



**HAL**  
open science

## A manual categorization of new quality issues on automatically-generated tests

Geraldine Galindo-Gutierrez, Narea Maxilimiliano, Blanco Alison Fernandez,  
Nicolas Anquetil, Alcocer Juan Pablo Sandoval

### ► To cite this version:

Geraldine Galindo-Gutierrez, Narea Maxilimiliano, Blanco Alison Fernandez, Nicolas Anquetil, Alcocer Juan Pablo Sandoval. A manual categorization of new quality issues on automatically-generated tests. Proceedings of the 39st IEEE International Conference on Software Maintenance and Evolution (ICSME'23), IEEE computer society, Oct 2023, Bogota, Colombia. 10.1109/IC-SME58846.2023.00035 . hal-04344531

**HAL Id: hal-04344531**

**<https://hal.science/hal-04344531>**

Submitted on 14 Dec 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# A manual categorization of new quality issues on automatically-generated tests

Geraldine Galindo-Gutierrez  
Bolivian Catholic University, Bolivia

Maxilimiliano Narea  
Pontificia Universidad Católica de Chile, Chile

Alison Fernandez Blanco  
Pontificia Universidad Católica de Chile, Chile

Nicolas Anquetil  
Université de Lille, France  
0000-0003-1486-8399

Juan Pablo Sandoval Alcocer  
Pontificia Universidad Católica de Chile, Chile  
0000-0002-8335-4351

**Abstract**—Diverse studies have analyzed the quality of automatically generated test cases by using test smells as the main quality attribute. But recent work reported that generated tests may suffer a number of quality issues not necessarily considered in previous studies. Little is known about these issues and their frequency within generated tests. In this paper, we report on a manual analysis of an external dataset consisting of 2,340 automatically generated tests. This analysis aimed at detecting new quality issues, not covered by past recognized test smells. We use thematic analysis to group and categorize the new quality issues found. As a result, we propose a taxonomy of 13 new quality issues grouped in four categories. We also report on the frequency of these new quality issues within the dataset and present eight recommendations that test generators may consider to improve the quality and usefulness of the automatically generated tests.

**Index Terms**—

## I. INTRODUCTION

Unit tests are a critical asset in software development, as they help identify and resolve issues early in the development cycle. By verifying the behavior of individual program units in isolation, unit tests help ensure that each unit of code performs as intended, saving developers significant time and resources. Unit tests can be created either manually or automatically. Manually created tests can be time-consuming and may only cover a limited set of test scenarios that developers or testers consider. For this reason, diverse test strategies have been proposed to automatically generate tests and reduce manual effort. The latter can complement the manually created test battery, increasing coverage and fault detection.

To ensure the effectiveness of unit tests, they must meet several quality attributes, such as coverage, readability, maintainability, isolation, and speed. Previous studies have used test smells to detect generated tests that contain code pieces that may compromise their quality [24], [39], [45]. For example, it has been found that automatically generated tests often exhibit typical test smells, such as Eager Testing or Assertion Roulette. However, Panichella et al. [39] recently identified additional issues not fully captured by test smells. For instance, tests that lack assertions and tests that involve a substantial

Juan Pablo Sandoval Alcocer thanks ANID FONDECYT Iniciación Folio 11220885 for supporting this article.

```
1 public void test28() throws Throwable {  
2     XmlCalendar xmlCalendar0 = new XmlCalendar();  
3     CronText cronText0 = new CronText(xmlCalendar0);  
4     cronText0.setStartDate(xmlCalendar0);  
5     assertEquals(1, cronText0.getFrequency());  
6     assertEquals(0, cronText0.getInterval());  
7 }
```

Fig. 1. Example of test with inconsistent assertions.

amount of setup but still result in exceptions. Consider an example we found in our manual inspection, as shown in Figure 1. In this case, the test assertions (lines 5–6) are only related to the object initialization (lines 2–3), and the method call `setStartDate` (line 4) is unrelated to the rest of the test. These issues are partly associated with the unique nature of automatically generated tests and reflect a disconnect between the optimization metric of "coverage" and the real-world validity of test cases. While previous studies reported three of these issues, it remains unclear whether there are additional similar issues and how widespread they are in automatically generated tests.

In this paper, we manually review 2,340 automatically generated tests provided in the curated dataset of Panichella and colleagues [40] to detect quality issues that do not fully match previously studied test smells. Similar to previous work, we use the term "issue" to refer to portions of code that reflect symptoms or concerns related to a broader quality problem [39]. We refer to our proposition as *quality issues*, whereas we use the phrase *test smells* for smells that have been published previously.

We investigated the quality issues through an empirical study. First (Section III-B), two of the authors manually reviewed the curated dataset to locate (already known) test smells and ensure that the (new) quality issues were not already captured by any test smells. Then (Section III-C), they detected and categorized "issues" by conducting a thematic analysis and holding four meetings to refine the codes and themes generated. The remaining authors held three meetings to discuss disagreements to check the consistency of the thematic analysis. Finally, we manually analyzed the "issues" and counted the frequency of issues for projects, classes, and

methods within the curated dataset. As a result, we developed a taxonomy with 13 categories and eight recommendations that test generator maintainers and researchers can consider to improve or use test generator tools. The paper’s main contribution is the design and execution of an empirical study that analyzes quality issues not fully covered by current test smells.

## II. RELATED WORK

This section summarizes the literature related to (i) test smells, (ii) automatic test generation, (iii) test smells in generated tests. It also points out the limitations of the prior work.

**Test Smells.** Test smells extend the definition of code smells [8], [17], [37] to tests (i.e., structures or patterns that indicate lousy test code [15]). Deursen et al. [15] presented the first set of test smells in 2001, and posterior studies involving other authors [34], [44] expanded the collection. By 2018, Garousi et al. [23] found eight formally published sources presenting new test smells [1], [10], [15], [16], [22], [32], [33], [36]. Aljeedaani et al. [3], [7], [14], [26]–[29], [45] studied test smells detection tools and found that nine presented new test smell definitions. Many empirical studies [5], [53] have investigated the prevalence of test smells and their effects and found them to be detrimental to program comprehension and maintenance. As a result, some of these test scents are used for analyzing automatically generated tests (see Section II-B) and for improving the quality of generated tests.

### A. Automatically-Generated Tests

The goal of automatic test generation is to save effort and time by generating test suites with minimal human intervention [2], [21], [30], [46]. Among various generation tools [11], [35], [52], our focus in this study is on the EvoSuite Java tool [18], [40]. EvoSuite is designed to achieve the maximum possible coverage on its generated test suites [19] using generation algorithms that consider the test generation as a multi-objective problem [41], [42]. With its latest algorithm, DynaMOSA, EvoSuite produces shorter tests with higher coverage [9], [39], [42]. In order to improve EvoSuite, a number of studies have focused on: (i) readability [12], [13], [25], [31], [43], [47] and (ii) scenario generation [48], [50].

### B. Test Smells in Automatically-Generated Tests

Previous studies have explored test smells in generated tests using detection tools (e.g., tsDetect [45], JNose [54]) and manual revision [40]. Palomba et al. [38] used the Test Smell Detector [4] to investigate the frequency of the Deursen et al. test smells [15] on the SF110 dataset [20]. The results showed that 83% of the classes were smelly. Grano et al. conducted a similar study with a smaller but significant sample of SF110 [24]. The results showed a high proportion of tests generated with Assertion Roulette and Eager Test smells. Panichella et al. [39] presented a new review of the topic analyzing the Deursen smells [15], using the same sample as Grano et al. [24], but generating tests with DynaMOSA algorithm [41] and extended parameters in EvoSuite. As a

result, Panichella et al. found that such test smells were less extensive than thought when tests were generated with better configurations.

**Limitations in the Prior Work.** Most previous studies examine the quality of generated tests by focusing on test smells using detection tools and manual inspection. However, these studies voluntarily ignore other problems present in generated tests. Recently, Panichella and colleagues [39] reported “issues” that are not necessarily fully captured by current test smells. For example, they found tests with too many assertions, and tests that required significant setup that ended up throwing exceptions indicating failed setup. There is no thorough investigation of these problems or atypical code patterns within the generated tests. Therefore, little is known about their frequency, impact, or the existence of other problems.

Compared to most of the literature, our study is focused on locating, exploring, and quantifying “issues” that (i) affect the quality of automatically generated tests [6], [39] and (ii) are not necessarily captured by the current test smells [39].

## III. METHODOLOGY

To identify and quantify quality issues [6] that do not fully match with existing test smells [39] in generated tests, we focus on the following research questions:

- *RQ<sub>1</sub> – What issues or atypical patterns are present in automatically-generated tests?*
- *RQ<sub>2</sub> – How widespread are these issues in automatically-generated tests?*

**RQ1** centers on discovering and categorizing quality issues in automatically-generated tests. These issues (i) affect the quality of generated tests [6] and (ii) are not necessarily represented by current test smells [39]. To answer **RQ1**, we opted for an approach of two steps: (i) collecting existing test smells in automatically-generated tests (described in Section III-B), and (ii) manually locating issues and generating a classification scheme based on the detected issues (described in Section III-C). We also provide recommendations to improve the quality and usefulness of automatically generated tests based on the detected issues. Finally, **RQ2** focuses on quantifying the identified issues. To answer **RQ2**, we manually analyze the issues and collect the frequency of issues for projects, classes, and methods.

The following subsections detail the steps in our methodology.

### A. Dataset Selection

To detect and quantify the quality issues in automatically-generated tests, we chose the curated dataset of Panichella and colleagues [39]. We opted for this dataset because (i) it contains non-trivial classes sample suitable for studies centered on understanding and characterizing automatically-generated tests [24], [49], and (ii) it involves tests generated with better configurations using the DynaMOSA algorithm [41], and enhanced hyperparameters.

## B. Test Smells Review

To identify quality issues that may not be captured by the existing test smells definitions, we conducted a thorough review of the definitions of test smells considered in previous studies focused primarily on automatically generated tests [15], [25], [39], [44]. We used this information as a starting point for our manual inspection process, which allowed us to analyze the code and detect any issues not covered by the current test smells.

## C. Quality Issues Identification and Classification

We performed a thematic analysis [51] to detect and classify quality issues. To locate quality issues, the reviewer should detect an issue that affects a test case's quality and examine that this issue is not captured by a test smell. Consequently, the current test smells could not be reported as new quality issues.

In order to create a classification scheme and locate quality issues in automatically-generated tests, two authors conducted a thematic analysis following a number of specific steps detailed in the next paragraphs.

**Familiarization.** The two authors independently read and reread each test case to have an overview of the structure and the information of the test cases in the dataset.

**Generating codes.** The two authors independently determine if any quality issue in each test case was not captured by existing test smells in Section III-B. If an author considers a quality issue in the source code, reviews the test smells' definition, and notes that the located quality issue does not match any current test smell. Then, the author makes an annotation where he/she assigns codes that reflect the relevant features of the detected quality issue and mentions the variable or method involved with the issue. For example, one author assigned the code *duplicated\_setup* to indicate that in a test case, two or more code lines are equally exact from the setup in common. Additionally, the two authors conducted two continuous reviews to refine codes and determine if they were assigned correctly. The latter requires comparing two test cases assigned to the same code to inspect if they reflect the same issue.

**Constructing initial themes.** To generate coherent groups, the two authors independently compiled the assigned codes along their associated test cases. These groups allow identifying initial themes (broader patterns) that help address the research questions. The codes that did not belong to a specific theme were grouped as miscellaneous and analyzed in the next step.

**Reviewing themes.** The two authors held two meetings to check the initial themes created previously against the associated test case (e.g., lines of code with the issue and annotations). Then, they refined the initial themes to create a final set of themes.

**Defining and naming themes.** Each theme of the final set was defined with a description and an informative name based on the identified patterns. For example, the theme *Redundant code*

contains quality issues involved with duplicated, redundant, and/or unnecessary test code that can be extracted to avoid redundancy or erased for simplicity. Note that this theme is related to the test smell duplicated code. However, each quality issue presents features that do not precisely match the exact definition of the test smell duplicated code. Consequently, our quality issues illustrate this difference in the interpretation rules and corresponding examples.

Finally, to ensure the consistency of the process, the remaining authors checked the consistency of codes and themes against the associated data. They examined if the themes were created to respond to the research questions. Three meetings were held to discuss the disagreements or potential issues of the generated codes and themes. As a consequence, we minimized potential inconsistencies in the coding process.

Note that the identified quality issues can be represented as patterns in test cases. Therefore, they can be used to improve the quality of test generators.

## IV. RESULTS

We manually review 2,340 automatically generated tests to detect and quantify quality issues. A total of two meetings were held to identify test smells, and seven meetings were carried out to classify and locate quality issues. The results of our empirical study are detailed in the following subsections.

### A. Quality Issues (RQ1)

After manually inspecting and categorizing the data using thematic analysis, we identified a catalog of 13 quality issues that we grouped into four categories. Table I summarizes these categories, their rationale, as well as the rules (symptoms) we used to manually detect and collect the frequency. This section details each of the categories we found, providing examples and a number of recommendations to improve the test generator tools.

1) *Act-Assert Mismatch:* The *Arrange-Act-Assert* (AAA) pattern is a standard structure for unit tests. In this pattern, the *Arrange* step involves setting up the necessary objects and data to perform the test, the *Act* step calls the stimulus (i.e., the methods or actions being tested), and the *Assert* step verifies the output of the stimulus. During our manual analysis, we found that some tests included method calls unrelated to the assertions made.

**Not Asserted Side Effect.** This category group tests that contain calls to void methods that cause a side-effect (i.e., they change the object's state), but the assertions within the test do not verify these effects. For example, consider the generated test for the class `SubstringLabeler` shown in Figure 2. This test only calls the void method `addDataSourceListener`, which adds a `TextViewer` instance to an internal collection of listeners within the `SubstringLabeler` object. However, the assertions within the test only verify attributes such as `matchAttributeName` and `customName`, which are unrelated to the void method call. Although the void method call increases test coverage, the test itself does not help to verify the correct behavior of the `addDataSourceListener` method.

TABLE I  
OBSERVED QUALITY ISSUES IN AUTOMATICALLY GENERATED TESTS

Act-Assert Mismatch		
Not asserted side effects (NASE)	<b>Symptom.</b> At least one void method call causes a side effect, but no assertion verifies the side effect or references the method call.	<b>Rationale.</b> Although a void method call may increase coverage if no assertion is associated, the test reduces its ability to detect future change behavior rather than exceptions.
Not asserted return values (NARV)	<b>Symptom.</b> At least one method call returns a value that is not checked by any assertion or used later in the test.	<b>Rationale.</b> If a test excludes method call return values from its assertions, it misses the opportunity to verify that the behavior of the method under test is correct.
Assertions with not related parent class method (ARPM)	<b>Symptom.</b> At least one assertion calls an inherited method that checks a value unrelated to the method under test.	<b>Rationale.</b> Asserting methods that are not tested is not very useful. Also, the peculiarity of asserting parent class methods may fit as a special case of indirect testing.
Redundant Code		
Asserting object initialization multiple times (OIMT)	<b>Symptom.</b> Two or more tests in which the assertions check for values that are set in the constructor, or check for default values.	<b>Rationale.</b> A limited number of tests should be dedicated to checking the correct initialization of an object. However, this becomes redundant when repeated throughout the suite.
Duplicated Setup (DS)	<b>Symptom.</b> Two or more tests sharing at least two equal lines of setup code.	<b>Rationale.</b> Test maintenance and analysis becomes difficult due to duplicate code.
Testing the same exception scenario (TSES)	<b>Symptom.</b> Two or more tests with the same exceptions and console output calling different methods with the same setup.	<b>Rationale.</b> If there are two tests that evaluate the same exception scenario, one of them can be considered to be semantically redundant.
Testing the same void method (TSVM)	<b>Symptom.</b> Two or more tests that call the same void method and that report instances of side effects that are not asserted.	<b>Rationale.</b> It is important to have two or more tests of the same method to cover different branches. But their usefulness for finding bugs is limited if the assertions are unrelated to the behavioral differences in the method.
Redundant not null Assertion (NNA)	<b>Symptom.</b> At least one <code>assertNotNull</code> is redundant. It is redundant (i) after initializing an object, (ii) if some other <code>assert</code> checks the value content, and (iii) if the method being tested does not return null.	<b>Rationale.</b> In order to reduce the size of the test and the amount of time that developers spend analyzing it, it may be useful to remove redundant statements.
Failed Setup		
Exceptions due to null arguments (EDNA)	<b>Symptom.</b> Tests handling <code>NullPointerException</code> caused by null arguments passed directly to constructors or methods.	<b>Rationale.</b> In some scenarios, testing for exceptions caused by null arguments can be valuable. However, if the class is designed to assume that no null arguments are sent, testing will not cover new branches, and the result is predictable.
Exceptions due to external dependencies (EDED)	<b>Symptom.</b> Tests that handle <code>HeadlessException</code> , <code>SQLException</code> or <code>NotYetConnectedException</code> .	<b>Rationale.</b> Typically, these exceptions occur when a test uses external resources. Thus, these tests are no longer self-contained.
Exceptions due to incomplete setup (EDIS)	<b>Symptom.</b> Tests that construct an object, call a method and handle an exception thrown by an uninitialized value.	<b>Rationale.</b> In these cases, much of the generated suite contains exceptions caused by incomplete test setup steps. Rather than interesting test scenarios, these tests reflect the difficulty of the generator tool.
Testing only Field Accessors and Constants		
Testing only field accessors (TOFA)	<b>Symptom.</b> Tests containing only object initialization and field accessor (getter/setter) assertions.	<b>Rationale.</b> Field accessors are straightforward and predictable operations; we expect an obvious result upon their assertion.
Asserting constants (AC)	<b>Symptom.</b> At least one assertion verifies the value of a variable declared as a static final on its class.	<b>Rationale.</b> Asserting constants is uncommon because the <code>assert</code> won't fail unless the value is manually changed.

We found 328 (14%) tests that contain this code pattern. One question that arises for these tests is: *Is it possible to verify these side effects?* After manually reviewing the classes under test, we realized that 31% of the side-effects can be verified by calling another method within the class, but 69% of the classes do not provide any public method (i.e., accessors) to help verify these side-effects (e.g., printing values in a terminal). We found only one test that contains both cases.

**Not Asserted Return Values.** This category group tests that contain a method call that returns a value not verified by the test's assertions. For instance, consider the generated

test for the class `ExternalXid` (Figure 3), where none of its assertions verify the returned value of the `equals` method call. We inspected the `equals` method and confirmed that: (i) it returns a `boolean`, (ii) it does not cause any side effects, and (iii) it has no relation with the `formatId` (the only assertion in the test). Consequently, we concluded that although this increases the coverage, it does not help verify this method's behavior, as it is the only method being tested. We found 242 (10.3%) tests in this category.

**Assertions with not related parent class method.** We found that 121 (5.1%) tests contain assertions that use methods

```

1 public void test79() throws Throwable {
2     SubstringLabeler substringLabeler0 = new
3         SubstringLabeler();
4     TextView textView0 = new TextView();
5     substringLabeler0.addDataSourceListener(textView0);
6     assertEquals("Match",
7         substringLabeler0.getMatchAttributeName());
8     assertEquals("SubstringLabeler",
9         substringLabeler0.getCustomName());
10 }

```

Fig. 2. Not asserted side effect.

```

1 public void test22() throws Throwable {
2     byte[] byteArray0 = new byte[6];
3     ExternalXid externalXid0 = new ExternalXid((-1), ...);
4     Object object0 = new Object();
5     externalXid0.equals(object0);
6     assertEquals((-1), externalXid0.getFormatId());
7 }

```

Fig. 3. Example of not asserted return values.

from the parent class of the class under test. However, these methods are not related to the methods being tested. For instance, consider the following generated test that asserts the method `isFocusTraversalPolicySet` defined in the `Container` class, although the class under test is `Field` (Figure 4). The `Container` class is extended by the `JComponent` class, which is then extended by the `JPanel` class, and finally by the `Field` class. It is also important to note that we verified that the method `isFocusTraversalPolicySet` is not overridden by the `Field` class. In this example, the `paintComponent` method defined in `Field` does not modify the parent class attribute `focusTraversalPolicy`. Therefore, the assertion that calls the method `isFocusTraversalPolicySet` is not related to the `paintComponent` method, which is the only operation called after object initialization.

This category is similar to the previous one, except that the methods used in the asserts belong to a parent class of the class under test. We have separated this category as we believe it may be useful for tool maintainers.

**Recommendation 1.** Test generator tools should prioritize the relationship between the act and assert steps in the generation algorithm to enhance the effectiveness of generated tests.

```

1 public void test12() throws Throwable {
2     ...
3     Field field0 = new Field(handballModel0, colorModel0);
4     field0.paintComponent(graphics2D0);
5     assertFalse(field0.isFocusTraversalPolicySet());
6 }
7 /* Field class parent*/
8 class Container{
9     public boolean isFocusTraversalPolicySet() {
10         return (focusTraversalPolicy != null);
11     }
12 }

```

Fig. 4. Example of assertions with not related parent class method.

```

1 public void test07() throws Throwable {
2     ScriptOrFnScope scriptOrFnScope0 =
3         new ScriptOrFnScope((-806), (ScriptOrFnScope) null
4             );
5     ScriptOrFnScope scriptOrFnScope1 =
6         new ScriptOrFnScope((-330), scriptOrFnScope0);
7     scriptOrFnScope1.preventMunging();
8     ...
9 }
10 public void test08() throws Throwable {
11     ScriptOrFnScope scriptOrFnScope0 =
12         new ScriptOrFnScope((-806), (ScriptOrFnScope) null
13             );
14     ScriptOrFnScope scriptOrFnScope1 =
15         new ScriptOrFnScope((-330), scriptOrFnScope0);
16     scriptOrFnScope1.declareIdentifier("");
17     ...
18 }
19 public void test09() throws Throwable {
20     ScriptOrFnScope scriptOrFnScope0 =
21         new ScriptOrFnScope((-806), (ScriptOrFnScope) null
22             );
23     ScriptOrFnScope scriptOrFnScope1 =
24         new ScriptOrFnScope((-330), scriptOrFnScope0);
25     scriptOrFnScope1.munge();
26     ...
27 }

```

Fig. 5. Example of test with duplicated setup.

2) *Redundant Code:* We discovered pieces of duplicate, redundant, or unnecessary test code that could be eliminated for simplicity or extracted to prevent redundancy.

**Duplicated Setup.** We observed that several tests within the same test suite initialize the same group of objects at the beginning. Although not all tests in the suite have the same initialization, some tests contain exactly the same statements at the start. For instance, `test07`, `test08`, and `test09` (shown in Figure 5) are generated tests that demonstrate this duplication, as each of them initializes the variables `scriptOrFnScope0` and `scriptOrFnScope1` using the same constructor and arguments. Therefore, the duplicated code could be extracted into a `setUp` method to avoid redundancy.

We found 498 (21.2%) test methods with at least two identical setup statements. These test methods belong to 35 test classes. It is important to note that *EvoSuite* generates a test class for each class under test, which means that there are at least 14 test methods per test class on average.

**Recommendation 2.** Test generator tools should identify and refactor tests with duplicated setups into a separate class with a common setup method. The latter will reduce code duplication and improve code quality.

**Testing the same exception scenario.** We found 166 (7.1%) tests that test the same exception scenario as another test in the battery, with minor differences. These tests contain mostly the same statements, differing only in the method call that triggers the exception. However, the root cause of the exception is either related to the shared statements or another common factor. For example, two generated tests (`test04` and `test08`) differ only in the method called (`initialize` and `getClusterNodeAddresses`), but both test the same exception

```

1 public void test04() throws Throwable {
2     ClusterExecutorUtil clusterExecutorUtil0 =
3     new ClusterExecutorUtil();
4     ClusterExecutorImpl clusterExecutorImpl0 =
5     new ClusterExecutorImpl();
6     clusterExecutorUtil0.setClusterExecutor(
7     clusterExecutorImpl0);
8     // Undeclared exception!
9     try {
10        ClusterExecutorUtil.initialize();
11        fail("Expecting exception: NoClassDefFoundError");
12    } catch (NoClassDefFoundError e) {
13        // Could not initialize class com.liferay.portal.
14        util.PropsValues
15        verifyException("com.liferay.portal.cluster.
16        ClusterBase", e);
17    }
18 }
19 public void test08() throws Throwable {
20     ClusterExecutorUtil clusterExecutorUtil0 =
21     new ClusterExecutorUtil();
22     ClusterExecutorImpl clusterExecutorImpl0 =
23     new ClusterExecutorImpl();
24     clusterExecutorUtil0.setClusterExecutor(
25     clusterExecutorImpl0);
26     // Undeclared exception!
27     try {
28        ClusterExecutorUtil.getClusterNodeAddresses();
29        fail("Expecting exception: NoClassDefFoundError");
30    } catch (NoClassDefFoundError e) {
31        // Could not initialize class com.liferay.portal.
32        util.PropsValues
33        verifyException(
34        "com.liferay.portal.cluster.ClusterBase", e);
35    }
36 }

```

Fig. 6. Testing the same exception scenario.

scenario where PropsValues cannot be found (shown in Figure 6). As such, we consider one of these tests redundant. This category of tests handles various exception kinds, such as NullPointerException, RuntimeException, and ClassCastException.

**Recommendation 3.** To improve the quality of the test battery and reduce its size, test generator tools should detect and eliminate redundant tests using inter-test analysis, especially for exception scenarios.

**Testing the same void method.** We found 117 (5%) tests that evaluate the same void method with small variations. While these variations may help cover new branches of the class under test, the assertions of these tests are unrelated to the branch under test, limiting their usefulness for bug fault detection. For example, consider two generated tests (Figure 7). In both tests, the release method is called from DbConnectionBroker, which affects the maxConnections attribute. However, the tests assert a different attribute (maxReached) and do not verify the method's side effects. This category is similar to the previous category, "not asserted side-effect," but with two or more tests executing different branches of the same void method, where the assertions are unrelated to this method, further limiting their usefulness.

```

1 public void test00() throws Throwable {
2     DbConnectionBroker dbConnectionBroker0 = new
3     DbConnectionBroker();
4     DbConnectionAttributes dbConnectionAttributes0 = new
5     DbConnectionAttributes(10);
6     dbConnectionBroker0.release(dbConnectionAttributes0);
7     assertEquals(0, dbConnectionBroker0.getMaxReached());
8 }
9 public void test11() throws Throwable {
10    DbConnectionBroker dbConnectionBroker0 = new
11    DbConnectionBroker();
12    assertEquals(30, dbConnectionBroker0.getMax());
13    DbConnectionAttributes dbConnectionAttributes0 = new
14    DbConnectionAttributes(15);
15    dbConnectionBroker0.release(dbConnectionAttributes0);
16    assertEquals(0, dbConnectionBroker0.getMaxReached());
17 }

```

Fig. 7. Multiple calls to the same void method.

**Recommendation 4.** Test generator tools should generate assertions that verify the behavior associated with different branches of the same method already covered by the test battery.

**Asserting object initialization multiple times.** This category includes tests with assertions that verify behavior related to the object initialization during the test setup phase. While testing for proper initialization is valid, some tests in the same battery focus solely on initialization instead of the tested method. These tests often have similar assertions. For example, test14 and test27 in Figure 8 have identical assertions for the PhotoController class that verify values set during initialization. Note that although these two tests differ slightly, with test14 calling the method setLens and test27 calling the method getContentManager, there are no assertions related to these methods. We found 550 (23,5%) tests that assert the object initialization multiple times. These tests belong to 50 test suites.

**Recommendation 5.** Test generator tools should avoid redundant tests that verify the same object initialization in multiple tests and instead focus on testing the particular behavior of each test. This approach can reduce the size of the test suite and ensure that each test contributes uniquely to code coverage and error detection.

**Redundant not Null Assertion.** We found 77 (3.3%) tests with redundant not null assertions. In particular, we identify two cases: (i) *When the not null assertion is used to verify a recently created object* – For example, consider the test generated for the class HookHotDeployListener in Figure 9. The test initializes an instance of the class using the new operator and immediately follows it with an assertNotNull statement over the newly created object. However, since an error would be triggered if the object were not created, verifying its existence with an assertNotNull statement is unnecessary.

(ii) *When another assertion indirectly verifies that a given variable is not null, and deleting will not alter the test*

```

1 public void test14() throws Throwable {
2     Home home0 = new Home();
3     SwingViewFactory swingViewFactory0 = new
4     SwingViewFactory();
5     PhotoController photoController0 = new
6     PhotoController(home0, (UserPreferences) null,
7     (View) null, swingViewFactory0, (ContentManager
8     ) null);
9     Camera.Lens camera_Lens0 = Camera.Lens.PINHOLE;
10    photoController0.setLens(camera_Lens0);
11    assertEquals(0, photoController0.getQuality());
12    assertEquals(300, photoController0.getHeight());
13    assertEquals(400, photoController0.getWidth());
14    assertEquals(1.0F, photoController0.
15    get3DViewAspectRatio(), 0.01F);
16    assertEquals(1392409281320L, photoController0.
17    getTime());
18    assertEquals(13684944, photoController0.
19    getCeilingLightColor());
20 }
21 public void test27() throws Throwable {
22     Home home0 = new Home();
23     SwingViewFactory swingViewFactory0 = new
24     SwingViewFactory();
25     PhotoController photoController0 = new
26     PhotoController(home0, (UserPreferences) null(
27     View) null, swingViewFactory0, (ContentManager
28     ) null);
29    photoController0.getContentManager();
30    assertEquals(13684944, photoController0.
31    getCeilingLightColor());
32    assertEquals(300, photoController0.getHeight());
33    assertEquals(0, photoController0.getQuality());
34    assertEquals(400, photoController0.getWidth());
35    assertEquals(1392409281320L, photoController0.
36    getTime());
37    assertEquals(1.0F, photoController0.
38    get3DViewAspectRatio(), 0.01F);
39 }

```

Fig. 8. Asserting object initialization multiple times.

```

1 public void test00() throws Throwable {
2     HookHotDeployListener hookHotDeployListener0 = new
3     HookHotDeployListener();
4     assertNotNull(hookHotDeployListener0);
5 }

```

Fig. 9. Redundant not null assertion that verify a recently created object

*semantics*. – This happens when a test method contains another assertion that indirectly verifies that a given variable is not null. Therefore the `assertNotNull` could be deleted without changing the test behavior<sup>1</sup>. To illustrate, consider the following generated test for the class `SubstringLabeler`, where there is an `assertNotNull` for the variable `string0` and an `assertEquals` that checks that the variable `string0` is equal to the string "trainingSet" (Figure 10). In this test, the `assertEquals` indirectly verifies that the variable `string0` is not null since it does not contain the expected value.

We found 47 tests that contained an `assertNotNull` after initializing an object and 33 tests that presented another assertion that indirectly verify the content of a given variable besides the `assertNotNull`.

<sup>1</sup>It is important to mention that there are cases where even though another assertion indirectly verifies a variable is not null, the not null assertion is not redundant.

```

1 public void test74() throws Throwable {
2     ...
3     String string0 = substringLabeler_Match0.getMatch();
4     assertNotNull(string0);
5     assertEquals("trainingSet", string0);
6     ...
7 }

```

Fig. 10. Redundant not null assertion since the `assertEquals` indirectly verifies that the variable is not null

```

1 public void test02() throws Throwable {
2     PeersItem peersItem0 = new PeersItem();
3     // Undeclared exception!
4     try {
5         peersItem0.refresh((TableCell) null);
6         fail("Expecting exception: NullPointerException");
7     } catch (NullPointerException e) {
8         // no message in exception (getMessage() returned
9         // null)
10        verifyException("org.gudy.azureus2.ui.swt.views.
11        tableitems.tracker.PeersItem", e);
12    }
13 }

```

Fig. 11. Exceptions due to null arguments.

**Recommendation 6.** To improve the effectiveness of test generation, it is recommended that test generators analyze the need for not-null assertions and avoid generating redundant assertions.

3) *Failed Setup*: Previous studies show that many tests require significant setup but often fail, causing exceptions. If the test generator tool struggles with initializing objects, it can result in a large number of tests with failed setups. We identified three types of failed setups during our manual inspection, including eight test suites with all test methods handling exceptions due to failed setups.

*Exceptions due to null arguments.* We observed that 300 (12.8%) tests trigger a `NullPointerException` because null arguments are sent to a method call or class constructor. For example, consider the `test02` in Figure 11, where a null value is sent to the `refresh` method. The exception is directly linked to the null value and is triggered when the `refresh` method attempts to execute the `getDataSource` method on the null value. Note that the `refresh` method does not handle null arguments. While sending a null argument may help cover new branches or test exception scenarios, these tests have limited usefulness as the exception triggered is directly linked to the null value rather than any new branch of the tested method.

*Exceptions due external dependencies.* This category involves tests that handle exceptions triggered by external dependencies or resources rather than due to issues in the behavior of the method being called. For example, consider the following generated test for the `EntryListView` class shown in Figure 12, which includes a try-catch statement that catches a `HeadlessException`. This exception type is linked to code that relies on keyboard, display, or mouse events in an environment that does not support them. We found similar exceptions

```

1 public void test00() throws Throwable {
2     try {
3         EntryListView.showView();
4         fail("Expecting exception: HeadlessException");
5     } catch (HeadlessException e) {
6         // no message in exception (getMessage() returned null
7         verifyException("java.awt.GraphicsEnvironment",
8             e);
9     }
10 }

```

Fig. 12. Exceptions due to external dependencies.

```

1 public void test00() throws Throwable {
2     AlphabeticTokenizer alphabeticTokenizer0 = new
3     AlphabeticTokenizer();
4     // Undeclared exception!
5     try {
6         alphabeticTokenizer0.nextElement();
7         fail("Expecting exception: NullPointerException");
8     } catch (NullPointerException e) {
9         // no message in exception (getMessage() returned
10        null)
11        verifyException("weka.core.tokenizers.
12        AlphabeticTokenizer", e);
13    }
14 }
15 /* Class under test*/
16 class AlphabeticTokenizer{
17     protected int m_CurrentPos;
18     public Object nextElement() {
19         int beginpos, endpos;
20         beginpos = m_CurrentPos;
21         while ( (beginpos < m_Str.length) && ((m_Str[
22             beginpos] < 'a') && (m_Str[beginpos] > 'z'))
23             && ((m_Str[beginpos] < 'A') && (m_Str[
24             beginpos] > 'Z')) ) {
25             beginpos++;
26         }
27         ...
28         return s;
29     }
30 }

```

Fig. 13. Exceptions due to incomplete call sequence.

such as `SQLException` and `NotYetConnectedException`. We found 52 (2.2%) test methods that handle exceptions due to external dependencies. Multiple tests in the same test class have similar exception handling code and the same root cause, which is a failed setup.

**Exceptions due to incomplete setup.** This category contains tests that handle exceptions triggered by an incomplete setup. We analyzed the identified cases and discovered that uninitialized values or uncalled methods triggered the exceptions. To illustrate, consider the following generated test for the `AlphabeticTokenizer` class shown in Figure 13. To successfully call the `nextElement` method, the `m_CurrentPos` attribute needs to be initialized. Otherwise, the method will attempt to access a null index inside an array. In `test00`, the default constructor of `AlphabeticTokenizer0` does not initialize `m_CurrentPos`. Therefore, when calling `nextElement`, a `NullPointerException` is thrown. We found 41 (1.7%) test methods that handle exceptions due to external dependencies.

```

1 public void test05() throws Throwable {
2     DirEntry dirEntry0 = new DirEntry();
3     dirEntry0.setModtime("");
4     String string0 = dirEntry0.getModtime();
5     assertEquals("", string0);
6 }

```

Fig. 14. Testing only field accessors.

```

1 public void test06() throws Throwable {
2     ...
3     connectionConsumer0.connectionError("summa.persistent")
4     ;
5     assertEquals(0, ConnectionConsumer.
6     DEFAULT_SUBSEQUENT_RETRIES);
7 }

```

Fig. 15. Asserting constants.

**Recommendation 7.** Test generator tools should be able to detect situations where the object initialization is unsuccessful because this can lead to multiple tests with similar exception scenarios.

4) *Testing Field Accessors, and Constants:* Tests, where the behavior under test is rather simple, belong to this category. Apart from improving code coverage and preventing exceptions, the behavior under test is predictable and straightforward.

**Testing only field accessors.** We identified 139 (5.9%) tests in which, apart from object initialization, the behavior under test solely consisted of field accessors such as getters and setters. An example of this is the generated test for the `DirEntry` class shown in Figure 14, which solely reads and writes the `modtime` field using the `getModtime` and `setModtime` methods. Although calling these methods increases coverage for the class under test, we believe that testing only field accessors that do not add any additional behavior beyond exposing a class field are strong candidates for test reduction or prioritization.

**Asserting Constants.** We observed that 20 (0.8%) tests only check the value of a constant using assertions. While testing constant values can be valuable in certain cases, these tests also tend to call different methods and only assert unrelated constant values. For example, consider the generated test for the `ConnectionConsumer` class shown in Figure 15. This test initializes an object of the class and invokes the `connectionError` method, but the only assertion in the test checks the `DEFAULT_SUBSEQUENT_RETRIES` constant, which is always initialized to zero unless manually modified. Moreover, this constant is irrelevant to the `connectionError` method. Tests that only check constant values in this manner do not significantly contribute to the test suite, as the outcome of such assertions is always obvious.

**Recommendation 8.** Tests that solely use accessors or validate constant values might not substantially impact the test suite. Therefore, test generator tools should identify such tests as potential candidates for prioritization or removal to reduce the total number of tests to execute or analyze.

### B. Quality Issue Distribution (RQ2)

The frequency of the categories above at the test method, class, and project levels can be found in Table IV-B. Our results indicate that the reported issues were present across various classes and projects within the data set. The latter suggests that the issues are not specific to any particular project or class. Although the percentage of test methods with at least one issue is small (less than 20%), the number of classes and test suites containing at least one issue is considerable. For instance, the issue “not asserted side effects” appear at least once in 57 out of 100 generated test suites.

TABLE II  
FINE-GRAINED QUALITY ISSUES: CATEGORIES AND FREQUENCY

	% Test Methods	# Test Methods	# Test Classes	# Projects
<b>Act-Assert Mismatch</b>				
Not asserted side effects	14.0%	328	49	26
Not asserted return values	10.3%	242	30	19
Assertions w/ not related parent class method	5.1%	121	13	9
<b>Redundant Code</b>				
Asserting object initialization multiple times	23.5%	550	50	24
Duplicated setup	21.2%	498	35	18
Testing the same exception scenario	7.1%	166	24	16
Testing the same void method	5.0%	117	20	15
Redundant not null assertion	3.3%	77	12	12
<b>Failed Setup</b>				
Exceptions due to null arguments	12.8%	300	64	30
Exceptions due to incomplete setup	2.2%	52	19	12
Exceptions due to external dependencies	1.7%	41	20	12
<b>Testing Field Accessors and Constants</b>				
Testing only field accessors	5.9%	139	17	13
Asserting constants	0.8%	20	11	6

**Relation between Quality Issues.** The heat map shown in Figure 16 shows the relationships between the identified issues, with the color of each cell indicating the number of generated tests that have both of these issues. The two most closely related quality issues are “Not asserted side effects” (NASE) and “Assertion with no related parent class” (ARPM). Therefore, a significant proportion of tests that do not assert side effects also contain an assertion with a method call to a parent class unrelated to the under-test behavior. Two other related issues are: “Exceptions due to null arguments” (TWNA) and “Testing the same exception scenario” (TSES). Although the tests with these issues have slightly different codes, in many cases, the root cause of the exception is null arguments. It is worth noting that the exception scenarios are not exclusively related to null values, which is why the color box is not dark blue. Overall, Figure 16 shows that no two quality issues overlap entirely in the same group of

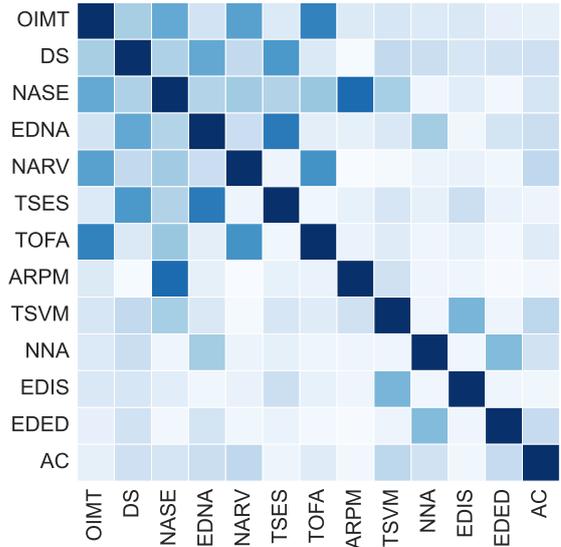


Fig. 16. Relation between identified quality issues

tests, suggesting that each issue has unique characteristics not covered by other issues.

## V. THREATS TO VALIDITY

This study presents a manual categorization of a particular sample of automatically generated tests, and like any other study, it is subject to a number of threats to validity.

**Constructor Validity.** In order to search for code parts or code anti-patterns that do not match the current test smell definitions, we use a set of test smells as a basis for the analysis. The fact that not all relevant test smells proposed by previous literature are covered may be a threat to the validity of this study. Consequently, a quality issue may partly overlap the definition of a test smell not included in our set of test smells. To reduce this threat, we utilized the test smells considered by Panichella and colleagues [39], since this study collects the test smells proposed by formally-published sources [15], [25], [44] which can affect tests generated by EvoSuite.

**Internal Validity.** We performed a thematic analysis to answer our research questions. Two authors performed the systematic process to conduct the analysis. Since this process includes generating codes and defining themes, these aspects vary depending on the coder’s experience, point of view, and level of abstraction. For example, to respond to RQ1, we detected quality issues involved with duplicated, redundant, and/or unnecessary test code. These quality issues are related to the code smell of duplicated code. However, each quality issue shows particularities associated to the nature of generated tests. We tried to reduce this threat by checking the consistency of the process. The remaining authors examined the description of themes, generated codes, and the test cases involved with codes. We held three discussion meetings to diagnose the

codes and the themes generated. As a result, we solve any discrepancies.

**External Validity.** In our study, we identified and categorized quality issues in a sample of 2,340 tests generated by EvoSuite. However, besides involving a significant manual effort, the generalizability of our findings is limited. Additional issues may be present in tests generated by other tools or in different projects. Nevertheless, this paper presents an initial taxonomy of quality issues highlighting the discrepancy between generated and manually-written tests. The latter confirms that test smells alone do not capture all the quality issues present in generated tests.

## VI. CONCLUSION

This paper presents a manual categorization of quality issues for a data set of 2,340 generated tests, resulting in a taxonomy of 13 issues that reflect atypical patterns in automatically generated tests. Our study shows that the reported issues are not specific to a particular class or project. Furthermore, our issue catalog highlights the importance of considering the code being tested when analyzing the test quality. For example, methods with side effects and complex object initialization may trigger a number of issues in the generated tests. Our findings are consistent with previous work in revealing a discrepancy between the optimization metric used by test generator tools and the usefulness of the generated tests. Based on the issues we identified, we provide eight recommendations that demonstrate their applicability and usefulness. As future work, we plan to extend the external validity of the issues found by automatically searching the reported issues in other projects, and generator tools.

## REFERENCES

- [1] Jennitta Andrea, Gerard Meszaros, and Shaun Smith. Catalog of xp project ‘smells’. In *2002 International Conference on Extreme Programming and Agile Processes in Software Engineering*, 06 2002.
- [2] Alberto Bacchelli, Paolo Ciancarini, and Davide Rossi. On the effectiveness of manual and automatic unit test generation. In *2008 The Third International Conference on Software Engineering Advances*, pages 252–257, 2008.
- [3] P. Baker, D. Evans, J. Grabowski, H. Neukirchen, and B. Zeiss. Trex - the refactoring and metrics tool for ttcn-3 test specifications. In *Testing: Academic & Industrial Conference - Practice And Research Techniques (TAIC PART’06)*, pages 90–94, 2006.
- [4] Gabriele Bavota, Abdallah Qusef, Rocco Oliveto, Andrea De Lucia, and Dave Binkley. Are test smells really harmful? an empirical study. *Empirical Software Engineering*, 20:1052–1094, 2015.
- [5] Gabriele Bavota, Abdallah Qusef, Rocco Oliveto, Andrea De Lucia, and David Binkley. An empirical analysis of the distribution of unit test smells and their impact on software maintenance. In *2012 28th IEEE International Conference on Software Maintenance (ICSM)*, pages 56–65, 2012.
- [6] Carolin Brandt and Andy Zaidman. Developer-centric test amplification: The interplay between automatic generation human exploration. *Empirical Softw. Engg.*, 27(4), jul 2022.
- [7] Manuel Breugelmans and Bart Van Rompaey. Testq: Exploring structural and maintenance characteristics of unit test suites. In *WASDeTT-1: 1st International Workshop on Advanced Software Development Tools and Techniques*. Citeseer, 2008.
- [8] Aloisio S Cairo, Glauco de F Carneiro, and Miguel P Monteiro. The impact of code smells on software bugs: A systematic literature review. *Information*, 9(11):273, 2018.
- [9] José Campos, Yan Ge, Nasser Albulian, Gordon Fraser, Marcelo Eler, and Andrea Arcuri. An empirical evaluation of evolutionary algorithms for unit test suite generation. *Information and Software Technology*, 104:207–235, 2018.
- [10] Woei-Kae Chen and Jung-Chi Wang. Bad smells and refactoring methods for gui test scripts. In *2012 13th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing*, pages 289–294, 2012.
- [11] Christoph Csallner and Yannis Smaragdakis. Jcrasher: An automatic robustness tester for java. *Softw. Pract. Exper.*, 34(11):1025–1050, sep 2004.
- [12] Ermira Daka, José Campos, Gordon Fraser, Jonathan Dorn, and Westley Weimer. Modeling readability to improve unit tests. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, page 107–118, New York, NY, USA, 2015. Association for Computing Machinery.
- [13] Ermira Daka, José Miguel Rojas, and Gordon Fraser. Generating unit tests with descriptive names or: Would you name your children thing1 and thing2? In *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 57–67, 2017.
- [14] Julien Delplanque, Stéphane Ducasse, Guillermo Politto, Andrew P. Black, and Anne Etien. Rotten green tests. In *Proceedings of the 41st International Conference on Software Engineering, ICSE ’19*, page 500–511. IEEE Press, 2019.
- [15] Arie Deursen, Leon M.F. Moonen, A. Bergh, and Gerard Kok. Refactoring test code. Technical report, Delft University of Technology, NLD, 2001.
- [16] Bill Dudley and Jonathan Lehr. *Jakarta Pitfalls: Time-Saving Solutions for Struts, Ant, JUnit, and Cactus (Java Open Source Library)*. John Wiley & Sons, 2003.
- [17] Martin Fowler. *Refactoring: Improving the Design of Existing Code*. Addison-Wesley Professional, 1st edition, 1999.
- [18] Gordon Fraser and Andrea 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, ESEC/FSE ’11*, page 416–419, New York, NY, USA, 2011. Association for Computing Machinery.
- [19] Gordon Fraser and Andrea Arcuri. Whole test suite generation. *IEEE Transactions on Software Engineering*, 39(2):276–291, 2012.
- [20] Gordon Fraser and Andrea Arcuri. A large-scale evaluation of automated unit test generation using evosuite. *ACM Trans. Softw. Eng. Methodol.*, 24(2), dec 2014.
- [21] Gordon Fraser, Matt Staats, Phil McMinn, Andrea Arcuri, and Frank Padberg. Does automated unit test generation really help software testers? a controlled empirical study. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 24(4):1–49, 2015.
- [22] Steve Freeman and Nat Pryce. *Growing Object-Oriented Software, Guided by Tests*. Addison-Wesley Professional, 1st edition, 2009.
- [23] Vahid Garousi and Barış Küçük. Smells in software test code: A survey of knowledge in industry and academia. *Journal of Systems and Software*, 138:52–81, 2018.
- [24] Giovanni Grano, Fabio Palomba, Dario Di Nucci, Andrea De Lucia, and Harald C Gall. Scented since the beginning: On the diffuseness of test smells in automatically generated test code. *Journal of Systems and Software*, 156:312–327, 2019.
- [25] Giovanni Grano, Simone Scalabrino, Harald C. Gall, and Rocco Oliveto. An empirical investigation on the readability of manual and generated test cases. In *Proceedings of the 26th Conference on Program Comprehension, ICPC ’18*, page 348–351, New York, NY, USA, 2018. Association for Computing Machinery.
- [26] Michaela Greiler, Arie Van Deursen, and Margaret-Anne Storey. Automated detection of test fixture strategies and smells. In *2013 IEEE Sixth International Conference on Software Testing, Verification and Validation*, pages 322–331. IEEE, 2013.
- [27] Alex Gyori, August Shi, Farah Hariri, and Darko Marinov. Reliable testing: Detecting state-polluting tests to prevent test dependency. In *Proceedings of the 2015 international symposium on software testing and analysis*, pages 223–233, 2015.
- [28] Chen Huo and James Clause. Improving oracle quality by detecting brittle assertions and unused inputs in tests. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 621–631, 2014.
- [29] Negar Koochakzadeh and Vahid Garousi. Tcrevis: a tool for test coverage and test redundancy visualization. In *Testing-Practice and*

- Research Techniques: 5th International Academic and Industrial Conference, TAIC PART 2010, Windsor, UK, September 3-5, 2010. Proceedings*, pages 129–136. Springer, 2010.
- [30] Jeshua S. Kracht, Jacob Z. Petrovic, and Kristen R. Walcott-Justice. Empirically evaluating the quality of automatically generated and manually written test suites. In *2014 14th International Conference on Quality Software*, pages 256–265, 2014.
- [31] Bin Lin, Csaba Nagy, Gabriele Bavota, Andrian Marcus, and Michele Lanza. On the quality of identifiers in test code. In *2019 19th International Working Conference on Source Code Analysis and Manipulation (SCAM)*, pages 204–215, 2019.
- [32] Mikael Lundin. *Testing with F#*. Packt Publishing, 2015.
- [33] Gerard Meszaros. *xUnit Test Patterns*. Addison-Wesley, 2007.
- [34] Gerard Meszaros, Shaun M. Smith, and Jennitta Andrea. The test automation manifesto. In Frank Maurer and Don Wells, editors, *Extreme Programming and Agile Methods - XP/Agile Universe 2003*, pages 73–81, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [35] Carlos Pacheco, Shuvendu K. Lahiri, Michael D. Ernst, and Thomas Ball. Feedback-directed random test generation. In *Proceedings of the 29th International Conference on Software Engineering, ICSE '07*, page 75–84, USA, 2007. IEEE Computer Society.
- [36] Alan Page and Ken Johnston. *How We Test Software at Microsoft*. Microsoft Press, USA, 2008.
- [37] Fabio Palomba, Gabriele Bavota, Massimiliano Di Penta, Fausto Fasano, Rocco Oliveto, and Andrea De Lucia. On the diffuseness and the impact on maintainability of code smells: a large scale empirical investigation. In *Proceedings of the 40th International Conference on Software Engineering*, pages 482–482, 2018.
- [38] Fabio Palomba, Dario Di Nucci, Annibale Panichella, Rocco Oliveto, and Andrea 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*, pages 5–14, 2016.
- [39] A. Panichella, Sebastiano Panichella, Gordon Fraser, Anand Ashok Sawant, and Vincent Hellendoorn. Test smells 20 years later: Detectability, validity, and reliability. *Empirical Software Engineering*, 27(7), 2022.
- [40] Annibale Panichella, José Campos, and Gordon Fraser. Evosuite at the sbst 2020 tool competition. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops, ICSEW'20*, page 549–552, New York, NY, USA, 2020. Association for Computing Machinery.
- [41] Annibale Panichella, Fitsum Meshesha Kifetew, and Paolo Tonella. Automated test case generation as a many-objective optimisation problem with dynamic selection of the targets. *IEEE Transactions on Software Engineering*, 44(2):122–158, 2017.
- [42] Annibale Panichella, Fitsum Meshesha Kifetew, and Paolo Tonella. Automated test case generation as a many-objective optimisation problem with dynamic selection of the targets. *IEEE Transactions on Software Engineering*, 44:122–158, 2018.
- [43] Sebastiano Panichella, Annibale Panichella, Moritz Beller, Andy Zaidman, and Harald C Gall. The impact of test case summaries on bug fixing performance: An empirical investigation. In *Proceedings of the 38th international conference on software engineering*, pages 547–558, 2016.
- [44] Anthony Peruma, Khalid Almalki, Christian D. Newman, Mohamed Wiem Mkaouer, Ali Ouni, and Fabio 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, CASCON '19*, page 193–202, USA, Nov. 2019. IBM Corp.
- [45] Anthony Peruma, Khalid Almalki, Christian D Newman, Mohamed Wiem Mkaouer, Ali Ouni, and Fabio 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*, pages 1650–1654, 2020.
- [46] José Miguel Rojas, Gordon Fraser, and Andrea Arcuri. Automated unit test generation during software development: A controlled experiment and think-aloud observations. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis, ISSTA 2015*, page 338–349, New York, NY, USA, 2015. Association for Computing Machinery.
- [47] Devjeet Roy, Ziyi Zhang, Maggie Ma, Venera Arnaoudova, Annibale Panichella, Sebastiano Panichella, Danielle Gonzalez, and Mehdi Mirakhorli. Deeptc-enhancer: Improving the readability of automatically generated tests. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering, ASE '20*, page 287–298, New York, NY, USA, 2021. Association for Computing Machinery.
- [48] Domenico Serra, Giovanni Grano, Fabio Palomba, Filomena Ferrucci, Harald C. Gall, and Alberto Bacchelli. On the effectiveness of manual and automatic unit test generation: Ten years later. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*, pages 121–125, 2019.
- [49] Sina Shamshiri, José Miguel Rojas, Gordon Fraser, and Phil McMinn. Random or genetic algorithm search for object-oriented test suite generation? In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation, GECCO '15*, page 1367–1374, New York, NY, USA, 2015. Association for Computing Machinery.
- [50] Sina Shamshiri, José Rojas, Luca Gazzola, Gordon Fraser, P.S. McMinn, Leonardo Mariani, and Andrea Arcuri. Random or evolutionary search for object-oriented test suite generation? *Software Testing, Verification & Reliability*, 2017.
- [51] Gareth Terry, Nikki Hayfield, Victoria Clarke, and Virginia Braun. Thematic analysis. *The Sage handbook of qualitative research in psychology*, pages 17–37, 2017.
- [52] Paolo Tonella. Evolutionary testing of classes. In *Proceedings of the 2004 ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA '04*, page 119–128, New York, NY, USA, 2004. Association for Computing Machinery.
- [53] Michele Tufano, Fabio Palomba, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, Andrea De Lucia, and Denys Poshyvanyk. An empirical investigation into the nature of test smells. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE '16*, page 4–15, New York, NY, USA, 2016. Association for Computing Machinery.
- [54] Tássio Virgínio, Luana Martins, Larissa Rocha, Railana Santana, Adriana Cruz, Heitor Costa, and Ivan Machado. Inose: Java test smell detector. In *Proceedings of the XXXIV Brazilian Symposium on Software Engineering, SBES '20*, page 564–569, New York, NY, USA, 2020. Association for Computing Machinery.