

Generalizing Test Cases for Comprehensive Test Scenario Coverage

BINHANG QI, National University of Singapore, Singapore

YUN LIN*, Shanghai Jiao Tong University, China

XINYI WENG, Shanghai Jiao Tong University, China

CHENYAN LIU, National University of Singapore, Singapore

HAILONG SUN, Beihang University, China

GORDON FRASER, University of Passau, Germany

JIN SONG DONG, National University of Singapore, Singapore

Test cases are essential for software development and maintenance. In practice, developers derive multiple test cases from an implicit pattern based on their understanding of requirements and inference of diverse test scenarios, each validating a specific behavior of the focal method. However, producing comprehensive tests is time-consuming and error-prone: many important tests that should have accompanied the initial test are added only after a significant delay, sometimes only after bugs are triggered.

Existing automated test generation techniques largely focus on code coverage. Yet in real projects, practical tests are seldom driven by code coverage alone, since test scenarios do not necessarily align with control-flow branches. Instead, test scenarios originate from requirements, which are often undocumented and implicitly embedded in a project's design and implementation. However, developer-written tests are frequently treated as executable specifications; thus, even a single initial test that reflects the developer's intent can reveal the underlying requirement and the diverse scenarios that should be validated.

In this work, we propose TestGeneralizer, a framework for generalizing test cases to comprehensively cover test scenarios. TestGeneralizer orchestrates three stages: (1) enhancing the understanding of the requirement and scenario behind the focal method and initial test; (2) generating a test scenario template and crystallizing it into various test scenario instances; and (3) generating and refining executable test cases from these instances. To ensure accuracy and completeness, TestGeneralizer combines rule-based prompts, automatically optimized via a prompt auto-tuning technique, with crucial project knowledge retrieved through program analysis. We evaluate TestGeneralizer against three state-of-the-art baselines (EvoSuite, gpt-o4-mini, and ChatTester) on 12 open-source Java projects, covering 506 multi-test focal methods and 1,637 test scenarios. TestGeneralizer achieves significant improvements: +57.67% and +59.62% over EvoSuite, +37.44% and +32.82% over gpt-o4-mini, and +31.66% and +23.08% over ChatTester, in mutation-based and LLM-assessed scenario coverage, respectively. In a field study, we submitted 27 generalized tests overlooked by developers; 16 were accepted and merged into official repositories, demonstrating the practical usefulness of TestGeneralizer.

CCS Concepts: • **Software and its engineering** → **Software testing and debugging**.

Additional Key Words and Phrases: Test Generation, Software Testing, Large Language Models

*Corresponding author.

Authors' Contact Information: Binhang Qi, qibh@nus.edu.sg, National University of Singapore, Singapore; Yun Lin, lin_yun@sjtu.edu.cn, Shanghai Jiao Tong University, China; Xinyi Weng, nanakusa@sjtu.edu.cn, Shanghai Jiao Tong University, China; Chenyan Liu, chenyang@u.nus.edu, National University of Singapore, Singapore; Hailong Sun, sunhl@buaa.edu.cn, Beihang University, China; Gordon Fraser, Gordon.Fraser@uni-passau.de, University of Passau, Germany; Jin Song Dong, dcsdjs@nus.edu.sg, National University of Singapore, Singapore.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

ACM 2994-970X/2026/7-ARTFSE209

<https://doi.org/10.1145/3808216>

ACM Reference Format:

Binhang Qi, Yun Lin, Xinyi Weng, Chenyan Liu, Hailong Sun, Gordon Fraser, and Jin Song Dong. 2026. Generalizing Test Cases for Comprehensive Test Scenario Coverage. *Proc. ACM Softw. Eng.* 3, FSE, Article FSE209 (July 2026), 22 pages. <https://doi.org/10.1145/3808216>

1 Introduction

Software testing plays a crucial role in ensuring the quality of both open-source and industrial software by verifying whether implementations satisfy user requirements. A large body of research has been devoted to automating test generation, which has been predominantly **code coverage-driven**. Classical approaches [6, 14, 20, 30] such as EvoSuite [12] frame test generation as the problem of maximizing structural coverage (e.g., branches or paths). This formulation reduces test generation to a constraint-solving problem, giving rise to techniques including symbolic execution [4–6, 14, 36] and search-based testing [2, 4, 12, 15, 18, 20, 21, 30]. With the rise of large language models (LLMs), recent work [9, 16, 18, 24, 25, 40, 46] recasts test generation as a special case of code generation. Given a focal method and an instruction prompt, these approaches leverage LLMs (e.g., ChatGPT) to inductively produce test code by translating the focal method, the prompt, or both into executable test cases. Yet despite their different form, these LLM-based approaches remain code coverage-driven: their primary goal is to improve code coverage, whether by optimizing prompts and pipelines [24, 46] or by integrating LLMs with classical techniques [18, 45].

While existing test generation approaches have shown promising results, practical test cases are rarely written solely to maximize coverage metrics or exercise focal code. In practice, developers design test cases around diverse **test scenarios** that validate whether a focal method satisfies specific requirements. A test scenario is derived from a requirement, aiming to validate whether a specific behavior of a focal method is expected by the requirement. For example, in the `ofdrw` project [28], the requirement (obtained by comments) for the focal method `setPaint(Paint paint)` is: “To generate OFD through graphic drawing, it must be possible to set the brush color for both filling and stroking.” To validate this requirement, developers wrote multiple tests to cover various test scenarios, including setting the color as a solid, a multi-stop linear gradient, and a radial gradient. Designing high-quality tests to comprehensively cover such scenarios is important, yet time-consuming and error-prone. It is often the case that developers realize important test scenarios only later, sometimes after failures occur. In this example, the test `linearGradientPaint()` was first added in commit 91af2eb [26] (February 1, 2023), while the test `setPaintRadialGradientPaint()` was only introduced in commit 9f82f37 [27] (March 6, 2023). Notably, we verified that the latter test also executes successfully on the earlier version, indicating that it should have been included from the initial commit. A similar case arises in the `cron-utils` project, where a developer supplements an additional test case [8] to validate an overlooked test scenario: “infinite loop when daylight savings time starts at midnight,” raised through an issue report [7]. These examples highlight the importance and the difficulty of deriving comprehensive test scenarios. This motivates the following research question: *Given a focal method and a developed test case, how can we generalize it into additional test cases that comprehensively cover the test scenarios intended by the developer?*

Apparently, code coverage-driven solutions, whether classical [6, 14, 20, 30] or LLM-based [9, 16, 18, 24, 25, 40, 45, 46], are not well-suited for this task, since test scenarios do not necessarily align with control-flow branches (e.g., `setPaint()` contains no branches). The key challenge lies in **Implicit Functionality Requirement**: complete and well-maintained requirements are rare, especially under agile development. Instead, requirements are often only implicitly embedded in the project’s functional design and developer intent. Without understanding such implicit requirements, it is difficult to infer comprehensive test scenarios without omissions or redundancies. This setting also differs from established paradigms like property-based testing [22, 38, 39, 41], which rely on

explicit properties. In practice, behavioral intent is often implicitly encoded in developer-written tests. We hypothesize that such tests contain a latent scenario template, and that a single initial test can provide sufficient information to be generalized into diverse, behaviorally consistent tests.

In this work, we propose TestGeneralizer, a three-stage framework for generalizing test cases to cover comprehensive test scenarios. Across these stages, TestGeneralizer addresses the implicit requirement challenge through *project knowledge collection* and *prompt auto-tuning*, enabling accurate generalization from a single test case. Specifically, given a focal method and one initial test case, TestGeneralizer proceeds as follows. In Stage ❶ (Enhancing Test Scenario Understanding), TestGeneralizer transforms the initial test into multiple-choice “exams” by mutating its oracles. The LLM is asked to select the correct oracle, providing a measurable check of its understanding of the requirement and the current test scenario. In Stage ❷ (Test Scenario Generalization), TestGeneralizer prompts the LLM to generate a *test scenario template*, a semi-structured plan describing how to validate the focal method under diverse test scenarios. The template consists of step-by-step actions with *variation points* (VPs), whose different settings derive distinct intended test scenarios. By fixing the VP settings, the template is crystallized into concrete *test scenario instances*, each validating a meaningful behavior of the focal method. Each instance contains both *primary oracles* (deduced from project design and implementation) and *alternative oracles* (inferred from common requirement knowledge). In practice, alternative oracles can alert developers to potential design flaws or implementation bugs. In Stage ❸ (Test Generation), each test scenario instance serves as guidance for the LLM to generate a concrete test. Generated tests are iteratively refined with the help of project knowledge, resolving compilation errors, execution errors, and assertion failures until the test passes or a maximum iteration limit is reached. To make this pipeline effective, TestGeneralizer overcomes two primary challenges. **(1) Implicit Pattern Recognition:** distinguishing true scenario VPs from noisy code elements is highly complex. TestGeneralizer addresses this through a prompt auto-tuning technique (in Stage ❷) that automatically learns precise rules for VP identification. **(2) Reasonable Scenario Crystallization:** determining semantically valid and project-specific settings for these VPs is error-prone. TestGeneralizer overcomes this by proactively retrieving crucial project knowledge (in Stage ❶ – ❸) via program analysis, grounding generalized scenarios in the actual project context rather than relying on LLM hallucination.

We extensively evaluate TestGeneralizer on 506 focal methods from 12 open-source projects, involving 1,637 test scenarios. Compared to state-of-the-art approaches (e.g., gpt-o4-mini [29], ChatTester [46] and EvoSuite [12]), TestGeneralizer demonstrates clear advantages. In particular, compared to ChatTester, TestGeneralizer achieves: **(1) Higher Scenario Coverage:** TestGeneralizer generalizes more comprehensive test scenarios, improving scenario coverage by 31.66% and 23.08% in mutation-based (i.e., overlap between mutants killed by generated and ground-truth tests) and LLM-assessed scenario coverages, respectively. Also, the effectiveness of TestGeneralizer is consistent on both commercial LLMs (e.g., ChatGPT) and open-source LLMs (e.g., DeepSeek-V3.1). **(2) Practical Impact:** In a field study, we submitted pull requests with tests generalized by TestGeneralizer but overlooked by developers. Of the 27 submitted tests, 16 were accepted and merged into official repositories, demonstrating the practical usefulness of TestGeneralizer.

In summary, this work makes the following contributions:

- **Methodology.** To the best of our knowledge, TestGeneralizer is the first approach to generalize test cases. We introduce techniques for enhancing requirement understanding and inferring test scenarios, laying the groundwork for deriving tests directly from requirements.
- **Dataset and Evaluation.** We release a benchmark for test generalization, including focal code, test code, and high-quality test scenario templates. The benchmark consists of 506 multi-test focal methods from 12 open-source projects, involving 1,637 test scenarios, curated to ensure quality

```

void linearGradientPaint() throws Exception {
  ① Path dst = Paths.get("linearGradPaint.ofd");
  ② try (OFDDoc doc = new OFDDoc(dst)) {
    OFDPPage2D g = doc.newPage(null);

    Point2D start = new Point2D.Float(0, 0);
    Point2D end = new Point2D.Float(50, 50);
    float[] dist = {0.0f, 0.2f, 1.0f};
    ③ Color[] colors = {Color.RED, ...};
    LinearGradientPaint p =
      new LinearGradientPaint(...);
    g.setPaint(p);

    ④ g.fillRect(0, 0, 50, 50);
  }
  ⑤ System.out.println(">> " + dst.toPath());
}

void radialGradientPaint() throws Exception {
  Path dst = Paths.get("radialGradPaint.ofd");
  try (OFDDoc doc = new OFDDoc(dst)) {
    OFDPPage2D g = doc.newPage(500, 500);

    Color[] colors = {Color.red, ...};
    float[] dist = {0.0f, 0.5f, 1.0f};
    Point2D center = new Point2D.Float(...);

    RadialGradientPaint p =
      new RadialGradientPaint(...);
    g.setPaint(p);

    g.fillRect(0, 0, 500, 500);
  }
  System.out.println(">> " + dst.toPath());
}

void fillOval() throws Exception {
  Path dst = Paths.get("fillOval.ofd");
  try (OFDDoc doc = new OFDDoc(dst)) {
    OFDPPage2D g = doc.newPage(null);

    g.setPaint(Color.RED);

    g.fillOval(25, 25, 120, 60);
  }
  System.out.println(">> " + dst.toPath());
}

```

Fig. 2. Test cases for the focal method `setPaint()`, covering three test scenarios. The tests follow a common pattern consisting of five parts (highlighted by colored backgrounds), which collectively validate the requirement behind the focal method.

and practice relevance. Using this benchmark, we extensively evaluate TestGeneralizer against state-of-the-art baselines. Results show that TestGeneralizer generalizes more comprehensive test scenarios and consistently outperforms baselines in scenario coverage.

- **Field Study.** To assess practical usefulness, we conduct a field study by submitting pull requests with test cases generalized by TestGeneralizer but overlooked by project developers. Of the 27 submitted tests, 16 were accepted and merged into official repositories.

All source code, benchmark, experimental results, and field study data are available at [32].

2 Motivation Example

In practice, when developers write test cases to validate a requirement, they often follow a recurring pattern to crystallize multiple test scenarios from this pattern and derive corresponding test cases.

For example, Figure 1 shows a focal method in `ofdrw` [28], which allows users to set the brush color for both filling and stroking when generating OFD through graphic drawing. Figure 2 presents the corresponding tests written by developers to validate this requirement.

```

/* set brush color for filling and stroking */
public void setPaint(Paint paint) {
  this.drawParam.setColor(paint);
}

```

Fig. 1. Focal method in `ofdrw` project [28].

By comparing these tests, we observe that they follow a shared pattern consisting of five parts, highlighted in the figure. If such a pattern is already identified and crucial project knowledge (e.g., `LinearGradientPaint`, `RadialGradientPaint`, `fillRect`, and `fillOval`) is already known, generalizing the test scenarios and implementing the tests is straightforward. However, when only a single test case is available as a reference, generalizing the remaining test scenarios and writing their corresponding tests becomes non-trivial. This challenge arises for several reasons:

- **Implicit Pattern Recognition.** Focal methods often provide little guidance on which test scenarios are intended, since scenarios do not necessarily correspond to control-flow branches and branchless methods such as `setPaint()` are common. Even with an initial test as a reference, it is still difficult to infer the underlying pattern and variation points (VPs) regarding test scenarios from numerous changeable code elements. For example, with `linearGradientPaint()` as the initial test, one can observe many changeable code elements across the five parts, such as the path in part ①, the parameter of `newPage()` in part ②, the parameters of `LinearGradientPaint()` and `setPaint()` in part ③, and the filling method and its parameters in part ④. Yet test scenarios involve three VPs: the canvas setting, the paint style, and the drawing shape. These correspond to only three code elements, including the parameters of `newPage()` and `setPaint()`, and the choice of filling method. Other code elements are either trivial (e.g., the parameters of `fillRect()`) or dependent on these VPs (e.g., the path). Thus, inferring the underlying pattern with intended VPs is highly challenging.

```

1 void testSetPaintWithBasicColor() {
2     DummyOFDPPageGraphics2D g = new DummyOFDPPageGraphics2D();
3     Color testColor = Color.BLUE;
4     g.setPaint(testColor);
5     assertEquals(testColor, g.getPaint(), "The paint should be updated to Color.BLUE");
6 }
7
8 void testChangingPaintMultipleTimes() { X Redundant test scenario
9     DummyOFDPPageGraphics2D g = new DummyOFDPPageGraphics2D();
10    Color firstColor = Color.MAGENTA;
11    g.setPaint(firstColor);
12    assertEquals(firstColor, g.getPaint(), "Paint should be updated to Color.MAGENTA");
13    // ...
14    LinearGradientPaint secondPaint = new LinearGradientPaint(start, end, dist, colors);
15    g.setPaint(secondPaint);
16    assertEquals(secondPaint, g.getPaint(), "Paint should now be updated to the new GradientPaint");
17 }
18
19 void testSetPaintWithNull() { X Redundant test scenario
20    DummyOFDPPageGraphics2D g = new DummyOFDPPageGraphics2D();
21    g.setPaint(null);
22    assertEquals(null, g.getPaint(), "Paint should be null after setting it to null");
23 }
24
25 public void linearGradientPaintWithTransform() throws Exception { X Redundant test scenario
26     final Path dst = Paths.get("target/linearGradientPaintWithTransform.ofd");
27     try (OFDGraphicsDocument doc = new OFDGraphicsDocument(dst)) {
28         OFDPageGraphics2D g = doc.newPage(500, 500);
29         g.translate(250, 250);
30         g.rotate(Math.toRadians(30));
31         g.scale(1.2, 1.2);
32         // ...
33         LinearGradientPaint p = new LinearGradientPaint(start, end, dist, colors);
34         g.setPaint(p);
35     }
36     System.out.println(">> " + dst.toAbsolutePath());
37 }

```

Listing 1. ChatTester-generated tests, with one missing and three redundant cases relative to ground truth.

- **Reasonable Scenario Crystallization.** Even having accurate VPs, determining which test scenarios are reasonable to crystallize remains difficult. Feasible choices for VP settings may be unknown, or syntactically valid but semantically unreasonable with respect to requirements. For example, deriving these tests based on the pattern requires several project-specific knowledge: `fillOval()` is an alternative to `fillRect()` with similar functionality, while `RadialGradientPaint()` and `Color` are feasible parameters for `setPaint()`. Without such knowledge, guessing appropriate APIs or parameter values is difficult and prone to error.

Facing these challenges, code coverage-driven test generation often produces redundant test cases and misses important scenarios, leading to inefficiency and incomplete validation.

Coverage-driven Test Generators. Traditional coverage-driven approaches rely solely on control-flow branches, without awareness of requirements or test scenarios. Since test scenarios do not necessarily correspond to branches, these approaches cannot capture the intended behaviors of the focal method. In this example, coverage-driven software testing tools such as EvoSuite [12] stop immediately after generating a single executable test case.

LLM-based Test Generators. Listing 1 shows the results of ChatTester [46], a state-of-the-art test generation solution built on GPT-o4-mini. We adapt ChatTester for test generalization by additionally providing it with an initial test (i.e., `linearGradientPaint`) and modifying its prompt from “generate one test case” to “generate more test cases with reference to the initial test case,” while keeping all other settings as default. As shown in Listing 1, the generalized tests produced by ChatTester remain considerably distant from the ground-truth tests, due to two major issues: (1) **Failure to capture the underlying pattern.** The generated tests miss key parts of the pattern and the associated VPs. For instance, `testSetPaintWithBasicColor()` misses the first part of setting a path and the final part of invoking a filling method. (2) **Failure to infer reasonable test scenarios.** While ChatTester does infer a scenario involving `Color` (corresponding to the ground-truth test `fillOval()`), likely guided by common knowledge and hints in the initial test (where `Color` is already used), it fails to produce the important scenario `radialGradientPaint()`.

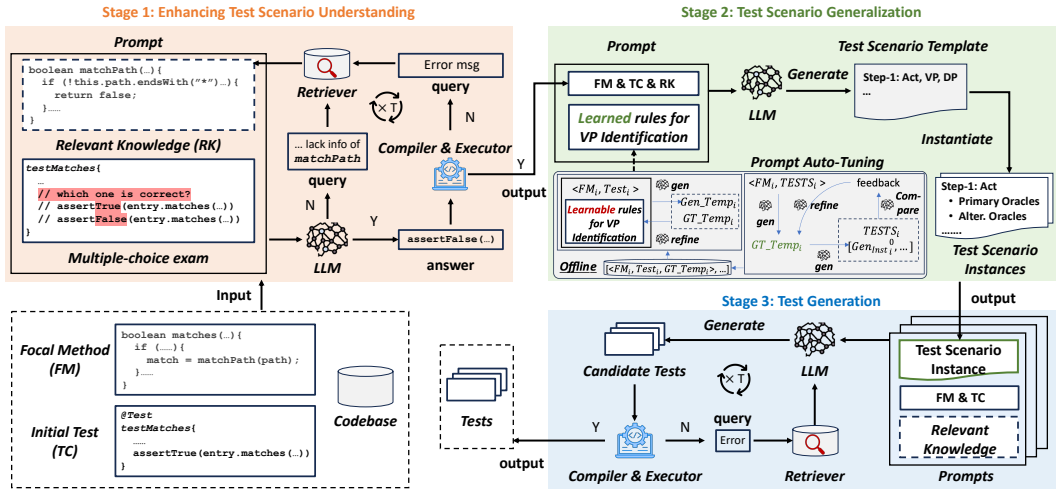


Fig. 3. Overview of TestGeneralizer: Given a focal method, an initial test, and the codebase, TestGeneralizer generalizes the initial test into additional tests that comprehensively cover intended scenarios. The retriever is implemented using predefined CodeQL and JDtLS script templates instantiated with LLM-generated queries.

Table 1. Definitions of core terminology used in TestGeneralizer.

Term	Definition
Relevant project knowledge	Project-specific facts retrieved from the codebase, such as method bodies, class relationships, and field definitions, that help understand requirements, identify variation points, determine feasible settings, or repair generated tests.
Variation point (VP)	A scenario-relevant factor, either a code element or high-level testing choice, whose alternative settings produce distinct intended behaviors of the focal method
Test scenario template	A general, semi-structured test plan that abstracts the developer’s intended testing pattern into steps, variation points, and dependencies among steps.
Test scenario instance	A concrete instantiation of a test scenario template, obtained by fixing all variation-point settings and determining the corresponding oracle(s).

To address these challenges, we introduce TestGeneralizer, a framework for accurately generalizing comprehensive test scenarios. TestGeneralizer first constructs a test scenario template with precise VPs, guided by a high-quality prompt with well-designed rules obtained through prompt auto-tuning. It then retrieves crucial project knowledge to strengthen the understanding of requirements and test scenarios, enabling the crystallization of reasonable test scenario instances from the template. From these instances, it generates corresponding test cases, achieving comprehensive coverage of the test scenarios targeted by ground-truth tests. Moreover, TestGeneralizer produces additional tests beyond the ground-truth; several of these were submitted as pull requests and successfully merged by project maintainers, as detailed in our field study (Section 4.5).

3 Approach

Figure 3 illustrates the workflow of TestGeneralizer, which consists of the following three stages. Table 1 summarizes the core terminology used throughout the approach.

- Stage 1 (Enhancing Test Scenario Understanding):** Given a focal method fm and a developer-written test tc , TestGeneralizer infers the underlying requirement and test scenario reflected by tc . To achieve this, TestGeneralizer introduces **Oracle-Mutation-Guided Scenario Understanding (OMSU)**, which validates and enhances the LLM’s requirement and scenario understanding

by probing it with multiple-choice “exams” constructed from oracle mutations and triggering retrieval of relevant project knowledge from K when needed.

- **Stage 2 (Test Scenario Generalization):** Building on the understanding from Stage 1, TestGeneralizer produces a *test scenario template* so_{temp} , a concise step-by-step plan describing how to test fm and what variation points (VPs) are. The template is then crystallized into a set of *test scenario instances* $\{so_{ins}^m\}_{m=1}^M$, each with concrete VP settings and corresponding oracles. To improve the comprehensiveness of $\{so_{ins}^m\}_{m=1}^M$, TestGeneralizer (1) prompts the LLM to proactively query relevant knowledge K to recover potential VPs absent from the current context, and (2) applies prompt auto-tuning to optimize VP identification.
- **Stage 3 (Test Generation):** Finally, TestGeneralizer generates executable tests from $\{so_{ins}^m\}_{m=1}^M$. For each instance, the LLM is given fm , tc , K , and so_{ins}^m to produce a test tc_{gen} . It iteratively repairs failing tests using compiler and runtime feedback together with retrieved knowledge.

3.1 Enhancing Test Scenario Understanding

Understanding the underlying requirements and test scenarios is the foundation of test generalization. Since requirements and scenarios are often implicitly embedded in project design, collecting relevant project knowledge is essential for improving the understanding. However, even with a developer-written test for reference, two challenges remain: (1) determining whether the current context already contains sufficient knowledge, and (2) identifying which pieces of knowledge are important amid the abundance of potentially noisy information. Relying solely on the LLM to judge sufficiency is unreliable due to hallucination. To address this, we propose **Oracle-Mutation-Guided Scenario Understanding (OMSU)**, which examines the LLM using multiple-choice exams derived by mutating the oracles in the initial test case. These exams provide a more deterministic measure of the LLM’s understanding while also guiding the retrieval of relevant knowledge.

Exam Generation. Given fm and tc with assertions, TestGeneralizer prompts the LLM to generate up to Q incorrect oracles (e.g., $Q = 10$) for each assertion in tc . The key instruction is: “1. Identify the assertions and oracles used in the test case. 2. Come up with up to 10 WRONG oracles for each identified oracle. NOTE: The oracle is wrong, but the assertion statement should still be compilable, which means you cannot pinch APIs or fields out of thin air.” The LLM generates wrong oracles by, for example, altering expected results or modifying parameter values in the invocation of fm . As shown in Stage 1 of Figure 3, one wrong oracle changes `assertTrue` to `assertFalse`. TestGeneralizer filters out wrong oracles that cause compilation or execution errors, retaining only those that trigger assertion failures. If no valid wrong oracles remain, the LLM is instructed to revise its output.

Examination and Knowledge Collection. TestGeneralizer then examines the LLM on each exam. When answering, the LLM is allowed to query any project knowledge it deems necessary. The core instruction is: “... decide which one among the list of alternative assertion statements is equipped with a correct oracle. You must be very confident in your answer. You CANNOT guess or assume the existence or values for APIs and Fields. If you think the context lacks some important information for making the decision, please do not make a choice and output only a list of the required information ...” To support such queries, TestGeneralizer utilizes CodeQL to collect, offline, the definitions and invocations of all classes, constructors, methods, and fields in the project. During examination, retrieval is restricted to symbols directly referenced by the focal method or the initial test (e.g., parent classes and overridden methods of the referenced symbols), ensuring semantic completeness while controlling context size. For example, as shown in Figure 3, the LLM queries the definition of `matchPath()`, because understanding its functionality is essential for selecting the correct oracle. Queries can be raised proactively by the LLM or enforced by TestGeneralizer when the LLM selects an incorrect option, which indicates insufficient semantic understanding of fm . The examination

Table 2. Prompt template for generating test scenario templates. Green-highlighted contents are case-specific and are to be filled in accordingly. Gray-highlighted contents are predefined definitions and instructions.

#Instruction: Given the Java Focal Method, generate a Test Scenario Template for the Focal Method with reference to the provided Test Case. [Definition of a test scenario template.]

#Focal Method: [Code of focal method]
#Focal Method Context: [Skeleton of focal file]
#Initial Test Case: [Developer-written test]
#Project Knowledge: [List of collected knowledge]

#Rules for Variation Point Identification: [Auto-tuned rules]
#Requirements:
 ① Your response must first output the analysis and thinking: [Example analysis format]
 ② Based on your analysis and thinking, if you require more Relevant Project Knowledge (e.g., constructor, method, and field) for accurately identifying variation points, output your query in the following format without any commentary: [Example query format]
 ③ If no more information is required, output Test Scenario Template in the following format: [Example template format]

follows an iterative loop of (i) answering, (ii) knowledge retrieval if needed, and (iii) re-evaluation, until the LLM passes the exams or a maximum number (e.g., 3) of iterations is reached.

By examining the LLM, TestGeneralizer reliably collects relevant project knowledge and strengthens the understanding of requirements and test scenarios, thereby improving subsequent test generalization (a concrete example illustrating the knowledge retrieval process is presented in Section 4.3.2). The collected knowledge is passed to Stage 2 for test scenario generalization. All prompts are available at [32].

3.2 Test Scenario Generalization

Given the focal method fm , initial test tc , and the project knowledge K (if any) obtained from Stage 1, TestGeneralizer generalizes test scenarios by first generating a test scenario template so_{temp} and then crystallizing it into a set of test scenario instances $\{so_{ins}^m\}_{m=1}^M$ (see Section 3.2.1). To ensure comprehensive scenario coverage, the prompt for generating so_{temp} is optimized using our prompt auto-tuning technique (see Section 3.2.2).

3.2.1 Test Scenario Template and Instance Generation. Test scenario generalization is the core of TestGeneralizer. It first generates a test scenario template and then crystallizes the template into a comprehensive set of test scenario instances.

Test Scenario Template Generation. Given fm , tc , and K , TestGeneralizer orchestrates a learnable prompt (shown in Table 2) to guide the LLM in producing a test scenario template. A template is defined as a general test plan that captures the developer’s intent to validate the fm across diverse scenarios. It consists of concise steps expressed in natural language, optionally supplemented with minimal program elements. Each step includes: (1) *Action (ACT)*: an imperative instruction for a tester. (2) *Variation Points (VPs)*: factors that can vary, either concrete code elements or abstract objects in the high-level scenario. VPs can be instantiated into specific settings to produce meaningful behaviors of the focal method. (3) *Dependency (DP)*: optional links to VPs defined in earlier steps.

A high-quality template should abstract all developer-intended test scenarios and be capable of deriving them. For example, Figure 4 shows the template generated for the focal method `setPaint()` with `linearGradientPaint()` as the initial test. The template consists of five steps, corresponding to the underlying pattern illustrated in Figure 2. Among them, Steps 2–4 correctly capture the VPs and their reasonable candidate settings.

To generate such a high-quality template, the key challenge lies in accurately identifying VPs. To address this, the learnable prompt incorporates auto-tuned rules (“#Rules for Variation Point Identification” in Table 2), derived through our prompt auto-tuning technique (see Section 4.2). In addition, TestGeneralizer encourages the LLM to proactively query additional project knowledge

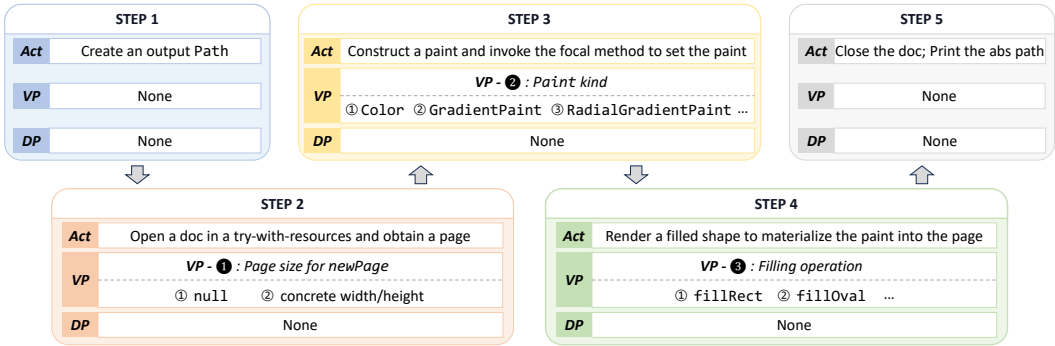


Fig. 4. Generated test scenario template for `setPaint()` using `LinearGradientPaint()` as the initial test.

beyond what was retrieved in Stage 1, since some information about VPs or their feasible settings may reside in other parts of the codebase. Such queries are resolved using the same offline CodeQL index constructed in Stage 1, and the retrieved definitions/usages are appended to the prompt before continuing template or instance generation.

Test Scenario Instance Generation. A test scenario instance is a concrete test plan for validating a specific behavior of fm in a particular scenario. It is derived from the test scenario template by fixing all VPs and determining the associated oracles. To instantiate a test scenario template, TestGeneralizer adopts a structured instantiation strategy that separates scenario-level decision making from natural-language instance generation. Rather than directly asking the LLM to generate full natural-language test scenario instances, TestGeneralizer first asks it to output structured settings, including the selected setting of each variation point, dependencies among these settings, and the corresponding primary oracle(s). A deterministic instantiation procedure then fills the original test scenario template with these settings to produce test scenario instances compatible with the downstream test-generation stage, while validating that all settings refer to declared variation points and valid template steps.

Considering that fm may be buggy in practice, TestGeneralizer does not assume the correctness when determining oracles. Instead, TestGeneralizer prompts the LLM to identify and list all reasonable oracles that a developer might intend, including both the *primary oracle* (deduced from the code implementation) and *alternative oracles* (inferred from requirement understanding and general knowledge). In practice, TestGeneralizer can optionally allow developers to select among alternative oracles; if no intervention occurs, the primary oracle is used.

The test scenario instances then serve as guidance for generating executable tests that comprehensively cover the developer’s intended scenarios (see Section 3.3).

3.2.2 Prompt Auto-Tuning. As discussed in Section 2 (Implicit Pattern Recognition), even with all necessary information available, accurately identifying VPs regarding test scenarios from changeable code elements remains challenging. Many elements may appear to be plausible VPs, but only a small subset is truly meaningful.

Thus, well-designed rules are essential for accurate identification. However, manually crafting such rules is often time-consuming, error-prone, and lacks generalizability across projects and programming languages. To address this, we propose a *prompt auto-tuning* technique, enabling the LLM itself to derive and optimize these rules. Algorithm 1 outlines the procedure. It takes as input an LLM LLM and a dataset $D = \{(x_i, y_i)\}$, where $x_i = (fm_i, tc_i^j, K_i)$ bundles a focal method fm_i , one of its initial tests tc_i^j , and the relevant knowledge K_i . The target y_i is the ground-truth test scenario template containing accurate VPs. Prompt auto-tuning derives a rule set S_{rule} and

Algorithm 1 Prompt Auto-Tuning for Variation Point Identification

```

1: Input: (1) Black-box large language model  $LLM$ ; (2) Dataset  $D = \{(x_i, y_i)\}_{i=1}^{|D|}$ , where  $x_i = (fm_i, tc_i, K_i)$  bundles a focal method  $fm_i$ , one initial test  $tc_i$ , and the relevant knowledge  $K_i$ ;  $y_i$  is the ground-truth test scenario template.
2: Output: Tuned prompt  $p^*$  with an optimal rule set  $S_{rule}$  for VP identification.
3: Randomly split  $D$  into training and test sets:  $D_{train}$  and  $D_{test}$ 
4: Randomly split  $D_{train}$  into batches:  $B = \{b_n\}_{n=1}^N$ , where  $b_n = \{(x^k, y^k)\}_{k=1}^{|b_n|}$ 
5: Initialize base prompt  $p_0$  with empty  $S_{rule}$ 
6: for epoch  $e = 1$  to  $E$  do
7:    $p_e \leftarrow p_{e-1}$ 
8:   for each batch  $b_n \in B$  do
9:      $T_n = \{t^k | t^k \leftarrow LLM(p_e, x^k)\}_{k=1}^{|b_n|}$   $\triangleright$  Generate scenario templates
10:     $F_n = \{f^k | f^k \leftarrow LLM(p_e, y^k, t^k)\}_{k=1}^{|b_n|}$   $\triangleright$  Generate feedback
11:     $p_e \leftarrow LLM(p_e, F_n)$   $\triangleright$  Optimize prompt based on feedback
12:   end for
13:   Record  $p_e$ 
14: end for
15:  $p^* = \arg \max_{p \in \{p_1, \dots, p_E\}}$  Evaluate( $p; D_{test}$ )  $\triangleright$  Select the optimal prompt
16: Return:  $p^*$ 

```

iteratively refines it on D , leveraging feedback from LLM by comparing generated templates against y_i to progressively improve VP identification.

Dataset Construction. We construct D in a semi-automated manner. First, we select a set of focal methods $\{fm_i\}_{i=1}^I$ from our dataset of tests (see Section 4.1.1). Each selected focal method is inspected to ensure that it is equipped with sufficient developer-written tests covering comprehensive scenarios. Given a focal method fm_i and its complete set of tests $\{tc_i^j\}_{j=1}^J$, we prompt the LLM to generate a test scenario template t_i . Since the LLM has access to all tests, this generation task is relatively straightforward: it can infer VPs by analyzing differences among the tests. To validate t_i , we then ask the LLM to instantiate it into test scenario instances and manually compare them against $\{tc_i^j\}_{j=1}^J$. If missing or redundant instances are observed—indicating misidentified VPs—we provide the LLM with feedback, including the missing or redundant instances and any manually collected relevant knowledge K_i , and ask the LLM to refine t_i . This process is repeated until t_i passes evaluation, at which point it is accepted as the ground-truth template.

Tuning Process. The dataset D is divided into a training set D_{train} and a test set D_{test} . D_{train} is used to optimize S_{rule} , while D_{test} is used to select the optimal version of S_{rule} across E epochs. Specifically, D_{train} is partitioned into batches $B = \{b_n\}_{n=1}^N$. The base prompt p_0 (shown in Table 2) is initialized with an empty S_{rule} . During the e -th epoch, the prompt is iteratively optimized across B : ❶ For each sample (x^k, y^k) in batch b_n , the current prompt p_e is used to generate a template t^k for x^k , producing a set of templates $T_n = \{t^k\}_{k=1}^{|b_n|}$. ❷ LLM evaluates t^k against the corresponding ground-truth template y^k , identifying missing or redundant VPs. ❸ For t^k , LLM generates feedback f^k describing how S_{rule} should be adjusted to fix the issues while preserving its generality, yielding a set of feedback $F_n = \{f^k\}_{k=1}^{|b_n|}$. ❹ To avoid overfitting to particular focal methods, LLM synthesizes general feedback from F_n by merging complementary suggestions and resolving conflicts (i.e., retaining the most general suggestion). This feedback is then used to refine S_{rule} , resulting in updated p_e . At the end of each epoch e , the updated p_e is recorded.

After all epochs, all prompts $\{p_e\}_{e=1}^E$ are evaluated on D_{test} using precision, recall, and F1-score metrics. Precision measures the proportion of VPs in the generated template that also appear in the ground-truth template, while recall measures the proportion of ground-truth VPs correctly

Table 3. Statistics of the dataset

Project	Commit Version	# Focal Methods	# Tests (Scenarios)	# Tests per FM		
				Min	Max	Avg
itext-java	9c895e8410	177	626	2	15	3.5
hutool	3e716cf0a	30	73	2	7	2.4
yavi	de9e9ab34d	74	167	2	8	2.3
lambda	d360ae809f	59	144	2	5	2.4
jlinstagram	82834b0ea2	29	90	2	13	3.1
truth	23171822b	63	308	2	15	4.9
cron-utils	d31697ec4c	19	88	2	13	4.6
imglib	49e238162c	19	49	2	4	2.6
ofdrw	beeeb31c4c	12	28	2	4	2.3
RocketMQC	6790ee5dfc	10	25	2	5	2.5
blade	ecf15b0664	2	4	2	2	2.0
spark	1973e402f5	12	35	2	9	2.9
Total		506	1637	2.0	8.3	3.0

identified in the generated template. Algorithm 1 returns the prompt with the highest F1-score on D_{test} , which is then used for test scenario template generation (see Section 3.2.1).

Due to space limitations, the prompts are presented at [32], and a concrete example illustrating the contribution of auto-tuned rules is presented in Section 4.3.2

3.3 Test Generation

For each generated test scenario instance, TestGeneralizer prompts the LLM to produce a corresponding test. The test is then compiled and executed, and feedback messages are collected from the compiler and runtime. If the test fails due to compilation errors, execution errors, or assertion failures, TestGeneralizer refines it based on the feedback: ❶ Extract error messages and filter out irrelevant messages that originate outside the project (e.g., third-party dependencies). ❷ Use regular expressions to extract program elements and their positions from the error messages. ❸ For each extracted program element, retrieve relevant project knowledge (e.g., the implementation of a method that failed during compilation or execution) via JDTLS [11], using the element’s position for lookup. ❹ TestGeneralizer then constructs a prompt that integrates the focal method, the generated test, error messages, and the retrieved knowledge. This prompt guides the LLM in refining the test to resolve the error. The refinement process repeats until the test passes or a maximum number of iterations is reached (e.g., three). The prompt details are available at [32].

4 Evaluation

We evaluate TestGeneralizer through the following research questions:

RQ1 (Overall Performance): How effective is TestGeneralizer in generalizing tests compared to state-of-the-art baselines?

RQ2 (Prompt Auto-Tuning Effectiveness): How much does prompt auto-tuning improve VP identification over standard prompting strategies?

RQ3 (Ablation Study): What are the contributions of project knowledge (i.e., retrieval) and rules (i.e., prompt auto-tuning) to TestGeneralizer’s overall performance?

RQ4 (Sensitivity Analysis): How robust is TestGeneralizer when the initial test is of low quality?

RQ5 (Field Study): Do the generalized tests provide practical value in real projects?

4.1 Overall Performance (RQ1)

4.1.1 Dataset. We curate focal methods, each with developer-written tests, from 12 diverse open-source projects. These projects span domains such as web development and image processing, each with over 100 GitHub stars and forks.

Criteria. A method is selected as a focal method if it satisfies the following criteria:

- **Multiple Developer-Written Tests.** The method must be exercised by at least two distinct developer-written tests. Our approach relies on scenario diversity to generalize developer intent.
- **Clear Test-Method Mapping.** The method must be explicitly invoked in the test body (directly or via a simple wrapper), ensuring unambiguous scenario attribution.
- **Meaningful Behavioral Logic.** We exclude trivial methods (e.g., simple getters/setters or one-line delegations) where scenario variation is inherently limited.
- **Interpretable Test Scenarios.** Four authors inspect the associated tests to ensure they encode identifiable behavioral scenarios (e.g., nominal, boundary, or exceptional cases). Ambiguous cases are resolved through discussion.

Statistics. Table 3 presents the commit versions of the projects in which the focal methods are selected, the number of focal methods, the total number of existing tests for these focal methods, and the minimal, maximal, and average number of existing tests per focal method. In total, the dataset contains 506 focal methods and 1,637 tests. Tests range from 13 to 140 lines of code (LOC; average 37), while focal methods range from 3 to 90 LOC (average 11). For RQ1, we evaluate TestGeneralizer on seven projects (see Table 4). The remaining five projects are exclusively used for prompt auto-tuning in RQ2.

4.1.2 Baselines. We use gpt-o4-mini (o4-mini-2025-04-16) as the default LLM for TestGeneralizer. To assess practicality with an open-source model, we also report TestGeneralizer powered by DeepSeek-V3.1 (denoted “Ours (ds)”). For comparison, we evaluate three representative baselines: ① **o4-mini**, a vanilla baseline, ② **ChatTester** [46], a state-of-the-art LLM-based test generator, and ③ **EvoSuite** [12], a state-of-the-art search-based test generator.

We use o4-mini as a vanilla baseline, as it is a straightforward tool for test generalization. ChatTester is chosen for its outstanding performance among LLM-based test generators. For fairness, we upgrade its backend from GPT-3.5 to o4-mini. To adapt ChatTester to the test generalization task, we provide it with an initial test and modify its prompt from “generate one test case” to “generate more test cases with reference to the initial test case”, leaving other settings unchanged.

EvoSuite is chosen as a representative search-based technique. As a coverage-driven generator, it naturally produces multiple tests for a given focal method, making it a suitable non-LLM baseline.

4.1.3 Metrics. We evaluate TestGeneralizer and the baselines using *scenario coverage*, which measures how well the generated tests cover the test scenarios targeted by developer-written tests. For a focal method $fm^n \in S_{fm}$ with generated tests S_{gen}^n and ground-truth tests S_{gt}^n , scenario coverage is defined as:

$$Cov_{so}^n = \frac{1}{|S_{gt}^n|} \sum_{tc_{gt} \in S_{gt}^n} \text{Coverage}(tc_{gt}, S_{gen}^n). \quad (1)$$

We adopt two complementary definitions of $\text{Coverage}(\cdot, \cdot)$:

- **Mutation-based Scenario Coverage.** The intuition is that a scenario targeted by a test can be characterized by the set of potential bugs it can expose. We utilize Pitest [31] to mutate the class containing fm^n , yielding a mutant set. For a test tc , let $S_{\mu}(tc)$ denote the mutants killed by tc . For a ground-truth test tc_{gt} and a generated test tc_{gen} , the pairwise coverage score is

$$s(tc_{gt}, tc_{gen}) = \frac{|S_{\mu}(tc_{gt}) \cap S_{\mu}(tc_{gen})|}{|S_{\mu}(tc_{gt})|}. \quad (2)$$

To ensure *order-independent* and prevent a single generated test from matching multiple ground-truth tests, we compute a maximum-weight bipartite matching between S_{gt}^n and S_{gen}^n with edge weights $s(\cdot, \cdot)$, and the matched scores (unmatched tc_{gt} contribute 0) are then used in Equation 1.

Table 4. Comparison of TestGeneralizer and baselines in test scenario coverage (in %).

Projects	# Scenarios	Mutation-based Scenario Coverage					LLM-Assessed Scenario Coverage				
		EvoSuite	Vanilla	ChatTester	Ours (ds)	Ours	EvoSuite	Vanilla	ChatTester	Ours (ds)	Ours
itext-java	627	8.53	31.91	30.88	70.57	65.52	19.28	33.86	32.36	56.14	57.95
hutool	73	4.39	65.33	51.19	81.79	84.07	4.88	37.67	44.75	69.97	70.21
lambda	144	n/a	65.86	76.89	96.05	95.49	n/a	57.99	65.74	78.76	77.89
cron-utils	88	16.32	26.81	29.38	72.46	64.75	15.79	23.81	35.83	58.16	57.91
ofdrw	28	n/a	12.21	14.89	75.28	42.14	n/a	26.79	23.57	47.62	64.88
RocketMQC	25	25.99	58.86	64.86	82.37	89.75	29.00	54.00	75.00	81.00	85.00
blade	4	42.67	17.67	51.00	84.33	99.00	0.00	50.00	75.00	100.00	100.00
Average	989	19.58	39.81	45.58	80.41	77.25	13.79	40.59	50.32	70.24	73.41

- LLM-Assessed Scenario Coverage.** Mutation-based scenario coverage can be overly strict and miss semantically valid matches. For example, testing a default behavior of a focal method either by omitting a setting or by explicitly resetting it may kill different mutants while targeting the same scenario. Moreover, a developer-written test may encode multiple scenarios, whereas generated tests target a single scenario. Therefore, LLM-assessed scenario coverage serves as a complementary metric. We prompt the LLM (same as that used for test generalization) to judge whether the test scenario(s) exercised by tc_{gt} can be fulfilled by any tests in S_{gen}^n . If so, $Coverage(tc_{gt}, S_{gen}^n) = 1$; otherwise it is 0. To avoid double counting, the tests decided to match tc_{gt} are removed from S_{gen}^n before evaluating the next ground-truth test.

4.1.4 Results. Table 4 reports the results of TestGeneralizer and baselines in terms of both mutation-based and LLM-assessed scenario coverage across seven projects and 989 test scenarios. On average, a ground-truth test scenario involves 82.1 mutants. The column “Ours (ds)” shows results when TestGeneralizer is powered by DeepSeek-V3.1, while the column “Ours” shows TestGeneralizer’s performance when powered by gpt-o4-mini. For projects of drw and lambda, the results of EvoSuite are marked as “n/a” because EvoSuite fails to generate tests for most focal methods in these projects, which make extensive use of generic Java types. The last row “Average” presents the average performance across all projects, with the projects where “n/a” values occur excluded for EvoSuite.

Baseline Comparison in Scenario Coverage. Overall, EvoSuite, despite its efficiency in achieving high code coverage, performs far worse than both LLM-based baselines and TestGeneralizer. For example, its mutation-based scenario coverage is lower by 20.23% than Vanilla and by 57.67% than TestGeneralizer. These results are expected: Code coverage-driven techniques rely entirely on control-flow branches in the focal method, but scenarios do not necessarily correspond to branches, and many focal methods contain no branches. For setPaint (introduced in Section 2), once a single test is generated, such approaches terminate, leaving many important scenarios uncovered. By contrast, LLM-based baselines (i.e., Vanilla and ChatTester) achieve much better performance. Their advantage comes from the ability of LLMs to infer or “imagine” plausible scenarios beyond code structural exploration. For example, as shown in Listing 1, ChatTester can imagine a test scenario where the paint is set to a basic color using a Color object.

However, imagination requires both constraint and stimulation; without them, it can lead to *hallucination* and *exhaustion*. Hallucination occurs when a tool treats unintended factors as VPs. For example, in project spark [3], ChatTester treats the parameters `StaticFilesConfiguration` and `ExceptionHandler` of focal method `EmbeddedJettyFactory.create()` as VPs, generating tests that vary them (e.g., setting both to null). Yet the initial test uses mock objects for these parameters, clearly indicating that developers do not intend their effects to be part of the test scenarios.

Exhaustion occurs when a tool fails to derive a comprehensive set of reasonable scenarios, even after correctly identifying VPs. As shown in Listing 1, ChatTester generated only one reasonable scenario (`testSetPaintWithBasicColor()`) regarding a VP (i.e., the parameter of `setPaint()`), while missing others intended by the developer, such as `setPaintRadialGradientPaint()`.

Table 5. Branch and line coverage achieved by TestGeneralizer (powered by gpt-o4-mini) and baselines.

Project	Branch Coverage				Line Coverage			
	EvoSuite	Vanilla	ChatTester	Ours	EvoSuite	Vanilla	ChatTester	Ours
itext-java	28.27	8.23	14.19	21.02	35.63	14.05	19.65	29.22
hutool	3.33	3.97	6.89	5.43	4.98	3.70	8.50	6.23
lambda	-	38.50	38.18	40.10	-	28.21	27.52	29.22
cron-utils	10.87	31.89	30.93	36.66	22.34	43.96	43.50	47.17
ofdrw	-	1.53	3.02	13.71	-	4.35	6.73	21.68
RocketMQC	1.24	1.22	1.26	1.85	5.08	0.73	4.87	5.93
blade	5.80	6.40	6.40	6.70	10.14	11.12	11.12	11.30
Average	9.90	13.11	14.41	17.92	15.63	15.16	17.41	21.54

These issues stem from the lack of explicit rules and project-specific knowledge to constrain VP identification and to stimulate exploration of reasonable settings. TestGeneralizer addresses both: it combines rule-based guidance with retrieved project knowledge and therefore consistently outperforms LLM-based baselines. Overall, TestGeneralizer achieves 77.25% mutation-based scenario coverage and 73.41% LLM-assessed scenario coverage, improving over Vanilla by 37.44% and 32.82%, and over ChatTester by 31.66% and 23.08%, respectively. A mutation-based scenario coverage close to 80% indicates that the generalized scenarios are sufficiently close to the ground truth, given the large number of mutants per ground-truth scenario on average (82.1).

Mutation-based vs. LLM-Assessed Scenario Coverage. Comparing mutation-based and LLM-assessed scenario coverage, we observe moderate differences between them. For example, mutation-based scenario coverage is often slightly higher than LLM-assessed scenario coverage (e.g., on project RocketMQC). This is because mutation-based coverage provides a continuous score in $[0, 1]$ that reflects partial overlap in killed mutants, whereas LLM-assessed coverage is binary (1 if any generated tests fulfill the scenario of the ground-truth test, 0 otherwise). Mutation-based coverage can capture partial alignment between two tests, while LLM-assessed scenario coverage cannot.

On the other hand, LLM-assessed coverage is more flexible when a ground-truth test encodes multiple scenarios. For example, in project ofdrw [28], the focal method setClip has a test clip() that exercises three distinct clipping scenarios: (1) a fill entirely inside the clip region, (2) a fill partially overlapping the region, and (3) a fill completely outside the region. TestGeneralizer correctly generalized these into three separate tests. In this case, LLM-assessed coverage recognizes that the three generated tests collectively fulfill the test scenario targeted by the ground-truth test (i.e., $\text{Coverage}(tc_{gt}, S_{gen}^n) = 1$ in Equation 1), while mutation-based coverage assigns only partial overlap ($\text{Coverage}(tc_{gt}, S_{gen}^n) = 0.33$).

Finally, the column “Ours (ds)” reports the performance of TestGeneralizer powered by DeepSeek-V3.1. As we can see, TestGeneralizer still consistently outperforms all baselines across all metrics, demonstrating its robustness and practicality for deployment in the industry.

Baseline Comparison in Code Coverage and Mutation Score. We also measure traditional code coverage and mutation score for a comprehensive comparison. Table 5 and Table 6 compare the total code coverage and mutation score achieved across the whole project with all test cases generated by each approach. TestGeneralizer achieves the highest average branch coverage, line coverage, and mutation score, with gains of 3.51%, 4.13%, and 9.87% over the best baseline, ChatTester. We observe

Table 6. Mutation score on the whole project achieved by TestGeneralizer (powered by gpt-o4-mini) and baselines.

Project	EvoSuite	Vanilla	ChatTester	Ours
itext-java	11.71	28.41	32.86	43.44
hutool	5.59	12.39	17.13	15.53
lambda	-	56.85	60.67	62.25
cron-utils	23.64	59.74	71.86	71.88
ofdrw	-	11.73	12.29	26.82
RocketMQC	16.37	27.19	45.03	51.46
blade	31.88	48.47	50.22	87.77
Average	17.84	34.97	41.44	51.31

that certain individual tests generated by baselines achieve higher code coverage or mutation score than TestGeneralizer. We manually inspected these cases and found that the baseline tests often exercise additional classes or methods that are unnecessary for testing the meaningful scenarios, thereby inflating project-level coverage.

Improvement Over Existing Tests. We analyze how generated tests improve existing test suites. Specifically, we measure the incremental mutation and coverage gains obtained by augmenting the original tests (T_{exist}) with generated tests. For each project, we compare T_{exist} with $T_{exist} + T_{ours}$ and compute the additional mutants killed and additional branch and line coverage achieved.

Across all projects, augmenting the existing tests with TestGeneralizer yields an average increase of 7.07% (57.93% to 65.00%) in mutation score, 2.46% (17.47% to 19.93%) in branch coverage, and 2.05% (21.45% to 23.50%) in line coverage. Although the improvements are moderate — since the focal methods in our benchmark are already supported by well-developed test suites and TestGeneralizer operates from a single initial test — these results indicate that TestGeneralizer can still introduce fault-detection capability and structural coverage beyond the original developer-written tests. This suggests that scenario-driven generalization complements existing suites by systematically exploring intention-aligned variations rather than merely duplicating already-tested behaviors.

Token and Time Consumption. We measure the token usage and wall-clock time required to generate one test. Across all generations, TestGeneralizer consumes on average 12K tokens per test. The average end-to-end processing time per test is 2 minutes. Future work could reduce cost by partially replacing LLM calls in Stage 1 with static or mutation-based program analysis.

4.2 Prompt Auto-Tuning Effectiveness (RQ2)

4.2.1 Baselines. Well-designed rules are crucial for generating high-quality test scenario templates with accurate variation points. TestGeneralizer derives these rules through a prompt auto-tuning technique. We evaluate its effectiveness by comparing it against three standard prompting strategies:

- **Zero-Shot Prompting:** The LLM generates templates using a prompt with only the task description, without rules or examples for guidance.
- **Few-Shot Prompting:** The LLM is provided with one to three in-context examples, each consisting of a focal method, an initial test, and the corresponding ground-truth template.
- **Hand-Crafted Prompt:** A prompt with manually designed rules, created by annotators based on their experience when constructing the dataset for prompt auto-tuning.

4.2.2 Dataset. We construct a dataset of 50 samples (see Section 4.2), each consisting of a focal method, its ground-truth tests, and the ground-truth template. These focal methods are selected from spark, imglib, truth, jInstagram, and yavi projects, which span diverse scales and domains and are not used in our field study. Each focal method is manually inspected to ensure that it is equipped with sufficient test cases covering comprehensive scenarios. On average, a focal method has 5.6 tests (ranging from 2 to 21), and a template has 3.2 variation points (ranging from 1 to 6).

4.2.3 Setup. We evaluate prompt auto-tuning and baselines using precision, recall, and F1-score for variation point identification. For zero-shot prompting and the hand-crafted prompt, all 50 samples are used for evaluation. For prompt auto-tuning, we apply the Leave-One-Project-Out strategy: for each fold, training on four projects and testing on the remaining one

Table 7. Comparison of prompt auto-tuning and baseline strategies for variation point identification (in %).

Method	Precision	Recall	F1
Zero-Shot Prompting	49.55	75.49	59.83
Few-Shot Prompting	51.64	70.24	59.52
Hand-Crafted Prompt	68.47	77.29	72.61
Auto-Tuned Prompt	81.21	86.39	83.72

Table 8. Ablation study results for TestGeneralizer (in %).

Projects	Mutation-based Scenario Coverage			LLM-Assessed Scenario Coverage		
	w/o knowledge	w/o rules	Ours	w/o knowledge	w/o rules	Ours
itext-java	64.78	57.17	65.52	54.32	53.21	57.95
hutool	81.20	77.05	84.07	67.43	63.51	70.21
yavi	91.63	83.09	96.92	77.84	70.34	81.69
lambda	93.97	95.83	95.49	71.06	78.18	77.89
jlinstagram	74.54	80.12	77.28	63.81	64.70	64.48
truth	62.27	41.53	66.25	59.94	43.69	63.24
cron-utils	59.42	55.73	64.75	54.89	43.65	57.91
imglib	94.76	90.83	94.76	93.20	87.41	93.20
ofdrw	32.76	33.31	42.14	39.21	41.94	64.88
RocketMQC	84.37	88.24	89.75	74.00	79.00	85.00
blade	83.00	98.67	99.00	100.00	100.00	100.00
spark	81.48	77.15	85.53	67.50	70.12	72.19
Average	75.35	73.23	80.12	68.60	66.31	74.05

project. We report the average evaluation results across the five folds. Regarding the tuning hyperparameters, we set the number of epochs to 3 and the batch size to 5 (results for batch sizes of 1, 3, and 7 are reported at [32]). For few-shot prompting, we randomly select 2 samples as in-context references (results for using 1 and 3 references are reported at [32]), and the rest of the samples are used for evaluation. To account for variability due to reference selection, we repeat the experiment three times with different references and report the average performance across runs.

4.2.4 Results. Table 7 reports the results for each prompting strategy. Zero-shot and few-shot prompting yield the worst performance across all metrics. Notably, few-shot prompting performs even worse than zero-shot prompting (see Recall and F1-score metrics), because a small number of references provides only partial guidance or causes overfitting, while adding more references can degrade performance due to context length constraints (the 3-reference setting performs even worse [32]). Hand-crafted prompting outperforms zero- and few-shot prompting but remains limited, as designing comprehensive and non-conflicting rules for complex scenarios is difficult for humans. In contrast, prompt auto-tuning can synthesize and reconcile rules from diverse samples into a unified, conflict-resolved rule set, achieving the best results across all metrics with improvements of 31.66%, 10.90%, and 23.89% in precision, recall, and F1-score over zero-shot prompting. These findings demonstrate the effectiveness of the proposed prompt auto-tuning technique.

4.3 Ablation Study (RQ3)

4.3.1 Setup. To assess the contribution of project knowledge (i.e., retrieval), we remove all retrieved knowledge from the prompts used in Stage 2 and Stage 3. We then re-run the full pipeline (*w/o knowledge*). As for the contribution of rules (i.e., prompt auto-tuning), we remove them from the prompt used in Stage 2 (i.e., template generation) and re-run the pipeline (*w/o rules*).

4.3.2 Results. Table 8 reports the ablation results. The columns “*w/o knowledge*” and “*w/o rules*” show the performance when project knowledge and auto-tuned rules are removed, respectively.

Contribution of Project Knowledge. As shown in the “*w/o knowledge*” column, compared to TestGeneralizer, scenario coverage drops by 4.77% and 5.45% on average in mutation-based and LLM-assessed metrics, respectively. This highlights the importance of project knowledge collected across the three stages.

Knowledge from Stage 1. In project spark [3], the focal method `matches(HttpMethod httpMethod, String path)` processes the parameter `httpMethod` directly but delegates `path` handling to `matchPath(String path)`, which distinguishes four patterns (i.e., `/path`, `/path/`, `/path/other`,

/path/*). Without the implementation of `matchPath` in context, the LLM cannot fully infer the functional requirement behind the focal method, and thus generalizes tests only over `httpMethod`, missing four test scenarios related to path. For this example, `TestGeneralizer` proactively queries the implementation of `matchPath` in Stage 1, as it recognizes that it is crucial for determining the oracle and answering the exam correctly.

Knowledge from Stage 2. In Stage 2, project knowledge again proves essential. For example, for `setPaint(Paint paint)` from `ofdrw` [28], the initial test `linearGradientPaint()` contains no assertions, so Stage 1 is skipped. `TestGeneralizer` prompts the LLM to proactively query project knowledge. Since the parameter `p` is of type `LinearGradientPaint`, a subtype of `Paint`, the LLM queries the family of related types and thus is aware of `Color`, `GradientPaint`, `RadialGradientPaint`, and `TexturePaint`. This knowledge guides the LLM to infer comprehensive test scenarios.

As for the project knowledge collected based on error messages, it is useful for correcting compilation and execution errors during test generation. In Stage 3, project knowledge collected from error messages is useful for resolving compilation and execution errors. Since this effectiveness is well established in prior work, we omit discussion here but provide illustrative examples at [32]. **Contribution of Auto-Tuned Rules.** Comparing the column “w/o rules” with the column “Test-Generalizer”, we observe that scenario coverage decreases by 6.89% and 7.74% on average in mutation-based and LLM-assessed metrics, respectively. This confirms the effectiveness of auto-tuned rules in test scenario generalization.

Without the rules, the LLM tends to indiscriminately treat changeable details as variation points, producing inferior test scenario templates where the valuable variation points are overwhelmed by the meaningless ones. On the one hand, redundant test scenarios derived from these spurious variation points generate noise, forcing developers to spend extra effort distinguishing useful tests. On the other hand, these redundant test scenarios hinder the derivation of valuable test scenarios, given the limited context length and the constrained attention capacity of the LLM.

For example, for `create(Route, StaticFilesConfiguration, ExceptionMapper, boolean)` from the spark project [3], the generated template incorrectly treats `StaticFilesConfiguration` and `ExceptionMapper` as variation points, even though the initial test mocks these parameters—explicitly signaling that their behavior is irrelevant to the intended scenarios. Similarly, trivial parameters such as connection timeouts or maximum connection counts, whose values are determined by earlier variation points, are wrongly considered independent variation points. As a result, the template includes eight meaningless variation points and only two actual ones, resulting in four redundant scenarios and missing one of the three intended scenarios.

4.4 Sensitivity Analysis (RQ4)

4.4.1 Setup. To evaluate the robustness of `TestGeneralizer` under degraded initial test quality, we progressively remove semantic information from the original initial test while preserving executability and the focal method invocation. We define three quality levels:

- **L0 (Original test).** The developer-written test is used as-is. This serves as the baseline.
- **L1 (Oracle-removed).** All oracle statements are removed via deterministic AST-based transformation. We delete assertion and verification calls (e.g., `assertEquals`, `assertTrue`, `assertThat`, `verify`), while preserving the test setup and focal method invocation.
- **L2 (Smoke test).** Starting from the L1 result, we use an LLM to further simplify the test into a minimal smoke test. The LLM is instructed to: (1) rename the test method to a generic name (e.g., `test_focal_method`), (2) remove scenario-specific setup code and retain only the minimal statements required to instantiate necessary objects and invoke the focal method once, and (3)

maintain executability. The LLM receives only the L1 test and the focal method signature as input (not the original L0 test), to avoid reintroducing removed oracle information.

4.4.2 Results. Table 9 reports the mutation-based and LLM-assessed scenario coverage across the seven projects evaluated in RQ1 for all three quality levels. On average, L0 achieves 77.2% / 73.4% scenario coverage. L1 (oracle removed) achieves 65.1% / 67.4%, corresponding to a decrease of 12.1% / 6.0%. L2 (smoke test) achieves 60.9% / 63.7%, corresponding to a decrease of 16.3% / 9.7% from L0. The degradation from L0 to L1 is moderate, indicating that TestGeneralizer does not rely solely on explicit oracle statements; substantial scenario signal remains embedded in the test setup and invocation structure. While the larger drop from L0 to L2 confirms that lower-quality inputs affect performance, it also demonstrates graceful degradation. Even when provided with a minimal smoke test, TestGeneralizer maintains over 60% scenario coverage, proving the pipeline remains highly effective without requiring rich initial tests.

4.5 Field Study (RQ5)

While the previous research questions evaluate TestGeneralizer on datasets of developer-written tests, its practical effectiveness in real projects remains to be assessed. In particular, some generalized tests extend beyond the ground-truth tests, raising the question of whether they are redundant or capture meaningful scenarios that developers intended but overlooked. To investigate, we conduct a field study by submitting generalized tests as pull requests to active GitHub projects and recording the number of accepted, rejected, and pending submissions.

Specifically, we selected four actively maintained projects—blade, ofdrw, cron-utils, and hutool—based on their recent activity (commits, accepted pull requests, or issue responses within six months prior to the time of submission). To minimize burden on maintainers, we choose the tests to be submitted based on the following criteria: (1) the focal method must already be covered by at least one existing test, indicating it is an important and stable API; and (2) when a focal method had more than two existing tests, our generalized tests typically included them, so only the additional ones were submitted. Moreover, we manually adjust the tests only for project conventions, including removing or adding comments and renaming variables (e.g., changing `plainPasswd` to `plain` for consistency with existing tests).

Table 10 reports the results of the field study, including the number of accepted, pending, and rejected submitted tests. In total, we submitted 5 pull requests containing 27 tests for 10 focal methods. Among them, 16 tests have been accepted, 11 tests remain pending, and none are rejected. As an example, in the motivating case of `setPaint(Paint paint)` from project ofdrw (Section 2), we submitted two tests: `setPaintGradientPaint()` and `setPaintTexturePaint()`. The former has been accepted, while the latter remains pending. Because the project currently does not provide complete functionality for texture paint, the pending test may have uncovered a new requirement, which may be under consideration by the maintainers. We also evaluated ChatTester on these

Table 9. Performance comparison of TestGeneralizer across three levels of initial test quality (in %).

Project	Mutation-Based			LLM-Assessed		
	L0	L1	L2	L0	L1	L2
itext-java	65.5	55.9	55.9	58.0	53.5	52.2
hutool	84.1	75.3	67.9	70.2	52.7	43.3
lambda	95.5	89.9	90.0	77.9	70.6	68.5
cron-utils	64.8	56.0	50.4	57.9	52.1	48.8
ofdrw	42.1	23.6	9.1	64.9	63.1	60.1
RocketMQC	89.8	80.2	76.9	85.0	80.0	73.0
blade	99.0	74.7	76.3	100	100	100
Average	77.2	65.1	60.9	73.4	67.4	63.7

Table 10. Results from a field study of five pull requests, comprising 27 tests for 10 focal methods.

Project	# Accepted	# Pending	# Rejected
hutool	15	0	0
cron-utils	0	9	0
ofdrw	1	1	0
blade	0	1	0
Total	16	11	0

focal methods. ChatTester failed to generate 37.0% (10/27) of the generalized tests. Importantly, among the 16 tests accepted by developers, ChatTester missed 37.5% (6/16). This confirms that TestGeneralizer produces valuable tests SOTA approaches miss.

5 Discussion

5.1 Limitations and Future Work

5.1.1 Quality of Generated Oracles. Oracle generation has long been an open problem in test generation. The quality of oracles generated by TestGeneralizer is potentially affected by bugs in the projects under test. In our evaluation, this issue is minimized because the dataset consists of well-established projects where focal methods are already equipped with multiple developer-written tests. However, in practice, developers cannot always assume correctness. If a focal method or related project knowledge contains bugs, then the oracles deduced from them may also be incorrect.

To mitigate this, TestGeneralizer (as described in Section 3.2.1) generates not only a primary oracle—derived from the focal method implementation and collected knowledge—but also alternative oracles inferred from the common knowledge embedded in LLMs. Developers can intervene in the generalization process and confirm or refine the intended oracle when alternatives are presented. In a preliminary experiment, we found that TestGeneralizer has a 72.98% probability of generating at least one alternative oracle capable of revealing the same synthetic bugs as the ground-truth tests. The detailed experimental setup and results are provided at [32]. Future work will further explore strategies to improve the quality and reliability of generated oracles.

5.1.2 Completeness of Knowledge Collection. TestGeneralizer prompts the LLM to proactively query crucial knowledge based on its analysis and reasoning. In Stage 1, knowledge is collected via examination: the LLM answers multiple-choice exams to determine whether it requires relevant project knowledge to correctly understand the test scenarios. A limitation of this design is that the LLM may guess the correct oracle, causing important knowledge to be missed. Future work will investigate more robust strategies, such as exploring the project's knowledge graph, to improve the completeness of knowledge collection and reduce reliance on chance.

5.2 Threats to Validity

Language generalizability. Evaluated only on Java, the transferability of auto-tuned rules to other languages remains uncertain. However, TestGeneralizer's pipeline is language-agnostic, and the involved tools like CodeQL and LSP naturally support multiple programming languages. Adapting to a new language mainly requires re-running prompt auto-tuning on a target-language dataset.

LLM-assessed scenario coverage. The nondeterministic nature of LLMs may affect the reliability of the LLM-assessed scenario coverage metric. To alleviate this, we set the temperature to zero, constraining randomness and ensuring more deterministic responses. We also repeat the evaluation three times to confirm its reliability and report the averaged results.

Emphasis on developer-written tests. Our primary metric (i.e., scenario coverage) measures how well generalized tests cover the scenarios targeted by developer-written ground-truth tests. Generalized tests that go beyond the ground truth may capture additional valuable scenarios and should not be deemed redundant by default. Assessing the value of such extra scenarios requires feedback from project maintainers, making a fully automated, large-scale evaluation infeasible. To mitigate this, we conducted a field study, which demonstrated the practical value of such tests.

Initial-test assumption. TestGeneralizer assumes an initial test as input for generalization, representing the first test a developer writes for a focal method. This makes TestGeneralizer less fully automated than end-to-end test generation approaches. However, our goal is to augment developer workflows by generalizing an existing test into broader intention-aligned scenarios,

rather than replacing fully automated tools. An initial test could also be generated by automated approaches, making TestGeneralizer naturally composable with existing test generation techniques.

6 Related Work

Coverage-driven software testing. Software testing has traditionally been treated as a constraint-solving problem, where the goal is to generate tests that cover targeted program branches. Representative approaches include symbolic execution (both dynamic and static) [4–6, 14, 36] and search-based testing [2, 4, 12, 15, 18, 20, 21, 30]. Although coverage-driven techniques can, to some extent, address test generalization, they cannot capture requirements or infer test scenarios, which do not necessarily correspond to control-flow branches. This limitation motivates our design of TestGeneralizer, which leverages LLMs for requirement understanding and scenario generalization.

LLM-based test generation. With the rise of LLMs, recent work applies LLMs to software testing [1, 10, 13, 17, 19, 23, 33–35, 37, 42–44, 47]. One closely related work is ChatTester [46], which generates a single test for a given focal method. In contrast, TestGeneralizer generalizes from an initial test to produce multiple tests that comprehensively cover developer-intended scenarios. Another relevant work is IntUT [24], which first generates test intentions for a focal method—each corresponding to a specific branch and specifying input parameters and expected outputs—and then generates test cases from these intentions to maximize coverage. Nevertheless, IntUT is essentially a coverage-driven approach that focuses on branch coverage. By contrast, TestGeneralizer targets scenario coverage, which is requirement-driven rather than code-coverage-driven.

Property-based test generation. Property-based testing [22, 41] and parameterized unit tests [38, 39] generalize inputs against developer-specified invariants, whereas TestGeneralizer infers behavioral intent directly from an example test without requiring explicit properties. For example, PROZE [38] derives parameterized tests from runtime data, primarily generalizing observed input values. In contrast, TestGeneralizer performs scenario-level generalization, identifying variation points that affect both inputs and oracle logic without relying on execution traces.

7 Conclusion

This work presents TestGeneralizer, a framework that generalizes developer-written tests to cover requirement-driven scenarios beyond traditional code coverage-driven approaches. Given a focal method with an initial test, TestGeneralizer generates scenario templates, instantiates them into concrete scenarios, and produces diverse and practical tests. Evaluation on 506 focal methods and 1,637 scenarios demonstrates that TestGeneralizer significantly outperforms state-of-the-art baselines. A field study further demonstrates its practical value.

8 Data Availability

All source code, benchmark, experimental results, and field study data are available at [32].

Acknowledgments

This research is supported in part by the National Natural Science Foundation of China (62572300), the Minister of Education, Singapore (MOE-T2EP20124-0017, MOET32020-0004), the National Research Foundation, Singapore and the Cyber Security Agency under its National Cybersecurity R&D Programme (NCRP25-P04-TAICeN), DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-GC-2023-008-1B), and Cyber Security Agency of Singapore under its National Cybersecurity R&D Programme and CyberSG R&D Cyber Research Programme Office. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore, Cyber Security Agency of Singapore as well as CyberSG R&D Programme Office, Singapore.

References

- [1] Nadia Alshahwan, Jubin Chheda, Anastasia Finogenova, Beliz Gokkaya, Mark Harman, Inna Harper, Alexandru Marginean, Shubho Sengupta, and Eddy Wang. 2024. Automated unit test improvement using large language models at meta. In *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering*. 185–196.
- [2] Andrea Arcuri and Xin Yao. 2008. Search based software testing of object-oriented containers. *Information Sciences* 178, 15 (2008), 3075–3095.
- [3] Spark authors. [n. d.]. Spark - a tiny web framework for Java 8. <https://github.com/perwendel/spark>.
- [4] Pietro Braione, Giovanni Denaro, Andrea Mattavelli, and Mauro Pezzè. 2017. Combining symbolic execution and search-based testing for programs with complex heap inputs. In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 90–101.
- [5] Pietro Braione, Giovanni Denaro, Andrea Mattavelli, and Mauro Pezzè. 2018. SUSHI: a test generator for programs with complex structured inputs. In 2018 IEEE/ACM 40th International Conference on Software Engineering: Companion (ICSE-Companion).
- [6] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. 2008. Klee: unassisted and automatic generation of high-coverage tests for complex systems programs.. In *OSDI*, Vol. 8. 209–224.
- [7] cron-utils authors. 2018. Infinite loop when daylight savings time starts at midnight #332. <https://github.com/jmrozanec/cron-utils/issues/332>.
- [8] cron-utils authors. 2018. Test Case for Issue #332. <https://github.com/jmrozanec/cron-utils/blob/master/src/test/java/com/cronutils/Issue332Test.java>.
- [9] Elizabeth Dinella, Gabriel Ryan, Todd Mytkowicz, and Shuvendu K Lahiri. 2022. Toga: A neural method for test oracle generation. In *Proceedings of the 44th International Conference on Software Engineering*. 2130–2141.
- [10] Chunhao Dong, Yanjie Jiang, Yuxia Zhang, Yang Zhang, and Liu Hui. 2025. ChatGPT-Based Test Generation for Refactoring Engines Enhanced by Feature Analysis on Examples . In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, 746–746. doi:10.1109/ICSE55347.2025.00210
- [11] Eclipse JDT Language Server Authors. 2025. Eclipse JDT Language Server. <https://github.com/eclipse-jdtls/eclipse.jdt.ls>.
- [12] 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*. 416–419.
- [13] Shuzheng Gao, Chaozheng Wang, Cuiyun Gao, Xiaoqian Jiao, Chun Yong Chong, Shan Gao, and Michael Lyu. 2025. The Prompt Alchemist: Automated LLM-Tailored Prompt Optimization for Test Case Generation. arXiv:2501.01329
- [14] Patrice Godefroid, Nils Klarlund, and Koushik Sen. 2005. DART: Directed automated random testing. In *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*. 213–223.
- [15] Javier Godoy, Juan Pablo Galeotti, Diego Garbervetsky, and Sebastián Uchitel. 2021. Enabledness-based testing of object protocols. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 30, 2 (2021), 1–36.
- [16] Sungmin Kang, Juyeon Yoon, and Shin Yoo. 2023. Large language models are few-shot testers: Exploring llm-based general bug reproduction. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2312–2323.
- [17] Myeongsoo Kim, Saurabh Sinha, and Alessandro Orso. 2025. LlamaRestTest: Effective REST API Testing with Small Language Models. *Proc. ACM Softw. Eng.* 2, FSE (2025), 24 pages. doi:10.1145/3715737
- [18] 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)*. IEEE, 919–931.
- [19] Tsz-On Li, Wenxi Zong, Yibo Wang, Haoye Tian, Ying Wang, Shing-Chi Cheung, and Jeff Kramer. 2023. Nuances are the key: Unlocking chatgpt to find failure-inducing tests with differential prompting. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 14–26.
- [20] Yun Lin, You Sheng Ong, Jun Sun, Gordon Fraser, and Jin Song Dong. 2021. Graph-based seed object synthesis for search-based unit testing. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1068–1080.
- [21] Yun Lin, Jun Sun, Gordon Fraser, Ziheng Xiu, Ting Liu, and Jin Song Dong. 2020. Recovering fitness gradients for interprocedural Boolean flags in search-based testing. In *Proceedings of the 29th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 440–451.
- [22] David R MacIver, Zac Hatfield-Dodds, et al. 2019. Hypothesis: A new approach to property-based testing. *Journal of Open Source Software* 4, 43 (2019), 1891.
- [23] Simone Mezzaro, Alessio Gambi, and Gordon Fraser. 2024. An empirical study on how large language models impact software testing learning. In *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering*. 555–564.

- [24] Zifan Nan, Zhaoqiang Guo, Kui Liu, and Xin Xia. 2025. Test Intention Guided LLM-Based Unit Test Generation. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*. 1026–1038.
- [25] Pengyu Nie, Rahul Banerjee, Junyi Jessy Li, Raymond J Mooney, and Milos Gligoric. 2023. Learning deep semantics for test completion. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2111–2123.
- [26] ofdrw authors. 2023. OFD Reader & Writer. Commit Version: 91af2eb. <https://github.com/ofdrw/ofdrw/commit/91af2ebc12cb5eb9a4a4a74b96b2a06504e21e7#diff-29573571d1af02d4502f7fc750f378c0ba807a539b9cdfac1899c03dd1c6c789R22>.
- [27] ofdrw authors. 2023. OFD Reader & Writer. Commit Version: 9f82f37. <https://github.com/ofdrw/ofdrw/commit/9f82f37022d9b0d020ca56f40832ebf876b786c4#diff-94ead98a7ae23957cc9c48bb19f5ab36ee48e0e478b75ed480e31b0b779b9a08R515>.
- [28] ofdrw authors. 2025. OFD Reader & Writer. <https://github.com/ofdrw/ofdrw>.
- [29] OpenAI. 2025. Introducing OpenAI o3 and o4-mini. <https://openai.com/index/introducing-o3-and-o4-mini/>.
- [30] Carlos Pacheco and Michael D Ernst. 2007. Randoop: feedback-directed random testing for Java. In *Companion to the 22nd ACM SIGPLAN conference on Object-oriented programming systems and applications companion*. 815–816.
- [31] pitest authors. 2025. State of the art mutation testing system for the JVM. <https://github.com/hcoles/pitest>.
- [32] Binhang Qi. 2026. Website of TestGeneralizer. <https://github.com/code-philias/TestGeneralizer>.
- [33] Binhang Qi, Yun Lin, Xinyi Weng, Yuhuan Huang, Chenyan Liu, Hailong Sun, Zhi Jin, and Jin Song Dong. 2025. Intention-driven generation of project-specific test cases. *arXiv preprint arXiv:2507.20619* (2025).
- [34] Binhang Qi, Hailong Sun, Wei Yuan, Hongyu Zhang, and Xiangxin Meng. 2021. Dreamloc: A deep relevance matching-based framework for bug localization. *IEEE Transactions on Reliability* 71, 1 (2021), 235–249.
- [35] 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* (2023).
- [36] Koushik Sen, Darko Marinov, and Gul Agha. 2005. CUTE: A concolic unit testing engine for C. *ACM SIGSOFT Software Engineering Notes* 30, 5 (2005), 263–272.
- [37] Jiho Shin, Sepehr Hashtroudi, Hadi Hemmati, and Song Wang. 2024. Domain Adaptation for Code Model-Based Unit Test Case Generation. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 1211–1222.
- [38] Deepika Tiwari, Yogya Gamage, Martin Monperrus, and Benoit Baudry. 2024. PROZE: Generating Parameterized Unit Tests Informed by Runtime Data. In *2024 IEEE International Conference on Source Code Analysis and Manipulation (SCAM)*. 166–176. doi:10.1109/SCAM63643.2024.00025
- [39] Deepika Tiwari, Long Zhang, Martin Monperrus, and Benoit Baudry. 2021. Production monitoring to improve test suites. *IEEE Transactions on Reliability* 71, 3 (2021), 1381–1397.
- [40] 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* (2020).
- [41] Vasudev Vikram, Caroline Lemieux, Joshua Sunshine, and Rohan Padhye. 2023. Can large language models write good property-based tests? *arXiv preprint arXiv:2307.04346* (2023).
- [42] Hongtai Wang, Ming Xu, Yanpei Guo, Weili Han, Hoon Wei Lim, and Jin Song Dong. 2025. RulePilot: An LLM-Powered Agent for Security Rule Generation. *arXiv:2511.12224* [cs.CR] <https://arxiv.org/abs/2511.12224>
- [43] Jin Wen, Qiang Hu, Yuejun Guo, Maxime Cordy, and Yves Le Traon. 2025. Variable Renaming-Based Adversarial Test Generation for Code Model: Benchmark and Enhancement. *ACM Transactions on Software Engineering and Methodology* (2025).
- [44] Chunqiu Steven Xia, Matteo Paltenghi, Jia Le Tian, Michael Pradel, and Lingming Zhang. 2024. Fuzz4all: Universal fuzzing with large language models. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
- [45] Chen Yang, Junjie Chen, Bin Lin, Jianyi Zhou, and Ziqi Wang. 2024. Enhancing LLM-based Test Generation for Hard-to-Cover Branches via Program Analysis. *CoRR* abs/2404.04966 (2024).
- [46] Zhiqiang Yuan, Mingwei Liu, Shiji Ding, Kaixin Wang, Yixuan Chen, Xin Peng, and Yiling Lou. 2024. Evaluating and Improving ChatGPT for Unit Test Generation. *Proc. ACM Softw. Eng.* 1, FSE, Article 76 (jul 2024), 24 pages. doi:10.1145/3660783
- [47] Junwei Zhang, Xing Hu, Shan Gao, Xin Xia, David Lo, and Shanping Li. 2025. Less Is More: On the Importance of Data Quality for Unit Test Generation. *Proc. ACM Softw. Eng.* 2, FSE, Article FSE059 (June 2025), 24 pages. doi:10.1145/3715778

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009