repository-level code generation. Firstly, given varying context levels (full, medium, small), which correspond to the information provided for dependencies—such as full implementation, docstrings, and function signatures—we observe that models achieve the best performance when given full context dependencies while also showing potential with smaller contexts. Secondly, pretrained LLMs and instruction-tuned LLMs exhibit different strengths: pretrained LLMs generally excel in pass@k metrics, while instruction-tuned LLMs perform better on the Dependency Invocation Rate (DIR). This suggests that while pretrained LLMs can generate runnable code, they often struggle with properly utilizing dependencies, which means the code may pass tests but might not be correctly implemented in a repository-level context. On the other hand, instruction-tuned LLMs make better use of contextual information. These findings highlight the importance of using both pass@1 and DIR for a more holistic evaluation, emphasizing the need to improve both functional correctness and context utilization in CodeLLMs.

Furthermore, we investigate two methods to enhance repository-level code generation: multi-round debugging and instruction tuning with dependency contexts. Multi-round debugging with

test execution, particularly with models like GPT-3.5 and WizardCoder, significantly boosts both pass@1 and DIR after three rounds. Additionally, fine-tuning on our dependency-enhanced dataset improves these metrics while reducing computational costs. These findings highlight the benefits of executable code testing, instruction tuning, and iterative debugging for improving CodeLLM performance and managing code dependencies more effectively. Besides, it also demonstrates that better context utilization leads to higher pass@k scores (detailed analysis provided in Sections 6.2, 6.3 and Appendix G). In summary, our contributions are as follows:

1. We introduce a novel *evaluation paradigm* for repository-level code generation, assessing both functional correctness and quality factors such as maintainability and adherence to clean code principles through efficient dependency use.
2. We present the Dependency Invocation Rate (DIR), a novel metric that measures the proportion of provided dependencies successfully incorporated into the generated code. This metric helps gauge the models’ understanding and utilization of dependencies and shows a strong correlation with functional correctness.
3. We introduce REPOEXEC, a novel benchmark that aligns with our evaluation paradigm, addressing the gaps in existing benchmarks and offering a comprehensive assessment of code generation quality. Additionally, an effective pipeline is introduced within REPOEXEC from dependency extraction, dynamically generate high-coverage test cases and automatic evaluation with execution and code dependencies. Our pipeline offers practical usage and scalability for the community.
4. We release a tool named `pydepcall` to extract dependencies for all functions within any repos-

itory, providing a practical use for advancing research in this domain.

5. Our experiments reveal key insights into CodeLLMs’ performance in repository-level code generation. While foundation models show high initial accuracy, instruction-tuned models excel in dependency management. Additionally, multi-round debugging tests further improve performance, enhancing dependency management. Notably, a strong correlation between  $\text{pass}@1$  and  $\text{DIR}$  underscores that better context utilization leads to more functionally correct code.

## 2 Related works

Coding-related tasks have been crucial for assessing the performance of Large Language Models (LLMs), with code generation emerging as a primary focus (Chen et al., 2021; Li et al., 2023; Jiang et al., 2024; Touvron et al., 2023; Roziere et al., 2023; Xu et al., 2022; Allal et al., 2023; Nijkamp et al., 2022; Phan et al., 2024a; To et al., 2023). Early benchmarks have been introduced to address this issue (Yin et al., 2018; Iyer et al., 2018; Nguyen et al., 2023; Chen et al., 2021; Austin et al., 2021; Hendrycks et al., 2021); however, they often had limited scope or employed weak evaluation approaches. Some benchmarks (Yin et al., 2018; Iyer et al., 2018; Nguyen et al., 2023) exhibit domain diversity akin to real-world applications; however, their evaluation methodologies are constrained to match-based metrics, thereby decreasing the reliability of these benchmarks (Chen et al., 2021). Meanwhile, benchmarks with reliable evaluation approaches such as HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021) and APPS (Hendrycks et al., 2021) often entail limitations to specific domains like competitive programming. Recently, there have been efforts to extend the domains of generation tasks in various benchmarks. For instance, ExeDS (Huang et al., 2022) focuses on data science code generation, while ODEX (Wang et al., 2022) serves as an open-domain benchmark for code generation. Besides, all the benchmarks mentioned primarily focus on standalone function generation, lacking consideration for cross-contextual and dependency invocation scenarios.

Several recent studies have introduced frameworks and benchmarks for repository-level code generation (Ding et al., 2024; Shrivastava et al., 2023; Ding et al., 2023; Liao et al., 2023; Liu et al., 2023b). These studies closely align with real-world

scenarios, underscoring the importance of cross-contextual considerations. However, these works are still limited to match-based evaluation methods. A recent study (Li et al., 2024) introduced the DevEval benchmark for code generation within repository contexts, evaluating performance based on extracted tests available in the repository. These benchmarks primarily focus on assessing the functional correctness of generated outputs but have not extensively investigated the correctness in conjunction with the ability to utilize dependency contexts.

## 3 Evaluation Paradigm

In this section, we outline our paradigm to achieve a more robust and comprehensive evaluation of repository-level code generation. Our paradigm encompasses two key attributes: Functional correctness and Dependency utilization.

**Functional correctness:** This evaluation criterion ensures that the code accurately performs its intended tasks and requirements. Specifically, it typically involves using test cases to compare the execution output of the generated code against the expected output for a given input. This criterion has been widely employed for evaluating code generation in numerous studies (Chen et al., 2021; Austin et al., 2021; Zhuo et al., 2024; Li et al., 2024). We follow the well-known automatic metric  $\text{Pass}@k$  (Chen et al., 2021) to measure the functional correctness of generation outputs.

**Dependency utilization:** Functional correctness alone cannot fully capture code quality, as it may overlook poor implementations or redundancy. Match-based metrics like BLEU and edit similarity evaluate alignment between generated and reference code at the token level. However, not all tokens impact code quality equally, and alternative implementations can maintain quality despite low similarity scores. Tokens representing called dependencies—such as packages, functions, variables, and classes within the repository—imply effective use of human intent for efficient implementation. Ignoring these dependencies may suggest workaround implementations misaligned with human intent, leading to increased verification and maintenance costs. To assess the models’ ability to utilize provided dependencies in accordance with human intent, we introduce the *Dependency Invocation Rate (DIR)*. This metric represents the percentage of invoked dependencies out of the total

number of dependencies provided. Let  $D_g$  denote the set of identifiers in the generated output, and  $D_s$  denote the set of provided dependencies extracted from the solution. The Dependency Invocation Rate (DIR) is calculated as follows:

$$\text{DIR} = \frac{|D_g \cap D_s|}{|D_s|}$$

A higher DIR indicates that the model successfully incorporates a larger proportion of the provided dependencies into the generated code, demonstrating a better understanding of the dependencies’ relevance and their intended usage. Conversely, a lower DIR suggests that the model may struggle to identify and utilize the appropriate dependencies, potentially generating code that is less aligned with the human-intent solution. In summary, achieving a high-quality solution requires the generated code to excel in both functional correctness and dependency utilization. Otherwise, it indicates a lack of ability to generate correct solutions or suggests poor implementation practices, including technical debt or code smell issues.

## 4 Data Collection Pipeline

Developing an executable benchmark within repository contexts is challenging due to complex setup requirements and frequent lack of clear installation guidelines in repositories. Previous studies (Ding et al., 2024; Shrivastava et al., 2023; Ding et al., 2023; Liao et al., 2023) have often relied on match-based metrics for evaluation, which may not fully capture code functionality. In addition, test cases are essential for assessing code functionality. However, extracting test cases from repositories (Zhang et al., 2023b; Li et al., 2024; Zhang et al., 2024) often relies on available functions and heuristic rules, limiting adaptability and excluding data without existing tests. For instance, Li et al. 2024 found that over 99% of functions were discarded due to the absence of suitable tests. We propose a dataset collection pipeline to ensure that repositories can build executable environments and that test cases are automatically generated. The complete data pipeline is illustrated in Figure 1, with details of the data sources provided in Appendix A.

### 4.1 Functions and Dependencies Extraction

**Function extraction:** We extract functions and their dependencies from repositories, considering only those suitable for function-level code generation, akin to benchmarks like HumanEval (Chen

	#No problem		#No testcase		Avg tokens	
	Cross-file	Total	Avg LC (%)	Avg	Prompt	Solution
Full	22.8%	355	96.25	99.45	362.92	78.46
Medium	-	-	-	-	253.05	-
Small	-	-	-	-	179.66	-

Table 1: Dataset attributes with different levels of contexts. **Cross-file** refers to the percentage of problems that involve cross-file dependencies. Tree-sitter is used for tokenization. **AVG LC**: Average Line Coverage.

et al., 2021) and MBPP (Austin et al., 2021). Using tree-sitter, we parse files into Abstract Syntax Trees (AST) to extract functions, focusing on those with comprehensive docstrings and excluding functions used as entry points or that do not produce verifiable outputs.

**Dependencies extraction:** We employ static analysis on call graph to identify dependencies, excluding identifiers that are function parameters or typing objects. Each dependency name is then mapped to its definition within the repository using a repository graph and static analysis tools. We release our tool pydepcall for community usage and provide a brief description in Appendix I.

For example in Figure 2, for the function camel\_case\_to\_snake in manipulation.py, our process identifies CAMEL\_CASE\_REPLACE\_RE as an in-file dependency and analyzes import statements to track cross-file dependencies from errors.py and validation.py. Dependencies are then parsed and incorporated into the input prompt to ensure comprehensive context for code generation.

### 4.2 Test case generation

To overcome the limitation of requiring available tests in the repository for evaluation as in previous works, we leverage large language models (LLMs) to generate test cases automatically. Our proposed approach ensures the correctness of created test cases while also controlling and enhancing test coverage. To execute and validate the generated test cases, it is necessary to configure each repository to create an executable environment. Follow Lemieux et al. 2023, we use pipreqs<sup>1</sup> to identify each project’s dependency packages.

In our test generation process, we conduct two phases corresponding to **Correctness Control** and **Coverage Enhancement**. After the test generation process, we exclude samples with a line coverage lower than 40%, as they are insufficient to accurately assess the correctness of the generated code.

<sup>1</sup><https://github.com/bndr/pipreqs>

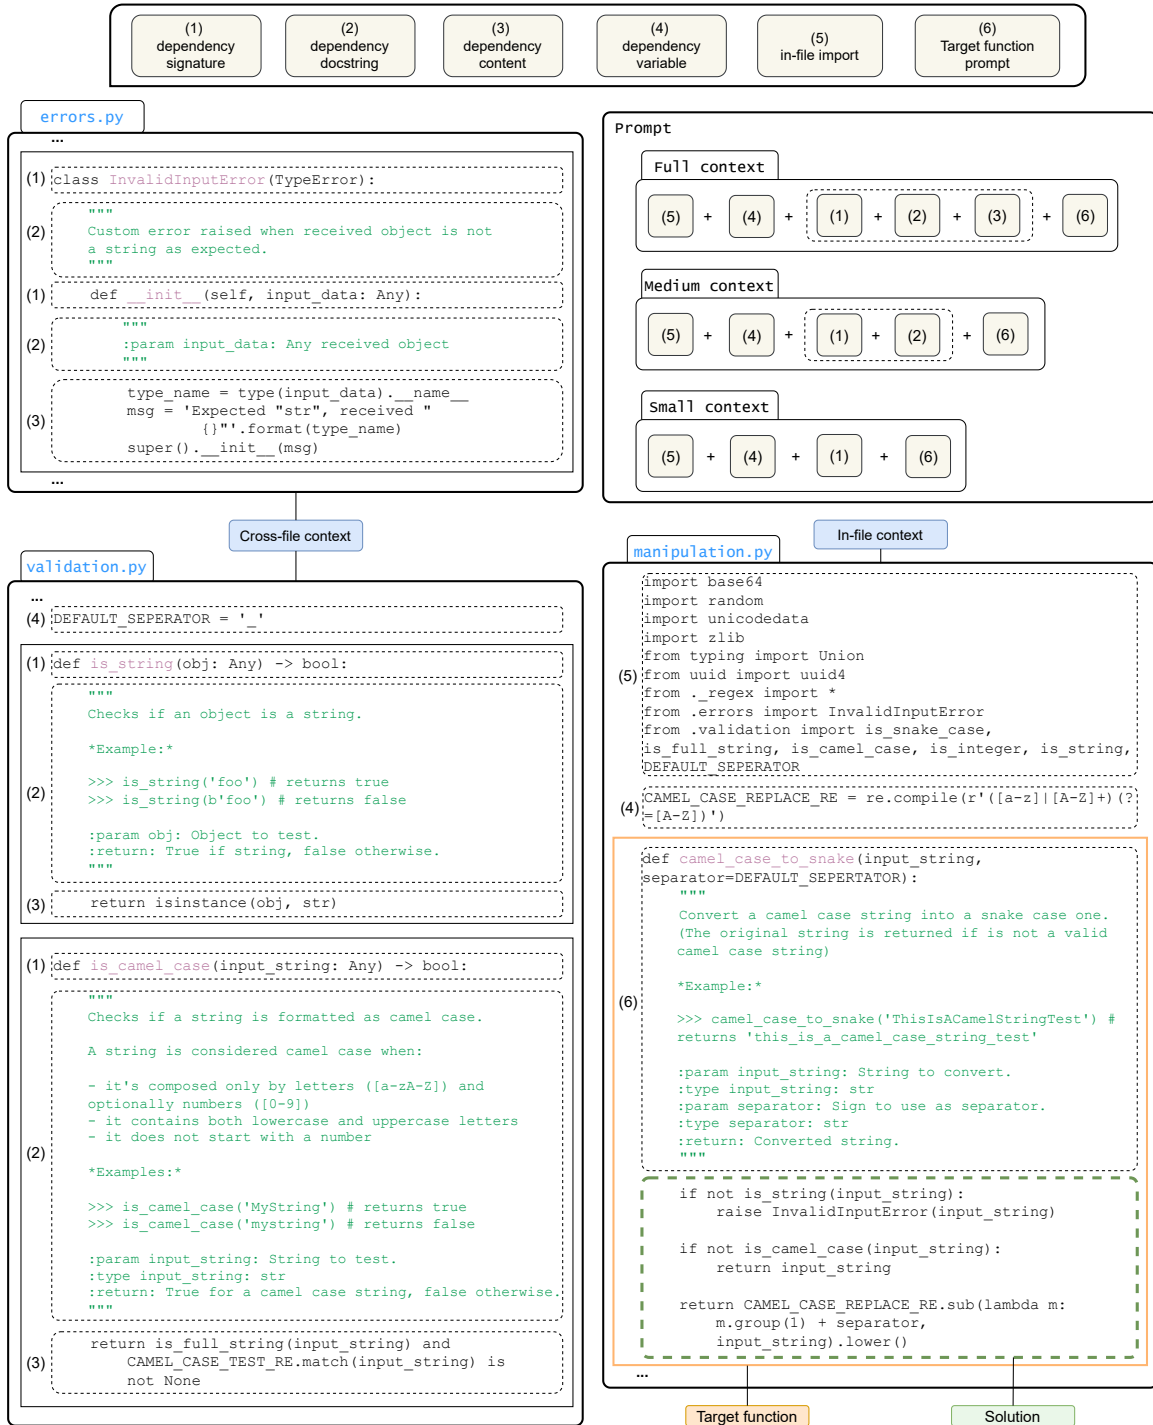


Figure 2: Illustration of a data instance in RepoExec. The target function signatures and their associated docstrings, which describe the functionality of the functions, are shown in (6). The in-file-imports and variable declarations are represented by (5) and (4), respectively. The remaining components, (1), (2), and (3), represent the function and class contexts. Specifically, (1) denotes the class or function signature, (2) may contain the description of the class, and (3) represents the function body of the cross-file function.

#### 4.2.1 Correctness Control

In this stage, we present the procedure for generating initial test cases and ensure the tests' correct-

ness. Specifically, we use CodeLlama-13b (Roziere et al., 2023) and provide the model with the prompt detailed in Appendix D. The first 20 assertions form

Pipeline	RepoContext Mining	Test Generation	Test Enhancement
RepoCoder	✓	✗	✗
R2E	✓	✓	✗
CodeBenchGen	✗	✓	✓
REPOEXEC	✓	✓	✓

Table 2: Comparison of data construction pipelines across different repo-level code generation benchmarks.

the test cases, which are then filtered via syntax and execution checks to ensure correctness.

**Syntax-based filter:** We filter out tests that present syntax errors during parsing into AST. Additionally, tests that do not invoke the function under test are excluded. For instance, while the statement “assert 1” may pass during execution, it is meaningless and negatively impacts the evaluation.

**Execution-based filter:** We use `pytest`<sup>2</sup> to run the generated tests, discarding those with errors except for `AssertionError`. If an assertion error occurs, the Assertion Fixer resolves it by executing the test with the target function, extracting the output, and updating the assertion. To handle complex return types (e.g., custom objects), we employ `pickle`<sup>3</sup> to store results. Additionally, to address flaky tests that are inconsistent due to randomness, each test is executed 10 times for comparison.

#### 4.2.2 Coverage Enhancement

Weak unit tests may allow incorrect implementations to pass (Liu et al., 2023a). To address this, we propose a strategy for enhancing test coverage using LLMs. Given the complexity of this task, requiring a strong understanding of code, we utilize GPT-3.5 to improve the quality of test cases. We provide GPT-3.5 with three prompts (see Appendix D) to handle challenging scenarios, including edge and corner cases. The initial generated tests serve as few-shot examples. We then extract the newly generated tests and ensure their correctness using our methodology from Section 4.2.1. Table 3 shows a line coverage improvement of about 4% after enhancement, reaching 96.25%. The performance gap has also increased to over 5% (Appendix C), indicating greater robustness.

## 5 Data Characteristics

### 5.1 Data Formatting

Figure 2 illustrates the input data format used in REPOEXEC. We retain import information and append dependencies in the order presented in the import statements, placing the target function signature and description at the end of the prompt. To evaluate the reasoning capability of CodeLLMs in repo-level code generation, we propose three prompt types with varying context lengths:

- **Full-size context:** All contexts, including cross-file and in-file contexts, are preserved to assess the model’s ability to navigate and utilize the complete information available in the repository.
- **Medium-size context:** Class and function bodies are removed, while their signatures and docstrings are retained. This tests the model’s ability to infer the functionality and usage of dependencies based on their interfaces and documentation, reducing the input context length.
- **Small-size context:** Only the signatures of the dependencies are retained. This tests if CodeLLMs can infer the usage of dependencies in the target function given only the function signatures without docstrings, representing the most challenging case with minimal information.

Evaluating the model’s performance across these context sizes provides insights into the trade-offs between input context length and the model’s reasoning capabilities, helping to understand the optimal balance between providing sufficient information and minimizing input size for effective code generation at the repository level.

We follow Muennighoff et al. 2023 to define 2 types of prompt formats in our evaluation of REPOEXEC across LLMs: (1) BasePrompt, which concatenates all contexts with the target functions (Figure 2), and (2) InstructPrompt, which includes specific instructions for the LLMs to follow, utilizing two variations as input formats (further details and examples in Appendix E).

### 5.2 Dataset Statistic

**Comparison to Existing Benchmarks:** Table 3 compares the details of REPOEXEC with existing code generation benchmarks. Benchmarks

<sup>2</sup><https://github.com/pytest-dev/pytest>

<sup>3</sup><https://docs.python.org/3/library/pickle.html>

Dataset	#Samples	RC	TC	LC (%)
CoNaLA (Yin et al., 2018)	500	✗	✗	-
CONCODE (Iyer et al., 2018)	2,000	✗	✗	-
HumanEval (Chen et al., 2021)	164	✗	✓, H	99.43
MBPP (Austin et al., 2021)	974	✗	✓, H	98.48
RepoCoder (Zhang et al., 2023b)	373	✓	✓, P	-
CrossCodeEval (Ding et al., 2023)	2,665	✓	✗	-
CoCoMIC (Ding et al., 2024)	6,888	✓	✗	-
DevEval (Li et al., 2024)	1,874	✓	✓, P	-
REPOEXEC	355	✓	✓, A	92.46
+ coverage-enhancement				96.25

Table 3: The comparison between popular code generation benchmarks and REPOEXEC. For test case, we denote **H**: Human annotated, **P**: Pre-existing, **A**: Automated. **RC**: Repo-context utilization. **TC**: Test Cases. **LC**: Line Coverage

that exclude execution-based evaluation (Yin et al., 2018; Iyer et al., 2018; Ding et al., 2023, 2024) may gather substantial amounts of data; however, they are inadequate for assessing the quality of the generated code. For HumanEval and MBPP, the majority of problems involve standalone functions, which do not reflect real-world scenarios. Besides, benchmarks that rely on human-annotated and pre-existing test cases (Chen et al., 2021; Austin et al., 2021; Zhang et al., 2023a; Li et al., 2024) are challenging to scale and control the test coverage. Finally, repository-context benchmarks (Zhang et al., 2023a; Ding et al., 2023, 2024; Li et al., 2024) primarily focus on investigating retrieval modules. For example, RepoCoder analyzes a retriever using a sparse bag-of-words model. Similarly, CrossCodeEval experiments and reports performance using various retrievers such as BM25, UniXCoder, and OpenAI ada. CoCoMic proposed CCFINDER to retrieve cross-file context. DevEval is the closest to our work; however, they compare the performance of the generation model based on different given types of file-level context, which can also align with the different contexts provided by different retrievers. In contrast, our approach emphasizes the generation module’s ability to understand and utilize human-provided dependency contexts.

**Comparison to Existing Data Pipeline:** Several studies, including RepoCoder (Zhang et al., 2023a), R2E (Jain et al., 2024), and CodeBenchGen (Xie et al., 2024), have introduced data pipelines for repository-level code generation. Table 2 compares our pipeline with these approaches across three key aspects: RepoContext mining, Test Generation, and Test Enhancement. These components are crucial for robust evaluation, as missing any can compromise quality. RepoCoder extracts existing

test cases but lacks enhancement, limiting scalability. CodeBenchGen uses LLMs to synthesize function contexts, reducing practicality. Lastly, R2E is the most closely related to our approach, however, it similarly neglects test case enhancement like RepoCoder, which can undermine evaluation robustness (as detailed in Appendix C).

**Dataset attributes:** Table 1 outlines the characteristics of REPOEXEC, including the total number of examples, the average number of test cases, and the number of tokens in prompts and solutions.

## 6 Experiment

### 6.1 Evaluation Results

We evaluated 18 CodeLLMs on REPOEXEC and presented the results (pass@1, pass@5, and DIR) in Table 4. We use nucleus sampling with temperature set to 0.2, top-p to 0.95, and 10 outputs generated for all models. For evaluating super-large models requiring paid API access, we utilize greedy decoding and report pass@1 and DIR metrics under the Full Context setting. The results show that retaining the full context of dependencies yields the best performance across all models. Surprisingly, Small-size context proves to be more effective than Medium-size context, which we attribute to the context’s input format using BasePrompt, potentially misleading the model into interpreting the dependency functions as few-shot examples. Using the Small-size context results in a fair decrease in performance compared to the Full context while effectively reducing the input length.

DeepSeek-R3 achieves the highest pass@1 rate among the evaluated models. Additionally, BasePrompt and pretrained models exhibit superior effectiveness over instruction-tuned models in terms of functional correctness (pass@k) when comparing models of similar size on REPOEXEC. However, our analysis reveals limitations in both types of models. Additional discussions and examples are provided in Appendix F.

1. Instruction-tuned LLMs demonstrate a higher capacity for utilizing given dependencies than foundation LMs, sometimes enabling them to address edge or corner cases that foundation models have overlooked.
2. Despite the high DIR, instruction-tuning LMs may not utilize dependencies correctly and frequently produce overly complex code, leading to incorrect solutions.

Model		Full context			Medium context			Small context		
		pass@1	pass@5	DIR	pass@1	pass@5	DIR	pass@1	pass@5	DIR
<b>BasePrompt</b>										
Pre-models	CodeLlama-13b-Python (Roziere et al., 2023)	38.65	<u>43.24</u>	62.26	<u>32.96</u>	<u>38.33</u>	56.38	<u>35.66</u>	<u>42.41</u>	62.67
	StarCoder (Li et al., 2023)	28.08	33.95	58.75	22.54	31.83	50.74	25.54	31.45	56.67
	StarCoder2-15b (Lozhkov et al., 2024)	27.77	32.60	60.57	18.70	23.28	39.28	23.27	29.78	53.49
	Mixtral-8x7B-v0.1 (Jiang et al., 2024)	22.82	29.14	55.90	19.38	25.25	47.22	20.54	26.30	53.40
	Phi-2 (Javaheripi et al., 2023)	19.04	24.56	48.22	14.54	20.34	40.85	14.82	20.69	44.54
	Phi-1 (Gunasekar et al., 2023)	14.99	18.38	43.17	12.48	15.42	37.45	12.54	15.96	38.75
Inst-models	DeepSeek-R1 (Guo et al., 2025)	<b>42.57</b>	-	70.86	-	-	-	-	-	-
	DeepSeek-V3 (Liu et al., 2024a)	<u>42.00</u>	-	<u>80.35</u>	-	-	-	-	-	-
	Llama 3.1-405B-Instruct (Dubey et al., 2024)	34.86	-	78.81	-	-	-	-	-	-
	GPT-4o	37.14	-	<b>81.43</b>	-	-	-	-	-	-
	GPT-4o-mini	30.29	-	74.75	-	-	-	-	-	-
	CodeLlama-34b-Python (Roziere et al., 2023)	40.93	<b>49.54</b>	68.85	<b>35.92</b>	<b>42.95</b>	58.15	<b>39.80</b>	<b>45.79</b>	64.23
	WizardCoder-Python-13B-V1.0 (Luo et al., 2023)	34.31	40.06	62.90	30.99	36.75	59.50	32.54	38.34	64.67
	Phind-CodeLlama-34B-v2	30.08	33.49	59.47	25.25	29.40	50.73	27.55	31.85	58.93
	CodeLlama-13b-Instruct (Roziere et al., 2023)	28.56	32.67	57.09	26.25	30.72	49.48	26.73	33.50	54.53
	GPT-3.5	27.27	37.69	63.59	23.15	33.94	52.79	22.59	33.63	55.22
	DeepSeek-Coder-7b-Instruct (Guo et al., 2024)	25.18	29.91	58.50	20.23	26.02	45.76	22.20	27.74	56.69
	Mixtral-8x7B-Instruct-v0.1 (Jiang et al., 2024)	23.41	28.71	59.83	19.04	24.83	52.75	20.45	25.84	58.17
	<b>InstructPrompt</b>									
Inst-models	DeepSeek-R1 (Guo et al., 2025)	39.71	-	53.78	-	-	-	-	-	-
	DeepSeek-V3 (Liu et al., 2024a)	41.71	-	62.58	-	-	-	-	-	-
	Llama 3.1-405B-Instruct (Dubey et al., 2024)	39.43	-	73.26	-	-	-	-	-	-
	GPT-4o	38.00	-	68.15	-	-	-	-	-	-
	GPT-4o-mini	32.29	-	53.34	-	-	-	-	-	-
	WizardCoder-Python-13B-V1.0	26.20	30.68	67.32	24.56	30.25	<u>67.54</u>	24.70	29.56	<u>68.34</u>
	CodeLlama-13b-Instruct	25.66	30.82	62.04	27.44	34.11	63.33	26.73	32.43	64.85
	GPT-3.5	23.82	39.10	40.55	19.62	36.03	37.48	19.00	34.00	35.28
	Mixtral-8x7B-Instruct-v0.1	18.11	23.04	<u>67.73</u>	18.54	23.12	<b>69.66</b>	15.38	20.61	<b>68.86</b>

Table 4: Pass@k and DIR results of various LLMs on REPOEXEC. **Bold scores** indicate the highest, while Underlined scores denote the second highest. Pre- and Inst- denote Pretrained and Instruction-tuned, respectively. More results are presented in Table 8.

3. Pretrained LLMs generate functionally correct code but often fail to effectively utilize the provided dependencies, occasionally reimplementing dependencies already present in the given context. This leads to redundancy and potentially creates technical debt or code smell issues (Maldonado and Shihab, 2015; Sierra et al., 2019; Santos et al., 2018; Hai et al., 2024a).

These issues can lead to a low-quality codebase, requiring substantial human effort for reviewing and fixing, which may even exceed the effort needed to write the code from scratch. Moreover, the performance of current advanced LLMs highlights that DeepSeek models, particularly DeepSeek-R1 and DeepSeek-V3, demonstrate strong overall capabilities, effectively balancing accuracy and dependency management. Among smaller models, GPT-4o-mini achieves the highest DIR (74.75) but underperforms in pass@1 (30.29), emphasizing it often invokes dependencies incor-

rectly despite its emphasis on handling them. Models with high pass@k and DIR scores, such as DeepSeek-V3 (80.35) and Llama 3.1 (73.26), excel in accurately invoking dependencies, which is crucial for real-world code generation tasks. Finally, REPOEXEC remains challenging to existing models, as evidenced by consistently low performance, underscoring the need for more advanced systems.

## 6.2 Generation with Multi-round Debugging

In this section, we examine the models’ self-refinement capabilities in enhancing generation performance. We provide the models with error output logs and ask them to fix the errors. We experiment with WizardCoder, GPT3.5, and CodeLlama-13b-Python. The number of rounds to debug is set to 3 and the input template is presented in Appendix G. In this experiment, we employ a greedy search algorithm to generate only a single output.

Table 5 shows the improvement across three rounds of debugging in various models. We ob-



Round	GPT-3.5	WizardCoder	CodeLlama-13b-Python
0	27.04	34.37	39.44
1	36.34	40.85	39.44
2	40.00	41.69	39.44
3	41.97	42.54	39.44

Table 5: Pass@1 scores of various models across three rounds of debugging. Round 0 represents the initial generation stage.

serve that GPT-3.5 and WizardCoder demonstrate a high capacity for debugging with improvement of over 10% and 7% in pass@1, respectively, while CodeLlama fails to take advantage of this process. Additionally, the DIR has also shown a significant improvement (over 7%) after three rounds of debugging in these two instruction models (Figure 7). These findings indicate a promising approach using self-refinement with debugging for code generation, which can enhance both the correctness and the utilization of given dependencies.

### 6.3 Instruction-tuning with Code Dependencies

While the multi-round debugging experiment has demonstrated effectiveness in leveraging given dependencies to provide correct solutions (Section 6.2), it requires a strong model to generate good test cases and can be time-consuming due to the repeated generation and execution of code and test cases. To address these challenges, we propose an instruction-tuning dataset for fine-tuning base LLMs. We collected the 1,555 most-starred repositories from 2018 onward and extracted functions with their corresponding dependencies, following the procedure outlined in Section 4.1. We obtained 154,818 functions, of which 57,746 samples include docstrings. We use 50K samples with docstrings and applied instruction prompts, while 80K samples adhered to the raw code format (Full context). Recognizing the potential of the Small context format, we allocated the remaining 20K samples to follow this structure. We fine-tuned Phi-2, StarCoder, StarCoder2, and CodeLlama-13b-Python models for 5 epochs with LoRa (Hu et al., 2021) and used 10% of the training data as the validation set to select the best checkpoint.

Table 6 illustrates the efficacy of our training data. All 4 models demonstrate improvements in both Pass@1 and DIR after instruction tuning. Specifically, there is a slight increase of around 1% in Pass@k for all models, while DIR shows a significant improvement, reaching the highest

Model	Full context		Small context	
	Pass@1	DIR	Pass@1	DIR
phi-2	19.04	48.22	14.82	44.54
phi-2 <sub>DepIT</sub>	<b>20.20</b>	<b>61.66</b>	<b>20.31</b>	<b>70.30</b>
StarCoder	28.08	58.67	25.49	56.67
StarCoder <sub>DepIT</sub>	<b>29.43</b>	<b>69.80</b>	<b>28.73</b>	<b>71.48</b>
StarCoder2	27.77	60.57	23.27	53.49
StarCoder2 <sub>DepIT</sub>	<b>28.45</b>	<b>69.76</b>	<b>27.27</b>	<b>73.98</b>
CodeLlama	38.65	62.26	35.66	62.67
CodeLlama <sub>DepIT</sub>	<b>38.85</b>	<b>68.89</b>	<b>36.93</b>	<b>73.19</b>

Table 6: Comparison of the performance of several models on REPOEXEC after instruction tuning for dependency calls ( $_{DepIT}$ ) with their pre-trained versions.

scores (over 70%) compared to other models after tuning with our dataset. Notably, performance improves significantly with small context, matching results from full context, enabling more efficient processing and reducing computational costs. In summary, our instruction-tuning method enhances the model’s ability to utilize dependencies and ensure functional correctness. While multi-round debugging (Section 6.2) is more effective, instruction-tuned models rely on single-turn generation, making them more practical and efficient. *To facilitate open research in this domain, we will publicly release this dataset.*

## 7 Conclusion

We propose an evaluation approach for repository-level code generation that rethinks the limitations of prior methods by assessing not only the functional correctness aspect but also dependency utilization to ensure code quality. We introduce REPOEXEC, a novel Python code generation benchmark with executable capabilities, designed to evaluate the alignment of generated code with developer intent and correctness. Our experiments show that while pretrained LLMs excel in functional correctness, instruction-tuned models perform better in utilizing dependencies and debugging. However, existing models struggle to reuse provided dependencies, risking technical debt and code smells. We provide a comprehensive analysis of how LLMs leverage dependencies, revealing that context information significantly influences the results. Besides, we also introduce an instruction-tuning dataset that enhances dependency invocation accuracy and output correctness, even with limited context. Our contributions establish a foundation for future research in code generation, providing valuable evaluation techniques to drive the development of more capable and reliable models.

## 8 Limitations

In this work, we currently consider one level of dependency context, specifically the dependencies directly called from the target function. While this simplification facilitates manageable analysis and model development, it may not fully capture the valuable context necessary for leveraging models effectively. However, incorporating deep dependencies could significantly extend the input length, posing challenges in managing long context inputs and potentially exceeding the maximum input length. Our approach has demonstrated promising outcomes with the small-size context version, creating opportunities for integrating additional input context. Future research could explore incorporating multiple levels of dependencies, creating a more comprehensive graph that includes transitive dependencies, indirect calls, and broader contextual information. By doing so, we could enhance the model’s understanding of code interactions and improve its ability to handle intricate software execution scenarios.

## References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, et al. 2023. Santacoder: don’t reach for the stars! *arXiv preprint arXiv:2301.03988*.
- 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*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Nghi DQ Bui, Hung Le, Yue Wang, Junnan Li, Akhilesh Deepak Gotmare, and Steven CH Hoi. 2023. Codetf: One-stop transformer library for state-of-the-art code llm. *arXiv preprint arXiv:2306.00029*.
- 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*.
- Yangruibo Ding, Zijian Wang, Wasi U. Ahmad, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, Dan Roth, and Bing Xiang. 2024. CoCoMIC: Code completion by jointly modeling in-file and cross-file context. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*.
- Yangruibo Ding, Zijian Wang, Wasi Uddin Ahmad, Hantian Ding, Ming Tan, Nihal Jain, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, Dan Roth, and Bing Xiang. 2023. Crosscodeeval: A diverse and multilingual benchmark for cross-file code completion. In *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*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. Codebert: A pre-trained model for programming and natural languages. *arXiv preprint arXiv:2002.08155*.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. 2023. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y Wu, YK Li, et al. 2024. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*.
- Nam Le Hai, Anh MT Bui, Phuong T Nguyen, Davide Di Ruscio, and Rick Kazman. 2024a. Improving the detection of technical debt in java source code with an enriched dataset. *arXiv preprint arXiv:2411.05457*.

- Nam Le Hai, Trang Nguyen, Linh Ngo Van, Thien Huu Nguyen, and Khoat Than. 2024b. Continual variational dropout: a view of auxiliary local variables in continual learning. *Machine Learning*, 113(1):281–323.
- Rajarshi Haldar and Julia Hockenmaier. 2024. Analyzing the performance of large language models on code summarization. *arXiv preprint arXiv:2404.08018*.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. 2021. Measuring coding challenge competence with apps. *NeurIPS*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi 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, Jie M Zhang, Michael Luck, and Heming Cui. 2023. Agentcoder: Multi-agent-based code generation with iterative testing and optimisation. *arXiv preprint arXiv:2312.13010*.
- Junjie Huang, Chenglong Wang, Jipeng Zhang, Cong Yan, Haotian Cui, Jeevana Priya Inala, Colin Clement, Nan Duan, and Jianfeng Gao. 2022. Execution-based evaluation for data science code generation models. *arXiv preprint arXiv:2211.09374*.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, et al. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Maor Ivgi, Uri Shaham, and Jonathan Berant. 2023. Efficient long-text understanding with short-text models. *Transactions of the Association for Computational Linguistics*, 11:284–299.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2018. Mapping language to code in programmatic context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1643–1652.
- Naman Jain, Manish Shetty, Tianjun Zhang, King Han, Koushik Sen, and Ion Stoica. 2024. R2e: Turning any github repository into a programming agent environment. In *Forty-first International Conference on Machine Learning*.
- Mojan Javaheripi, Sébastien Bubeck, Marah Abdin, Jyoti Aneja, Sebastien Bubeck, Caio César Teodoro Mendes, Weizhu Chen, Allie Del Giorno, Ronen Eldan, Sivakanth Gopi, et al. 2023. Phi-2: The surprising power of small language models. *Microsoft Research Blog*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Caroline Lemieux, Jeevana Priya Inala, Shuvendu K Lahiri, and Siddhartha Sen. 2023. Codamosa: Escaping coverage plateaus in test generation with pre-trained large language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 919–931. IEEE.
- Jia Li, Ge Li, Yunfei Zhao, Yongmin Li, Huanyu Liu, Hao Zhu, Lecheng Wang, Kaibo Liu, Zheng Fang, Lanshen Wang, et al. 2024. Deveval: A manually-annotated code generation benchmark aligned with real-world code repositories. *arXiv preprint arXiv:2405.19856*.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161*.
- Dianshu Liao, Shidong Pan, Qing Huang, Xiaoxue Ren, Zhenchang Xing, Huan Jin, and Qinying Li. 2023. Context-aware code generation framework for code repositories: Local, global, and third-party library awareness. *arXiv preprint arXiv:2312.05772*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and LINGMING ZHANG. 2023a. 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*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024b. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Tianyang Liu, Canwen Xu, and Julian McAuley. 2023b. Repobench: Benchmarking repository-level code auto-completion systems. *arXiv preprint arXiv:2306.03091*.
- Zhijie Liu, Yutian Tang, Xiapu Luo, Yuming Zhou, and Liang Feng Zhang. 2024c. No need to lift a finger anymore? assessing the quality of code generation by chatgpt. *IEEE Transactions on Software Engineering*.
- Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, et al. 2024. Starcoder 2 and the stack v2: The next generation. *arXiv preprint arXiv:2402.19173*.

- Stephan Lukasczyk, Florian Kroiß, and Gordon Fraser. 2023. An empirical study of automated unit test generation for python. *Empirical Software Engineering*, 28(2):36.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*.
- Everton da S Maldonado and Emad Shihab. 2015. Detecting and quantifying different types of self-admitted technical debt. In *2015 IEEE 7th international workshop on managing technical debt (MTD)*, pages 9–15. IEEE.
- Dung Nguyen Manh, Thang Phan Chau, Nam Le Hai, Thong T Doan, Nam V Nguyen, Quang Pham, and Nghi DQ Bui. 2024. Codemmlu: A multi-task benchmark for assessing code understanding capabilities of codellms. *arXiv preprint arXiv:2410.01999*.
- Debanjan Mondal, Abhilasha Lodha, Ankita Sahoo, and Beena Kumari. 2023. Robust code summarization. In *Proceedings of the 1st GenBench Workshop on (Benchmarking) Generalisation in NLP*, pages 65–75, Singapore. Association for Computational Linguistics.
- 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*.
- Dung Nguyen, Le Nam, Anh Dau, Anh Nguyen, Khanh Nghiem, Jin Guo, and Nghi Bui. 2023. The vault: A comprehensive multilingual dataset for advancing code understanding and generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4763–4788, Singapore. Association for Computational Linguistics.
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. Codegen: An open large language model for code with multi-turn program synthesis. *arXiv preprint arXiv:2203.13474*.
- Huy N Phan, Hoang N Phan, Tien N Nguyen, and Nghi DQ Bui. 2024a. Repohyper: Better context retrieval is all you need for repository-level code completion. *arXiv preprint arXiv:2403.06095*.
- Huy Nhat Phan, Tien N Nguyen, Phong X Nguyen, and Nghi DQ Bui. 2024b. Hyperagent: Generalist software engineering agents to solve coding tasks at scale. *arXiv preprint arXiv:2409.16299*.
- Nikhil Pinnaparaju, Reshith Adithyan, Duy Phung, Jonathan Tow, James Baicoianu, Ashish Datta, Maksym Zhuravinskyi, Dakota Mahan, Marco Bellagente, Carlos Riquelme, et al. 2024. Stable code technical report. *arXiv preprint arXiv:2404.01226*.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- José Amancio M Santos, João B Rocha-Junior, Luciana Carla Lins Prates, Rogeres Santos Do Nascimento, Mydiã Falcão Freitas, and Manoel Gomes De Mendonça. 2018. A systematic review on the code smell effect. *Journal of Systems and Software*, 144:450–477.
- Max Schäfer, Sarah Nadi, Aryaz Eghbali, and Frank Tip. 2023. Adaptive test generation using a large language model. *arXiv e-prints*, pages arXiv–2302.
- Disha Shrivastava, Hugo Larochelle, and Daniel Tarlow. 2023. Repository-level prompt generation for large language models of code. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 31693–31715. PMLR.
- Giancarlo Sierra, Emad Shihab, and Yasutaka Kamei. 2019. A survey of self-admitted technical debt. *Journal of Systems and Software*, 152:70–82.
- Ankita Nandkishor Sontakke, Manasi Patwardhan, Lovekesh Vig, Raveendra Kumar Medicherla, Ravindra Naik, and Gautam Shroff. 2022. Code summarization: Do transformers really understand code? In *Deep Learning for Code Workshop*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Qwen Team. 2024. *Qwq: Reflect deeply on the boundaries of the unknown*.
- Hung Quoc To, Nghi DQ Bui, Jin Guo, and Tien N Nguyen. 2023. Better language models of code through self-improvement. *arXiv preprint arXiv:2304.01228*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Quyen Tran, Nguyen Thanh, Nguyen Anh, Nam Hai, Trung Le, Linh Ngo, and Thien Nguyen. 2024. Preserving generalization of language models in few-shot continual relation extraction. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13771–13784.

- Yue Wang, Hung Le, Akhilesh Deepak Gotmare, Nghi DQ Bui, Junnan Li, and Steven CH Hoi. 2023. Codet5+: Open code large language models for code understanding and generation. *arXiv preprint arXiv:2305.07922*.
- Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021. Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. *arXiv preprint arXiv:2109.00859*.
- Zhiruo Wang, Shuyan Zhou, Daniel Fried, and Graham Neubig. 2022. Execution-based evaluation for open-domain code generation. *arXiv preprint arXiv:2212.10481*.
- Yiqing Xie, Alex Xie, Divyanshu Sheth, Pengfei Liu, Daniel Fried, and Carolyn Rose. 2024. Codebenchgen: Creating scalable execution-based code generation benchmarks. *arXiv preprint arXiv:2404.00566*.
- Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. A systematic evaluation of large language models of code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, pages 1–10.
- Ruijie Xu, Zengzhi Wang, Run-Ze Fan, and Pengfei Liu. 2024. Benchmarking benchmark leakage in large language models. *arXiv preprint arXiv:2404.18824*.
- Prateek Yadav, Qing Sun, Hantian Ding, Xiaopeng Li, Dejjiao Zhang, Ming Tan, Parminder Bhatia, Xiaofei Ma, Ramesh Nallapati, Murali Krishna Ramanathan, Mohit Bansal, and Bing Xiang. 2023. [Exploring continual learning for code generation models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 782–792, Toronto, Canada. Association for Computational Linguistics.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. Learning to mine aligned code and natural language pairs from stack overflow. In *Proceedings of the 15th international conference on mining software repositories*, pages 476–486.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. 2023a. Repocoder: Repository-level code completion through iterative retrieval and generation. *arXiv preprint arXiv:2303.12570*.
- Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. 2023b. Repocoder: Repository-level code completion through iterative retrieval and generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 2471–2484. Association for Computational Linguistics.
- Kechi Zhang, Jia Li, Ge Li, Xianjie Shi, and Zhi Jin. 2024. Codeagent: Enhancing code generation with tool-integrated agent systems for real-world repo-level coding challenges. *arXiv preprint arXiv:2401.07339*.
- Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhu Chen, and Xiang Yue. 2024. Opencodeinterpreter: Integrating code generation with execution and refinement. *arXiv preprint arXiv:2402.14658*.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, et al. 2024. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*.

# Appendix

## A Data Source

Creating an executable benchmark within repository contexts poses significant challenges due to intricate setup requirements and the often inadequate installation instructions provided in repositories. Therefore, we select repositories that can be automatically set up using `pipreqs` to auto-detect the required packages. Additionally, to effectively utilize the test generation module of our pipeline, we select repositories that align with the literature on unit test generation. As a result, REPOEXEC, is based on repositories from the unit test generation studies (Schäfer et al., 2023; Lukasczyk et al., 2023; Lemieux et al., 2023; Liu et al., 2024c), which provide executable environments and robust test generation capabilities.

Some might claim that our dataset is constructed based on existing works, thus, this approach can potentially lead to data leakage, especially when modern models are trained on similar datasets. If the benchmark relies heavily on known works, there’s a risk that the model may inadvertently learn from specific patterns or features present in those works, compromising its generalization ability. Despite this concern, experimental results demonstrate that even modern models struggle to handle the challenges posed by REPOEXEC, indicating that the benchmark remains a valuable tool for assessing model performance. Besides, several models, including StarCoder2 and DeepSeek-Coder, have been pretrained using repository-level context. However, these models typically concatenate the contents of multiple files in a repository without filtering out irrelevant information or considering the selection of dependencies. This helps our dataset distinguish from the pretraining datasets of these models, thereby helping to mitigate the data leakage issue. To further measure the data leakage level compare to other benchmarks, we follow BenBench (Xu et al., 2024) to report the perplexity and n-gram metrics, against three datasets—HumanEval, RepoCoder, and DevEval—analyzing their levels of data leakage when tested on three powerful LLMs. As indicated in Table 7, while HumanEval exhibits a significant level of data leakage, as expected, RepoExec demonstrates notably low levels of leakage. This is reflected in its higher perplexity and reduced n-gram overlap in most cases. The extent of data leakage in RepoExec is comparable to that of RepoCoder, a recently developed benchmark based on repositories created after January 1, 2022, to mitigate the data leakage problem. Consequently, the leakage issue in RepoExec is minimal, aligning with benchmarks like RepoCoder that utilize the latest repositories. We believe that repository-level code generation could be a valuable contribution to more robust and high-quality evaluation for advanced agent systems and dynamic environments (Huang et al., 2023; Phan et al., 2024b; Yadav et al., 2023; Hai et al., 2024b; Tran et al., 2024). By assessing models in a more comprehensive coding context, this approach enables a deeper evaluation of their ability to understand, integrate, and adapt code components effectively.

Model	perplexity				n-gram			
	HumanEval	RepoCoder	DevEval	RepoExec	HumanEval	RepoCoder	DevEval	RepoExec
Mistral-7B-v0.3	2.10	<b>177.99</b>	170.67	120.77	0.35	0.30	0.26	<b>0.18</b>
CodeLlama-7B-Python	1.96	<b>90.71</b>	32.14	72.77	0.43	<b>0.30</b>	0.43	<b>0.30</b>
CodeLlama-13B-Python	1.96	72.66	23.83	<b>126.51</b>	0.48	<b>0.31</b>	0.41	<b>0.31</b>

Table 7: Comparison of data leakage level between REPOEXEC and other benchmarks.

## B Evaluation Metrics: Match-Based vs. Execution-Based

Several code generation benchmarks utilize match-based metrics like Edit similarity (ES), BLEU, and CodeBLEU for evaluation (Ding et al., 2024; Shrivastava et al., 2023; Ding et al., 2023; Liao et al., 2023; Yin et al., 2018; Iyer et al., 2018). These metrics are straightforward to apply and may exhibit a strong correlation with execution metrics such as Pass@k. However, they cannot accurately measure functional correctness. For instance, comparing two Python code snippets where the only difference is the ":" character could result in a good ES and BLEU score. Nevertheless, one snippet may contain a syntax error, highlighting a limitation in these metrics for assessing true functionality.




Figure 3: Correlation between Match-based metrics and Execution-based metric (Pass@1).

Figure 3 demonstrates that these metrics on average can achieve a strong correlation with Pass@k, with CodeBLEU showing the highest correlation with Pearson score of 0.92. However, upon closer inspection of the score distribution between correct and incorrect solutions (Figure 4), a considerable overlap becomes apparent. This underscores the limitation of match-based metrics in accurately measuring the correctness of code generation.




Figure 4: Match-based metric distributions between Correct and Incorrect solutions

## C Coverage Enhancement Effectiveness

Weak unit tests may inadvertently allow incorrect implementations to be determined as correct. Even with human-written tests, the overlooked coverage rates lead to evaluations that are incomplete and potentially misleading (Liu et al., 2023a). We present evidence supporting this argument, underscoring the limitations of prior work on code generation within repository-level contexts. As depicted in Figure 5, enhancing the number of test cases and coverage rates leads to a significant increase in the identification of incorrect generated solutions, causing the Pass@1 score to drop markedly (by over 5%). We investigated several solutions and found that most of the generated results did not fully utilize the given context (considered as human-provided). Instead, they primarily focused on addressing the problem described in the given natural language description. This indirectly overlooks the developer’s intentions, such as testing edge or corner cases, highlighting the limitations in following and understanding the provided intent and dependency context in these models. In summary, these findings underscore the effectiveness and importance of maintaining high-quality test cases for evaluation purposes.

## D Test case generation

**Generated test quality discussion:** To ensure generated tests align with the target function’s requirements, we employ several quality control measures, as discussed in Section 4.2. First, we provide the full target function in the LLM prompt to guide the generation of tests with the intended input and output formats. We then apply a Syntax-based filter to confirm that the generated tests correctly call the function with its signature and parameters. Additionally, to verify the accuracy of the tests, we execute them to confirm their pass status and measure their line coverage.

In some minor cases, weak tests or even wrong ones may still arise despite these measures. For instance, flaky tests could pass and achieve high coverage but produce inconsistent outputs due to randomness. To address this, we execute tests multiple times and exclude those with inconsistent results. Similarly, concerns about unintended function calls are mitigated by deduplicating test cases and ensuring diversity in how functions are invoked. This process yields an average of 99 test cases per problem (Table 1), ensuring robust evaluation even if some weaker tests are included. While we do not require such extensive coverage for correctness, this approach enhances the likelihood of strong test cases identifying implementation issues effectively. Furthermore, the test enhancement module is provided with a few examples of generated tests from the previous step to enhance coverage of edge cases and corner cases to ensure specific requirements.




Figure 5: Performance of various CodeLMs on REPOEXEC before (bf-) and after (af-) Coverage Enhancement (CovEn) stage.

**Test generation prompts:** Below, we present two prompts designed for generating initial unit tests and improving the quality of existing tests.

### Initial Test Generation Prompt

```
{function_under_test}
# test to check the correctness of "{function_name}" function
assert
```



## Coverage Enhancement Prompts

### **Prompt 1:**

Here are some Python unit test functions and the focal function that they test:

# Test functions:

{existing\_test\_functions}

# Focal function:

{function\_under\_test}

Write more unit test functions that will increase the test coverage of the function under test.

---

### **Prompt 2:**

Here are some Python unit test functions and the focal function that they test:

# Test functions:

{existing\_test\_functions}

# Focal function:

{function\_under\_test}

Write more unit test functions that will cover corner cases missed by the original and will increase the test coverage of the function under test.

---

### **Prompt 3:**

Here is a focal function under test:

{function\_under\_test}

This function under test can be tested with these Python unit test functions:

{existing\_test\_functions}

Here is an extended version of the unit test function that includes additional unit test cases that will cover methods, edge cases, corner cases, and other features of the function under test that were missed by the original unit test functions:

## E Data Formating

### E.1 BasePrompt

#### Example of Full Context

```
import base64
import random
import unicodedata
import zlib
from typing import Union
from uuid import uuid4
from ._regex import *
from .errors import InvalidInputError
from .validation import is_snake_case, is_full_string, is_camel_case, is_integer, is_string

CAMEL_CASE_REPLACE_RE = re.compile(r'([a-z]|[A-Z]+)(?=[A-Z])')

class InvalidInputError(TypeError):
    """
    Custom error raised when received object is not a string as expected.
    """

    def __init__(self, input_data: Any):
        """
        :param input_data: Any received object
        """
        type_name = type(input_data).__name__
        msg = 'Expected "str", received "{}'.format(type_name)
        super().__init__(msg)

def is_string(obj: Any) -> bool:
    """
    Checks if an object is a string.

    *Example:*

    >>> is_string('foo') # returns true
    >>> is_string(b'foo') # returns false

    :param obj: Object to test.
    :return: True if string, false otherwise.
    """
    return isinstance(obj, str)

def is_camel_case(input_string: Any) -> bool:
    """
    Checks if a string is formatted as camel case.

    A string is considered camel case when:

    - it's composed only by letters ([a-zA-Z]) and optionally numbers ([0-9])
    - it contains both lowercase and uppercase letters
    - it does not start with a number

    *Examples:*

    >>> is_camel_case('MyString') # returns true
    >>> is_camel_case('mystring') # returns false

    :param input_string: String to test.
    :type input_string: str
    :return: True for a camel case string, false otherwise.
    """
    return is_full_string(input_string) and CAMEL_CASE_TEST_RE.match(input_string) is not None

def camel_case_to_snake(input_string, separator='_'):
    """
    Convert a camel case string into a snake case one.
    (The original string is returned if is not a valid camel case string)

    *Example:*

    >>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

    :param input_string: String to convert.
    :type input_string: str
    :param separator: Sign to use as separator.
    :type separator: str
    :return: Converted string.
    """
```

## Example of Medium Context

```
import base64
import random
import unicodedata
import zlib
from typing import Union
from uuid import uuid4
from .regex import *
from .errors import InvalidInputError
from .validation import is_snake_case, is_full_string, is_camel_case, is_integer, is_string

CAMEL_CASE_REPLACE_RE = re.compile(r'([a-z]|[A-Z]+)(?=[A-Z])')

class InvalidInputError(TypeError):
    """
    Custom error raised when received object is not a string as expected.
    """

    def __init__(self, input_data: Any):
        """
        :param input_data: Any received object
        """

def is_string(obj: Any) -> bool:
    """
    Checks if an object is a string.

    *Example:*

    >>> is_string('foo') # returns true
    >>> is_string(b'foo') # returns false

    :param obj: Object to test.
    :return: True if string, false otherwise.
    """

def is_camel_case(input_string: Any) -> bool:
    """
    Checks if a string is formatted as camel case.

    A string is considered camel case when:

    - it's composed only by letters ([a-zA-Z]) and optionally numbers ([0-9])
    - it contains both lowercase and uppercase letters
    - it does not start with a number

    *Examples:*

    >>> is_camel_case('MyString') # returns true
    >>> is_camel_case('mystring') # returns false

    :param input_string: String to test.
    :type input_string: str
    :return: True for a camel case string, false otherwise.
    """

def camel_case_to_snake(input_string, separator='_'):
    """
    Convert a camel case string into a snake case one.
    (The original string is returned if is not a valid camel case string)

    *Example:*

    >>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

    :param input_string: String to convert.
    :type input_string: str
    :param separator: Sign to use as separator.
    :type separator: str
    :return: Converted string.
    """
```

## Example of Small Context

```
import base64
import random
import unicodedata
import zlib
from typing import Union
from uuid import uuid4
from ._regex import *
from .errors import InvalidInputError
from .validation import is_snake_case, is_full_string, is_camel_case, is_integer, is_string

CAMEL_CASE_REPLACE_RE = re.compile(r'([a-z]|[A-Z]+)(?=[A-Z])')

class InvalidInputError(TypeError):
    def __init__(self, input_data: Any):
        ...

def is_string(obj: Any) -> bool:
    ...

def is_camel_case(input_string: Any) -> bool:
    ...

def camel_case_to_snake(input_string, separator='_'):
    """
    Convert a camel case string into a snake case one.
    (The original string is returned if is not a valid camel case string)

    *Example:*

    >>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

    :param input_string: String to convert.
    :type input_string: str
    :param separator: Sign to use as separator.
    :type separator: str
    :return: Converted string.
    """
```

## E.2 InstructPrompt

### Instruction Prompt Templates

#### Prompt 1:

### Instruction:

Write a Python function `{target\_function\_signature}` to solve the following problem:  
{target\_function\_docstring}

### Response:

{BasePrompt}

---

#### Prompt 2:

### Instruction:

{dependency\_context}  
The provided code snippet includes necessary dependencies for implementing the `{target\_function\_name}` function. Write a Python function `{target\_function\_signature}` to solve the following problem:  
{target\_function\_docstring}

### Response:

{target\_function\_prompt}

## Example of Prompt 1 for Small Context

```
### Instruction:
Write a Python function `camel_case_to_snake(input_string, separator='_)` to solve the following
  problem:
"""
Convert a camel case string into a snake case one.
(The original string is returned if is not a valid camel case string)
*Example:*
>>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

:param input_string: String to convert.
:type input_string: str
:param separator: Sign to use as separator.
:type separator: str
:return: Converted string.
"""

### Response:
import base64
import random
import unicodedata
import zlib
from typing import Union
from uuid import uuid4
from ._regex import *
from .errors import InvalidInputError
from .validation import is_snake_case, is_full_string, is_camel_case, is_integer, is_string

CAMEL_CASE_REPLACE_RE = re.compile(r'([a-z]|[A-Z]+)(?=[A-Z])')

class InvalidInputError(TypeError):
    def __init__(self, input_data: Any):

def is_string(obj: Any) -> bool:

def is_camel_case(input_string: Any) -> bool:

def camel_case_to_snake(input_string, separator='_'):
    """
    Convert a camel case string into a snake case one.
    (The original string is returned if is not a valid camel case string)
    *Example:*
    >>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

    :param input_string: String to convert.
    :type input_string: str
    :param separator: Sign to use as separator.
    :type separator: str
    :return: Converted string.
    """
```

## Example of Prompt 2 for Small Context

```
### Instruction

import base64
import random
import unicodedata
import zlib
from typing import Union
from uuid import uuid4
from ._regex import *
from .errors import InvalidInputError
from .validation import is_snake_case, is_full_string, is_camel_case, is_integer, is_string

CAMEL_CASE_REPLACE_RE = re.compile(r'([a-z]|[A-Z]+)(?=[A-Z])')

class InvalidInputError(TypeError):

    def __init__(self, input_data: Any):

def is_string(obj: Any) -> bool:

def is_camel_case(input_string: Any) -> bool:

The provided code snippet includes necessary dependencies for implementing the `camel_case_to_snake`
function. Write a Python function `camel_case_to_snake(input_string, separator='_)` to solve
the following problem:
"""
Convert a camel case string into a snake case one.
(The original string is returned if is not a valid camel case string)

*Example:*

>>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

:param input_string: String to convert.
:type input_string: str
:param separator: Sign to use as separator.
:type separator: str
:return: Converted string.
"""

### Response:
def camel_case_to_snake(input_string, separator='_):
"""
Convert a camel case string into a snake case one.
(The original string is returned if is not a valid camel case string)

*Example:*

>>> camel_case_to_snake('ThisIsACamelStringTest') # returns 'this_is_a_camel_case_string_test'

:param input_string: String to convert.
:type input_string: str
:param separator: Sign to use as separator.
:type separator: str
:return: Converted string.
"""
```

## F Studied LLMs: Supplemental results

In this section, we offer supplementary results from the evaluation of LLMs on REPOEXEC. Table 4 presents the Dependency Invocation Rate (DIR) for the experimented LLMs. When comparing models of the same size, it is shown that instruction-tuned models more effectively follow human intent in utilizing the provided dependencies with InstructPrompt. For example, WizardCoder outperforms CodeLlama by 5%, and the instruction-tuned version of Mixtral-8x7B shows a 10% improvement over its foundation version. This highlights the strong capability of instruction-tuned models to follow the given context effectively. Besides, using the Medium context leads to a significant decline in both Pass@k and DIR using BasePrompt. This implies the generation of empty function bodies using this template.

Indeed, Figure 6 illustrates the proportion of generated functions that are empty for each LLM using BasePrompt. The findings indicate that utilizing Medium context results in a substantial number of empty functions. This may be due to the input format of the context when using BasePrompt, which can mislead the model into interpreting dependency functions as few-shot examples. In the Medium context, the function bodies of dependencies are removed, making their format identical to the target function

prompt. This similarity can mislead the LMs, resulting in empty solutions. Particularly, Starcoder-2 is heavily impacted by this issue, as over 31% of its generated results are empty functions, revealing a significant weakness of the model. Meanwhile, small context effectively decreases the occurrence of empty function generation by the model and, in certain instances, improves models' ability in dependency calls (e.g. CodeLlama-13b-Python, WizardCoder-Python-13B-V1.0, and Mixtral-8x7B-Instruct-v0.1 in Table 4). We believe that the following reasons could contribute to this observation. Firstly, employing small context reduces the input token count, preventing truncation when exceeding the maximum length limit, thus allowing uninterrupted solution generation by the model. This reduced context enables models to concentrate exclusively on dependency signatures, thereby enhancing the probability that generated solutions effectively utilize these dependency token names. Moreover, function names hold substantial semantic value by delineating the function's purpose. Many studies have underscored that code summarization heavily relies on extracting information from function names (Haldar and Hockenmaier, 2024; Mondal et al., 2023; Sontakke et al., 2022). Therefore, this concise representation of dependencies has the potential to improve how models utilize dependencies in generating code.

Additionally, Table 9 presents examples that support our findings in Section 6.1. For instance, in the first example, we observe that the instruction-tuned model can effectively utilize the given dependencies to manage edge cases to raise an error, whereas the pretrained model fails to do so. This supports our first findings. Meanwhile, the second and third examples demonstrate that pretrained models try to reimplement or devise workarounds instead of leveraging the available context. Additionally, in the second example, the instruction-tuned model successfully identifies the relevant dependencies; however, it complicates the situation and fails to produce the correct solution. This may suggest that hallucinations complicate the outputs generated by instruction-tuned models. In the third example, both types of models successfully pass the tests. However, although the context provides the dependency for creating a new Future object, these models attempt to reimplement the `_create_future` function but fail to optimize memory usage effectively. This observation implies the potential of code smell and technical debt in these generated codes.

Besides, we present supplementary results for both closed-source and open-source models in Table 8. Specifically, we employ greedy decoding and report pass@1 and DIR metrics, selecting the best-performing results for each model based on either BasePrompt or InstructPrompt with Full context. The results demonstrate strong performance in both pass@k and DIR metrics with advanced reasoning models like DeepSeek-R1 and GPT-4o, highlighting their reliability in code generation. However, all pass@1 results remain below 50%, underscoring the challenges posed by REPOEXEC.

## G Multi-round Debugging

### Prompt for Debugging

```
{dependency_context}
# The provided code snippet includes necessary dependencies for implementing
the `{target_function_name}` function. Write a Python function `{
  target_function_signature}` to solve the following problem:
{target_function_docstring}

# Here is the current solution.
{error_solution}

# When executing the below test case.
{failed_test_case}

# The provided python code solution fails the test with the following errors,
  please correct them.
{error_log}

# Please provide the modified code for me to review and provide feedback.
{target_function_prompt}
```

ID	Model	Size (B)	Pass@1	DIR
1	DeepSeek-R1 (Guo et al., 2025)	671	<b>42.57</b>	70.86
2	DeepSeek-V3 (Liu et al., 2024a)	671	<u>42.00</u>	<b>80.35</b>
3	Llama 3.1-Instruct (Dubey et al., 2024)	405	39.43	<u>73.26</u>
4	GPT-4o	-	38.00	68.15
5	QwQ-Preview (Team, 2024)	32	37.43	57.07
6	DeepSeek-Coder-Instruct (Guo et al., 2024)	33	35.71	67.02
7	Yi-1.5 (Young et al., 2024)	34	35.43	61.34
8	Qwen-2.5-Coder-Instruct (Hui et al., 2024)	14	34.28	65.68
9	CodeQwen1.5 (Bai et al., 2023)	7	32.29	59.48
10	OpenCodeInterpreter (Zheng et al., 2024)	33	31.42	61.28
11	Phi-4 (Abdin et al., 2024)	14	30.86	71.10
12	GPT-4o-mini	-	30.29	74.75
13	Llama 3.3-Instruct (Dubey et al., 2024)	70	29.43	67.20
14	DeepSeek-Coder (Guo et al., 2024)	6.7	28.86	57.44
15	Qwen-2.5-Instruct (Yang et al., 2024)	14	28.85	65.63
16	starcoder2 (Lozhkov et al., 2024)	15	28.57	60.97
17	Gemma2 (Team et al., 2024)	27	28.00	71.68
18	Llama-3.1-Instruct (Dubey et al., 2024)	70	25.14	66.95

Table 8: Performance of various LLMs on REPOEXEC.

We employ Multi-round debugging in code generation, which iteratively refines and improves the generated code through multiple cycles of debugging. Following the execution of unit tests on the generated functions, we extract the error log if the code fails to run. We employ the following prompt to utilize the model for bug fixing. This process is iterated multiple times until either the correct code is achieved or the maximum number of rounds is reached. Specifically, we set the maximum number of rounds to 3 and experimented on three models WizardCoder, GPT-3.5 and CodeLlama-13b-Python.

Table 5 shows the improvement across three rounds of debugging in various models. We observe that GPT-3.5 and WizardCoder demonstrate a high capacity for debugging with improvement of over 10% and 7% in Pass@1, respectively, while CodeLlama fails to take advantage of this process. Additionally, the DIR has also shown a significant improvement (over 7%) after three rounds of debugging in these two instruction models (Figure 7). These findings indicate a promising approach using self-refinement with debugging for code generation, which can enhance both the correctness and the utilization of given dependencies.

We also present data on the number of error types corrected in each round of WizardCoder, as illustrated in Figure 8. We can see that `AssertionError` makes up the majority of errors across all rounds. This error type indicates either incorrect outputs from the generated code or the presence of empty function bodies (return None). However, by incorporating the test output guide, the model effectively addressed most of these errors. Furthermore, fundamental issues like `SyntaxError` or `AttributeError` were promptly rectified during the initial round.

## H Context length, model size and families analysis

Long-context models can enhance the ability to comprehend and select relevant context from lengthy inputs to effectively solve the required task. In this section, we examine how the context length supported by each model affects their performance in REPOEXEC. Figure 9 demonstrates the relationship between support context length, model size, and model family in relation to pass@k and DIR scores.

We observe that context length has a weak correlation with model performance on our dataset. In contrast, model size (scaling law) and model family, which encompass different training methods (pretraining or instruction tuning), training datasets, and architectures, show a more significant impact. The weak correlation with context length can be explained by our approach’s ability to already capture relevant






Figure 6: Percentage of generated outputs that result in empty functions across various context types.

information (e.g. dependency) for each data sample, resulting in the pruning of context length for practical usage (363 tokens on average shown in Table 1). Meanwhile, models that support long contexts are often trained on data containing a mix of relevant and irrelevant information and are evaluated on their ability to retrieve the correct context in a needle-in-a-haystack scenario (Roziere et al., 2023; Ivy et al., 2023; Liu et al., 2024b). Therefore, models with varying context lengths might show a weak correlation to performance in our scenario.

## I Dependency Extraction Tool

We present `pydepcall`, a Python library designed to extract function dependencies from any repository. Our tool provides a quick, practical solution for scalability, requiring only three lines of code to extract dependencies for all functions in any repository. Specifically, our tool is based on two conventions for retrieving function dependencies of a programming language:

- **Import Convention (for cross-file dependencies):** The convention mentions how a programming language imports modules within a repository. This can be utilized to extract dependency names from different files. For instance, in Python, this can be done using statements such as `"from file import function/class/variable"` or `"from file1.file2 import function/class/variable."`
- **Calling Dependencies Convention:** This principle pertains to how a program invokes variables, functions, and classes. In Python, functions are typically called by specifying their names followed by parentheses (e.g., `function1(param1, param2)`), or by using a dot notation to access attributes or methods from a module (e.g., `file1.function1, class1.attribute1`).

Besides, `pydepcall` can extract dependencies at various depths, up to 100 levels deep, which can help the community explore the integration of a deeper context. We provide a brief overview of its usage in the following code snippet.




Figure 7: Improvement of instruction-tuning models on Pass@1 and DIR after 3-round debugging process.




Figure 8: Fixed error types of WizardCoder across 3 rounds of the debugging process.

```

pydepcall usage

from pydepcall import Extractor

# If you want to extract all module files in the repository
reposrc = YOUR_LOCAL_PATH_OF_REPO
extractor = Extractor(reposrc)
output = extractor.extract()

# If you want to extract a specific module file in the repository
reposrc = YOUR_LOCAL_PATH_OF_REPO
module_file = YOUR_LOCAL_PATH_OF_FILE_IN_REPO
extractor = Extractor(reposrc, module_file)
output = extractor.extract()

```




Figure 9: Correlation of context length to the model performance on REPOEXEC. The size of the dots indicates the model size, while the color represents the model family.

Input Prompt & Solution	Generation output
<pre># Prompt: import base64 import random import unicodedata ...  class InvalidInputError(TypeError):     ...  def is_string(obj: Any) -&gt; bool:     ...  def reverse(input_string: str) -&gt; str:     """     Returns the string with its chars reversed.      *Example:*      &gt;&gt;&gt; reverse('hello') # returns 'olleh'      :param input_string: String to revert.     :type input_string: str     :return: Reversed string.     """</pre> <hr/> <pre># Solution: def reverse(input_string: str) -&gt; str:     """     Returns the string with its chars reversed.      *Example:*      &gt;&gt;&gt; reverse('hello') # returns 'olleh'      :param input_string: String to revert.     :type input_string: str     :return: Reversed string.     """     if not is_string(input_string):         raise InvalidInputError(input_string)      return input_string[::-1]</pre>	<pre># Model type: Pretrained # Passed: No def reverse(input_string: str) -&gt; str:     """     Returns the string with its chars reversed.      *Example:*      &gt;&gt;&gt; reverse('hello') # returns 'olleh'      :param input_string: String to revert.     :type input_string: str     :return: Reversed string.     """     return input_string[::-1]</pre> <hr/> <pre># Model type: Instruction-tuned # Passed: Yes def reverse(input_string: str) -&gt; str:     """     Returns the string with its chars reversed.      *Example:*      &gt;&gt;&gt; reverse('hello') # returns 'olleh'      :param input_string: String to revert.     :type input_string: str     :return: Reversed string.     """     if not is_string(input_string):         raise InvalidInputError(input_string)     return input_string[::-1]</pre>

<pre> # Prompt: import random import luigi import luigi.configuration ...  class hadoopcli(luigi.Config):     command = ...     version = ...  def get_configured_hadoop_version():     """     CDH4 (hadoop 2+) has a slightly different     syntax for interacting with hdfs     via the command line.      The default version is CDH4, but one can     override     this setting with "cdh3" or "apache1" in the     hadoop section of the config     in order to use the old syntax.     """ </pre> <hr style="border-top: 1px dashed black;"/> <pre> # Solution: def get_configured_hadoop_version():     """     CDH4 (hadoop 2+) has a slightly different     syntax for interacting with hdfs     via the command line.      The default version is CDH4, but one can     override     this setting with "cdh3" or "apache1" in the     hadoop section of the config     in order to use the old syntax.     """     return hadoopcli().version.lower() </pre>	<pre> # Model type: Pretrained # Passed: Yes def get_configured_hdfs_client():     """     This is a helper that fetches the     configuration value for 'client' in     the [hdfs] section. It will return the     client that retains backwards     compatibility when 'client' isn't configured     .     """     return luigi.configuration.get_config().get(         'hdfs', 'client', 'hadoopcli') </pre> <hr style="border-top: 1px dashed black;"/> <pre> # Model type: Instruction-tuned # Passed: No def get_configured_hadoop_version():     """     CDH4 (hadoop 2+) has a slightly different     syntax for interacting with hdfs     via the command line.      The default version is CDH4, but one can     override     this setting with "cdh3" or "apache1" in the     hadoop section of the config     in order to use the old syntax.     """     hadoop_cli = hadoopcli().command.split()     hadoop_version = hadoopcli().version     if hadoop_version == "cdh3":         return hadoop_cli + ["fs"]     elif hadoop_version == "apache1":         return hadoop_cli + ["dfs"]     else:         return hadoop_cli </pre>
---	--

<pre> # Prompt: import asyncio import builtins import collections ...  def is_future(x: Any) -&gt; bool:     return isinstance(x, FUTURES)  def _create_future() -&gt; Future:     future = Future() # type: Future     # Fixup asyncio debug info by removing     # extraneous stack entries     source_traceback = getattr(future, "_source_traceback", ())     while source_traceback:         # Each traceback entry is equivalent to         # a         # (filename, self.lineno, self.name,         # self.line) tuple         filename = source_traceback[-1][0]         if filename == __file__:             del source_traceback[-1]         else:             break     return future  def maybe_future(x: Any) -&gt; Future:     """Converts ``x`` into a ``Future``.      If ``x`` is already a ``Future``, it is     simply returned; otherwise     it is wrapped in a new ``Future``. This is     suitable for use as     ``result = yield gen.maybe_future(f())``     when you don't know whether     ``f()`` returns a ``Future`` or not.      .. deprecated:: 4.3     This function only handles ``Futures``, not     other yieldable objects.     Instead of ``maybe_future``, check for the non     -future result types     you expect (often just ``None``), and ``     yield`` anything unknown.     """ </pre> <hr/> <pre> # Solution: def maybe_future(x: Any) -&gt; Future:     """Converts ``x`` into a ``Future``.      If ``x`` is already a ``Future``, it is     simply returned; otherwise     it is wrapped in a new ``Future``. This is     suitable for use as     ``result = yield gen.maybe_future(f())``     when you don't know whether     ``f()`` returns a ``Future`` or not.      .. deprecated:: 4.3     This function only handles ``Futures``,     not other yieldable objects.     Instead of ``maybe_future``, check for the     non-future result types     you expect (often just ``None``), and ``     yield`` anything unknown.     """     if is_future(x):         return x     else:         fut = _create_future()         fut.set_result(x)         return fut </pre>	<pre> # Model type: Pretrained # Passed: Yes def maybe_future(x: Any) -&gt; Future:     """Converts ``x`` into a ``Future``.      If ``x`` is already a ``Future``, it is     simply returned; otherwise     it is wrapped in a new ``Future``. This is     suitable for use as     ``result = yield gen.maybe_future(f())``     when you don't know whether     ``f()`` returns a ``Future`` or not.      .. deprecated:: 4.3     This function only handles ``Futures``, not     other yieldable objects.     Instead of ``maybe_future``, check for the non     -future result types     you expect (often just ``None``), and ``     yield`` anything unknown.     """     if is_future(x):         return x     future = Future() # type: Future     future.set_result(x)     return future </pre> <hr/> <pre> # Model type: Instruction-tuned # Passed: Yes def maybe_future(x: Any) -&gt; Future:     """Converts ``x`` into a ``Future``.      If ``x`` is already a ``Future``, it is     simply returned; otherwise     it is wrapped in a new ``Future``. This is     suitable for use as     ``result = yield gen.maybe_future(f())``     when you don't know whether     ``f()`` returns a ``Future`` or not.      .. deprecated:: 4.3     This function only handles ``Futures``, not     other yieldable objects.     Instead of ``maybe_future``, check for the non     -future result types     you expect (often just ``None``), and ``     yield`` anything unknown.     """     if isinstance(x, Future):         return x     elif isawaitable(x):         return asyncio.ensure_future(x)     else:         future = Future()         future.set_result(x)         return future </pre>
--	---

Table 9: Output samples of Pretrained and Instruction-tuned models.