

TESTEVAL: Benchmarking Large Language Models for Test Case Generation

Wenhan Wang^{1*}, Chenyuan Yang^{3*}, Zhijie Wang^{1*}, Yuheng Huang²,
Zhaoyang Chu⁴, Da Song¹, Lingming Zhang³, An Ran Chen¹, Lei Ma^{2,1}

¹University of Alberta, ²The University of Tokyo, ³University of Illinois Urbana-Champaign,

⁴Huazhong University of Science and Technology

wenhan12@ualberta.ca cy54@illinois.edu zhijie.wang@ualberta.ca
yuhenghuang42@g.ecc.u-tokyo.ac.jp chuzhaoyang@hust.edu.cn dsong4@ualberta.ca
lingming@illinois.edu anran6@ualberta.ca ma.lei@acm.org

Abstract

For program languages, testing plays a crucial role in the software development cycle, enabling the detection of bugs, vulnerabilities, and other undesirable behaviors. To perform software testing, testers need to write code snippets that execute the program under test. Recently, researchers have recognized the potential of large language models (LLMs) in software testing. However, there remains a lack of fair comparisons between different LLMs in terms of test case generation capabilities.

In this paper, we propose TESTEVAL, a novel benchmark for test case generation with LLMs. We collect 210 Python programs from an online programming platform, LeetCode, and design three different tasks: overall coverage, targeted line/branch coverage, and targeted path coverage. We further evaluate 17 popular LLMs, including both commercial and open-source ones, on TESTEVAL. We find that generating test cases to cover specific program lines/branches/paths is still challenging for current LLMs, indicating a lack of ability to comprehend program logic and execution paths. We have open-sourced our dataset and benchmark pipelines at <https://github.com/LLM4SoftwareTesting/TestEval>.

1 Introduction

Software testing is a crucial aspect of software development, allowing developers to identify potential bugs and check if the program behavior meets expectations. A key task in software testing is test case generation, which involves creating test inputs to cover different statements and branches in the program under test. Previous research indicates that test case generation can be extremely time-consuming, taking up over 15% of the time spent in software development (Daka and Fraser, 2014).

Therefore, automated test case generation has been a long-standing challenge in software engineering research. Various methods have been developed to address this issue, including symbolic execution testing (Chipounov et al., 2011; Cadar et al., 2011), search-based testing (Fraser and Arcuri, 2011; Baresi and Miraz, 2010; Fraser and Zeller, 2010), and deep learning-based approaches (Tufano et al., 2020). Recently, researchers have been exploring the potential of using LLMs to generate unit test cases (Lemieux et al., 2023; Xie et al., 2023; Yuan et al., 2023). However, despite the rapid development of LLM-based test case generation, there is still a lack of public benchmarks to evaluate different LLMs’ capabilities in this area. Hence, there is a need for a comprehensive analysis to determine whether current LLMs can (1) generate diverse test cases to achieve high coverage on a program under test, (2) generate test cases to cover a specific line or branch, and (3) generate test cases to cover a specific execution path by following the tester’s intent.

To bridge this gap, we present a new benchmark, TESTEVAL, which focuses on evaluating LLMs’ test case generation capabilities. The TESTEVAL dataset consists of 210 Python programs from the online coding platform LeetCode. We design three tasks to address the aforementioned challenges: (1) overall coverage, (2) targeted line/branch coverage, and (3) targeted path coverage.

Notably, unlike popular code generation benchmarks (Chen et al., 2021; Austin et al., 2021) or software testing datasets (Just et al., 2014; Lemieux et al., 2023), the tasks in our TESTEVAL benchmark require LLMs to reason about intricate program execution behaviors. To generate inputs that invoke specific branches or paths in the program under test, the LLM must have a comprehensive understanding of how to satisfy certain branch conditions during execution. Furthermore, our tasks emphasize program logic analysis rather than

*These authors contributed equally to this work.



Figure 1: The pipeline for running and evaluating LLMs for test case generation on TESTEVAL.

merely simulating numerical operations, as seen in benchmarks designed for predicting a program’s input/output (Gu et al., 2024).

We perform extensive experiments on TESTEVAL with both commercial and open-source LLMs. Our results indicate that while state-of-the-art LLMs can generate executable and diverse test cases, they struggle to identify which specific statements or branches need to be covered. For example, in targeted line coverage, 12 out of 16 LLMs’ performances are not significantly improved (improvements $\leq 5\%$) compared to the results when the target line information is even **not** given. Quantitative results show that commercial LLMs, such as GPT-4, generally outperform open-source LLMs in both overall coverage and targeted line/branch/path coverage. These findings suggest that future work on test case generation should focus on developing advanced LLM-based reasoning frameworks to enhance the understanding of program behaviors during testing.

Our work makes the following contributions:

- **Benchmark.** We propose TESTEVAL, a benchmark focused on evaluating LLMs’ capabilities in generating test cases for a given program under test, encompassing three different tasks.
- **Evaluation.** We design new metrics to measure the LLM’s test generation performance and conduct experiments with 17 popular LLMs.
- **Analysis.** We perform a systematic analysis of LLMs’ performance on TESTEVAL and discuss the challenges and opportunities in test case generation using LLMs.

2 Approach

In this section, we first introduce the tasks included in our benchmark (§ 2.1). Following that, we provide an overview of the dataset used (§ 2.2).

2.1 Task Description

Figure 1 shows the workflow of TESTEVAL. We propose three distinct tasks in our benchmark: (1) overall coverage, (2) targeted line and branch coverage, and (3) targeted path coverage. For each task, we prompt an LLM to generate test cases for a specified program based on the task description in natural language. Specifically, in each query round, we prompt the LLM to generate a testing function containing a single test case (see Appendix A for the complete prompt templates). Then, we filter out any non-code content that may have been generated outside the testing function, retaining only the first test case generated in each query round to ensure a fair comparison across different LLMs.

After generation, all test cases must undergo a correctness check, which consists of *syntactic correctness*, *execution correctness*, and *assertion correctness*. Syntactic correctness determines if the generated test case is free of syntax errors, while execution correctness evaluates if the test case can be executed successfully without any runtime errors. Assertion correctness evaluates whether the generated test case contains correct test assertions. Regarding execution correctness, we do not consider incorrect test assertion statements to be failed cases, since test cases with assertion errors can still cover the program under test. Finally, we evaluate coverage metrics on test cases that pass the correctness check. We now illustrate our three benchmark tasks in detail.

Algorithm 1: Computing the average line/branch $cov@k$ given a set of programs

Input: A set of programs under test $\mathcal{P} = \{p_1, p_2, \dots\}$, k
Output: The average $cov@k$ for all programs: $cov@k_{all}$

```

 $cov@k = [];$ 
for  $p_i$  in  $\mathcal{P}$  do
    Generate  $N$  test cases  $T_i = \{t_{i1}, t_{i2}, \dots, t_{iN}\};$ 
    Retain  $M$  executable test cases  $T_i = \{t_{i1}, t_{i2}, \dots, t_{iM}\};$ 
    if  $T_i = \emptyset$  then
         $cov@k.append(0);$ 
    else
        Randomly split  $T_i$  into  $\max(\lfloor M/k \rfloor, 1)$  groups, each
        group  $T_{ij}$  with  $k$  test cases;
         $cov_i = [];$ 
        for  $T_{ij}$  in  $\{T_{i1}, T_{i2}, \dots, T_{i\lfloor M/k \rfloor}\}$  do
            Compute line/branch coverage  $cov_{T_{ij}};$ 
             $cov_i.append(cov_{T_{ij}});$ 
        end
         $cov@k.append(avg(cov_i))$ 
    end
 $cov@k_{all} \leftarrow avg(cov@k);$ 
return  $cov@k_{all}$ 

```

Overall coverage. In this task, we query each LLM for N rounds given a program under test. During the i th ($1 < i \leq N$) round, we prompt the

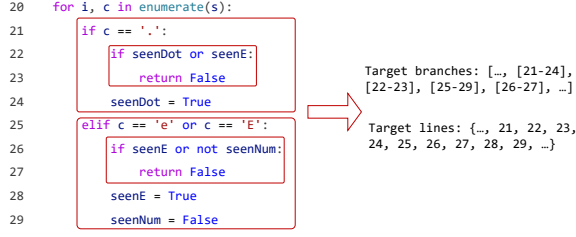


Figure 2: An example for selecting targeted lines and branches from programs under test.

LLM to generate a test method different from the j ($1 \leq j < i$)th rounds. After all rounds of query, we obtain N test cases for each program under test. The overall coverage for a program is computed by the proportion of lines/branches in the program that have been covered by at least one test case.

We further propose a new metric, $cov@k$, to measure the diversity of LLM’s generated test cases for a given program. Intuitively, $cov@k$ measures the line/branch coverage with a subset of the generated test cases with a size of k ($k < M$). To achieve this, we randomly split M executable test cases into $\max(\lfloor M/k \rfloor, 1)$ subsets. Then for each of these subsets, we calculate its overall line/branch coverage. In our experiments, we choose k as 1, 2, and 5. When k increases, the improvements of $cov@k$ can measure the diversity of the LLM’s generated test cases. We summarize the calculation of the average line/branch $cov@k$ for a set of programs under test in Algorithm 1.

Targeted line and branch coverage. Different from overall coverage, *targeted line and branch coverage* requires the LLM to generate test cases that could cover a specific branch, or a line inside this branch. This simulates the scenario in which a human tester is asked to craft test cases to cover a specific part of the program. Figure 2 shows an example of targeted branches and lines in a given program. To measure the targeted line/path coverage, we prompt the LLMs by including the line number(s) in the instruction (see prompt templates in Appendix A). For each targeted line/branch, an LLM is prompted to generate one test case.

Targeted path coverage. In real-world software development, testers sometimes need to craft test cases to cover a specific execution path that includes multiple branches. We refer to this task as the *target path coverage*. We show an example program in Figure 3 to demonstrate the importance of the target path coverage.

In Line 6, a bug (divided by zero) will occur only if branches “condition 1” and “condition 2” are both **not** executed. In this case, only covering the two conditional branches is not sufficient. By contrast, if we can cover all three paths (Figure 3), we can successfully detect the “divided by zero” bug. To obtain the target path coverage, we prompt an LLM by including a specific execution path (see Appendix A for the prompt template). For each path, an LLM is queried to generate one test case.

We further propose two metrics to evaluate the performance of target path coverage. First, for a given target path, we measure whether the generated test case covers the target path completely. Second, we measure the similarity between the given target path $PATH_{tgt}$ and the execution path of the LLM’s generated test case $PATH_{gen}$ by Eq. 1:

$$sim(PATH_{gen}, PATH_{tgt}) = \frac{lcs(PATH_{gen}, PATH_{tgt})}{len(PATH_{tgt})} \quad (1)$$

where $lcs()$ calculates the length of the longest contiguous common sub-sequence between two paths and $len()$ calculates the length of a path.

2.2 Benchmark Dataset

Data collection. To construct our benchmark dataset, we first collect solution programs of LeetCode, an online platform for evaluating a programmer’s coding performance. We choose LeetCode as our data source since it has a clear task description and input constraint for each programming task. We first select all publicly available tasks up to Apr. 2024. Then, we collect its Python solution for each task from a GitHub repository¹. At this stage, we collect 3,123 programs under test.

The main goal of our benchmark is to evaluate LLMs’ capability to generate test cases that cover specific statements/branches. Therefore, we filter out programs that are too simple (e.g., programs that only have one branch) according to the cyclomatic complexity (McCabe, 1976). Given the control flow graph of a program, the cyclomatic complexity V of this program is measured by: $V = e - n + p$, where e is the number of edges in the graph, n is the number of nodes, and p is the number of connected components. The cyclomatic complexity is positively correlated with the number of branches/loops in a program. In this work, we consider programs with the cyclomatic complexity

¹<https://github.com/walkccc/LeetCode>. The repository is under MIT license.

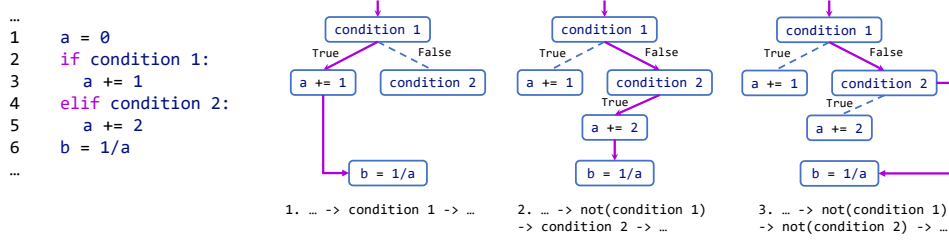


Figure 3: A motivating example showing the importance of path coverage (left), and examples of execution paths extracted from this program (right).

≥ 10 . This filters down the sample size from 3,123 to 216. We further check these 216 problems and remove similar problems with identical solutions. Finally, we collect 210 Python programs for our benchmark, consisting of 9 easy problems, 100 medium problems, and 101 hard problems according to LeetCode’s difficulty label. Each program under test is also paired with its task description in natural language. Note that most programs already have test cases in their task descriptions. We remove these cases to prevent LLMs from directly copying these test cases.

For each program, we perform the following preprocessing steps:

- We add all necessary import statements for the packages required by the Python solution.
- Python programmers often split long statements into multiple physical lines. For all statements split into multiple lines, we rewrite them in a single line. This ensures each statement only corresponds to one line when measuring line coverage.
- We reformat the in-line conditional statements (e.g., the ternary conditional operator) into multi-line blocks. This ensures that each line of the program is one statement that belongs to one specific branch.
- We remove all natural language comments.

Targeted line/branch/path identification. To obtain targeted lines/branches, we first extract all *conditional* branches of a given program based on its abstract syntax tree (AST). Since loop branches (i.e., for/while loops) are usually easy to cover, we only consider conditional branches in our task. Specifically, we extract all if, elif, and else branches. We refer to these branches as our targeted branches. Then, we consider all statements within

these targeted branches as our targeted lines (see Figure 2 for the example). Overall, we identified 983 target branches in 210 programs under test (4.7 target branches per program on average). The total number of target lines in 210 programs is 1,312 (6.2 target lines per program on average). The detailed algorithm for extracting target lines/branches can be found in Appendix B.

In a given program under test, certain branches could be hard to cover without carefully crafting the test cases. Therefore, we label each targeted branch in a program as easy, medium, or hard according to the average coverage after executing 100 randomly generated inputs. For each problem in TESTEVAL, we construct a *random input generator* to assess the difficulties of covering a specific branch. Each generator is a Python program that *uniformly* samples a valid test input for the LeetCode problem according to its constraint description. We leverage GPT-4 to generate these generators from the constraints in LeetCode problem descriptions. Then, we perform manual inspection and correction to ensure they adhere to the problem’s constraints. An example of an input generator is shown in Figure 4. These generators are then used to sample 100 executable test cases for each problem. Branch difficulty is determined by the frequency at which a branch is covered across these 100 sampled test inputs. We categorize branches as follows: easy (covered by [99%, 100%] of test cases), medium (covered by [40%, 99%] of test cases), and hard (covered by [0, 40%) of test cases). This partitioning ensures that easy branches do not significantly outnumber other categories, and promotes a balanced distribution between medium and hard branches. The number of easy, medium, and hard target branches are 498, 225, and 260.

For the targeted path coverage task, as the number of execution paths in a program can be enormous or even undecidable, it is impossible to col-

```

Constraints:
* `nums1.length == m`
* `nums2.length == n`
* `0 <= m <= 1000`
* `0 <= n <= 1000`
* `1 <= m + n <= 2000`
* `-10^6 <= nums1[i], nums2[i] <= 10^6`

def generate_random_input():
    # Define the constraints
    min_len = 1
    max_len = 1000
    min_val = -pow(10, 6)
    max_val = pow(10, 6)
    # Generate random length for two arrays within the constraints
    len_nums1 = random.randint(min_len, max_len)
    len_nums2 = random.randint(min_len, max_len)
    # Make sure the total length is within the limits
    while len_nums1 + len_nums2 > 2000:
        len_nums1 = random.randint(min_len, max_len)
        len_nums2 = random.randint(min_len, max_len)
    # Create two lists with random int within the given value constraints
    nums1 = sorted([random.randint(min_val, max_val) for _ in range(len_nums1)])
    nums2 = sorted([random.randint(min_val, max_val) for _ in range(len_nums2)])
    # Return the resulting lists as a tuple
    return nums1, nums2

```

Figure 4: The input constraints for a LeetCode problem (left) and its random input generator for TESTEVAL (right).

lect all execution paths. Instead, we collect the target execution paths from the example test cases given by LeetCode problem descriptions. For each example test case, we execute it and record its execution path using all the condition/loop branches it executed. The complete execution path would be too long and difficult for LLMs to understand, so we perform clipping after obtaining full paths. For each execution path, we randomly sample two sub-paths with lengths of 5 consecutive branches taken. We further remove duplicated sampled paths, resulting in an average of 4.1 target paths per problem.

3 Evaluation

3.1 Experiment Setup

We evaluate 17 popular instruction-following LLMs, including both commercial and open-source ones. The parameter sizes of open-source models range from 1.3B to 34B. The temperature is set to 0 or 1e-5 (for models on Huggingface that do not support temperature=0) to ensure that the evaluation results can be reproduced. All experiments on open-source LLMs are run on two NVIDIA A6000 GPUs. We set the length limit of outputs to 256 tokens. We use the `pytest-cov` (`pytest cov`) to measure the code coverage.

3.2 Overall Coverage

In this experiment, we query every model 20 rounds ($N = 20$) to generate test cases (one test case per round) for each program under test. Table 1 shows the evaluation results on the overall coverage task.

Regarding correctness metrics, we observe that most models can achieve high syntactical and acceptable execution correctness, but all models have much lower assertion correctness. For test cases that do not pass the execution correctness check, we perform a preliminary study in the Appendix C. Regarding the coverage performance, most of

the LLMs are able to generate test cases that cover over 80% lines/branches per program under test. Notably, the latest GPT-4o achieves the best overall line (98.65%) and branch (97.16%) coverage. We also notice that the open-source model, DeepSeek-coder-33b, outperforms the commercial LLM, Gemini-1.0-pro, on both overall line and branch coverage.

We further use $cov@k$ to measure the diversity of each LLM’s generated test cases. Similar to the overall coverage results, GPT-4o has the best line and branch $cov@1$, demonstrating its ability to craft complex test cases that are able to cover most of the program branches within a single attempt. We also find all LLMs have a higher $cov@2$ and $cov@5$ compared with $cov@1$. This indicates that the LLMs are able to generate different test cases. Gemma-7b shows the most significant improvements in the line (+9.67%) and branch (+13.14%) $cov@5$ compared with its line and branch $cov@1$. We also notice that Starcoder-2-Instruct has the least improvement on $cov@5$ compared with $cov@1$ (+0.47% and +0.70% for line and branch coverage, respectively). By manually checking the test cases generated by Starcoder-2-Instruct, we find that it frequently repeats previously generated cases despite being instructed to generate different ones.

3.3 Targeted Line and Branch Coverage

Table 2 and Table 3 show the evaluation results for the targeted line and branch coverage, respectively. For each subject LLM, we also include a baseline by excluding the information about the targeted lines/branches in the text prompt. For each program under test, we reuse the first test case generated for the overall coverage task and measure its coverage accuracy on the targeted lines/branches. The intuition is that, if an LLM could not outperform the baseline, it might be struggling with identifying the line/branch that is expected to cover when

Table 1: Result on the overall coverage task. The results in parenthesis are the improvements over $cov@1$.

Model	Size	Correctness			Overall coverage		Line $cov@k$			Branch $cov@k$		
		syntax	execution	assertion	line	branch	$k = 1$	$k = 2$	$k = 5$	$k = 1$	$k = 2$	$k = 5$
GPT-3.5-turbo	N/A	100	97.43	40.40	96.27	93.65	88.35	90.02 (1.67)	92.14 (3.79)	81.87	84.32 (2.45)	87.55 (5.68)
GPT-4	N/A	100	92.33	54.16	94.94	92.81	85.65	87.77 (2.12)	90.04 (4.39)	78.89	81.93 (3.04)	85.39 (6.50)
GPT-4-turbo	N/A	100	94.79	56.24	96.08	94.81	85.46	87.87 (2.41)	90.81 (5.35)	78.62	82.06 (3.44)	86.64 (8.02)
GPT-4o	N/A	99.59	98.30	52.99	98.65	97.16	90.23	92.16 (1.93)	94.33 (4.10)	84.05	86.89 (2.84)	90.31 (6.26)
GPT-4o-mini	N/A	100	99.92	43.86	98.76	97.58	88.06	90.33 (2.27)	93.51 (5.45)	81.64	85.03 (3.39)	89.60 (7.96)
Gemini-1.0-pro	N/A	93.05	71.93	35.31	93.01	90.66	84.48	86.60 (2.12)	88.47 (3.99)	78.35	81.29 (2.94)	84.11 (5.76)
CodeLlama	7b	99.52	73.86	31.07	86.09	81.56	79.46	80.72 (1.26)	82.04 (2.58)	72.28	73.96 (1.68)	75.90 (3.62)
	13b	67.55	50.40	25.28	85.66	80.55	80.49	82.26 (1.77)	83.44 (2.95)	73.21	75.54 (2.33)	77.13 (3.92)
	34b	66.33	46.86	40.32	87.96	83.74	78.83	81.25 (2.42)	83.71 (4.88)	71.37	74.50 (3.13)	77.80 (6.43)
Llama3	8b	99.25	82.24	44.61	90.98	89.02	77.40	80.08 (2.68)	84.42 (7.02)	69.47	73.37 (3.90)	79.22 (9.75)
Llama3.1	8b	98.69	94.69	50.00	88.94	85.79	74.42	77.49 (3.07)	82.07 (7.65)	65.65	69.92 (4.27)	76.16 (10.51)
Gemma	7b	98.98	64.64	35.30	93.16	91.46	76.23	80.54 (4.31)	85.90 (9.67)	67.15	72.94 (5.79)	80.29 (13.14)
StarCoder-2-Instruct	15b	97.07	94.07	54.11	89.84	84.41	88.03	88.22 (0.19)	88.50 (0.47)	81.80	82.09 (0.29)	82.50 (0.70)
DeepSeek-coder	1.3b	96.05	82.48	38.66	81.22	75.99	75.89	76.50 (0.61)	77.09 (1.20)	69.06	69.90 (0.84)	70.70 (1.64)
	6.7b	97.42	82.43	40.43	93.48	91.61	82.40	84.74 (2.34)	87.97 (5.57)	75.29	78.73 (3.44)	83.46 (8.17)
	33b	99.21	83.57	50.75	94.86	91.92	85.47	87.38 (1.91)	90.30 (4.83)	78.49	81.23 (2.74)	85.12 (6.63)
CodeQwen	7b	100	84.26	46.36	90.73	86.90	84.53	85.33 (0.80)	86.71 (2.18)	77.66	78.94 (1.28)	80.95 (3.29)

Table 2: Results for targeted line coverage. Results in parenthesis are the improvements over baselines.

Model	Size	Targeted line				Baseline: no targeted line		
		syntax	execution	assertion	cov. Recall	Syntax	execution	cov. Recall
GPT-3.5-turbo	N/A	99.40	95.67	41.93	67.76 (-1.27)	100	100	69.03
GPT-4	N/A	100	98.81	61.22	78.20 (10.14)	100	99.52	68.06
GPT-4-turbo	N/A	99.20	98.73	67.29	80.52 (11.64)	100	100	68.88
GPT-4o	N/A	99.63	98.96	67.52	80.97 (9.48)	100	100	71.49
GPT-4o-mini	N/A	100	99.92	56.02	76.94 (8.73)	100	100	68.21
Gemini-1.0-pro	N/A	100	96.04	53.37	70.75 (4.93)	100	95.71	65.82
CodeLlama	7b	99.85	90.97	34.04	58.13 (0.89)	99.52	93.81	57.24
	13b	99.63	85.22	48.42	54.63 (-4.03)	99.05	94.76	58.66
	34b	98.66	90.60	44.34	59.48 (-0.29)	100	96.19	59.77
Llama3	8b	98.96	85.52	37.08	60.22 (-0.60)	99.52	95.24	60.82
Llama3.1	8b	99.25	98.43	48.56	56.49 (-3.36)	99.05	88.10	59.85
Gemma	7b	99.78	88.21	33.28	62.91 (4.92)	99.52	89.52	57.99
StarCoder-2-Instruct	15b	98.36	92.84	57.14	64.40 (-2.39)	100	99.05	66.79
DeepSeek-coder	1.3b	98.81	91.04	41.05	58.81 (2.69)	94.76	90.0	56.12
	6.7b	94.78	92.99	45.09	65.60 (3.81)	99.05	96.67	61.79
	33b	99.63	97.61	59.29	70.52 (2.09)	100	99.52	68.43
CodeQwen	7b	94.78	92.99	61.05	65.60 (3.81)	99.05	96.67	61.79

generating the test case.

Regarding the targeted line coverage, we find that the GPT-4 series has the best performance improvement (around 10%) over their baselines. The best-performing LLM is GPT-4o, reaching a coverage accuracy of 80.97% on average. We also find that six out of seventeen LLMs do not improve over their baseline and seven LLMs only have marginal improvement (less than 5%). These results suggest that most LLMs may have trouble with multi-step reasoning. Specifically, to reach a specific line inside a branch, the LLM needs first to identify which branch the targeted line belongs to and then generate a valid test input to invoke this branch.

We observe a similar trend in the targeted branch coverage (Table 3). Specifically, the GPT-4 series has the best performance improvement (12%~15%) over their baselines. GPT-4o is the best-performing LLM, which can cover 80.87% branches, respec-

tively. By contrast, eight LLMs only exhibit marginal improvements and four LLMs do not improve compared with the baselines. Regarding branches with different difficulties, we find that branches more likely to be covered by random test cases are also easier for LLMs to cover (recall we use random testing to label each branch’s difficulty level in § 2.2). The GPT-4 series shows the smallest performance gap between branches with different difficulty levels. We also notice that twelve out of sixteen LLMs show the largest performance improvements over the baselines on “hard” branches. These results indicate that providing target branch information can indeed help us to cover branches that are hard to reach by random testing.

3.4 Targeted Path Coverage

Table 4 presents the results of the targeted path coverage task. We adopt a similar baseline as in

Table 3: Results for targeted branch coverage. Results in parenthesis are the improvements over the baseline. We omit the correctness metrics of the baseline because they are the same as the targeted line coverage task.

Model	Size	Targeted branch								Baseline: no targeted branch			
		Correctness			Coverage				Coverage				
		syntax	execution	assertion	total		easy	medium	hard	total	easy	medium	hard
GPT-3.5-turbo	N/A	100	98.78	47.62	70.40 (4.38)	82.93 (0.40)	65.33 (1.77)	50.77 (14.23)		66.02	82.53	63.56	36.54
GPT-4	N/A	100	98.17	61.88	78.23 (13.33)	86.14 (4.41)	79.56 (16.89)	61.92 (27.30)		64.90	81.73	62.67	34.62
GPT-4-turbo	N/A	100	98.67	67.60	80.77 (15.15)	88.35 (5.42)	79.11 (16.00)	67.69 (33.07)		65.62	82.93	63.11	34.62
GPT-4o	N/A	100	99.08	68.74	80.87 (12.61)	87.55 (3.21)	83.11 (11.55)	66.15 (31.53)		68.26	84.34	71.56	34.62
GPT-4o-mini	N/A	100	99.39	57.65	78.13 (12.01)	87.35 (4.02)	77.33 (13.33)	61.15 (26.15)		66.12	83.33	64.00	35.00
Gemini-1.0-pro	N/A	100	97.04	55.43	68.97 (5.80)	82.13 (2.21)	69.78 (6.22)	43.08 (12.31)		63.17	79.92	63.56	30.77
CodeLlama	7b	100	81.99	40.57	50.97 (-4.17)	64.25 (-8.04)	51.11 (-3.56)	25.38 (2.69)		55.14	72.29	54.67	22.69
	13b	99.29	82.91	52.86	51.58 (-4.68)	64.86 (-7.83)	46.67 (-11.55)	30.39 (5.78)		56.26	72.69	58.22	24.61
	34b	99.39	95.02	42.12	63.17 (5.69)	78.51 (2.20)	60.44 (6.22)	36.15 (11.92)		57.48	76.31	54.22	24.23
Llama3	8b	98.88	84.94	37.93	58.39 (-0.31)	73.09 (-0.61)	59.11 (0.89)	29.26 (-1.12)		58.70	73.70	58.22	30.38
Llama3.1	8b	99.49	85.86	48.20	58.09 (-0.71)	69.08 (-6.42)	57.33 (2.22)	37.69 (7.69)		58.80	75.50	55.11	30.00
Gemma	7b	99.59	85.35	37.47	56.15 (1.11)	71.89 (2.01)	49.78 (0.45)	31.54 (0.00)		55.04	69.88	49.33	31.54
StarCoder-2-Instruct	15b	98.68	95.42	64.63	64.19 (-0.41)	78.71 (0.20)	63.56 (-1.33)	36.92 (-0.77)		64.60	78.51	64.89	37.69
DeepSeek-coder	1.3b	97.05	89.32	41.11	54.22 (0.81)	68.67 (1.20)	52.89 (-1.78)	27.69 (2.31)		53.41	67.47	54.67	25.38
	6.7b	96.74	93.79	43.91	66.43 (7.22)	77.11 (5.62)	69.33 (4.89)	43.46 (12.31)		59.21	71.49	64.44	31.15
	33b	100	97.05	55.43	68.46 (2.54)	80.12 (-2.21)	66.22 (4.00)	48.08 (10.39)		65.92	82.33	62.22	37.69
CodeQwen	7b	99.49	95.02	61.76	65.82 (0.51)	81.12 (1.60)	63.56 (-3.55)	38.46 (1.92)		65.31	79.52	67.11	36.54

our targeted line/branch coverage tasks by excluding the targeted path in the text prompts. Overall, GPT-4o and Gemini-1.0-pro have the best performance on the path coverage, reaching 56.67% and 56.09% on average, respectively. However, they do not outperform their baselines. Generally, we do not find any LLMs that show obvious performance improvement (more than 5%) on the path coverage compared with the baselines. Nine out of sixteen LLMs do not outperform the baselines. Regarding the path similarity, we also do not find any LLMs exhibiting large performance improvement compared with the baselines. These results suggest that comprehending the program logic and identifying a specific execution path is still a challenging task for the current LLMs.

Targeted path coverage is considerably more complicated compared with overall coverage and targeted line/branch coverage. Specifically, the LLM needs to identify a sequence of multiple branches, and create a test input that can execute these branches following a certain order, which is challenging even for human programmers.

3.5 Advanced Prompting

Advanced prompting techniques, such as in-context learning (Brown, 2020) and chain-of-thought (COT) (Wei et al., 2022), can improve the performance of LLMs on language understanding and generation. We further conduct a study on the influence of different prompting strategies on TESTEVAL. In this advanced prompt setting, we adopt

an explicit two-step COT for the targeted line coverage task. LLMs are first asked to generate the conditions that need to be satisfied when the target line is executed. Then, we ask LLMs to generate a test case that satisfies these conditions. We provide a one-shot example of the reasoning process, which is created from the solution of a LeetCode easy-level problem (not included in the TESTEVAL dataset). The complete prompt template for this setting is shown in Appendix A.5.

Table 5 shows the results of our COT prompting on the targeted line coverage task. Because the cost of COT is significantly higher than basic prompting, we only run experiments on several cost-efficient models, and omitted expensive proprietary models or large open-source models. For most models (except GPT-4o-mini and DeepSeek-coder 6.7b), COT can improve the performances on target line coverage. This suggests that building more complex LLM pipelines or agents for test case generation is worth investigating in the future.

With the two-step COT setting, we can have a detailed analysis of the reason behind failures in generated test cases. Figure 5 demonstrates a test case generated by GPT-4o that failed to cover the target line: line 33. We find that although the LLM is capable of generating correct conditions (Figure 5 (b)) for covering the target line, the generated test case did not satisfy those conditions, suggesting that the LLM’s code generation ability needs further improvement. In this case, the generated test case (Figure 5 (c)) does not satisfy the condi-

Table 4: Results for targeted path coverage. Results in parenthesis are the improvements over the baseline.

Model	Size	Given target path					Baseline: no target path	
		syntax	execution	assertion	path cov	path similarity	path cov	path similarity
GPT-3.5-turbo	N/A	99.88	98.95	49.58	49.30 (-5.97)	77.35 (-2.39)	55.27	79.74
GPT-4	N/A	100	99.18	61.71	54.10 (-0.94)	80.77 (3.23)	55.04	77.54
GPT-4-turbo	N/A	100	99.41	63.00	50.47 (-3.74)	79.82 (1.08)	54.21	78.74
GPT-4o	N/A	100	99.53	70.62	56.67 (-1.76)	82.35 (1.29)	58.43	81.06
GPT-4o-mini	N/A	100	99.77	52.93	51.87 (-2.58)	80.09 (0.67)	54.45	79.42
Gemini-1.0-pro	N/A	100	96.02	54.51	56.09 (0.70)	77.59 (-0.23)	55.39	77.82
CodeLlama	7b	99.76	90.98	39.63	41.57 (-1.05)	67.66 (0.14)	42.62	67.52
	13b	99.18	94.15	44.80	40.28 (-4.10)	64.63 (-4.98)	44.38	69.61
	34b	98.95	96.25	41.11	48.01 (2.93)	72.33 (2.69)	45.08	69.64
Llama3	8b	98.24	89.46	33.72	41.92 (1.29)	68.03 (0.40)	40.63	67.63
Llama3.1	8b	99.88	95.78	39.64	44.02 (-1.41)	72.51 (4.24)	45.43	68.27
Gemma	7b	100	88.06	29.99	37.11 (4.09)	64.54 (2.29)	33.02	62.25
StarCoder-2-Instruct	15b	96.83	90.28	47.86	48.48 (-5.38)	70.91 (-6.78)	53.86	77.69
DeepSeek-coder	1.3b	97.89	88.99	40.27	40.16 (0.46)	64.91 (0.67)	39.70	64.24
	6.7b	99.06	95.90	43.49	53.04 (0.23)	76.77 (1.56)	52.81	75.21
	33b	100	96.49	63.28	54.10 (-4.33)	77.99 (-2.73)	58.43	80.72
CodeQwen	7b	99.77	94.96	62.87	55.97 (-3.16)	77.46 (-2.67)	59.13	80.13

Table 5: Results for two-step COT prompting on targeted line coverage. The results in parenthesis are the improvements over the basic prompting setting.

Model	Size	syntax	execution	assertion	line coverage
GPT-3.5-turbo	N/A	99.70	98.13	47.86	71.79 (4.03)
GPT-4o	N/A	100	98.66	63.37	84.85 (3.88)
GPT-4o-mini	N/A	100	97.84	49.73	76.72 (-0.22)
Llama3	8b	99.93	87.99	35.37	62.54 (2.32)
Llama3.1	8b	99.70	97.24	38.14	64.48 (7.99)
DeepSeek-coder	6.7b	99.92	96.79	47.82	65.07 (-0.53)

```

import math
import itertools
import numpy
import collections
import sys
import string
import random
import unittest

from typing import List, Dict, Tuple, Iterator

class Solution:
    def twoSum(self, nums: List[int]) -> List[int]:
        for i in range(0, len(nums)-1):
            return i, j

ans = []

def test():
    for i in range(len(nums)-2):
        for j in range(i+1, len(nums)-1):
            continue
            i = i + 1
            j = j + 1
            while i < j:
                sum = nums[i] + nums[j]
                if sum == target:
                    ans.append([nums[i], nums[j], i, j])
                    i = i + 1
                    j = j - 1
                elif sum < target:
                    i = i + 1
                else:
                    j = j - 1
            return ans

return ans

```

Figure 5: Example of a generated test case that failed to cover the target line. (a): the program under test. (b): LLM-generated reasoning steps. (c): LLM-generated test cases based on reasoning steps.

tion 'nums[1] == nums[1 - 1]'.

4 Related Work

Code-related Benchmarks for LLMs. In recent years, researchers have endeavored to develop more rigorous and comprehensive evaluation frameworks for LLMs on coding abilities from various perspectives. One of the earliest attempts is HumanEval (Chen et al., 2021), which consists of

164 hand-craft programming challenges that evaluate LLMs’ ability to understand natural language descriptions and generate the corresponding functional correct code. Since then, there have been several studies attempting to construct benchmarks with more diverse problems (Austin et al., 2021), more rigorous evaluations (Liu et al., 2024a), and more complex scenarios (Lai et al., 2023; Zheng et al., 2023; Li et al., 2024b). Beyond these established code-generation scenarios, numerous studies are expanding their focus to include a broader range of real-world applications, such as reviewing code (Li et al., 2022), performing repo-level code completion (Liu et al., 2023; Zhang et al., 2023a; Guo et al., 2023; Ding et al., 2024), and resolving GitHub issues (Jimenez et al., 2023). While all the aforementioned studies examine the coding abilities of LLMs from different perspectives, none specifically target test case generation, a crucial phase in the software engineering lifecycle. The most relevant study is DevBench (Li et al., 2024a), which evaluates LLMs across software development stages, including testing. Unlike DevBench, our benchmark provides more comprehensive evaluations specifically tailored to test case generation using coverage-guided tasks and includes a broader range of studied models.

LLMs for Software Testing. Recent studies have extensively utilized LLMs to develop efficient and effective testing pipelines for various software applications (Xia et al., 2023; Wang et al., 2024a). Unit test case generation (Schäfer et al., 2023), which aims to test individual software components independently, is the primary focus of current LLM-

aided software testing. One line of research tries to pre-train/fine-tune LLMs on focal methods and related assertion statements to enhance their test-generation capabilities (Alagarsamy et al., 2023; Hashtroudi et al., 2023; Rao et al., 2023; Steenhoek et al., 2023). Although effective, these methods can be cost-intensive and challenging to scale. Alternatively, some researchers focus on crafting effective prompts that instruct LLMs to analyze relevant information (Yuan et al., 2023; Xie et al., 2023; Zhang et al., 2023b; Li and Doiron, 2023; Dakhel et al., 2024; Ryan et al., 2024; Liu et al., 2024b; Pizzorno and Berger, 2024; Wang et al., 2024b) or documentation (Vikram et al., 2023; Plein et al., 2024), or integrate LLMs with traditional software testing tools (Lemieux et al., 2023).

5 Conclusion

We present TESTEVAL, a novel benchmark for evaluating automated test case generation with LLMs for Python programs. Based on this dataset, we propose three different tasks and standardized evaluation pipelines. Our targeted coverage tasks enable the assessment on the LLM’s capabilities in comprehending complex program logic and execution path and generating test cases following the tester’s intent, which is not considered in previous works on either code generation or test case generation with LLMs.

We further conduct extensive experiments with seventeen popular LLMs on TESTEVAL. We find that although LLMs can achieve high overall coverage by generating diverse test cases, generating test cases to cover a specific element is still challenging. Our results reveal that there is a common lack of abilities in comprehending program logic among current LLMs, despite their promising performance on other code-related tasks (e.g., code generation).

Limitations

As a pioneering work of benchmarking LLM-based test case generation, our TESTEVAL still has a few limitations. Here, we will discuss these limitations and how we addressed them in our work (or in the future).

First, the current TESTEVAL is only limited to Python. Although the solutions in LeetCode are written in multiple languages, we find that their adopted algorithm and logical structures are largely the same. We believe the behaviors of LLMs in other languages of LeetCode solutions will be sim-

ilar to those of Python, which we aim to verify in the future.

The second limitation is that our TESTEVAL dataset is created from online programming problems, which may be different from real-world scenarios. We argue that at the current stage of LLM for test case generation, datasets from programming problems are still important. First, many LLMs in real-world test case generation still struggle with the correctness problem (whether the generated test case can be executed), which makes it too early to consider the coverage problem. For example, in (Yuan et al., 2023), ChatGPT only achieves 42.1% success rate in compilation and 24.8% in execution. In contrast, on TESTEVAL, proprietary LLMs such as GPT-4 can achieve near 100% accuracies in execution (although some open-source LLMs still have difficulties in generating correctly formatted test cases), which allows researchers to focus on how to improve test coverage. Second, compared to real-world software testing datasets, programs in TESTEVAL have more complex control flow structures, which allow us to have a deeper study on how LLMs can reason about branches/loops in programs. For example, the real-world Python test case generation dataset CodaMOSA (Lemieux et al., 2023) has an average cyclomatic complexity of 5.85, while the average complexity of our TESTEVAL dataset is 13.35.

Ethical Discussion

Regarding the dataset, our dataset is built upon user-written solutions for LeetCode problems. These solutions are stored in a GitHub repository licensed with the MIT license, so we are granted permission to create our own dataset from this repository.

Regarding the use of automated systems, the automatic tools we used to create our dataset are all rule-based tools with no bias introduced.

The research of LLMs for test case generation may encourage the software development industry to use LLMs instead of human developers for software testing. However, our findings in the paper suggest that existing LLMs still encounter various difficulties in generating correct test cases with accurate target test coverage. As software testing in real-world practice may introduce new questions not discussed in this paper, thus the impact of this paper on the industry community is still limited and not likely to cause major concerns.

Also, using LLMs for automated software test-

ing may raise security concerns. As our dataset only consists of self-contained, single-file programs, there are no security vulnerabilities in our dataset that can be exploited by LLM-generated test cases. However, if we extend the scope of our dataset to real-world software in the future, the security of experiments should be carefully considered.

Acknowledgments

This work was supported in part by Canada CIFAR AI Chairs Program, Natural Sciences and Engineering Research Council of Canada; JST CRONOS Grant (No.JPMJCS24K8), JSPS KAKENHI Grant (No.JP21H04877, No.JP23H03372, and No.JP24K02920); and the Autoware Foundation.

References

- Saranya Alagarsamy, Chakkrit Tantithamthavorn, and Aldeida Aleti. 2023. A3test: Assertion-augmented automated test case generation. *arXiv preprint arXiv:2302.10352*.
- 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*.
- Luciano Baresi and Matteo Miraz. 2010. Testful: Automatic unit-test generation for java classes. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 2*, pages 281–284.
- Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Cristian Cadar, Patrice Godefroid, Sarfraz Khurshid, Corina S Păsăreanu, Koushik Sen, Nikolai Tillmann, and Willem Visser. 2011. Symbolic execution for software testing in practice: preliminary assessment. In *Proceedings of the 33rd International Conference on Software Engineering*, pages 1066–1071.
- 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*.
- Vitaly Chipounov, Volodymyr Kuznetsov, and George Candea. 2011. S2e: A platform for in-vivo multi-path analysis of software systems. *Acm Sigplan Notices*, 46(3):265–278.
- Ermira Daka and Gordon Fraser. 2014. A survey on unit testing practices and problems. In *2014 IEEE 25th International Symposium on Software Reliability Engineering*, pages 201–211. IEEE.
- Arghavan Moradi Dakhel, Amin Nikanjam, Vahid Majdinasab, Foutse Khomh, and Michel C Desmarais. 2024. Effective test generation using pre-trained large language models and mutation testing. *Information and Software Technology*, 171:107468.
- Yangruibo Ding, Zijian Wang, Wasi Ahmad, Hantian Ding, Ming Tan, Nihal Jain, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, Dan Roth, et al. 2024. Crosscodeeval: A diverse and multilingual benchmark for cross-file code completion. *Advances in Neural Information Processing Systems*, 36.
- Gordon Fraser and Andrea Arcuri. 2011. Evosuite: automatic test suite generation for object-oriented software. In *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*, pages 416–419.
- Gordon Fraser and Andreas Zeller. 2010. Mutation-driven generation of unit tests and oracles. In *Proceedings of the 19th international symposium on Software testing and analysis*, pages 147–158.
- Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I Wang. 2024. Cruxeval: A benchmark for code reasoning, understanding and execution. *arXiv preprint arXiv:2401.03065*.
- Daya Guo, Canwen Xu, Nan Duan, Jian Yin, and Julian McAuley. 2023. Longcoder: A long-range pre-trained language model for code completion. In *International Conference on Machine Learning*, pages 12098–12107. PMLR.
- Sepehr Hashtroudi, Jiho Shin, Hadi Hemmati, and Song Wang. 2023. Automated test case generation using code models and domain adaptation. *arXiv preprint arXiv:2308.08033*.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. 2023. Swe-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*.
- René Just, Darioush Jalali, and Michael D Ernst. 2014. Defects4j: A database of existing faults to enable controlled testing studies for java programs. In *Proceedings of the 2014 international symposium on software testing and analysis*, pages 437–440.
- Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2023. Ds-1000: A natural and reliable benchmark for data science code generation. In *International Conference on Machine Learning*, pages 18319–18345. PMLR.

- 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.
- Bowen Li, Wenhan Wu, Ziwei Tang, Lin Shi, John Yang, Jinyang Li, Shunyu Yao, Chen Qian, Binyuan Hui, Qicheng Zhang, et al. 2024a. Devbench: A comprehensive benchmark for software development. *arXiv preprint arXiv:2403.08604*.
- Jia Li, Ge Li, Yunfei Zhao, Yongmin Li, Huanyu Liu, Hao Zhu, Lecheng Wang, Kaibo Liu, Zheng Fang, Lanshen Wang, Jiazheng Ding, Xuanming Zhang, Yuqi Zhu, Yihong Dong, Zhi Jin, Binhua Li, Fei Huang, Yongbin Li, Bin Gu, and Mengfei Yang. 2024b. [DevEval: A manually-annotated code generation benchmark aligned with real-world code repositories](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3603–3614, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Vincent Li and Nick Doiron. 2023. Prompting code interpreter to write better unit tests on quixbugs functions. *arXiv preprint arXiv:2310.00483*.
- Zhiyu Li, Shuai Lu, Daya Guo, Nan Duan, Shailesh Jannu, Grant Jenks, Deep Majumder, Jared Green, Alexey Svyatkovskiy, Shengyu Fu, et al. 2022. Automating code review activities by large-scale pre-training. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1035–1047.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2024a. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems*, 36.
- Kaibo Liu, Yiyang Liu, Zhenpeng Chen, Jie M Zhang, Yudong Han, Yun Ma, Ge Li, and Gang Huang. 2024b. Llm-powered test case generation for detecting tricky bugs. *arXiv preprint arXiv:2404.10304*.
- Tianyang Liu, Canwen Xu, and Julian McAuley. 2023. Repobench: Benchmarking repository-level code auto-completion systems. *arXiv preprint arXiv:2306.03091*.
- Thomas J McCabe. 1976. A complexity measure. *IEEE Transactions on software Engineering*, (4):308–320.
- Juan Altmayer Pizzorno and Emery D Berger. 2024. Coverup: Coverage-guided llm-based test generation. *arXiv preprint arXiv:2403.16218*.
- Laura Plein, Wendkûuni C Ouédraogo, Jacques Klein, and Tegawendé F Bissyandé. 2024. Automatic generation of test cases based on bug reports: a feasibility study with large language models. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*, pages 360–361.
- pytest cov. <https://github.com/pytest-dev/pytest-cov>.
- Nikitha Rao, Kush Jain, Uri Alon, Claire Le Goues, and Vincent J Hellendoorn. 2023. Cat-lm training language models on aligned code and tests. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 409–420. IEEE.
- Gabriel Ryan, Siddhartha Jain, Mingyue Shang, Shiqi Wang, Xiaofei Ma, Murali Krishna Ramanathan, and Baishakhi Ray. 2024. Code-aware prompting: A study of coverage-guided test generation in regression setting using llm. *Proceedings of the ACM on Software Engineering*, 1(FSE):951–971.
- Max Schäfer, Sarah Nadi, Aryaz Eghbali, and Frank Tip. 2023. An empirical evaluation of using large language models for automated unit test generation. *IEEE Transactions on Software Engineering*.
- Benjamin Steenhoeck, Michele Tufano, Neel Sundaresan, and Alexey Svyatkovskiy. 2023. Reinforcement learning from automatic feedback for high-quality unit test generation. *arXiv preprint arXiv:2310.02368*.
- Michele Tufano, Dawn Drain, Alexey Svyatkovskiy, Shao Kun Deng, and Neel Sundaresan. 2020. Unit test case generation with transformers and focal context. *arXiv preprint arXiv:2009.05617*.
- Vasudev Vikram, Caroline Lemieux, and Rohan Padhye. 2023. Can large language models write good property-based tests? *arXiv preprint arXiv:2307.04346*.
- Junjie Wang, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. 2024a. Software testing with large language models: Survey, landscape, and vision. *IEEE Transactions on Software Engineering*.
- Wenhan Wang, Kaibo Liu, An Ran Chen, Ge Li, Zhi Jin, Gang Huang, and Lei Ma. 2024b. Python symbolic execution with llm-powered code generation. *arXiv preprint arXiv:2409.09271*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Chunqiu Steven Xia, Matteo Paltenghi, Jia Le Tian, Michael Pradel, and Lingming Zhang. 2023. Universal fuzzing via large language models. *arXiv preprint arXiv:2308.04748*.
- Zhuokui Xie, Yinghao Chen, Chen Zhi, Shuiguang Deng, and Jianwei Yin. 2023. Chatunitest: a chatgpt-based automated unit test generation tool. *arXiv preprint arXiv:2305.04764*.

Zhiqiang Yuan, Yiling Lou, Mingwei Liu, Shiji Ding, Kaixin Wang, Yixuan Chen, and Xin Peng. 2023. No more manual tests? evaluating and improving chatgpt for unit test generation. *arXiv preprint arXiv:2305.04207*.

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](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2471–2484, Singapore. Association for Computational Linguistics.

Ying Zhang, Wenjia Song, Zhengjie Ji, Na Meng, et al. 2023b. How well does llm generate security tests? *arXiv preprint arXiv:2310.00710*.

Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, et al. 2023. Codegeex: A pre-trained model for code generation with multilingual evaluations on humaneval-x. *arXiv preprint arXiv:2303.17568*.

A Prompt Templates

The prompt templates for TESTEVAL tasks are shown as follows. For the targeted line/branch/path coverage tasks, we add line numbers to both the program under test and the target information in order to accurately locate the position of the target line/branch/path. Notice that our prompt template is only a primary setting without advanced prompting techniques such as few-shot examples or chain-of-thought reasoning, and we encourage future researchers to design more advanced prompts for TESTEVAL.

A.1 Prompt Template for Overall Coverage

Please write a test method for the function ‘{func_name}’ given the following program under test and function description. Your answer should only contain one test input.

Program under test:

—
{program}

—
Function description for ‘{func_name}’:

—
description

—
Your test method should begin with:

```
def test_func_name():  
    solution=Solution()
```

Prompt for generating the next test case:

Generate another test method for the function under test. Your answer must be different from previously-generated test cases, and should cover different statements and branches.

A.2 Prompt Template for Targeted Line Coverage

Please write a test method for the function ‘{func_name}’ given the following program under test and function description. Your answer should only contain one test input.

Program under test:

—
{program}

—
Function description for ‘{func_name}’:

—
description

—
Your test case must cover line {target_line}.

Your test method should begin with:

```
def test_func_name():  
    solution=Solution()
```

A.3 Prompt Template for Targeted Branch Coverage

Please write a test method for the function ‘{func_name}’ given the following program under test and function description. Your answer should only contain one test input.

Program under test:

—
{program}

—
Function description for ‘{func_name}’:

—
description

—
Your test case must cover the branch {target_branch}.

Your test method should begin with:

```
def test_func_name():  
    solution=Solution()
```

A.4 Prompt Template for Targeted Path Coverage

Please write a test method for the function ‘{func_name}’ given the following program under test and function description. Your answer should only contain one test input.

Program under test:

—
{program}

—
Function description for ‘{func_name}’:

—
description —

Your test case must cover the following execution path in function {func_name}. The path is a sequence of branch conditions. When executing your test case, each branch condition in the target execution path must be satisfied sequentially.

Target execution path: {target_path}

—
Your test method should begin with:

```
def test_func_name():  
    solution=Solution()
```

A.5 Prompt Template for Two-step COT

Prompt template for generating conditions in the two-step COT

Given a Python code snippet and a target line number, you are asked to generate reasoning steps to satisfy a specific line to be executed.

[Example]

Given the following code snippet:

```
```Python
class Solution: #1
 def twoSum(self, nums: List[int], target: int) -> List[int]: #2
 numMap = #3
 n = len(nums) #4
 #5
 for i in range(n): #6
 numMap[nums[i]] = i #7
 #8
 for i in range(n): #9
 complement = target - nums[i] #10
 if complement in numMap and numMap[complement] != i: #11
 return [i, numMap[complement]] #12
 #13
 return [] #14
```
```

Identify when executing function twoSum, what conditions need to be satisfied if line 12 is to be executed.

Answer:

<cond>

Step 1: If line 12 'return [i, numMap[complement]]' is executed, then the 'if' condition '(complement in numMap and numMap[complement] != i)' in line 11 should be true.

Step 2: If condition 'complement in numMap' is true, at least one 'target - nums[i]' in line 10 equals an element in nums, which means there exists two elements in 'nums' that their sum is equal to 'target'.

Step 3: If condition 'numMap[complement] != i' is true, then 'numMap[target - nums[i]] != i', meaning that the index of 'target - nums[i]' is not equal to 'i'.

Step 4: Generally, to execute line 12, we need to ensure that there exists two different elements in 'nums' that their sum is equal to 'target'.

<\cond>

[Example]

In a similar fashion, identify the conditions that need to be satisfied when line targetline is to be executed for the following Python code.

```
```Python
{program}
```
```

Surround your answer with <cond> and <\cond>.

Prompt template for generating test case in the two-step COT

For the given code snippet and a list of conditions need to be satisfied, generate a test case that will satisfy these conditions. Here is an example:

[Example]

Code:

```
```Python
class Solution: #1
 def twoSum(self, nums: List[int], target: int) -> List[int]: #2
 numMap = #3
 n = len(nums) #4
 #5
 for i in range(n): #6
 numMap[nums[i]] = i #7
 #8
 for i in range(n): #9
 complement = target - nums[i] #10
 if complement in numMap and numMap[complement] != i: #11
 return [i, numMap[complement]] #12
 #13
 return [] #14
```
```

Conditions:

Step 1: If line 12 'return [i, numMap[complement]]' is executed, then the 'if' condition '(complement in numMap and numMap[complement] != i)' in line 11 should be true.

Step 2: If condition 'complement in numMap' is true, at least one 'target - nums[i]' in line 10 equals an element in nums, which means there exists two elements in 'nums' that their sum is equal to 'target'.

Step 3: If condition 'numMap[complement] != i' is true, then 'numMap[target - nums[i]] != i', meaning that the index of 'target - nums[i]' is not equal to 'i'.

Step 4: Generally, to execute line 12, we need to ensure that there exists two different elements in 'nums' that their sum is equal to 'target'.

Generated test case:

```
```Python
def test_twoSum():
 solution = Solution()
 assert solution.twoSum([2,7,11,15], 9) == [0, 1]
```
```

[Example]

In a similar fashion, generate a test case for the following code snippet and conditions. Your test function should be named 'test_func_name'. Code:

```
```Python
{program}
```
```

Conditions:

{conditions}

You should only generate the test case, without any additional explanation.

B Targeted Line/Branch Identification

The complete algorithm for extracting targeted lines/branches from a program under test is shown in Algorithm 2. At a high level, we first extract all conditional branches by locating the branches starting with conditional operators (i.e., ‘if’, ‘elif’, and ‘else’) through parsing the program’s abstract syntax tree. For each branch, we record the line numbers of its first and last lines (e.g., Lines 1:5) as one targeted branch. Then, we record the line numbers of all lines (except the line that only includes the ‘else’ operator) within this branch as the targeted lines (e.g., [1, 2, 3, 4, 5]). We repeat this process until finishing parsing all branches of a program.

Algorithm 2: Targeted Line/Branch Identification.

Input: Program with L lines:

$$p = \{s_1, s_2, \dots, s_L\}$$

Output: Target lines ls , target branches bs

$ls = [], bs = [], i = 1;$

while $i \leq L$ **do**

if s_i starts with ‘if’, ‘elif’, or ‘else’

then

$curent_branch = []$;

$j = i$;

repeat

$curent_branch.append(j)$;

$j = j + 1$

until s_j not in this branch;

$bs.append(curent_branch)$;

end

end

for $target_branch$ in bs **do**

for line s_i in $target_branch$ **do**

if s_i is inside a branch and not s_i

 starts with ‘else’ **then**

$ls.append(i)$;

end

end

end

return ls, bs

C Error Analysis

For the failure of LLMs in generating test cases that failed to execute, we choose Llama3 as the example. Figure 6 shows several examples of failed test cases generated by Llama 3. Figure 6 (a) shows an example with a slight syntax error: it generated

```
(a): format error      def test_isMatch():
                        solution=Solution()
                        assert not solution.isMatch("\"aa\"", \"a*\"),

(b): overlength test case  def test_getSubarrayBeauty():
                        solution=Solution()
                        nums = [-1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
-1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1,
```

Figure 6: Examples of erroneous test cases generated by LLMs.

a redundant comma at the end of the last statement. Figure 6 (b) is another common type of error: the LLM generates an endless statement by repeating a simple pattern. In our post-processing statements, we remove the last statement if it is uncompileable. These erroneous statements are removed and result in empty test cases, which are counted as execution errors. We find that all execution errors in Llama 3-8b for targeted line coverage are made up of these two types of errors.

D Data Leakage Analysis

We choose GPT-4o as an example to study the potential of data leakage. The training data of GPT-4o covers up to October 2023, so we filter the problems from our dataset released after Oct 2023, which results in a total of 21 problems. Correspondingly, we also create a subset with 21 oldest problems which are released before Oct 2023.

For the problems released after Oct 2023, in their 49 official test cases, we found none of them appeared in the generated test cases. On the contrary, for the 21 problems before Oct 2023, 35 out of 52 official test cases have been found in the generated test cases. However, as the LLM has generated 20 different test cases for each problem (which means 420 test cases for 21 problems), the issue of copying official test cases is minor. We further measure the overall coverage for all problems before/after Oct 2023, the results are shown in Table 6.

Table 6: Coverage metrics of the overall coverage task with data source before and after Oct 2023.

| Model | Before Oct 2023 | | After Oct 2023 | |
|--------|-----------------|--------|----------------|--------|
| | line | branch | line | branch |
| GPT-4o | 98.74 | 97.24 | 97.79 | 96.38 |

We can see that the coverage metrics before/after Oct 2023 are similar, indicating that potential data leakage is not a major concern of TestEval.