A System for Automated Unit Test Generation Using Large Language Models and Assessment of Generated Test Suites

Andrea Lops
Polytechnic University of Bari & Wideverse
Bari, Italy
andrea.lops@poliba.it

Fedelucio Narducci
Polytechnic University of Bari
Bari, Italy
fedelucio.narducci@poliba.it

Azzurra Ragone
University of Bari
Bari, Italy
azzurra.ragone@uniba.it

Michelantonio Trizio
Wideverse
Bari, Italy
michelantonio.trizio@wideverse.com

Claudio Bartolini
Independent Researcher
claudioluisbartolini@gmail.com

Abstract—Unit tests represent the most basic level of testing within the software testing lifecycle and are crucial to ensuring software correctness. Designing and creating unit tests is a costly and labor-intensive process that is ripe for automation. Recently, Large Language Models (LLMs) have been applied to various aspects of software development, including unit test generation. Although several empirical studies evaluating LLMs' capabilities in test code generation exist, they primarily focus on simple scenarios, such as the straightforward generation of unit tests for individual methods. These evaluations often involve independent and small-scale test units, providing a limited view of LLMs' performance in real-world software development scenarios. Moreover, previous studies do not approach the problem at a suitable scale for real-life applications. Generated unit tests are often evaluated via manual integration into the original projects, a process that limits the number of tests executed and reduces overall efficiency. To address these gaps, we have developed an approach for generating and evaluating more real-life complexity test suites. Our approach focuses on class-level test code generation and automates the entire process from test generation to test assessment. In this work, we present AGONETEST: an automated system for generating test suites for Java projects and a comprehensive and principled methodology for evaluating the generated test suites. Starting from a state-of-the-art dataset (i.e., METHODS2TEST), we built a new dataset for comparing human-written tests with those generated by LLMs. Our key contributions include a scalable automated software system, a new dataset, and a detailed methodology for evaluating test quality.

Index Terms—Software Testing, Large Language Model, Automatic Assessment

I. INTRODUCTION

Software testing is a critical step in the software development lifecycle, essential for ensuring code correctness and reliability. Within it, unit testing is the stage concerned with verifying the proper functioning of individual code units. Designing and building unit tests is a costly and labor-intensive process that requires significant time and specialized skills. Automating this process represents a promising area for research and development.

Automated tools for generating unit tests can reduce test engineers' and software developers' workload. These tools typically use static code analysis methods to generate test suites. For example, EvoSuite [1], a popular tool that combines static code analysis with evolutionary search, has been demonstrated to achieve adequate coverage.

Large Language Models (LLMs), efficiently exploited in various aspects of software development, could also handle the automatic generation of unit tests. Several empirical studies on LLMs have highlighted their ability to generate tests for simple scenarios, often limited to single methods [2]–[5]. Though directionally useful, these explorations focus on independent and small-scale test units, providing a limited view of LLMs' performance in real-world software development scenarios [6]. Moreover, previous studies do not approach the problem at a suitable scale for real-life examples. Generated unit tests are often evaluated via manual integration into the original projects, a process that limits the number of tests executed and reduces overall efficiency.

To address these gaps, we have developed an approach for generating and evaluating test suites that are more representative of real-life complex software projects. Our approach focuses on class-level test code generation and automates the entire process from test generation to test assessment.

In this work, we introduce AGONETEST, an automated system designed to generate test suites for Java projects, accompanied by a rigorous and systematic methodology to evaluate these generated test suites. Leveraging the METHODS2TEST dataset [7], we developed a new dataset specifically aimed at comparing human-written tests with those produced by LLMs. We integrate libraries such as JaCoCo, PITest, and TsDetect to compute the metrics for test evaluation.

The main contributions of our work are as follows:

- AGONETEST: we designed and developed a closed-loop, highly automated software system supporting the process of generation and assessment of unit tests, working at scale. This initial incarnation of the system works on real-life open-source Java projects integrating essential libraries like JaCoCo, PITest, and TsDetect;
- A methodology, underpinned and embodied into AGONETEST, for comprehensive evaluation of a variety of LLMs and relative prompting techniques and prompt schemata in the task of developing unit tests, and a set of metrics and test smells to assess the quality of the generated test suites;
- CLASSES2TEST[1]: An annotated open source Java project dataset extending METHODS2TEST [7], which maps focal classes to their related test classes. This extended dataset makes it possible to assess the test performance of an LLM on the entire class, rather than on a single method.

The paper is organized as follows. Section II sets the background and highlights differences between our work and related work. Section III gives an overview of AGONETEST and its modules, detailing their functional scope. Then, Section IV showcases how AGONETEST is applied in practice through an end-to-end example. Section V addresses key research questions, presenting a first evaluation of the framework, while Section VI highlights insights and lessons learned from our experiments. Section VII discusses the limitations of our approach and Section VIII concludes the paper, outlining potential directions for future work.

II. BACKGROUND AND RELATED WORK

A. Unit Test Generation

Unit test generation is the automated process of creating test cases for individual software components, such as functions, methods, or modules. These test cases are used to independently verify the correct functioning of each unit.

Present techniques employ randomness-based [8], [9], constraint-based [10], [11], or search-based approaches [12], [13]. The core idea behind these methods is to transform the problem into one that can be solved mathematically. For example, search-based techniques convert testing into an optimization problem, to generate unit test cases [14]. Consequently, the objective of these techniques is to generate all potential solutions and then select those that achieve better code coverage. EvoSuite [1] works by accepting a Java class or method as input and applying search-based algorithms to generate a test suite that meets coverage criteria such as code or branch coverage. EvoSuite assesses test fitness using iterative processes of variation, selection, and optimization. Not only does it generate JUnit test cases, but it also provides a comprehensive report produced by inspecting the efficiency of the created test suite, based on metrics such as code coverage and mutation score. One limitation of EvoSuite is that it often produces tests that lack clarity and readability

[15]. Additionally, EvoSuite can only be used on projects using Java 9 or lower, which limits its applicability to more modern Java projects (the last Java version, at the present time, is 22). Unlike EvoSuite, AGONETEST incorporates advanced evaluation metrics and test-smell recognition, providing a more comprehensive assessment of the quality of generated test suites and ensuring readability by leveraging human-like LLM-generated code. Moreover, AGONETEST supports all Java LTS versions, allowing projects built on newer versions to be tested as well, overcoming the compatibility limitations of EvoSuite.

B. Large Language Models for Test Generation

Since the emergence of LLMs, they have been used for test suite generation. The first techniques exploiting LLMs were thought of as solutions to neural machine translation problems [16], [17]. Such approaches work by translating from the primary method to the appropriate test prefix or test assertion while also fine-tuning the LLMs using the test generation dataset. For instance, AthenaTest [17] optimizes BART [18] using a test generation dataset in which the source is the primary method along with its corresponding code context, and the result is the complete test case. AthenaTest focuses mainly on generating method-level tests by fine-tuning a single model, while AGONETEST shifts the focus to the generation of class-level tests. Our approach makes it possible to use up-to-date LLMs and not constrain prompt design, thereby handling more complex, real-world scenarios. In light of the rapid evolution of instruction-tuned LLMs, the proliferation of methods for generating tests is on the rise, exploiting guided LLMs through appropriate prompts, as opposed to model fine-tuning [19], [20]. Several proposals for evaluating LLMs in test suite generation have emerged. For example, CHATTESTER [5] proposes a tool for evaluating and improving LLM-generated tests based on ChatGPT. ChatTester focuses on improving and evaluating tests generated by a specific LLM (ChatGPT), but requires human intervention to evaluate the generated code and does not provide an evaluation of class-level tests on multiple LLMs. AGONETEST provides support instead for a variety of LLMs and evaluates each LLM's performance on a wide range of real-life Java projects. TESTPILOT [3] is also focused on generating and improving tests using LLMs on JavaScript code. Although TestPilot performs an automated evaluation, it lacks wider applicability to projects other than the 25 repositories it considers in the work provided as reference here. AGONETEST offers far broader applicability by using a dataset of 9,410 Github repositories, and automatically integrating test libraries into them. CEDAR [21] instead proposes a prompt construction strategy based on *few-shot learning* [22] and the Codex model[2] to generate tests. Cedar uses a specific prompt construction strategy, but it does not incorporate a structured mechanism to evaluate multiple LLMs and prompt techniques in a unified framework. AGONETEST provides this by allowing the integration and evaluation of various

---

[1]https://anonymous.4open.science/r/classes2test

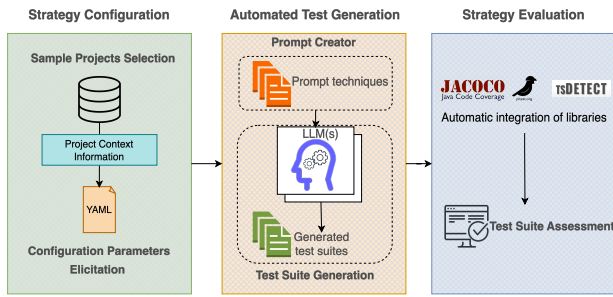[2]https://openai.com/index/openai-codex/

Fig. 1. Overview of AGONETEST framework

prompt engineering techniques and LLMs, offering a more holistic approach to test generation. Guilherme and Vincenzi [2] use gpt-3.5-turbo in analyzing the impact of variation in model hyperparameters. The study by Guilherme and Vincenzi presents an initial assessment but lacks automation in evaluating comprehensive test quality metrics like mutation coverage and test smells. AGONETEST goes a step further by automating these evaluations, integrating advanced metrics to provide a deeper analysis of the generated tests. Siddiq et al. [4] offer a new proposal for evaluating tests generated using common datasets and experimenting with the use of new metrics [23]. Although Siddiq et al. use Test Correctness (but not mutation coverage) on top of all the metrics that AGONETEST uses, their approach does not fully automate the test generation-execution-evaluation loop or focus on class-level tests. AGONETEST fills this gap by providing end-to-end automation and focusing on generating and evaluating complex, class-level test suites.

C. Limits of Current Approaches in Applying LLMs to Unit Test Generation

While promising, current approaches in applying LLMs to unit test generation exhibit several limitations:

a) Limited Scope: Current methods for evaluating how useful LLMs are in test code generation are mostly limited to generating code segments, rather than whole modules or components (e.g., whole classes in Java). Consequently, the research community lacks dedicated datasets for evaluating class-level test generation. To the best of our knowledge, studies often provide only punctual and anecdotal evaluations of the generated results [3], [5], [17].

b) Lack of Automation: No work has emerged in the literature that fully automates the test generation-execution-assessment loop, which is crucial for comprehensive and scalable testing [3], [19], [20]

c) Subjective Choice of Prompts: In most cases, the choice of prompts to get LLMs to generate testing code remains subjective. There is no thorough evaluation of alternate prompting techniques compared to those initially proposed, leaving room for further exploration and optimization in prompt engineering. [4], [20], [21].

III. OVERVIEW OF AGONETEST

The term *agone*, originating from ancient Greece and Rome, signified a contest wherein philosophers debated their ideas, with the audience determining the victor. We adopt the term *agone* metaphorically to represent the competitive evaluation of LLMs and their respective prompting strategies within an arena aimed at generating optimal unit test suites. AGONETEST determines the optimal strategies based on standard test quality metrics, which we elaborate on in subsequent sections.

AGONETEST is designed to provide software testers with a system for generating and assessing unit tests. This assessment focuses on key metrics such as code coverage and the presence of known test smells, thereby offering a comprehensive assessment of test suite quality.

AGONETEST operates on the principle that the evaluation of LLMs in the task of generating high-quality unit tests can be performed through the collaboration of test engineers and data scientists (or prompt engineers). However, in practice, a single experienced test engineer familiar with generative AI can perform both roles, allowing the focus to be only on defining new prompt techniques and the comparison of LLMs. This is the persona that we evoke when we refer to the AGONETEST user (alternatively "the test engineer") in the remainder of this paper.

The system helps test engineers through the following phases:

- **Strategy Configuration**
- **Automated Test Generation**
- **Strategy Evaluation**

Figure 1 provides a high-level diagram of the architecture of AGONETEST, showing the operating modules that streamline the test generation and evaluation process. These modules are described as follows:

Sample Projects Selection (Strategy Configuration - I): As an initial configuration step, the user chooses which repositories to generate test suites for. This initial phase leverages a comprehensive dataset of annotated open-source Java repositories, which we contribute to the community. It involves preparing, loading, and managing the repositories to be tested by the system.

Configuration Parameters Elicitation (Strategy Configuration - II): In this phase, configuration parameters are elicited from the selected repositories (e.g., the project java version, used testing framework, etc.) and processed to create prompts templates for the LLMs.

Prompt Creation (Automated Test Generation - I): During this phase, the prompt templates used in the previous phases are fully instantiated and then used to generate unit test suites in the next step.

Test Suite Generation (Automated Test Generation - II): Here, AGONETEST orchestrates the interaction with the selected LLMs, feeding them the instantiated prompts to produce the unit test code. Each LLM generates test classes that are then integrated into the project structure.

Characteristic	Value
Total Test Classes	147,473
Total Unique Repositories	9,410
Total Lines of Code	173,736,517
Total Cyclomatic Complexity	81,566,509

Test Suite Assessment (Strategy Evaluation): This phase assesses the quality of the test suites computing various metrics and identifying test smells.

This assessment enables a detailed analysis of the effectiveness and quality of the automated test generation strategies.

In the following, we describe each phase of the process in detail.

A. Sample Repositories Selection

We contribute a comprehensive, annotated dataset of open-source Java repositories from GitHub[3], which we leverage in this phase. Unlike popular datasets in the literature, our dataset enables the generation and validation of unit tests at the Java class level, rather than at the individual method level. To collate and annotate our dataset, we built on METHODS2TEST [7]. METHODS2TEST is a supervised dataset consisting of several Java methods, called *focal methods*, univocally mapped to their respective test methods. It is a real-world dataset built using data from 9,410 open-source git repositories selected from an initial analysis of 91,385 repositories [4]. The dataset comprises 780,944 instances. We chose METHODS2TEST as the starting point, as it contains not only the test methods to be tested, but also the corresponding test methods written and validated by humans. Human-written tests are a valid evaluation benchmark to evaluate the effectiveness of different LLMs in building a test suite.

We developed CLASSES2TEST to enable automatic performance evaluation of LLMs in generating unit tests for entire Java classes, rather than individual methods, since this scenario is more representative of real-world applications. To accomplish this, we extracted all references to the open-source repositories present in METHODS2TEST in order to map the Java classes, referred to as *focal classes*, to their corresponding test classes.

Here is the process we followed to create CLASSES2TEST:
1) Extract the repository reference, Github URL and selected branch;
2) Select the classes considered in METHODS2TEST;
3) Clone the repository and save the commit hash;
4) Map and save the focal classes along with their respective test classes.

The resulting dataset contains 147,473 test classes extracted from 9,410 unique repositories. A summary of the dataset's characteristics is shown in Table I.

---

[3]https://github.com
[4]We point out that the repositories have received updates in the last five years and are not forks, ensuring that they are actively maintained and representative of real-world software development practices.

B. Configuration Parameters Elicitation

Before unit test generation can begin, the system extracts some parameters from the projects selected in the previous step. These parameters are then fed into the module that selects prompts and LLMs. To query the model under examination, various prompting techniques are available and can be chosen [24].

The configuration parameters include:
- **focal_class**: This variable contains the Java class for which the test suite must be generated;
- **testing_framework**: This variable provides the name and version of the project's testing framework (e.g., JUnit 4), directly extracted from the project during execution;
- **java_version**: This variable allows you to retrieve the version of Java that the project uses.
- **example_focal_class & example_test_class**: These variables contain an example focal class and the corresponding test class extracted from a reference repository, useful to provide an example to the LLM if one wants to use the few-shot prompting technique;
- **example_testing_framework & example_java_version**: These variables provide the information about the example repo.

See Section IV-A for an example of a real implementation.

C. Prompt Creation

In this phase, the prompt templates described in the previous phases are fully instantiated to create viable prompts to guide the LLM in generating unit tests. We populate the user-supplied prompt structures by replacing the variables outlined in Section III-B.

It has to be noted that, in order to make sure our experiments and findings are reproducible, we prepared CLASSES2TEST by saving the commit hashes of the repositories used as sources. This allows AGONETEST to consistently extract information such as the Java version used, the type of test framework (e.g., jUnit), and its version.

Unlike previous approaches to creating unit testing with LLMs that require human intervention to input context information [2], [5], AGONETEST automates the process to a far greater degree. AGONETEST employs ElementTree [25] and a parser to read and modify the Maven and Gradle build (see Section III-E3). It analyzes the libraries present and the Java version used in each build system. This method, along with the ability to use examples, offers users a versatile system for generating prompts.

D. Test Suite Generation

At this point in the process, we have everything we need for the selected LLMs to generate test suites for each focal class of the project. To ensure each model has an appropriate number of tokens, we use tiktoken[5], a BPE tokenizer [26], to evaluate the token count in the prompt. If the limit is

---

[5]https://github.com/openai/tiktoken

exceeded, AGONETEST returns an error to the user, specifying the number of tokens exceeded.

We remark that AGONETEST allows users to evaluate a wide range of LLMs automatically. This capability is provided by the open-source LiteLLM library[6], which facilitates communication with more than 100 models[7] using a standard interaction based on the OpenAI API format[8]. Integration is made easier by LiteLLM, which translates inputs to satisfy the unique endpoint needs of each provider. This is crucial in today's environment, where the absence of standard API specifications for LLM providers makes it challenging to incorporate several LLMs into projects.

After invoking the LLM, AGONETEST selects relevant information from the LLM's answer (i.e., the generated test class). This step is crucial for automating the entire process, since LLMs can provide detailed descriptions or explain how the code should be structured without actually generating it. [27]. In this component, AGONETEST removes unnecessary parts (like outline descriptions) and creates a new file to integrate the test class into the project.

E. Test Suite Assessment

Here we evaluate the quality of the test suite according to the quality metrics and the test smells described below. The actual determination of metrics and test smells is done via library integration, allowing for fully automated test suite assessment. It is important to note that this component is separate from the experimental evaluation discussed later. Instead, it serves as an additional tool provided by AGONETEST to assist engineers in assessing the quality of the generated tests.

1) Coverage Metrics:

- **Line coverage** [28]: This metric measures the percentage of lines of code executed during the testing process. A 100% line coverage means that every line of code in the software has been run at least once during testing. We selected it because it provides direct visibility over the portion of the source code that is being tested.
- **Method coverage** [28]: Similar to line coverage, this metric focuses on the specific methods or functions in the code. A 100% method coverage score means that all methods have been run at least once during testing. This metric is useful to identify methods that may not have been adequately tested.
- **Branch coverage** [28]: This metric calculates the percentage of decision points (such as `if` or `switch` statements) that have been executed in tests. It ensures that all possible paths in the code are tested, which can uncover defects that might be missed by line or method coverage alone.
- **Instruction coverage** [9]: This metric calculates the number of Java bytecode instructions executed during test-

ing. It is a detailed metric, unaffected by source code formatting, and can be determined even without debug information in the class files. This helps in pinpointing the smallest code fragments not covered by tests.
- **Mutation coverage** [28]: This metric evaluates the effectiveness of tests in identifying deliberately introduced changes (mutations) in the code, such as modifying an arithmetic operation or reversing a condition. If the tests detect all mutations (i.e., identify all changes), the mutation coverage score is 100%. This metric was chosen because it measures the robustness of the test suite.

2) Test Smells [23]: These are indicators of inefficient or problem patterns that could negatively affect the maintainability and effectiveness of the test code. Identifying test smells helps improve the quality of the test code over time and raises awareness of potential issues in test design. AGONETEST determines whether the following test smells are present in the code:

- Assertion Roulette (AR) [29]: indicate the number of test methods containing more than one assertion statement without an explanation/message (parameter in the assertion method);
- Conditional Test Logic (CTL) [30]: indicate the number of test methods that contain one or more control statements (i.e., `if`, `switch`, conditional expression, `for`, `foreach` and `while` statement);
- Constructor Initialization (CI) [31]: indicate if the test class contains a constructor declaration;
- Default Test: indicate if the test class is named either 'ExampleUnitTest' or 'ExampleInstrumentedTest';
- Duplicate Assert (DA) [31]: indicate the number of test methods that contain more than one assertion statement with the same parameters;
- Eager Test (EA) [29]: indicate the number of test methods containing multiple calls to multiple production methods;
- Empty Test (EM) [31]: indicate the number of test methods that do not contain a single executable statement;
- Exception Handling (EH) [31]: indicate the number of test methods that contain either a throw statement or a catch clause;
- General Fixture: is 1 if not all fields instantiated within the `setUp` method of a test class are utilized by all test methods in the same test class;
- Ignored Test (IT) [31]: indicate the number of tests methods that contains the `@Ignore` annotation;
- Lazy Test (LT) [29]: indicate the number of test methods calling the same production method;
- Magic Number Test (MNT) [30]: indicate the number of test methods that contain a numeric literal as an argument;
- Mystery Guest: indicate the number of test methods containing object instances of files and databases classes;
- Redundant Print (RP) [31]: indicate the number of tests methods that invokes either the `print`, `println`, `printf` or `write` method of the System class;
- Redundant Assertion (RA) [31]: indicate the number of

test methods that contain an assertion statement in which the expected and actual parameters are the same;

- Resource Optimism (RO) [31]: indicate the number of tests methods utilize an instance of a File class without calling the `exists()`, `isFile()` or `notExists()` methods of the object;
- Sensitive Equality (SE) [29]: indicate the number of tests methods that invokes the `toString()` method of an object;
- Sleepy Test: indicate the number of tests methods that invokes the `Thread.sleep()` method;
- Unknown Test (UT) [31]: indicates the number of test methods that do not contain a single assertion statement and `@Test(expected)` annotation parameter.

3) Library integration: We utilized the following libraries to compute the metrics:

- JaCoCo[10]: JaCoCo is a free Java library used to measure code coverage in test suite execution. It helps developers identify which parts of their code base have been thoroughly tested and which have not, facilitating a better understanding of the test coverage within the project. We selected JaCoCo because of its widespread adoption, ease of integration with build tools, and report-generation features, which are essential for metric evaluation.
- PiTest [32]: PiTest is a mutation testing system for Java and JVM-based systems. It goes beyond traditional line and statement coverage metrics in that it offers more concrete insights into the robustness of a test suite. PiTest introduces minor changes, or *mutations*, into the source code and then re-runs the tests to determine whether these changes are detected. We chose PiTest because it provides a more granular and realistic view of the actual behavior and response of the system under test compared to traditional coverage tools.
- TSDETECT [33]: TSDETECT is a library that focuses on the automatic detection of test smells in software projects. Test smells refer to patterns in test code that may indicate design or implementation issues, leading to less maintainable tests and potentially hindering code comprehension. TSDETECT was chosen for its capability to identify these smells and provide actionable guidelines for code improvement.

In this phase, AGONETEST automatically includes these libraries into the project. For each run, AGONETEST checks the configuration files of the supported build systems (Maven and Gradle, Section III-C) to determine if the necessary libraries are already present. If they are not, it modifies the configuration to add the required dependencies.

AGONETEST demonstrates a high degree of automation, as illustrated by its handling of the PiTest library. Specifically, if the repo uses the JUnit 5 test framework, an additional library, *"pitest-junit5-plugin"*, is required. Utilizing information extracted from the repo in the Prompt Creation module (Section III-C), AGONETEST automatically identifies the test

framework in use and adds this dependency without any human intervention.

4) Automate Test Suite assessment: After adding the necessary libraries, AGONETEST runs a build and test to ensure there are no compilation errors. The test suite assessment phase of our process presents a high degree of automation, as we describe below.

AGONETEST generates a report with the results of the test smells and metrics computed for the LLM-generated tests. To achieve this, the tool automatically retrieves detailed information from the reports produced by the libraries, compiling these data for each class within each project.

This extensive computation process enables a detailed analysis of the generated test suites. By contrasting the results, the module helps identify specific strengths and weaknesses associated with each LLM and prompt configuration. It provides insights into areas where the LLMs excel and highlight potential gaps where improvements are needed.

Furthermore, this comparison facilitates a clear understanding of the nuances in how different LLMs and prompts impact the quality of test generation. It supports the identification of optimal configurations for generating high-quality tests. This detailed analysis is crucial for refining LLMs and enhancing their capabilities in automated test generation.

By providing such in-depth evaluations, AGONETEST serves as a valuable tool for researchers and developers. It aids in the continuous improvement of LLMs and contributes to advancements in the field of automated testing. Ultimately, it can ensure that the tests generated are robust and reliable, improving the effectiveness of automated testing solutions.

IV. AGONETEST IN PRACTICE

In this section, we will demonstrate how AGONETEST operates in practice by describing an end-to-end run of a practical example.

We will skip the repository selection phase in our account and move straight to the configuration phase, which concerns LLM selection and prompt specification. Then, we will exemplify how the results are presented back to the user for further analysis.

A. Configuration

As described in Section III-B, AgoneTest utilizes a YAML file as input, where it is possible to specify information related to two elements: `llms` and `prompts`. The YAML file represented in the Listing 1 declares usage of 'gpt-4'[11] and 'gpt-3.5 turbo'[12] models, both provided by OpenAI[13].

Listing 1. Setup of the YAML configuration file: setting of variables for two different LLMs and two different prompts.
```
llms:
- model: gpt-4-1106-preview
  temperature: 0
- model: gpt-3.5-turbo
  temperature: 0
```

```
prompts:
- name: zero-shot
  value:
  - role: system
    content: You are provided with Java class.
      Create a test class that fully tests the
      proposed Java class using the project
      information for imports. Reply with code
      only, do not add other text that is not code
      .
  - role: user
    content: "The project uses {{testing_framework}}
      and Java {{java_version}} and Java class is
      : \n<code>\n {{focal_class}}\n</code>"

- name: few-shot
  value:
  - role: system
    content: You are provided with an example with a
      Java class and its test class. You are then
      provided with a new Java class. Take a cue
      from the example and create a test class
      that fully tests the new proposed Java class
      . Reply with code only, do not add other
      text that is not code.
  - role: user
    content: "#Example:\nThe example Java class is:\
      n<code>\n {{example_java_class}} \n</code>\
      nThe example test class is: \n<code>\n {{
      example_test_class}} \n</code>.\nThe Java
      class you must create the test for is: \n<
      code>\n{{focal_class}}\n</code>"
```

Prompt specification declaration consists of two sections: `name` and `value`. `name` is an identifier for labeling the type of prompt (zero-shot, few-shots, etc.). In contrast, `value` is an array of message elements of type OpenAI [14]. Each individual message includes a `role` and a `content`.

`role` can be of the following types:

- `system`: This is used to instruct the model on the behavior it should adopt.
- `user`: This is used to indicate the request for the generation of the test class.

In the YAML configuration file, two types of prompts are specified: zero-shot and few-shot.

1) Zero-shot: Zero-shot refers to presenting the model with a single instance of a request or task without any previous examples for the model to draw upon [34]. This method emphasizes the model's ability to comprehend and accurately execute the given task.

2) Few-shot: Unlike zero-shot prompting, few-shot prompting involves providing the model with examples demonstrating the expected inputs and outputs [22]. This technique aids in contextual learning by including examples in the prompt, thereby guiding the model towards improved performance. These examples serve as a conditioning for the actual request, helping the model generate more accurate and relevant responses.

This configuration file will instruct AGONETEST to perform the steps described in Section III-C: it will fully instantiate template variables, including focal class, test frameworks used

[14]https://platform.openai.com/docs/api-reference/chat/create#chat-create-messages

by the repos (and versions thereof), and version of the JDK used in the repos.

B. Results presentation

After running the generation phase, AGONETEST generates a CSV file including, for each LLM selected and each prompting technique, the metrics computed for the focal classes as well as the results about test smells. As a way of example, Table II displays an extract of this file containing as well the results for the human-written tests, as they were present in the CLASSES2TEST dataset.

By examining this file, users can gain valuable insight into the strengths and weaknesses of each LLM and the prompt combination. Plus, software testers can accurately assess the effectiveness of the LLM in creating usable and effective class-level tests. How this is done is made clear in the following section, where we describe our experimental setup for validation of AGONETEST and discuss some results.

V. EVALUATION

In this experimental evaluation, we aim to address the following research questions:

- **RQ1:To what extent is it possible to implement an automated end-to-end process for generating test suites?** We analyze the degree of automation of the framework and the points (if any) where we need the human-in-the-loop.
- **RQ2: Can the quality of test suites automatically generated by different LLMs and prompt strategies be effectively assessed?** We investigate whether the framework can provide information about the quality of the test suite in terms of efficiency and robustness and help identify strengths, weaknesses, and potential improvements.

A. Dataset

In our experiment, we randomly selected 10 repositories from our dataset CLASSES2TEST. These repositories contain a total of 94 focal classes of various lengths and complexity, as shown in Table III. The size of the sample of randomly selected repositories is chosen to be representative enough of the variability encountered in real-world projects (usually comprising of one to a handful of co-dependent repositories), while ensuring that is tractable by our system in terms of scale.

B. LLMs and prompts configuration

For our experiment, we selected two LLMs from the models supported by LiteLLM[15]. We have chosen the 'gpt-4'[16] and 'gpt-3.5 turbo'[17] models. The 'gpt-4' model was selected for its outstanding performance on the HumanEval benchmark [35], while 'gpt-3.5 turbo' was chosen as an earlier generation model, allowing a meaningful comparison.

[15]https://docs.litellm.ai/docs/providers
[16]https://openai.com/index/gpt-4/
[17]https://platform.openai.com/docs/models/gpt-3-5-turbo

TABLE II

CSV FILE EXTRACT OF AGONETEST RESULTS. THE TEST SMELLS ARE FROM COLUMN 10 ONWARDS. PLEASE REFER TO SECTION III-E FOR ACRONYMS

model	prompt name	Project	Focal Class	instruction coverage	branch coverage	line coverage	method coverage	mutation coverage	AR	CTL	CI	DA	EA	EM	EH	IT	LT	MNT	RP	RA	RO	SE	UT
gpt-3.5-turbo	zero-shot	8313187	LDMLPredicateParser	0,72	0,5	0,67	0,5	0,5	6	1	0	0	0	0	1	0	11	11	10	0	0	0	0
gpt-3.5-turbo	few-shot	8313187	ConfigResourceBundleParser	0,53	1	0,67	0,5	0	8	0	0	0	0	1	1	0	10	10	13	0	0	0	0
gpt-3.5-turbo	few-shot	8313187	LDMLPredicateParser	0,72	0,5	0,67	0,5	0,5	5	1	0	0	0	1	0	0	11	12	1	0	0	0	0
gpt-4	zero-shot	8313187	ResourceResolutionContext	0,85	0,66	0,83	1	0,66	2	0	0	0	0	0	0	0	9	9	5	0	0	0	0
gpt-4	zero-shot	8313187	LDMLPredicateParser	0,88	1	0,83	0,5	1	3	0	0	0	0	0	0	0	9	9	1	0	0	0	0
gpt-4	few-shot	8313187	ResourceResolutionContext	0,79	0,64	0,78	0,96	0,6	1	0	0	0	0	0	0	0	9	9	5	0	0	0	0
gpt-4	few-shot	8313187	LDMLPredicateParser	0,72	0,5	0,67	0,5	0,5	2	1	0	0	0	0	0	0	9	9	1	0	0	0	0
human	-	8313187	ConfigResourceBundleParser	1	1	1	1	1	8	0	0	0	0	0	0	0	15	10	0	0	0	0	0
human	-	8313187	RefreshableResources	0,87	0,68	0,87	0,93	0,66	1	0	0	0	0	0	0	0	4	4	0	0	0	4	0
human	-	8313187	ResourceResolutionContext	0,8	0,72	0,81	0,87	0,69	6	0	0	0	0	0	0	0	15	11	0	0	0	0	0
human	-	8313187	LDMLPredicateParser	0,88	1	0,83	0,5	1	5	0	0	0	0	0	0	0	4	4	0	0	0	1	0
...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...

TABLE III

CODE CHARACTERISTICS OF THE REPOSITORIES SELECTED FOR THE EXPERIMENT

Characteristics	Value
Java Repositories	10
Total Focal Classes	94
Total Lines of Code	189,703
Total Cyclomatic Complexity	85,811

In addition to choosing LLMs, adjusting their temperature parameter is crucial. This feature allows users to control the level of randomness and creativity in the generated text. By varying the temperature, users can influence the breadth and exploration of the results. High-temperature settings, such as 1.0, introduce greater randomness, resulting in more creative and diverse outcomes. Conversely, a lower temperature setting, like 0.2, produces more focused and deterministic results, leading to predictable and cautious outcomes [2]. By carefully adjusting the temperature parameter, users can balance inno- vation and coherence in text generation, ensuring the output aligns with their specific task or application requirements. As shown in Listing 1, we set the temperature to 0 in our experiment to increase the level of coherence in text generation (and to decrease the level of randomness) and make the diverse test suite generated comparable. Regarding prompt types, we decided to experiment with two of the most popular techniques: zero-shot and few-shot. This allows us to show the flexibility of AGONETEST and further explore the use of the contextual variables seen in Section III-B. In our experiment, the example provided in the few-shot prompt is always the same across all LLMs to ensure uniformity. The example consists of a focal class and a test class extracted from an open sample repository[18].

C. Data collection and analysis

Not all test suites generated from our sample repositories were valid (i.e. tests did not pass, and in some cases, the code did not even compile). Similarly to [2], we encountered some issues: some tests failed to build due to syntax errors or incorrect or non-existent imports. To move to the assessment phase, the system automatically removes test classes with errors in order to proceed with the compilation. Once all errors

are removed, the system conducts a compile-and-run test to eliminate all classes that are not *green suite*, i.e. where all methods are executed without failures.

This is done for two reasons: first, because after performing root cause analysis, we found that the failing test classes did so because they were incorrectly specified, rather than for the presence of bugs. For example, tests failed because they invoked private or outdated methods of libraries. Or they failed because they made calls to non-existent APIs. The second reason is that having a *green suite* is a necessary condition for PiTest to calculate mutation coverage [19]. After completing this clean-up phase, the system performs a final compilation with the execution of the tests and the library to collect data for evaluation.

TABLE IV

NUMBER OF CLASSES PER EXPERIMENT THAT COMPILE AND FOR WHICH ALL TESTS ARE PASSED

model	prompt name	Build	Pass	Total Rejected
gpt-3.5-turbo	zero-shot	64 (68.08%)	36 (38.29%)	58 (61.70%)
gpt-3.5-turbo	few-shot	65 (69.14%)	35 (37.23%)	59 (62.76%)
gpt-4	zero-shot	76 (80.85%)	29 (30.85%)	65 (69.14%)
gpt-4	few-shot	76 (80.85%)	28 (29.78%)	66 (70.21%)
human	-	94 (100%)	94 (100%)	94 (100%)

Table IV shows the percentage of classes that compile and for which all tests are passed. On average, we have seen in our experiment that:

- 75% of the generated test classes compile successfully;
- 34% of generated test classes are *green suite* and have calculable mutation coverage.

Table V shows a comparative analysis of the performance of the different combinations of LLMs and prompt techniques with respect to the metrics computed by AGONETEST. We also report the metrics computed for the human-written tests, to have a benchmark.

VI. DISCUSSION AND LESSONS LEARNED

In this section, we will answer the research questions previously defined and discuss the lesson learned together with possible future research directions.

RQ1: To what extent is it possible to implement an automated end-to-end process for generating test suites?

---

[18]https://github.com/junit-team/junit5-samples

[19]https://pitest.org/faq/

model	prompt name	instruction coverage	branch coverage	line coverage	method coverage	mutation coverage
gpt-3.5-turbo	zero-shot	0,756923077	0,706363636	0,776923077	0,848217345	0,546923077
gpt-3.5-turbo	few-shot	0,813333333	0,681111111	0,775555556	0,832222222	0,455555556
gpt-4	zero-shot	**0,879090909**	0,776923077	**0,866363636**	**0,854545455**	0,546363636
gpt-4	few-shot	0,753333333	0,775555556	0,781666667	0,836923077	0,461666667
human	-	0,783611111	**0,808947368**	0,766388889	0,698055556	**0,690555556**

AGONETEST provides an end-to-end automated process to generate and evaluate test suites without human intervention. However, there are two points requiring attention that our experiments underline:

- The compilation success rate of the generated test classes shows room for improvement (in our experiment ranged between 64% and 76%);
- The percentage of tests passed was relatively low (between 30% and 38%).

Further analysis revealed that many generated tests failed due to incorrect imports or syntax errors and not because they discovered previously undetected bugs. To improve on these results and increase such percentages, there are different paths to explore. One is human-in-the-loop: where human intervention might include manually fixing code errors, adjusting settings, or installing required libraries for successful execution.

On the other hand, a good result, in light of automation, is having automated the process of extracting contextual information from project configuration files (such as Maven or Gradle). This minimizes the need for manual intervention and enhances the accuracy of the generated prompts and tests.

RQ2: Can the quality of test suites automatically generated by different LLMs and prompt strategies be effectively assessed? AGONETEST gives relevant information about the quality of the test suite generated, in terms of code coverage, robustness of the test suite to artificially injected bugs (i.e., mutation coverage), and test smells. Indeed, the presence of test smells indicates potential issues with test design and maintainability.

In Table V we can see that LLMs-written test suites already have good quality in terms of coverage, but should improve in terms of robustness. The benchmark of human-written tests shows always better results for mutation coverage. Comparing the output of our LLMs-powered system against tests written by actual test engineers gave us valuable insight into each model-setup pair's capabilities and limitations. Our experiment (Table V) demonstrates that the performance of various LLMs varied significantly depending on the prompting technique used. Surprisingly, we found better results for zero-shot prompts for 'gpt-4' than for the few-shot one. However, it is important to note that this does not mean that 'gpt-4' is the best model for all scenarios, and this is neither the objective of our experiment (i.e., to find the best model). Since the performance of a model can vary significantly depending

on the specifics of the context or the structure of the prompts. For this reason, we created AGONETEST: to provide users with a system that lets them experiment with various combinations to find the LLM and prompt configuration that best fits their specific requirements.

A. Lessons Learned

Throughout the development and evaluation of AGONETEST, we gathered several key insights that will guide future improvements in the framework. These lessons are crucial for refining the system and improving the efficacy of LLM-generated tests. Each subsection below highlights a specific challenge encountered and outlines a potential solution.

1) Compilation and Test Pass Rate: Our experiments show compilation and test pass rates that could be improved in light of the pursuit of full automation. The causes of these are diverse (e.g., import classes that do not exist or are missing). Automating the correction of these recurring problems is possible [36] and will increase the success rate of the generated tests. One promising approach involves asking the LLM itself to analyze errors in the generated test code and provide fixes. By supplying the identified errors as feedback, the LLM can generate corrected and functional code, thereby enhancing the initial output. Additionally, enhancing the robustness of the generated tests by incorporating context-aware validation and fixing mechanisms will ensure that the test suites align closely with the project's specific structures. This integrated approach not only automates error correction but also enhances the overall reliability and effectiveness of the test generation process, moving closer to the goal of fully automated, high-quality test suite production.

2) Performance in Mutation Testing: Human-written tests consistently outperformed LLM-generated tests in terms of mutation coverage, indicating that manually written tests are more effective at identifying code changes introduced through mutation. To address this, we should focus on improving the robustness of the generated test suites by refining the prompting algorithms and incorporating mutation-aware test generation techniques.

3) Scalability and Resource Management: Automating the entire pipeline-downloading projects, generating test suites, integrating libraries, and performing evaluations proved to be resource intensive. Efficiently managing and parallelizing these tasks can alleviate computational overhead and improve

scalability, allowing AGONETEST to handle larger datasets and codebases more effectively.

4) Impact of Prompting Techniques: The choice of prompting technique significantly impacts the quality of the generated tests. Our experiments showed that zero-shot prompting with gpt-4 yielded the best results, but performance varied across different combinations of LLMs and prompts. Systematically exploring and evaluating different prompting strategies will help identify the most effective configurations for various scenarios.

5) Automated Context Extraction: Providing the LLMs with accurate context information, such as the testing framework and the Java version, is essential for generating correct test classes. Automating the extraction of this context information reduces the need for manual intervention and improves the quality of generated prompts and tests. Enhancing the automation of context extraction by developing more sophisticated parsers and context inference algorithms will dynamically adapt to various project configurations.

6) Real-world Applicability: Building the dataset from actual open-source Java repositories on GitHub ensured that AGONETEST operates in real-world scenarios. However, ensuring that the dataset is representative of real-life situations across different types of repos and codebases remains an ongoing goal. To maintain and improve real-world applicability, we should continuously upgrade and update our dataset to include a broader range of real-world repositories and project structures, ensuring that the evaluation remains relevant and comprehensive.

These lessons direct us towards further improvements in AGONETEST. By implementing these improvements, we aim to develop a more robust, efficient, and reliable framework for automated unit test generation.

VII. LIMITATIONS

Although AGONETEST presents an innovative framework for automating the generation and evaluation of unit test suites using LLMs, several limitations should be acknowledged regarding its current implementation and the first experimental results.

A. *Dataset and Generalization*

For our evaluation, we relied on the newly created CLASSES2TEST dataset, derived from METHODS2TEST. Although this dataset is designed to evaluate class-level test generation, its scope is limited to Java projects. This makes our findings hardly generalizable to different programming languages. Moreover, the repositories included in CLASSES2TEST were selected based on their ability to compile without errors, potentially introducing a bias towards well-structured codebases.

B. *Model and Prompt Variability*

Limited Number of LLMs and Prompts Tested: Although AGONETEST supports various LLMs and prompting techniques, our initial experimental setup involved only two models (gpt-4 and gpt-3.5 turbo) and two prompt types (zero-shot

and few-shot). As LLMs and prompt engineering techniques continue to evolve, the results might vary significantly with newer models and advanced prompts. The limited scope of our initial tests could thus restrict the breadth of our conclusions.

Temperature Setting Restrictions: In our experiment, we set the temperature parameter to 0 to ensure consistency and reproducibility. Although this reduces randomness and increases coherence, it may inadvertently limit the creativity and diversity of the generated test cases. Different temperature settings could yield more varied results, which were not explored in this study.

C. *Compilation and Execution Failures*

A notable limitation observed during our experiments was the non-compilation and execution failures of some generated test classes. Approximately 66% of the test classes generated were either rejected during the compilation phase or failed to contribute positively to the metrics due to inherent errors. This reveals the current inability of some LLMs to consistently generate syntactically and semantically correct test code, affecting the overall evaluation.

D. *Evaluation Metrics*

While we employed a comprehensive set of metrics and test smell indicators, these metrics alone may not fully capture the quality of the test suite.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we present AGONETEST, a comprehensive framework for automating the generation and assessment of unit test suites using LLMs. This framework focuses on generating complex, class-level test suites while automating the entire testing process from test generation to integration and evaluation.

The results of our experiments demonstrate that AGONETEST can produce and evaluate unit tests across various real-world projects, offering detailed insights into the performance of different LLMs and prompting techniques. While the initial findings are promising, they also highlight challenges emphasizing the need for further refinement.

The automation of unit test generation using LLMs is a promising field. While current capabilities do not yet match those of human engineers for some tasks (such as mutation coverage), promising results in instruction, line, and method coverage indicate that further research and refinement can bridge this gap. Future work should focus on systematic research into the most effective LLMs and prompts, coupled with continuous improvements in automated correction mechanisms for recurrent issues.

REFERENCES

[1] G. Fraser and A. Arcuri, "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*, 2011, pp. 416–419.

[2] V. Guilherme and A. Vincenzi, "An initial investigation of chatgpt unit test generation capability," in *Proceedings of the 8th Brazilian Symposium on Systematic and Automated Software Testing*, 2023, pp. 15–24.

[3] M. Schäfer, S. Nadi, A. Eghbali, and F. Tip, "An empirical evaluation of using large language models for automated unit test generation," *IEEE Transactions on Software Engineering*, 2023.

[4] M. L. Siddiq, J. C. Da Silva Santos, R. H. Tanvir, N. Ulfat, F. Al Rifat, and V. Carvalho Lopes, "Using large language models to generate junit tests: An empirical study," in *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering*, 2024, pp. 313–322.

[5] Z. Yuan, Y. Lou, M. Liu, S. Ding, K. Wang, Y. Chen, and X. Peng, "No more manual tests? evaluating and improving chatgpt for unit test generation," *arXiv preprint arXiv:2305.04207*, 2023.

[6] Y. Tang, Z. Liu, Z. Zhou, and X. Luo, "Chatgpt vs sbst: A comparative assessment of unit test suite generation," *IEEE Transactions on Software Engineering*, 2024.

[7] M. Tufano, S. K. Deng, N. Sundaresan, and A. Svyatkovskiy, "Methods2test: A dataset of focal methods mapped to test cases," in *Proceedings of the 19th International Conference on Mining Software Repositories*, 2022, pp. 299–303.

[8] C. Csallner and Y. Smaragdakis, "Jcrasher: an automatic robustness tester for java," *Software: Practice and Experience*, vol. 34, no. 11, pp. 1025–1050, 2004.

[9] C. Pacheco, S. K. Lahiri, M. D. Ernst, and T. Ball, "Feedback-directed random test generation," in *29th International Conference on Software Engineering (ICSE'07)*. IEEE, 2007, pp. 75–84.

[10] L. Ma, C. Artho, C. Zhang, H. Sato, J. Gmeiner, and R. Ramler, "Grt: Program-analysis-guided random testing (t)," in *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2015, pp. 212–223.

[11] A. Sakti, G. Pesant, and Y.-G. Guéhéneuc, "Instance generator and problem representation to improve object oriented code coverage," *IEEE Transactions on Software Engineering*, vol. 41, no. 3, pp. 294–313, 2014.

[12] J. H. Andrews, T. Menzies, and F. C. Li, "Genetic algorithms for randomized unit testing," *Ieee transactions on software engineering*, vol. 37, no. 1, pp. 80–94, 2011.

[13] P. Derakhshanfar, X. Devroey, and A. Zaidman, "Basic block coverage for search-based unit testing and crash reproduction," *Empirical Software Engineering*, vol. 27, no. 7, p. 192, 2022.

[14] P. Tonella, "Evolutionary testing of classes," *ACM SIGSOFT Software Engineering Notes*, vol. 29, no. 4, pp. 119–128, 2004.

[15] G. Grano, S. Scalabrino, H. C. Gall, and R. Oliveto, "An empirical investigation on the readability of manual and generated test cases," in *Proceedings of the 26th Conference on Program Comprehension*, 2018, pp. 348–351.

[16] P. Nie, R. Banerjee, J. J. Li, R. J. Mooney, and M. Gligoric, "Learning deep semantics for test completion," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2111–2123.

[17] M. Tufano, D. Drain, A. Svyatkovskiy, S. K. Deng, and N. Sundaresan, "Unit test case generation with transformers and focal context," *arXiv preprint arXiv:2009.05617*, 2020.

[18] H. A. Chipman, E. I. George, and R. E. McCulloch, "Bart: Bayesian additive regression trees," *The Annals of Applied Statistics*, 2010.

[19] Y. Deng, C. S. Xia, H. Peng, C. Yang, and L. Zhang, "Large language models are zero-shot fuzzers: Fuzzing deep-learning libraries via large language models," in *Proceedings of the 32nd ACM SIGSOFT international symposium on software testing and analysis*, 2023, pp. 423–435.

[20] C. S. Xia, M. Paltenghi, J. Le Tian, M. Pradel, and L. Zhang, "Fuzz4all: Universal fuzzing with large language models," in *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, 2024, pp. 1–13.

[21] N. Nashid, M. Sintaha, and A. Mesbah, "Retrieval-based prompt selection for code-related few-shot learning," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 2450–2462.

[22] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.

[23] F. Palomba, D. Di Nucci, A. Panichella, R. Oliveto, and A. De Lucia, "On the diffusion of test smells in automatically generated test code: An empirical study," in *Proceedings of the 9th international workshop on search-based software testing*, 2016, pp. 5–14.

[24] P. Sahoo, A. K. Singh, S. Saha, V. Jain, S. Mondal, and A. Chadha, "A systematic survey of prompt engineering in large language models: Techniques and applications," *arXiv preprint arXiv:2402.07927*, 2024.

[25] R. Garabík, "Processing xml text with python and elementtree–a practical experience," *Bratislava, L'. Štúr Institute of Linguistics*, 2005.

[26] M. Berglund and B. van der Merwe, "Formalizing bpe tokenization," *arXiv preprint arXiv:2309.08715*, 2023.

[27] Y. Chen, Z. Hu, C. Zhi, J. Han, S. Deng, and J. Yin, "Chatunitest: A framework for llm-based test generation," in *Companion Proceedings of the 32nd ACM International Conference on the Foundations of Software Engineering*, 2024, pp. 572–576.

[28] M. Aniche, *Effective Software Testing: A developer's guide*. Simon and Schuster, 2022.

[29] A. Van Deursen, L. Moonen, A. Van Den Bergh, and G. Kok, "Refactoring test code," in *Proceedings of the 2nd international conference on extreme programming and flexible processes in software engineering (XP2001)*. Citeseer, 2001, pp. 92–95.

[30] G. Meszaros, S. M. Smith, and J. Andrea, "The test automation manifesto," in *Conference on extreme programming and agile methods*. Springer, 2003, pp. 73–81.

[31] A. Peruma, K. Almalki, C. D. Newman, M. W. Mkaouer, A. Ouni, and F. Palomba, "On the distribution of test smells in open source android applications: an exploratory study," in *Proceedings of the 29th Annual International Conference on Computer Science and Software Engineering*, 2019, pp. 193–202.

[32] F. Mariya and D. Barkhas, "A comparative analysis of mutation testing tools for java," in *2016 IEEE East-West Design & Test Symposium (EWDTS)*. IEEE, 2016, pp. 1–3.

[33] A. Peruma, K. Almalki, C. D. Newman, M. W. Mkaouer, A. Ouni, and F. Palomba, "Tsdetect: An open source test smells detection tool," in *Proceedings of the 28th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*, 2020, pp. 1650–1654.

[34] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.

[35] Z. Luo, C. Xu, P. Zhao, Q. Sun, X. Geng, W. Hu, C. Tao, J. Ma, Q. Lin, and D. Jiang, "Wizardcoder: Empowering code large language models with evol-instruct," *arXiv preprint arXiv:2306.08568*, 2023.

[36] S. Kang, J. Yoon, and S. Yoo, "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, 2023, pp. 2312–2323.