



Exploring the Interaction of Code Coverage and Non-Coverage Objectives in Search-Based Test Generation

Downloaded from: <https://research.chalmers.se>, 2025-10-15 13:39 UTC

Citation for the original published paper (version of record):

Fontes, A., Gay, G., Feldt, R. (2025). Exploring the Interaction of Code Coverage and Non-Coverage Objectives in Search-Based Test Generation. *Software Testing Verification and Reliability*, 35(6-7).
<http://dx.doi.org/10.1002/stvr.70009>

N.B. When citing this work, cite the original published paper.

RESEARCH ARTICLE OPEN ACCESS

Exploring the Interaction of Code Coverage and Non-Coverage Objectives in Search-Based Test Generation

Afonso Fontes | Gregory Gay  | Robert Feldt

Department of Computer Science and Engineering, Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden

Correspondence: Gregory Gay (greg@greggay.com)

Received: 10 June 2024 | **Revised:** 30 June 2025 | **Accepted:** 2 September 2025

Funding: This research was supported by Vetenskapsrådet grants 2019-05275 and 2020-05272. Computing resources were provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by Vetenskapsrådet grant agreement 2022-06725.

Keywords: adequacy criteria | automated test generation | branch coverage | coverage criteria | search-based test generation

ABSTRACT

Search-based test generation typically targets structural coverage of source code. Past research suggests that targeting coverage alone is insufficient to yield tests that achieve common testing goals (e.g., discovering situations where a class-under-test throws exceptions) or detect faults. A suggested alternative is to perform multi-objective optimization targeting both coverage and additional objectives directly related to the goals of interest. However, it is not fully clear how coverage and goal-based objectives interact during the generation process and what effects this interaction will have on the generated test suites. In this study, we assess five hypotheses about multi-objective test generation and the relationships between coverage-based and goal-based objectives, focusing on the effects on coverage, goal attainment, fault detection, test suite size, test case length and the impact of the search budget. We generate test suites using the EvoSuite framework targeting Branch Coverage, three testing goals—Exception Count, Output Coverage and Execution Time—and combinations of coverage and goal-based objectives. Ultimately, we find that targeting multiple objectives does not reduce code coverage, yields no or minor reductions in goal attainment, but—at the same time—detects more faults compared with single-target configurations. In addition, it produces larger test suites, but test case length is not increased. The benefits of multi-objective optimization are often more limited than hypothesized in past research, but improved fault detection is still sufficient to recommend multi-objective optimization over targeting coverage or testing goals alone. Our study offers insights and guidance into how coverage and goal-based objectives interact during multi-objective test generation.

1 | Introduction

Structural coverage criteria measure the percentage of the source code that has been executed according to a set of criterion-specific rules regarding (a) which code structures should be executed, and (b) how those structures should be executed [1–3]. Two of the most common criteria are Statement Coverage—which mandates that all code statements be executed, but places no constraints on how they are executed—and Branch Coverage—which mandates that all control-diverting

statements (e.g., if, case and loop conditions) evaluate to each of their possible outcomes [4].

Coverage measurement is a common advisory activity for testers [1]. The current percentage of coverage attained can serve as an approximation of ‘how much testing’ has been conducted, and missed coverage goals can serve as the targets of additional test cases. Because the attainment of most coverage criteria can be automatically measured through programme instrumentation and execution analysis, such criteria have also become the de

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2025 The Author(s). Software Testing, Verification and Reliability published by John Wiley & Sons Ltd.

facto basis of automated test generation—especially techniques such as search-based test generation, fuzzing, and symbolic or concolic execution [5, 6].

Consider, for example, search-based test generation. In search-based test generation, metaheuristic optimization algorithms sample from the space of possible test inputs to identify those that maximize or minimize *fitness functions*—numeric scoring functions representing properties of interest [6]. Coverage criteria serve as natural fitness functions, often associating each code structure of interest with a score representing how close an execution came to executing that structure in the manner prescribed by the criterion—e.g., how much x would need to change for the condition ($x = 0$) to evaluate to true within a particular control structure [7].

Coverage-directed testing is ubiquitous in automated test generation because structural coverage is easy to measure, easy to translate into an optimization target, and is hypothesized to have a correlation to the probability of fault detection [4]. However, concerns have been raised about its use as the primary target of automated generation [3, 8, 9]. We have previously conducted large-scale case studies on coverage-directed test generation, focusing on model and search-based test generation [3, 10, 11]. These studies have yielded important observations about the efficiency and effectiveness of coverage-directed test generation—at least, in the manner it is generally employed.

First, we have observed that achieving structural coverage is a reasonable *starting point* for effective automated test generation. For example, we observed that coverage was the single strongest predictor of the likelihood of fault discovery [11]. That is, if we want to detect potential faults, *we must execute the code*. The same basic observation holds for many other goals a tester may have. If we want to expose situations where the code can crash, *we must execute the code*. If we want to show that performance or reliability targets are met, *we must execute the code*. Other testing goals—e.g., diversity, exposing interaction faults, and more—similarly benefit from exploration of the codebase. Targeting code coverage during search-based test generation is an effective and efficient method of exploring a wide range of programme behaviours [11]. Therefore, even if a tester's goals lie beyond code coverage, *coverage is generally required to achieve those goals*.

However, we also observed that code coverage *alone* is a poor basis for producing test suites that meet these goals. In our past work, coverage only had a moderate correlation to the likelihood of fault detection [11], and was often weaker than random generation at detecting code mutations [3]. *Many different inputs can generally cover the same coverage goals*. While some coverage criteria are stricter than others, the majority impose few or no constraints on how code is executed [3, 11–13].

“How” is important. Testers rarely design tests for the sole purpose of attaining coverage [8, 14]. In practice, tests are designed around specifications, and coverage is used to identify clear

weaknesses in the suite [1]. That is—coverage serves an advisory role for testers, rather than the primary basis of test design. If we want to expose crashing code, we select input with a high probability of triggering a crash. If we want to violate performance requirements, we select input with a high probability of slowing programme execution. If there are multiple bugs in a branch, we typically need diverse inputs to uncover them all, as well as to cover the specification [15, 16]. In other words, while research in automated test generation has predominantly focused on code coverage, *coverage alone is not enough to ensure that testing goals are met*.

Search-based test generation offers potential solutions to this challenge. First, rather than optimizing fitness functions related to code coverage, one could attempt to optimize fitness functions based directly on goals of interest—for example, there are fitness functions that directly reward discovery of crashes or that assess performance. Second, rather than choosing coverage or goal optimization, one could attempt *multi-objective optimization*, where the combination of structural coverage and additional goal-based fitness functions are simultaneously targeted.

Multi-objective optimization is a particularly compelling solution, as each fitness function optimized shapes the resulting test suite. Our past research suggests that such a pairing can lead to better test suites than targeting coverage *or* a goal of interest alone [10, 11]. Consider a common testing goal—identifying situations where the system-under-test (SUT) throws an exception. This is a non-trivial goal, as we rarely know up front which exceptions could be thrown. Targeting coverage may not satisfy this goal, as exception-triggering input will only be chosen if it uniquely enhances coverage. We could alternatively try to directly maximize the number of exceptions thrown. However, this count offers no feedback; more exceptions will only be discovered by random chance. We observed situations where targeting both offered feedback missing when targeting either alone—with the exception count biasing the input used to attain coverage, and branch coverage offering a means to explore the code base.

These observations suggest the potential benefit of blending code coverage and goal-based fitness functions. While multi-objective test generation has been previously proposed and attempted (e.g., [11, 17, 18]), the interaction between objectives during this optimization—in particular, the interaction between coverage and goal-based fitness functions—has not been investigated previously. Understanding this interaction is important, as this understanding can influence the selection of optimization targets, the design of new fitness functions, and the development of new test generation tools.

Therefore, in this study, our goal is to assess and explore five hypotheses about this interaction:

Hypothesis 1. *The inclusion of goal-based fitness functions as additional generation targets will not have an impact on the attainment of code coverage, as compared with targeting coverage alone.*

That is, the hypothesis that targeting multiple objectives will not affect the evolution of code coverage during the test generation process—raising or lowering the final quantity of coverage attained, changing the specific coverage goals covered or affecting the rate at which coverage is attained during the generation process.

Hypothesis 2. *Targeting both coverage and a goal-based fitness function will not have an impact on the attainment of goal-based fitness functions, as compared with targeting coverage or a goal-based fitness function alone.*

Similar to Hypothesis 1, this hypothesis states that targeting multiple objectives will not affect the final fitness values of the goal-based objectives—raising or lowering goal attainment when compared with targeting a goal-based or a coverage-based objective alone.

Hypothesis 3. *Targeting both coverage and a goal-based fitness function will not have an impact on the fault detection of generated test suites, as compared with targeting coverage or a goal-based fitness function alone.*

This hypothesis states that targeting multiple objectives will not increase or decrease the likelihood that the generated test suites detect faults or the number of tests that fail when a fault is detected.

Hypothesis 4. *Targeting both coverage and a goal-based fitness function will not have an impact on the size of the test suite and the average test length, as compared with targeting coverage or a goal-based fitness function alone.*

Targeting multiple objectives could increase the number of test cases in the generated suites or increase the number of interactions in individual test cases, as each targeted objective adds additional obligations that the test suite must achieve. These obligations each may require distinct test input and setup to achieve, leading to the need for more or longer test cases.

Hypothesis 5. *An increase in the search budget will not lead to increased attainment of each objective.*

This hypothesis considers the effect of the search budget on fitness attainment. We further hypothesize that the effects that we observed when exploring the previous hypotheses will hold at higher search budgets. For example, if multi-objective optimization leads to higher fault detection at a limited search budget than single-objective optimization, we hypothesize that it will also do so at a higher search budget.

To assess these hypotheses, we target Branch Coverage—the most common structural coverage criterion [1]—as well as three specific testing goals:

- We further explore the goal of discovering situations where the SUT can crash.
- The discovery of situations that could violate performance goals—based on the maximization of execution time [19].
- Ensuring that test suites maximize coverage of diverse behaviours [16], specifically output diversity of the tested functions, which has been hypothesized to lead to faster coverage attainment and higher likelihood of fault detection [20].

Our study offers insight into how coverage and goal-based objectives interact during multi-objective test generation, with a focus on how this interaction affects code coverage, goal attainment, fault detection, the size of the test suite and the length of test cases. This research offers a starting point for exploring how search-based test generation can be adapted for particular goals, product domains, execution scenarios or code structures, enables guidance on how to use test generation to meet tester goals and can influence the creation of more efficient and effective test generation techniques and tools.

2 | Background

2.1 | Unit Testing

Testing can be performed at various levels of granularity. In this research, we focus on *unit testing*, where test cases target small segments of code that can be tested in isolation [21]. Unit tests are written as executable code, which can be re-executed on demand by developers. We refer to a purposefully grouped set of test cases as a *test suite*. Unit testing frameworks exist for many programming languages, such as JUnit for Java, and are integrated into most development environments.

An example of a unit test, written in JUnit, is shown in Figure 1. A unit test consists of a *test sequence (or procedure)*—a series of method calls to the class-under-test (CUT)—with *test input* provided to each method. Then, the test case will validate the output of the called methods and the class variables against a set of encoded expectations—the *test oracle*—to determine whether the test passes or fails. In a unit test, the oracle is typically formulated as a series of assertions on the values of method output and class attributes [22]. In the example in Figure 1, the *test input* consists of passing a string to the constructor of the `StringUtils` class, then calling its `removeWhitespace()` and

```

1      @Test
2      public void testPrintMessage() {
3          String str = "Test_Message";
4          StringUtils tCase = new StringUtils(str);
5          tCase.removeWhitespace();
6          assertEquals("TestMessage", tCase.getString());
7      }

```

FIGURE 1 | Example of a unit test case written using the JUnit notation for Java.

getString() methods. We use an assertion to ensure that a space is removed from the input.

2.2 | Adequacy (Coverage) Criteria

When testing, developers must judge both whether the tests they have written are effective and whether they have created enough test cases. Adequacy criteria have been developed to provide developers with guidance regarding these topics [1].

Each adequacy criterion prescribes, for a given programme, a set of goals—referred to as *test obligations*. Each obligation represents one constraint that the produced test suite must fulfil. If tests fulfil the test obligations, then testing is deemed ‘adequate’ with respect to faults that manifest through the structures of interest to the criterion. Most adequacy criteria are based on the execution of structural elements of the software. In such cases, an obligation may be expressed as the selection of an individual element of the source code—e.g., a statement, a branch of the software’s control flow or a boolean expression—and the conditions under which that element must be executed—e.g., the chosen expression evaluates to true or false [3, 21]. However, there are also adequacy criteria based, e.g., on coverage of formal requirements through test cases [23] or detection of synthetic faults (mutants) planted in the source code [10].

Adequacy criteria have seen widespread use in software development. Structural coverage is routinely measured as part of automated build processes [12] (for example, see <https://codecov.io/>) and is often mandated by safety standards in critical domains such as automotive [24] and avionics [25]. It is easy to understand the appeal of adequacy criteria. They offer clear checklists of testing goals that can be objectively evaluated and automatically measured through programme instrumentation and execution analysis [12]. These same qualities make adequacy criteria ideal for use as automated test generation targets [7].

One of the most common adequacy criteria is *Branch Coverage*. A branch refers to an outcome of any programme statement that can cause programme execution to diverge down a particular control flow path, such as the conditions in if, case or loop definitions. Branch Coverage requires that all outcomes of all control-diverging statements are executed at least once by the test suite under assessment.

To give an example, consider the removeWhitespace() method being tested in Figure 1, whose code is depicted in Figure 2. In this method, there are two programme statements that affect the control flow—the loop condition on Line 7 and the if-condition on Line 9. To achieve branch coverage over this method, both conditions must evaluate to true and false at least once when the test suite is executed. In other words, there are four test obligations that must be fulfilled.

By default, coverage obligations are formulated over the source code. However, in Java, test obligations are often instead formulated and measured over the bytecode representation as this form is easier and more efficient to instrument and monitor. The bytecode representation of removeWhitespace() is shown in Figure 3. The same control-altering expressions are present on Lines 11 and 19. Branch Coverage requires that both lines evaluate to true and false.

Branch Coverage is arguably the most commonly used coverage criterion, with ample tool support and industrial adoption [26]. For example, branch coverage measurement is built into the popular IntelliJ IDEA development environment. Therefore, we focus on Branch Coverage in this study as a representation of structural coverage criteria.

2.3 | Search-Based Test Generation

Manual creation of a large volume of test cases can be tedious and expensive. Automation of aspects of test creation, such as test input selection, can reduce and focus the required manual effort [5]. Search-based test generation frames input selection as a search problem, where metaheuristic optimization algorithms attempt to identify test input that best embody properties that testers seek in their test cases [5, 6].

These properties are assessed using one or more fitness functions—numeric scoring functions. The metaheuristic embeds a strategy for sampling solutions from the space of possible inputs, often based on a natural process such as evolution or swarm behaviour [27]. In test generation, a ‘solution’ is often either a single test case or a full test suite. The metaheuristic uses the selected fitness functions to assess solution quality, offering feedback to guide the selection and improvement of solutions over a series of generations. Search-based test generation has proven to

```

1      public class StringUtils{
2          private String str;
3          ...
4          public void removeWhitespace(){
5              String modified = "";
6
7              for (int index = 0; index < this.str.length(); index++) {
8                  char ch = this.str.charAt(index);
9                  if (!Character.isWhitespace(ch)) {
10                     modified += ch;
11                 }
12             }
13
14             this.str = modified;
15         }
16         ...
17     }

```

FIGURE 2 | Subset of the class-under-test in Figure 1.

```

1  public void removeWhitespace();
2  Code:
3      0: ldc          #7          // String
4      2: astore_1
5      3: iconst_0
6      4: istore_2
7      5: iload_2
8      6: aload_0
9      7: getfield      #9          // Field str:Ljava/lang/String;
10     10: invokevirtual #15         // Method java/lang/String.length:()I
11     13: if_icmpge     46
12     16: aload_0
13     17: getfield      #9          // Field str:Ljava/lang/String;
14     20: iload_2
15     21: invokevirtual #21         // Method java/lang/String.charAt:(I)C
16     24: istore_3
17     25: iload_3
18     26: invokestatic  #25         // Method java/lang/Character.isWhitespace:(C)Z
19     29: ifne         40
20     32: aload_1
21     33: iload_3
22     34: invokedynamic #31, 0      // InvokeDynamic #0:makeConcatWithConstants:
23                               // (Ljava/lang/String;C)Ljava/lang/String;
24     39: astore_1
25     40: iinc         2, 1
26     43: goto        5
27     46: aload_0
28     47: aload_1
29     48: putfield      #9          // Field str:Ljava/lang/String;
30     51: return

```

FIGURE 3 | Java Bytecode of the removeWhitespace() method from Figure 2.

be a flexible [28], scalable [29] and competitive [11, 18] method of automated test generation.

The most common metaheuristics for search-based test generation are genetic algorithms, which are modelled after the natural evolution of a population [30]. While specific aspects vary, a 'typical' test generation follows these steps:

- An initial population of solutions is randomly generated. Each solution represents a test suite, containing test cases.
- Each generation, the fitness score of each solution is calculated and a new population is created. This population is formed through four sources of solutions:
 - One of the best solutions may be carried over to the new population intact.
 - At a certain probability, elements of two solutions will be combined to create two 'children' (crossover). For example, the children may blend test cases from the parents.
 - At a certain probability, a solution can be mutated—e.g., a test case may be modified.
 - At a certain probability, a new randomly generated solution will be added to the population to maintain diversity.
- When the search budget—typically expressed in time or number of generations—expires, the best solution is returned.

When multiple fitness functions are targeted by an optimization algorithm, each fitness function and its individual sub-goals collectively influence the final solution generated. The interaction between fitness functions—or even between the sub-goals of one or more fitness functions—can be either positive or negative.

In the positive case, one or more fitness functions or sub-goals can provide missing feedback that is needed to optimize other fitness functions or sub-goals. This is common in situations where the fitness landscape is relatively flat, as is the case for functions based on simple counts—e.g., the number of exceptions thrown. Such fitness functions offer little feedback for improving solutions, e.g., for detecting new exceptions. However, in previous work, we observed cases where pairing such functions with more informative fitness functions (e.g., ones based on distance measurements) offered the missing feedback needed to identify additional exceptions [10, 11]. Even informative fitness functions can have a positive symbiotic relationship. For example, there are multiple methods of formulating a fitness function for Branch Coverage, each with differing fitness landscapes. We previously observed situations where targeting multiple forms of Branch Coverage simultaneously yielded improved performance over targeting any one form [31].

In the negative case, there are situations where optimization of a subset of fitness functions or sub-goals can lower the attainment of other functions or sub-goals—i.e., there are implicit trade-offs between these functions or sub-goals. As an example, EvoSuite offers a fitness function that rewards executing methods without an exception being triggered. This function directly contradicts the earlier-mentioned function that rewards suites where more exceptions are thrown. It is possible to optimize both functions by constructing a test suite where these functions are optimized in *separate* test cases [10, 11]. However, overcoming the contradiction between functions requires investing additional time in the search process and can result in larger and more cumbersome test suites.

This theoretical relationship between fitness functions and sub-goals of fitness functions is not well understood in the context of

search-based test generation. While we have discussed examples observed in prior work [10, 11, 32], the exact nature of this relationship has not been deeply explored. One of the goals of this study is to further empirically investigate this relationship.

3 | Related Work

Multi and many-objective optimization algorithms have become increasingly common in search-based test generation [33]. Even if the goal of the test generation process is solely code coverage, coverage can quickly be gained by representing each test obligation as an independent objective and applying multi or many-objective optimization [34]. Multiple studies have compared different algorithms for multi and many-objective optimization in terms of coverage achieved (e.g., [35–38]). However, these studies focused solely on coverage-based fitness and have not examined the interaction between coverage-based and goal-based fitness functions.

The number of exceptions or crashes discovered is a common secondary objective in search-based test generation, optimized in conjunction with coverage-based fitness functions (e.g., [11, 39–41]). Others have explored combinations of coverage criteria with non-functional criteria during test generation or test suite minimization, such as memory consumption [42] or execution time [43]. While these represent multi-objective optimization of coverage and goal-based fitness functions, these studies do not examine how these fitness functions interact, e.g., how the combinations affect coverage or fault detection.

Rojas et al. examined multi-objective optimization of Line Coverage—a structural coverage criterion—and additional fitness functions [41]. Relevant to our work, they also include two of the same goal-based objectives that we focus on, exception count and output diversity. They found that adding additional fitness functions led to only a minimal loss in the final percentage of Line Coverage achieved. They also found that coverage of secondary criteria increased over when Line Coverage was targeted alone. Therefore, there is a partial overlap in our focus. However, they only examined the final level of coverage and focused on different aspects of test generation. We address a broader set of hypotheses and examine coverage attainment more deeply.

Palomba et al. examine optimization of Branch Coverage and fitness functions intended to improve test quality based on the cohesion and coupling of test cases [17]. They found that targeting these quality objectives could increase code coverage over targeting coverage alone.

Weiglhofer et al. showed that coverage-directed test generation can be used to complement test generation based on testing goals [44]. In their approach, humans develop ‘test purposes’, specifications used in conjunction with formal models to generate test cases. Coverage-directed testing is then used to generate tests for parts of systems not covered by the test purposes. They do not apply multi-objective optimization, but the core concept is similar.

We previously examined the likelihood of fault detection of test suites generated targeting various fitness functions [10, 11].

Much of this research focused on single-objective optimization. However, we did find that some combinations of objectives, such as Branch Coverage and the exception count, had a higher likelihood of fault detection than targeting Branch Coverage or exception count alone [11]. This past work partially addresses one of our hypotheses, but we replicate this work and examine that hypothesis more closely in this study.

Zhou et al. propose an approach, ‘smart selection’, for selecting a subset of test obligations when targeting multiple coverage-based fitness functions for test generation [45]. Their approach reduces redundancy between fitness functions and eases optimization difficulty. McMinn et al. have also proposed using search techniques to evolve new coverage criteria that combine features of existing criteria [46]. In previous work, we also used reinforcement learning to dynamically select the fitness functions targeted during multi-objective test generation [47]. We demonstrated that fitness functions could be identified that increased attainment of common testing goals for particular classes-under-test. However, these studies do not examine the interaction between coverage and goal-based fitness functions.

4 | Methods

Our aim in this research is to examine the interaction between coverage-based and goal-based fitness functions during multi-objective test generation. In Section 1, we raised five informal hypotheses about how these objectives could interact. We assess those hypotheses by addressing the following specific research questions:

- Research Question 1.** How is the Branch Coverage of generated test suites influenced by targeting additional goal-based fitness functions compared with targeting Branch alone?
- Research Question 1.1.** How is the final percentage of attained Branch Coverage influenced?
- Research Question 1.2.** How is the set of satisfied Branch Coverage obligations influenced?
- Research Question 1.3.** How is the evolution of coverage attainment influenced?
- Research Question 2.** How is the attainment of testing goals by generated test suites influenced by targeting Branch Coverage in addition to a goal-based fitness function, compared with targeting coverage or a goal-based fitness function alone?
- Research Question 3.** How is the fault detection of generated test suites influenced by targeting Branch Coverage in addition to a goal-based fitness function, compared with targeting coverage or a goal-based fitness function alone?
- Research Question 4.** How is the suite size and test case length of generated test suites influenced by targeting Branch Coverage in addition to a goal-based fitness function, compared with targeting coverage or a goal-based fitness function alone?
- Research Question 5.** What influence does the search budget have on Branch Coverage, goal attainment, fault

detection, test suite size, and test case length attained by suites targeting different fitness function configurations?

As discussed in Section 1, we focus on three concrete testing goals: (1) discovery of scenarios where exceptions are thrown ('Exception Count'), (2) discovery of scenarios where the execution time may violate performance goals ('Execution Time') and (3) maximization of output diversity ('Output Coverage'). To address these research questions, we have performed the following experiment, targeting Branch Coverage, these three goals and combinations of Branch Coverage with each goal:

1. **Collected case examples:** We have selected 93 case examples from the Defects4J fault dataset, from 14 Java projects (Section 4.1).
2. **Defined test generation configurations:** We selected three single-objective configurations (Branch Coverage, Exception Count, Output Coverage) and three multi-objective configurations (Branch Coverage plus each testing goal listed above) and two search budgets (180 and 300s) to target in our experiments (Section 4.2).
3. **Generated test suites:** For each class modified by each case example, fitness function configuration and search budget, we generated 10 test suites using EvoSuite. We target the fixed version of each CUT (Section 4.2).
4. **Monitored coverage evolution:** We monitor how satisfaction of Branch Coverage obligations changes over the course of each invocation of EvoSuite (Section 4.3).
5. **Recorded generation statistics:** For each suite generated, at the end of the generation process, we record information on Branch Coverage obligation satisfaction, fitness values for each targeted function, test suite size and test case length (Section 4.3).
6. **Removed non-compiling and flaky tests:** Any tests that do not compile or that return inconsistent results are removed (Section 4.3).
7. **Assessed fault-finding effectiveness:** We measure the number of faults detected, the proportion of test suites that detect each fault to the number generated (likelihood of fault detection) and the number of failing tests in each suite (Section 4.3).
8. **Analysed the collected data:** We address the research questions using the data gathered above (Section 4.4).

We make a replication package available containing the data collected in this experiment: <https://doi.org/10.5281/zenodo.11047567>.

We also make available our modified version of EvoSuite:

- Code: <https://github.com/afonsohfontes/evosuite>
- Executable: https://github.com/afonsohfontes/defects4j/tree/master/framework/lib/test_generation/generation

4.1 | Case Example Selection

Defects4J is an extensible database of real faults extracted from Java projects [48] (available from <http://defects4j.org>). The current dataset, Version 2.0.1, consists of 835 faults from 17 Java projects. To control experiment costs, in this study, we aimed to select a sample of approximately 100 faults, chosen to reflect the proportion of faults-per-project in the full dataset. To select this sample, we initially selected **206** faults at random, sampled based on the number of faults-per-project in the full dataset. We then generated test suites targeting Branch Coverage and the three multi-objective configurations following the procedure described in Section 4.2, and omitted faults where either:

- Errors prevented the completion of 10 valid trials for all configurations, where a test suite was generated and all data collection completed successfully.
- Where the average Branch Coverage was below 5%—we judged that the research questions could not be reliably answered without a minimal level of coverage being reached over the classes-under-test.

This filtering process ultimately resulted in a set of **93** faults used in this study, listed in Table 1. Fifty-nine faults were excluded due to issues with the test generation or data collection processes, and 54 faults were excluded due to low Branch Coverage.

For each fault, Defects4J provides access to the faulty and fixed versions of the code, developer-written test cases that expose the

TABLE 1 | Subset of Defects4J faults selected for this study.

Project	Faults selected	Total
Chart	7, 6, 10, 8, 3, 5	6
Cli	27, 7, 29, 28, 1, 10, 3, 40, 2, 5	10
Closure	161, 74, 19, 154, 162, 164, 37, 55, 41, 70, 12, 71, 5	13
Codec	7, 6, 17, 1, 2	5
Collections	25, 26	2
Compress	2, 47, 46	3
Csv	1	1
Gson	3, 4, 5, 6	4
JacksonCore	11	1
JacksonDatabind	62, 93, 111, 112	3
Jsoup	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 55, 60, 77	16
Lang	4, 5, 6, 8, 9, 10, 11, 12, 41, 55, 64, 65	12
Math	95, 11, 87, 81, 100, 39, 90, 41, 3, 49, 40, 2	12
Mockito	6, 8, 37, 15, 2	5
	Total	93

fault, and a list of classes and lines of code modified by the patch that fixes the fault.

Each fault is required to meet three properties. First, a pair of code versions must exist that differ only by the minimum changes required to address the fault. The ‘fixed’ version must be explicitly labelled as a fix to an issue, and changes imposed by the fix must be to source code, not to other project artefacts such as the build system. Second, the fault must be reproducible—at least one test must pass on the fixed version and fail on the faulty version. Third, the fix must be isolated from unrelated code changes such as refactorings.

4.2 | Test Generation Configuration

In this study, we make use of the EvoSuite unit test generation framework for Java [49]. EvoSuite is mature, actively maintained, and has been successfully applied to a wide variety of projects [11, 18]—even winning multiple tool competitions (e.g., [50]). Specifically, we make use of a modified version of EvoSuite Version 1.2.1, where we have added an additional fitness function—Execution Time—as well as additional monitoring and data collection capabilities.

4.2.1 | Test Generation Algorithm

We make use of EvoSuite’s ‘whole test suite generation’ Genetic Algorithm [49]. This implementation of whole test suite generation has been replaced as the default optimization algorithm in EvoSuite by DynaMOSA, a many-objective optimization algorithm [51]. While DynaMOSA has been shown to achieve better coverage than whole test suite generation in some experiments [36], we use whole test suite generation to enable clearer comparison to our past research [10, 11]. Further, the DynaMOSA algorithm explicitly considers the code structure, and goal-based fitness functions cannot be decoupled from structural coverage. In this implementation of whole test suite generation, each solution represents a full test suite—in contrast to approaches where a solution represents a single test case. Then, rather than targeting one obligation (sub-goal) of each fitness function one-by-one, fitness is measured over all obligations of each fitness function at the same time.

In traditional multi-objective optimization algorithms, such as NSGA-II [52], an attempt is made to balance fitness function attainment, and each fitness function is treated as independent. In contrast, in this implementation of whole test suite generation, a single aggregate fitness score is calculated. The fitness for a test suite T over the CUT C is:

$$\text{fitness}(T, C) = \sum_{f \in F} \hat{f}(T, C) \quad (1)$$

That is, the aggregate fitness is the sum of the normalized score of each fitness function. EvoSuite treats all optimizations as *minimization* problems, where lower fitness scores represent better solutions.

4.2.2 | Fitness Function Configurations

We execute EvoSuite for each case example utilizing six fitness function configurations, representing three single-objective configurations (Branch Coverage, Exception Count and Output Coverage) and three multi-objective configurations (Branch Coverage with Exception Count, Output Coverage and Execution Time). The fitness functions are defined as follows:

Branch coverage: As defined in Section 2, Branch Coverage requires that all outcomes of all control-diverging statements are executed at least once by a test suite. For search-based test generation to be most effective, a fitness score should offer feedback to help guide the identification of better solutions. To that end, the most effective fitness functions tend to encode information about the *distance* to satisfying any unsatisfied goals. Therefore, rather than simply measuring whether each test obligation is covered or not, the fitness calculation for Branch Coverage instead embeds information—for each test obligation—on *how close execution came to satisfying that obligation*.

The branch coverage fitness function is a minimization of the following, where T refers to the test suite and B represents the set of test obligations. Each test obligation, $b \in B$, represents a control-diverging programme statement and a desired outcome for that statement (true or false).

$$\text{fitness}(T, B) = \sum_{b \in B} d(T, b) \quad (2)$$

where $d(T, b)$ is defined as:

$$d(T, b) = \begin{cases} 0 & \text{if } b \text{ has been satisfied} \\ \frac{0}{d_{\min}(t \in T, b)} & \text{if } b \text{ has been evaluated at least twice} \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

In the case where an obligation has not been satisfied, $\frac{0}{d_{\min}(t \in T, b)}$ represents the *branch distance*—the magnitude of change in execution that would be needed to achieve the targeted outcome for that control-diverging statement. The branch distance is determined based on how the condition has been formulated, following a standard set of formulae [7]. In this case, the minimal observed value of the branch distance is used in the fitness calculation, and is normalized to be between 0 and 1.

Exception count: This fitness function represents the goal of causing the CUT to throw as many exceptions as possible—either declared or undeclared. The fitness function is a minimization of the following formula, where T refers to the test suite, $E_{\text{discovered}}$ represents the number of exceptions discovered during the current generation process to date, and E_{thrown} represents the number thrown by the current solution T .

$$\text{fitness}(T) = 1 - \frac{E_{\text{thrown}}}{E_{\text{discovered}}} \quad (4)$$

$E_{discovered}$ is the number of unique exceptions thrown by *all* assessed solutions to date, while E_{thrown} is the number of unique exceptions thrown by the current test solution-under-assessment, which may not throw exceptions that were thrown in previous solutions. Therefore, $E_{thrown} \leq E_{discovered}$. As the number of possible exceptions that a class can throw cannot be known ahead of time, the number of test obligations may change each time EvoSuite is executed on a CUT.

Execution time: This fitness function represents a scenario where we seek test suites that could uncover potential violations of performance requirements, manifested as a test suite—containing individually short test cases—that take excessive time to execute. We have added a new fitness function to EvoSuite for this purpose. Due to technical details of its implementation, we are unable to target Execution Time without also targeting Branch Coverage. Therefore, Execution Time cannot currently be targeted as a single-objective configuration. The fitness function is calculated as follows:

$$\text{fitness}(T) = 1 - \frac{\text{Time}_{\text{current}}}{\text{Time}_{\text{max}}} + \text{penalty} \quad (5)$$

where $\text{Time}_{\text{current}}$ represents the execution time of the solution currently under assessment and Time_{max} is the largest execution time of any solution assessed during the current execution of the test generation framework.

One avenue to generate test suites with long execution times is simply to generate excessively long test cases that call many methods. Therefore, to prevent the generation of overly bloated test cases, the fitness calculation applies a penalty based on the average test case length within the suite:

$$\text{penalty} = 0.1 \times \frac{\text{Length}_{\text{current}}}{\text{Length}_{\text{max}}} \quad (6)$$

Output coverage: This configuration represents the goal of generating test suites that cover many different types of outcomes of the methods of the CUT. A tester may seek such diversity for two reasons. First, increased output coverage is hypothesized to lead to earlier and potentially higher code coverage [15, 16], and second, to potentially increase fault detection over pure white-box techniques [20]

Output coverage rewards diversity in the method output by mapping return types to a list of abstract values—Alshahwan and Harman provide a detailed explanation, including fitness formulae [20]. A test suite satisfies output coverage if, for each public method in the CUT that returns a data type covered by the fitness function, at least one test yields a concrete return value matching each abstract value. For numeric data types, distance functions similar to the branch distance offer feedback using the difference between the chosen value and the targeted abstract values.

We have selected these three goal-based objectives because they reflect three different, non-overlapping, goals that a tester may have. Exception Count reflects the desire of testers to discover situations where the software can crash, as unexpected crashes can indicate a lack of robustness in the code-under-test. Crashes

can occur in any software, can be detected without the need for specialized test oracles, and are the target of many test generation tools (e.g., fuzzers [53]). Output Coverage indicates that many different programme behaviours have been triggered, potentially ensuring that the actual requirements of the code have been thoroughly exercised [20]. Finally, the execution time represents an important non-functional property of software—its performance. Even if the correct execution outcome occurs, slow performance may degrade user satisfaction [43]. While other goal-based fitness functions exist, we believe that these three exemplify three important, common and distinct testing goals.

4.2.3 | Search Budgets

Two search budgets were used—180 and 300s per class. This allows us to examine how an increased search budget affects the test suites produced by each single- and multi-objective configuration.

4.2.4 | Generation Procedure

Test suites are generated individually for each of the classes modified to fix each fault chosen from Defects4J. We repeat generation a fixed number of times for each class, fitness function configuration, and search budget.

Test suites are generated targeting the fixed version of each CUT and applied to the faulty version to eliminate the oracle problem. EvoSuite generates assertion-based oracles. Generating oracles based on the fixed version of the class means that we can confirm that the fault is actually detected, and not just that there are coincidental differences in programme output. This translates to a regression testing scenario, where tests are generated using a version of the system understood to be ‘correct’ in order to guard against future issues. Tests that fail on the faulty version detect behavioural differences between the two versions.

Test suite generation and execution were performed on virtual machines, each configured with 4 vCPUs, 8 GB of RAM and 20 GB of storage, running a server version of Ubuntu 18.04.4 LTS. Each virtual machine was dedicated to executing experiments for a specific subset of faults and fitness function configurations, ensuring that experiments remained isolated and independent to ensure result reliability.

To control experiment cost, we deactivated assertion filtering—all possible regression assertions are included. We also disabled test suite reduction, an optional procedure that removes redundant test cases at the end of the generation process. We do this to maintain traceability between intermediate and final test suites during suite evolution. All other settings were kept at their default values. As results may vary, we performed 10 trials for each CUT, fitness function configuration, and search budget. This resulted in the generation of 11,160 test suites (2 budgets, 10 trials, 6 configurations, and 93 faults).

Generation tools may generate flaky (unstable) tests [18]. For example, a test case that makes assertions about the system time will only pass during generation. We automatically remove flaky

tests. First, all non-compiling test suites are removed. Then, each remaining test suite is executed on the fixed version of the CUT. If the test results are inconsistent, the test case is removed. This process is repeated until all tests pass five times in a row. On average, less than 1% of test cases were removed from each suite.

4.3 | Data Collection

To answer our research questions, we capture the following data during and after generation:

- **Final fitness function values:** For each test suite, we record the final fitness values for all four fitness functions considered in this experiment (Branch Coverage, Exception Count, Execution Time and Output Coverage).
- **Branch coverage obligation satisfaction:** Given a CUT, achieving Branch Coverage requires satisfying a set of test obligations, as defined in Section 2. We record information on the satisfaction of Branch Coverage obligations, including:
 - **Number of test obligations:** For each CUT, we record the number of Branch Coverage obligations.
 - **Percentage of obligations satisfied:** For each final test suite, we record the percentage of Branch Coverage obligations satisfied.
 - **Specific obligations satisfied:** For each final test suite, we record the specific obligations satisfied.
 - **Evolution of branch coverage during generation:** To understand the dynamic evolution of coverage over the course of each invocation of EvoSuite, we tracked the percentage of obligations and specific obligations covered by the best test suite in the population once per second during the generation process.
- **Fault detection:** To evaluate the fault-finding effectiveness of the generated test suites, we execute each test suite against the faulty version of each CUT. We then record the following:
 - **Likelihood of fault detection:** Across all trials for a particular fault, fitness function configuration, and search budget, we record the proportion of trials where the fault was detected to the total number of trials for that configuration.
 - **Number of failing tests:** For each test suite, we record the number of test cases that detect that fault (pass on the fixed version and fail on the faulty version).
- **Test suite size:** We recorded the number of tests in each test suite.
- **Average test case length:** Each test consists of one or more method calls, variable initializations and assertions. We record the average number of lines in each test case.

4.4 | Data Analysis

We answer each research question using the data gathered, comparing results attained by each fitness function configuration, split based on the search budget. To analyse the data, we employ a combination of descriptive statistics, distribution comparisons and effect size tests when distributions are found to differ.

Further explanation is provided in Section 5. Here, we provide a general overview of the data analysis procedure.

4.4.1 | Descriptive Statistics

Descriptive statistics provide an initial overview of the collected data.

1. **Data analysis:** Basic statistical measures such as the average, median, standard deviation, and percentiles are calculated for data appropriate for answering each research question. This provides an initial understanding of the data distribution and central tendencies [54].
2. **Data visualization:** We utilize box plots as a graphical representation to offer a visual insight into the result distribution across different configurations [55].

4.4.2 | Distribution Comparisons

We are interested in assessing whether the observed differences between two fitness function configurations at a particular search are significantly different.

For each research question, we select data relevant to that question (e.g., Branch Coverage attainment in Research Question 1). Then, for each pair of fitness function configurations, we formulate a hypothesis and a null hypothesis in the following format:

- H : Generated test suites have different distributions of X depending on the targeted fitness function configuration.
- H_0 : Observations of X for both configurations are drawn from the same distribution.

The informal hypotheses raised in Section 1 correspond to the null hypotheses used to answer each research question. Our observations for each of the collected data items defined above are drawn from an unknown distribution. To evaluate the null hypothesis without any assumptions on distribution, we use the Wilcoxon rank-sum test [56], a non-parametric test. We apply the test with $\alpha = 0.05$. A p -value less than α indicates a statistically significant difference [57].

4.4.3 | Effect Size

The Wilcoxon test determines if there are significant differences between the distributions of two configurations. To understand cases where there are differences, we use Cliff's delta to measure the effect size of these differences, providing a clearer understanding of their magnitude and practical significance [58]. We apply the standard interpretation of *Cliff's delta* (δ):

- $\delta > 0$ indicates that observations of configuration A are more likely to have higher values than observations of configuration B .
- $\delta < 0$ indicates that observations of configuration A are more likely to have lower values than observations of configuration B .

- The absolute value of δ is categorized as follows for further interpretation:
 - $|\delta| < 0.15$: Negligible effect
 - $0.15 \leq |\delta| < 0.33$: Small effect
 - $0.33 \leq |\delta| < 0.47$: Medium effect
 - $|\delta| \geq 0.47$: Large effect

5 | Results

5.1 | Effect on Structural Coverage (Research Question 1)

In this section, we address the following hypothesis:

Hypothesis 1. *The inclusion of goal-based fitness functions as additional generation targets will not have an impact on the attainment of code coverage, as compared with targeting coverage alone.*

Often, targeting multiple objectives can have *some* effect on each individual objective targeted, as compared with targeting a single

objective on its own. If objectives are contradictory, or if too many objectives are targeted at once, then the final attainment of each may be lowered [11, 41]. However, targeting one objective may also offer feedback that enhances attainment of another [10, 47]. Therefore, we wish to assess—first—the impact that targeting additional goal-based objectives has on code coverage-based objectives. We examine three aspects of coverage: (1) the final percentage of coverage attained, (2) the specific coverage obligations covered by the final test suites and (3) the evolution of coverage attainment over evolution. For this evaluation, we compare suite generated targeting Branch Coverage alone to suite targeting Branch Coverage and an additional goal-based fitness function.

5.1.1 | Attained Branch Coverage

Table 2 offers descriptive statistics on the final attainment of Branch Coverage by generated test suites. Figures 4 and 5 also depict the attained Branch Coverage overall and by project from Defects4J, respectively. Table 2 and Figure 4 do not demonstrate any clear differences between configurations with regard to the

TABLE 2 | Descriptive statistics of Branch Coverage across different test generation configurations and search budgets.

Configuration	Budget	Mean	Std Dev	Min	25th %	Median	75th %	Max
Branch	180	0.618	0.326	0.000	0.383	0.690	0.906	1.000
Branch and Exception	180	0.614	0.333	0.000	0.367	0.690	0.912	1.000
Branch and Execution Time	180	0.621	0.324	0.000	0.389	0.704	0.900	1.000
Branch and Output	180	0.620	0.325	0.000	0.383	0.704	0.900	1.000
Branch	300	0.665	0.317	0.000	0.484	0.726	0.929	1.000
Branch and Exception	300	0.665	0.325	0.000	0.507	0.749	0.936	1.000
Branch and Execution Time	300	0.659	0.322	0.000	0.484	0.742	0.932	1.000
Branch and Output	300	0.653	0.324	0.000	0.475	0.740	0.934	1.000

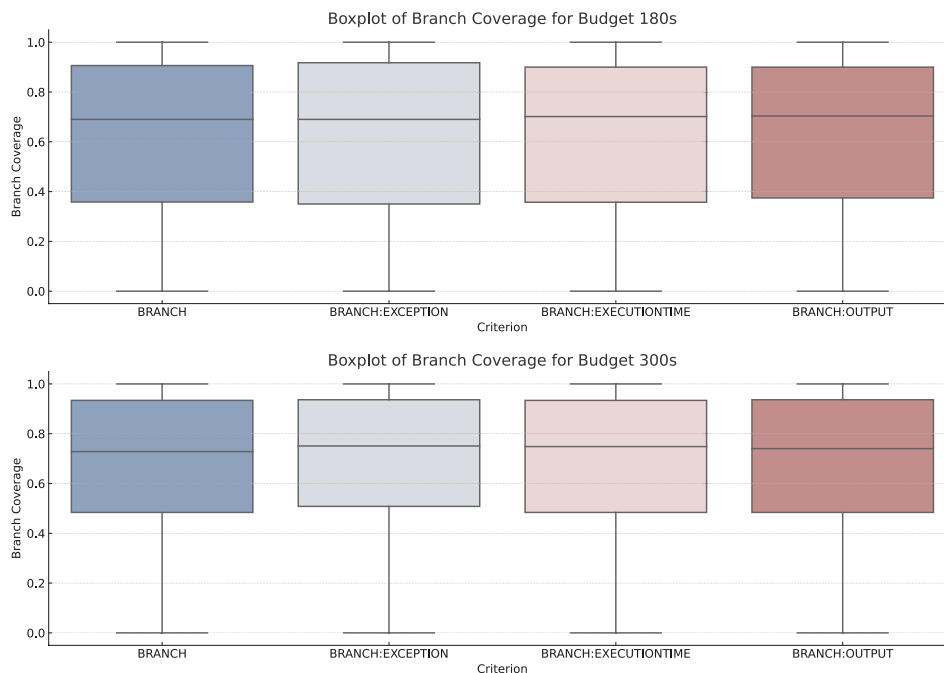


FIGURE 4 | Boxplots of the Branch Coverage attained by test suites, divided by budget.

final attained Branch Coverage. The distribution of results is visually similar for each configuration, and the mean and median Branch Coverage attained by each configuration are within a narrow range. An increase in search budget yields an increase in the average Branch Coverage, as well as less variance in the final results—seen in a rise in the 25th percentile. However, this improvement seems largely consistent across configurations.

In Figure 5, we do see some differences between configurations for particular projects. However, there are few clear trends, and only a small number of bugs were drawn from many of these projects. Still, we note some observations from the projects with over 10 included bugs. First, for project Cli, we see a higher median Branch Coverage for the combination of Branch Coverage and Execution Time at both search budgets. For

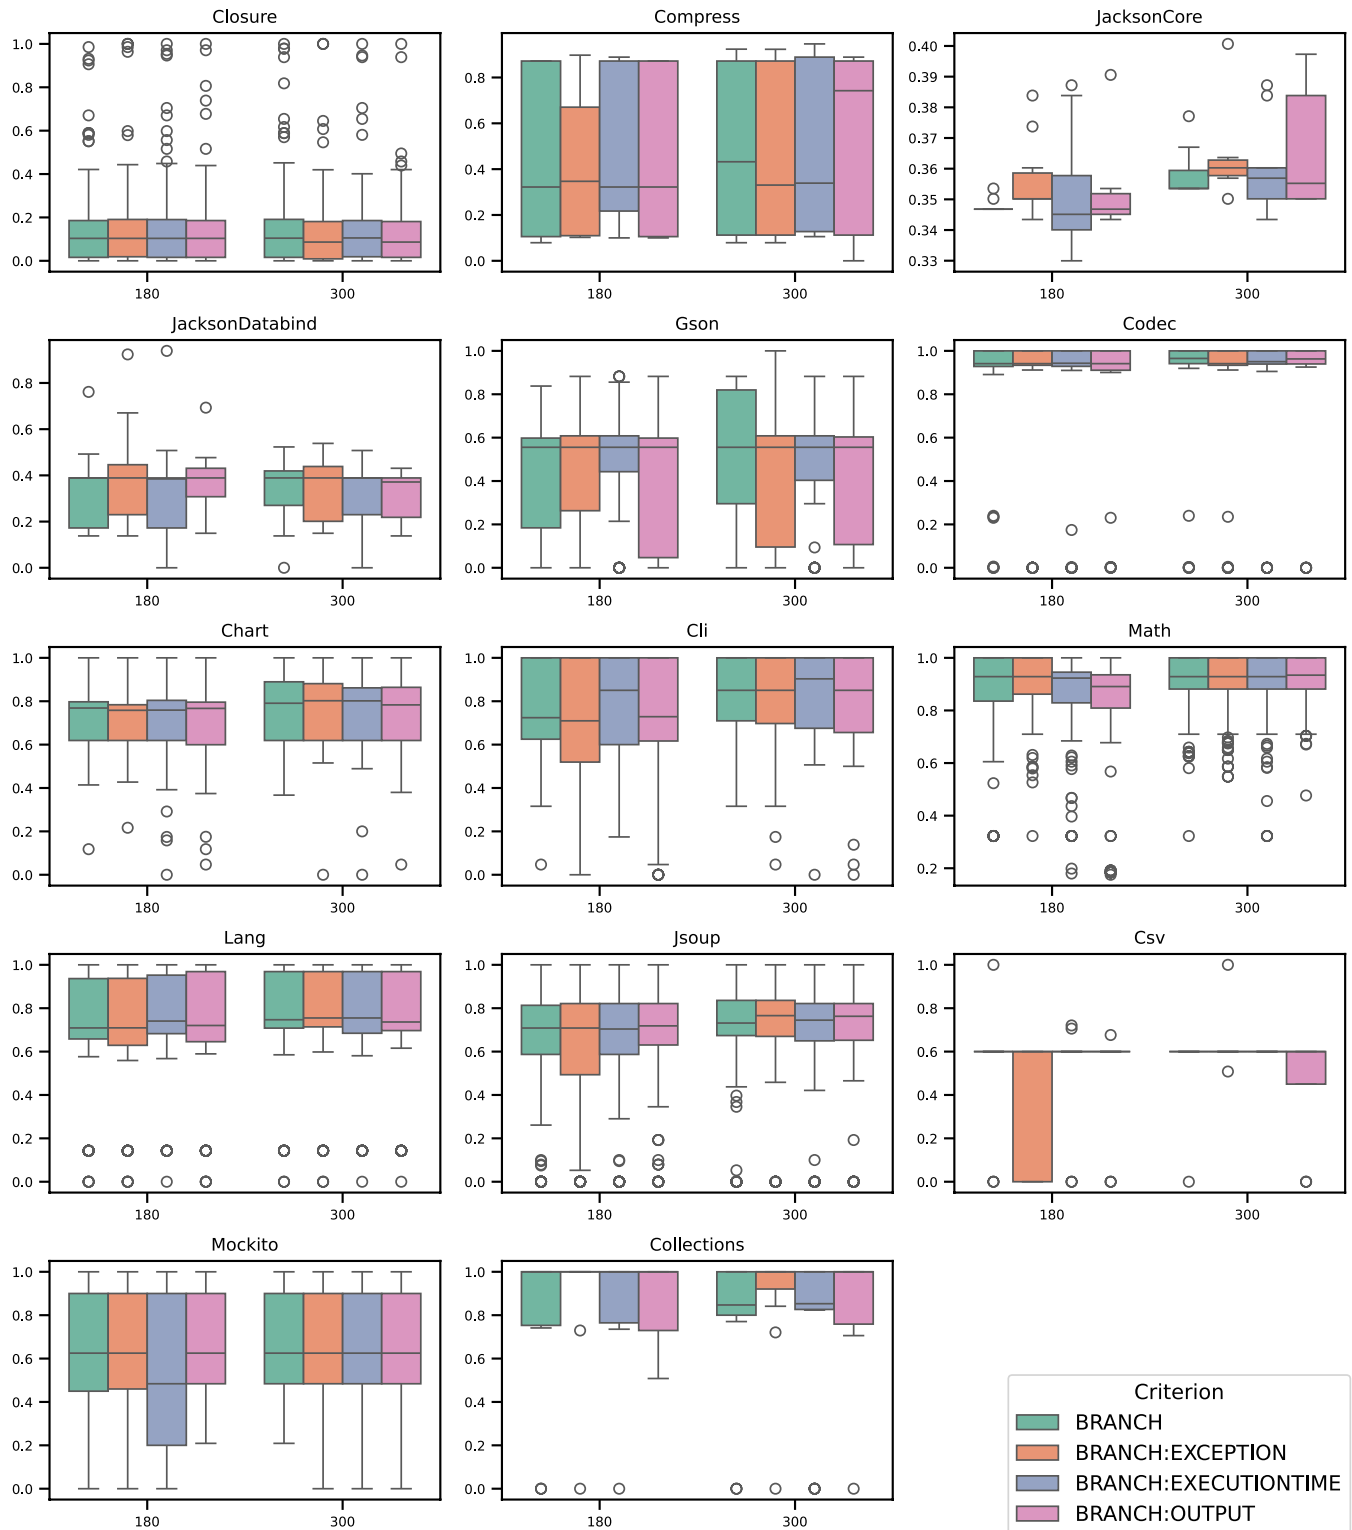
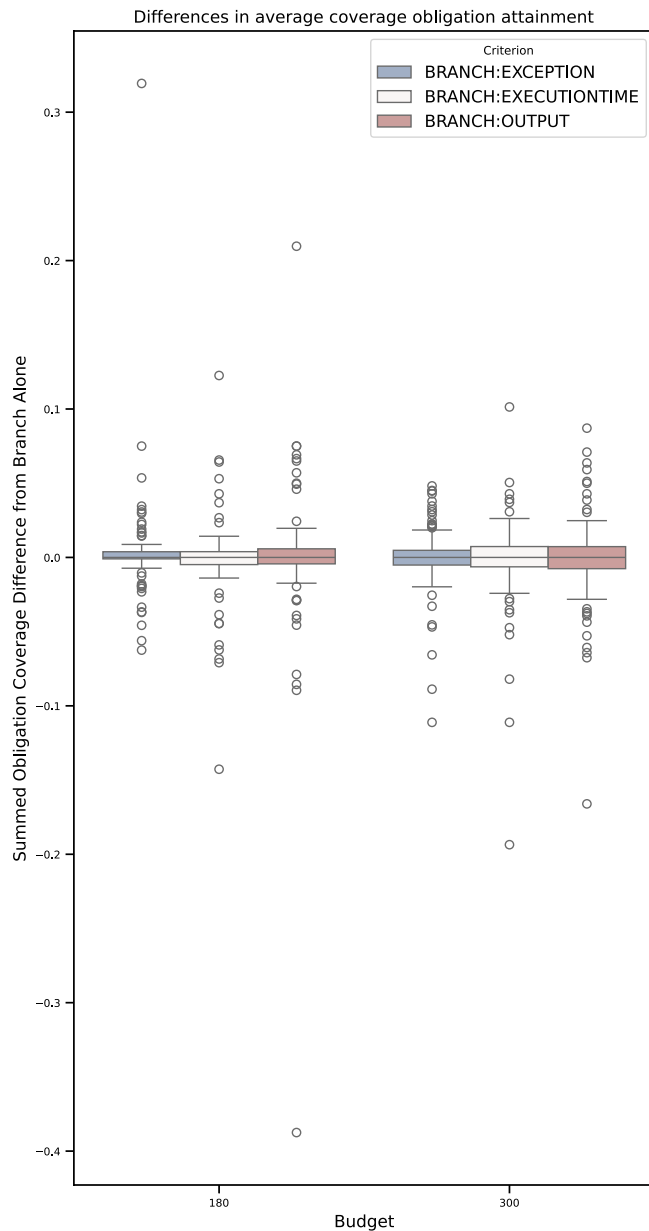


FIGURE 5 | Boxplots of the Branch Coverage, divided by both budget and project from Defects4J. The X-axis reports the budget and the Y-axis reports the Branch coverage.

TABLE 3 | Calculated p -values from comparisons of attained Branch Coverage by different configurations, split by search budget.

Comparison	Budget	p -value	Significant	Cliff's δ	Effect size
Branch vs. Branch and Execution Time	180	0.909	No	−0.005	Negligible
Branch vs. Branch and Exception	180	0.362	No	−0.007	Negligible
Branch vs. Branch and Output	180	0.647	No	0.005	Negligible
Branch vs. Branch and Exception	300	0.985	No	−0.017	Negligible
Branch vs. Branch and Execution Time	300	0.685	No	−0.000	Negligible
Branch vs. Branch and Output	300	0.347	No	0.023	Negligible

**FIGURE 6** | Difference in the average coverage of each obligation by Branch alone against Branch and Exception (green), Branch and Execution Time (blue) and Branch and Output (purple) for Chart-3, Class 1, when the budget is set to 180s.

the project JSoup, we see that the combination of Branch and Output Coverage yields a slightly higher median coverage at both search budgets. We will investigate both of these observations—as well as potential differences for other projects—more closely in future work.

To confirm our initial inspection, we performed a Wilcoxon rank-sum test to assess pairwise comparisons between different test generation configurations for the two search budgets. The results of this test are shown in Table 3, where we see that no comparison demonstrated statistically significant differences—that is, no p -value was below 0.05.

Research Question 1.1 (attained Branch Coverage): Optimizing a second goal-based fitness function does not have a significant impact on the final Branch Coverage attained by test suites.

5.1.2 | Attained Coverage Obligations

During and after the test generation process, we collected information on which specific Branch Coverage obligations were covered by generated test suites. To assess whether different configurations tend to cover distinct test obligations, we calculated the average coverage of each obligation for each CUT across all trials conducted for each configuration, search budget and bug. For example, if four of the 10 trials targeting Branch Coverage and Output Coverage for bug Chart-3 covered the first Branch Coverage obligation for the targeted class, then the average coverage of that obligation would be 0.40.

The resulting averages were then used to compare targeting Branch Coverage alone to targeting Branch Coverage as well as a second goal-based fitness function. For example, if there were four coverage obligations for a class:

- When targeting Branch Coverage alone, the average coverage of each obligation was 0.4, 0.7, 0.8 and 0.4.
- When targeting Branch and Exception Count, the average coverage of each obligation was 0.5, 0.7, 0.8 and 0.3.
- The resulting difference between the two would be −0.1, 0.0, 0.0 and 0.1.

Figure 6 visualizes the result of this comparison for one class from one bug and search budget, as an example of the differences that can emerge. In this example, positive spikes show cases where targeting Branch alone performed better for a particular obligation versus the compared configuration and vice versa.

To generalize the assessment across all bugs, we calculated the sum of these differences for each class for each bug. Figure 7 split the difference between each configuration across all bugs, split by search budget. Table 4 includes descriptive statistics on the difference in the average coverage of each obligation.

Figure 7 shows the vast majority of the summed differences is close to zero, with a median of 0.00 for all comparisons. This means that—in most cases—there are few major differences in

the obligations covered by each configuration. There are differences in the 25th and 75th percentiles between configurations, but relatively narrow ones.

Table 5 offers a complementary analysis, where we indicate the proportion of individual coverage obligations where one configuration outperformed the other in terms of the average coverage of each obligation (a ‘win’) or where their performance was identical (a ‘tie’). The results show that for the vast majority of obligations (88%–91%), the single- and multi-objective strategies perform identically, resulting in a tie. For the small fraction of branches where performance differs, there is no consistent winner—both Branch alone and multi-objective configurations win a similarly small percentage of the time. This reinforces that adding a secondary objective rarely changes which specific branches get covered.

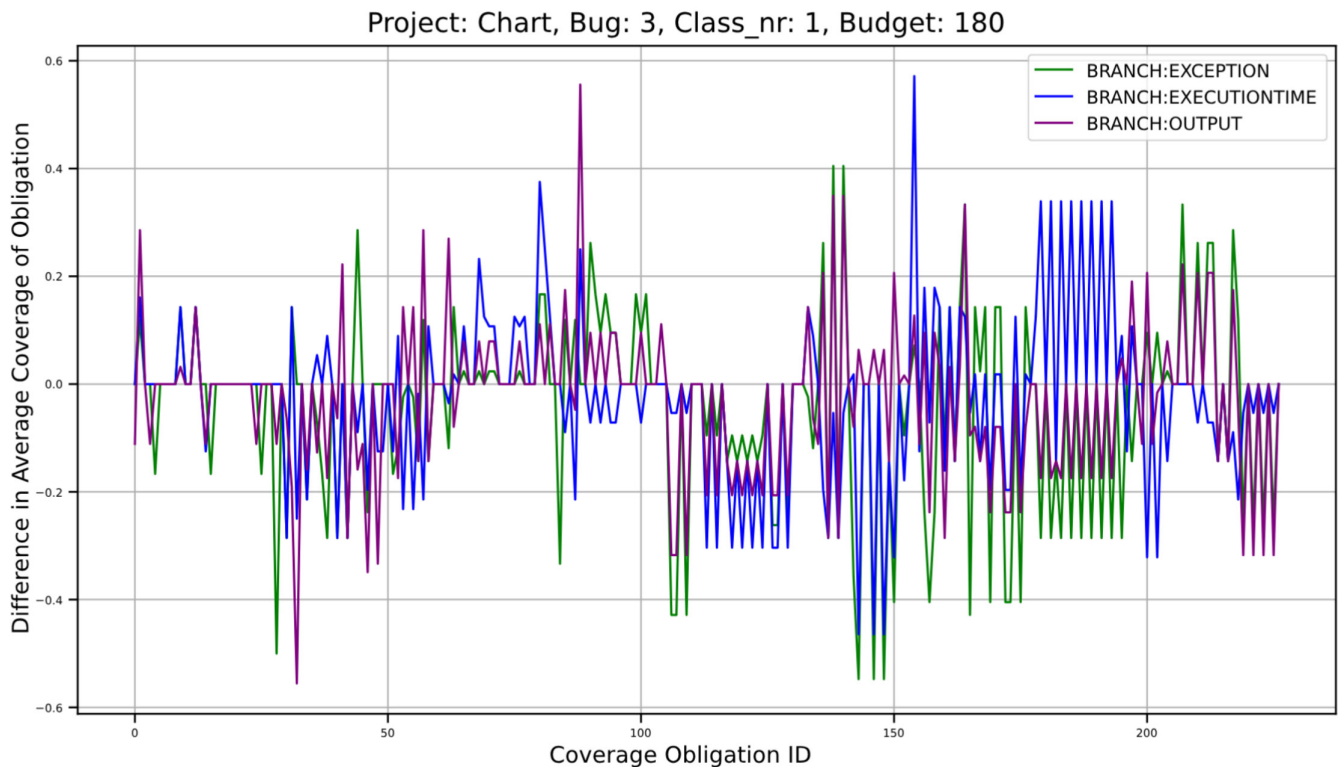


FIGURE 7 | Boxplots of the differences in covered obligations between targeting Branch Coverage alone versus Branch and goal-based fitness function.

TABLE 4 | Descriptive statistics of the difference in the average coverage of each obligation between branch and another configuration, split by configuration and budget.

Criterion	Budget	Mean	Std	Min	25%	50%	75%	Max
Branch and Exception	180	0.0033	0.0373	−0.0625	−0.0010	0.0000	0.0038	0.3194
Branch and Execution Time	180	−0.0015	0.0286	−0.1427	−0.0048	0.0000	0.0039	0.1226
Branch and Output	180	−0.0008	0.0527	−0.3875	−0.0043	0.0000	0.0058	0.2097
Branch and Exception	300	−0.0003	0.0232	−0.1111	−0.0051	0.0000	0.0048	0.0482
Branch and Execution Time	300	−0.0030	0.0311	−0.1936	−0.0063	0.0000	0.0073	0.1014
Branch and Output	300	−0.0011	0.0310	−0.1660	−0.0076	0.0000	0.0073	0.0871

TABLE 5 | Proportion of coverage obligations where one configuration outperforms another.

Comparison	Budget	Branch wins (%)	Branch + X wins (%)	Ties (%)
Branch vs. Branch and Exception	180	3.7%	5.1%	91.2%
Branch vs. Branch and Exec. Time	180	5.6%	6.4%	88.0%
Branch vs. Branch and Output	180	6.0%	5.7%	88.3%
Branch vs. Branch and Exception	300	4.6%	4.6%	90.8%
Branch vs. Branch and Exec. Time	300	5.5%	4.9%	89.7%
Branch vs. Branch and Output	300	5.2%	5.3%	89.5%

Research Question 1.2 (obligations covered): In the majority of cases, the addition of a goal-based fitness function does not change the likelihood of covering particular test obligations. Almost all differences in the average coverage of individual obligations are within 10% of when Branch Coverage is targeted alone.

Figure 7 shows that there are a number of outliers. Most outliers are clustered within -0.1 to 0.1 , i.e., within 10% difference in average coverage. However, it is possible that some of these outliers offer information that could improve the results of test generation. In particular, negative outliers are interesting, as they suggest cases where the addition of a goal-based fitness function improved Branch Coverage.

We inspected the negative outliers, with a particular focus on the six most extreme cases—Math-81 at the 180-s budget for both Branch and Execution Time and Branch and Output Coverage, JSoup-9 for Branch and Exception Count and Branch and Execution Time at the 300-s budget, Math-11 for Branch and Execution Time at the 300-s budget and Closure-164 for Branch and Output Coverage at the 300-s budget. Overall, there were few clear and actionable conclusions that we could draw from these outliers. However, we share some interesting observations.

The class modified in Math-81 for Branch and Output Coverage at the 180-s budget was the most extreme outlier, with an average difference of 38.75% in coverage. The explanation for this difference is relatively straightforward. Many of the methods

of the CUT have numeric return values. It is likely that Output Coverage, by placing emphasis on returning diverse results for these methods, helped to steer test generation towards covering Branch Coverage obligations in these methods and in methods indirectly called through these methods. This improvement disappears at the 300-s search budget, suggesting that Branch Coverage alone is eventually effective. However, the addition of Output Coverage speeds coverage attainment.

A similar observation can be made for Branch and Output Coverage for the CUT in Closure-164, where there was an average coverage difference of 16.60%. Most methods in this class return Boolean values. It is possible that the Output Coverage helped to encourage Branch Coverage by ensuring that the methods returned both possible values. Here, this difference increased with the search budget.

A potential explanation for the outliers for Branch and Execution Time is that the Execution Time fitness function encourages the generation of longer test cases, with more programme interactions. This function penalizes test cases that are too long, but the average test case length is still higher than when Branch Coverage is targeted alone. This may encourage improvement in coverage as well, in a small number of cases.

5.1.3 | Evolution of Coverage Attainment

We collected the evolution of coverage during the test generation process, based on the Branch Coverage achieved by the best test suite in the evolving population. A snapshot of coverage is captured each second during the generation process. This allows us to calculate the AUC (area under the curve) and the time the search took to achieve 25%, 50% and 75% coverage during each trial for each configuration.

Figure 8 and Table 6 show the statistics for AUC for each configuration and search budget. Higher AUC values indicate that Branch Coverage evolved early, while lower ones mean the search took more time to achieve coverage.

We observe that the median AUC is lower for both search budgets for Branch and Execution Time as well as for Branch and Output Coverage than for Branch alone, potentially indicating slightly slower coverage attainment. However, the 25th and 75th percentiles are similar. We also observe that the median AUC is slightly higher for Branch and Exception Coverage than for Branch alone at the 180-s budget, potentially indicating a slight improvement in the rate of coverage attainment.

In Table 7, we show the results of statistical testing on the AUC. All p -values are considerably above the 0.05 threshold. Figure 9 and Table 8 report the time needed to reach 25%, 50% and 75% Branch Coverage. The median time to reach each landmark is very similar across all configurations, regardless of search budget. The largest differences between configurations can be seen in the 75th percentile for each configuration. Here, we often see a higher 75th percentile for Branch and Output, as compared with the higher configurations, indicating again that coverage attainment may be slightly slowed with the inclusion of Output Coverage as a goal. However, there is no evidence that this effect is significant.

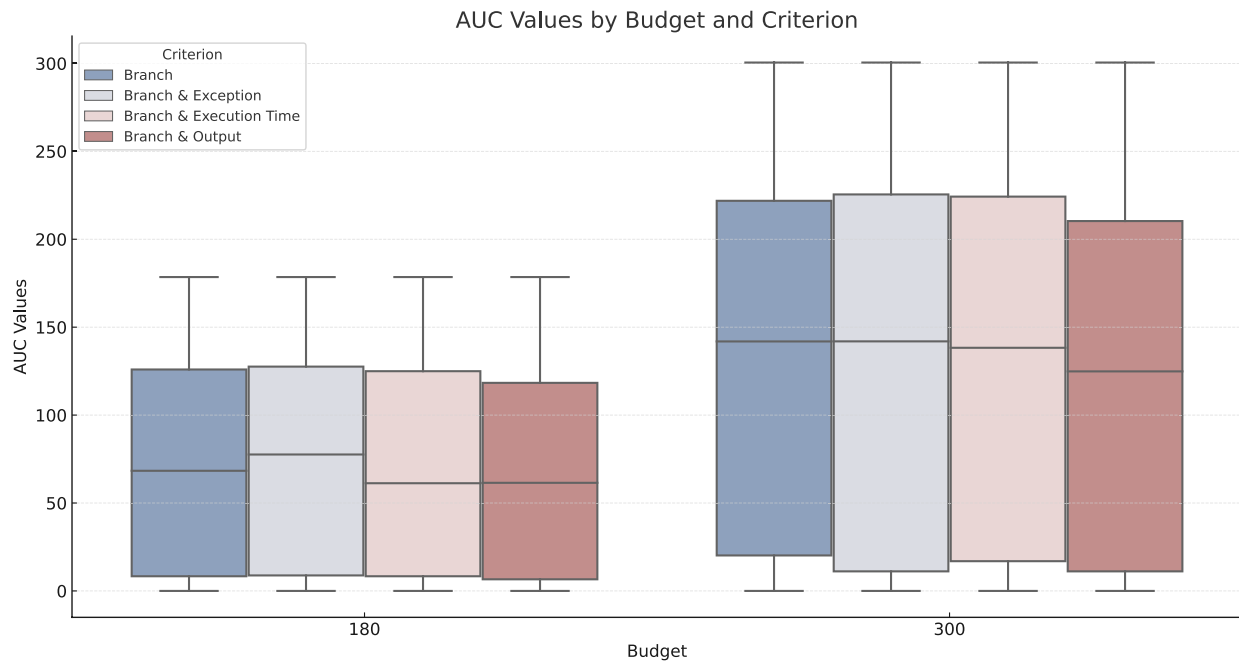


FIGURE 8 | Boxplot for aArea uUnder the cCurve (AUC) of the branch coverage evolution, split by budget and cCriterion.

TABLE 6 | Average and median values for AUC of coverage evolution for different configurations, split by budget.

Criterion	Budget	Average	Median
Branch	180	74.06	68.32
Branch and Exception	180	75.44	77.59
Branch and Execution Time	180	70.73	61.25
Branch and Output	180	70.30	61.48
Branch	300	134.19	141.89
Branch and Exception	300	135.65	141.93
Branch and Execution Time	300	130.55	138.27
Branch and Output	300	127.60	124.85

While the median remains relatively consistent across search budgets, the 75% percentile is often higher at the 300-s budget, especially at the 75% coverage threshold. The average also raises across search budgets. This is because a small number of additional test suites reach these thresholds later in the generation process under the higher budget. Again, the median is relatively consistent across all budgets and configurations.

Research Question 1.3 (coverage evolution): There are almost no significant differences between configurations with regard to the rate of attainment of Branch Coverage. Branch and Exception Coverage at a 180-s budget show statistically significant improvement, but with only a negligible effect size.

5.2 | Impact on Goal-Based Objectives (Research Question 2)

Our second hypothesis was the following:

Hypothesis 2. *Targeting both coverage and a goal-based fitness function will not have an impact on the attainment of goal-based fitness functions, as compared with targeting coverage or a goal-based fitness function alone.*

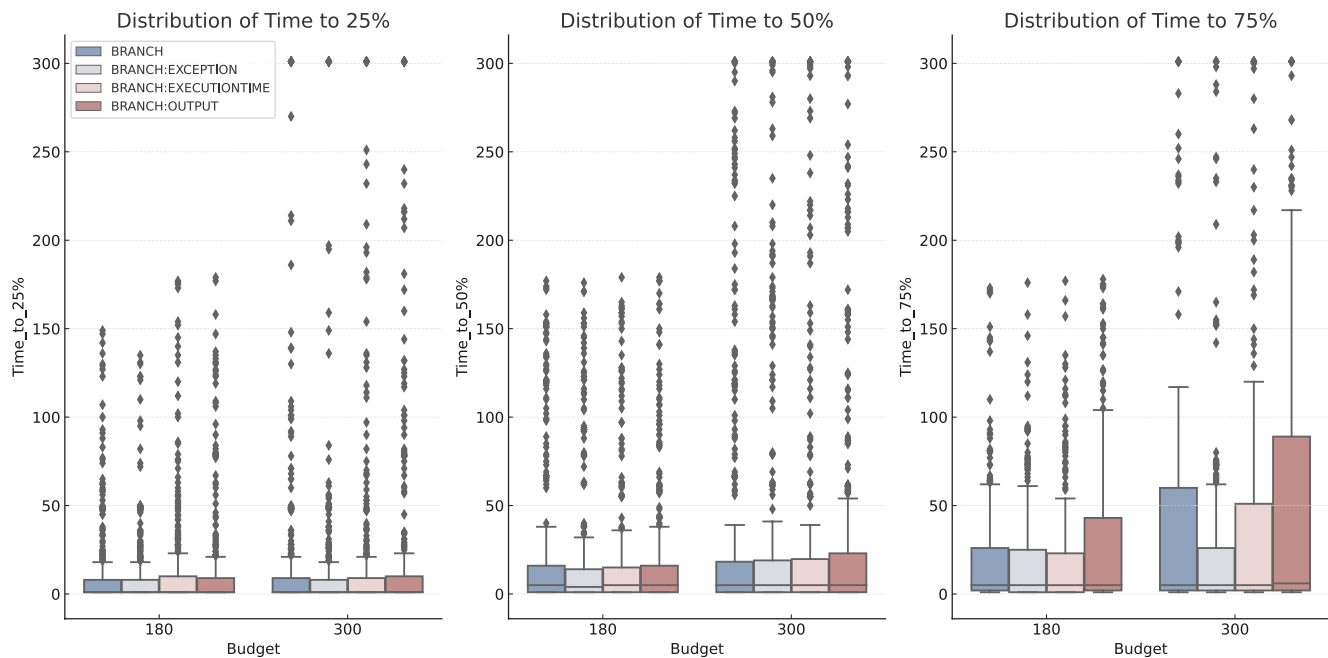
Similar to Research Question 1, targeting multiple objectives could affect the final fitness values of the goal-based fitness functions—e.g., raising or lowering goal attainment when compared with targeting a goal-based or a structure-based fitness function alone.

In addition to examining this hypothesis directly, there is a secondary hypothesis of interest. One of the reasons for the prevalence of structural coverage in search-based test generation is that structural coverage can be translated effectively into distance-based fitness functions, e.g., the branch distance used for optimizing Branch Coverage [7]. This means that tests can be efficiently generated that widely explore the codebase. Goal-based fitness functions often lack distance-based fitness functions [47]. Consequently, they may offer less feedback to the optimization process. As a result, targeting both coverage and goal-based objectives could potentially result in higher attainment of goal-based fitness by offering an additional feedback mechanism [11, 47]. Past research has not assessed this hypothesis. In this experiment, the Exception Count is one such example. In contrast, Output Coverage does have a distance-based fitness function, so such benefits may not be observed in this case.

During the experiment, we recorded the final attainment of each goal-based fitness function for all generated test suites. Note that

TABLE 7 | p -values and effect size on pairwise comparisons of AUC by different configurations, split by search budget.

Comparison	Budget	p -value	Significant	Cliff's δ	Effect size
Branch vs. Branch and Execution Time	180	0.909	No	-0.005	Negligible
Branch vs. Branch and Exception	180	0.362	No	-0.007	Negligible
Branch vs. Branch and Output	180	0.647	No	0.005	Negligible
Branch vs. Branch and Exception	300	0.985	No	-0.017	Negligible
Branch vs. Branch and Execution Time	300	0.685	No	-0.001	Negligible
Branch vs. Branch and Output	300	0.347	No	0.023	Negligible

**FIGURE 9** | Boxplots of the time taken to achieve 25%, 50%, and 75% Branch Coverage for each configuration, split by budget.

the Execution Time fitness function cannot be executed without also targeting Branch Coverage, so we were unable to generate test suites targeting Execution Time alone. In addition, note that the Exception Count is normalized between 0 and 1 for all bugs, based on the largest number of exceptions seen in any trial for that bug, as the number of possible exceptions differs between bugs. Figure 10 shows boxplots for Exception Count, Output Coverage and Execution Time for each fitness function configuration and search budget. Average and median values are reported in Table 9. Tables 10–12 report p -values and effect sizes for comparisons between single-objective generation versus multi-objective optimization.

First, we observe that no goal-based fitness function can serve as a proxy for another goal-based fitness function. Targeting Output Coverage yields a low Exception Count and Execution Time. Similarly, targeting Exception Coverage yields low Output Coverage and Execution Time. If one targets a goal-based fitness function *alone*, they should not expect high attainment of goals other than the one that function was designed for.

Targeting Branch Coverage *alone* yields better performance at each goal than targeting a fitness function designed for a

different goal, suggesting that coverage of the code base will always lead to *some* degree of goal attainment. However, as shown in Tables 10 and 11, these suites are also significantly worse at attaining Output Coverage or Exception Count than targeting either goal directly or targeting multiple objectives—with a large effect size for Exception Count at both search budgets, a small-medium effect size for Output Coverage at the 180-s search budget, and a medium-large effect size for Output Coverage at the 300-s budget. In other words, code coverage is a weak proxy for a goal-based fitness function—as noted in Section 1, coverage alone is not enough to ensure goal attainment.

Research Question 2 (goal coverage): Targeting code coverage *alone* leads to worse goal attainment than directly targeting a goal-based objective.

Next, we compare targeting a goal alone versus multi-objective optimization. Table 11 shows that, at both budgets, there is a significant difference in Output Coverage between targeting Output Coverage alone and targeting both Branch and Output Coverage, in favour of targeting Output Coverage alone. However, the effect

TABLE 8 | Descriptive statistics on the time (in seconds) to reach coverage thresholds, split by configuration and budget. ‘Count’ indicates the number of trials that reached this threshold.

Criterion	Budget	Count	Average	Median
Time to 25%				
Branch	180	544	10.33	1.00
Branch and Exception	180	527	9.12	1.00
Branch and Execution Time	180	501	12.06	1.00
Branch and Output	180	512	12.45	1.00
Branch	300	551	19.00	1.00
Branch and Exception	300	513	17.37	1.00
Branch and Execution Time	300	479	19.76	1.00
Branch and Output	300	492	21.34	1.00
Time to 50%				
Branch	180	433	22.55	5.00
Branch and Exception	180	416	18.70	4.00
Branch and Execution Time	180	376	19.67	5.00
Branch and Output	180	391	20.74	5.00
Branch	300	448	35.64	5.00
Branch and Exception	300	417	34.98	5.00
Branch and Execution Time	300	378	33.36	5.00
Branch and Output	300	391	37.39	5.00
Time to 75%				
Branch	180	233	24.29	5.00
Branch and Exception	180	236	22.08	5.00
Branch and Execution Time	180	209	23.15	5.00
Branch and Output	180	217	33.29	5.00

(Continues)

TABLE 8 | (Continued)

Criterion	Budget	Count	Average	Median
Branch	300	244	41.32	5.00
Branch and Exception	300	231	35.94	5.00
Branch and Execution Time	300	222	43.28	5.00
Branch and Output	300	229	57.98	6.00

size is only small at the 180-s budget and negligible at the 300-s budget. Similarly, Table 10 shows that, at both budgets, there is no significant difference between targeting Branch and Exception Count and targeting Exception Count alone.

In our past research, we observed situations where both Branch Coverage and the Exception Count offer the other function missing feedback—with the Exception Count steering Branch Coverage towards input that triggers exceptions and Branch Coverage offering feedback on how to further explore the code base [11]. Such cases are rare, but may explain the higher median seen for multi-objective generation in Figure 10.

We were unable to generate suites targeting Execution Time alone due to limitations. However, from Table 9, we can see that there is a slight improvement in the average (at both budgets) and the median (at the 180-s budget) Execution Time from targeting Branch and Execution Time simultaneously. However, there are no statistically significant differences between targeting both objectives versus targeting Branch Coverage alone (Table 12). No configuration was significantly better at yielding tests with high execution times. It is possible that the examples chosen from Defects4J had few or no performance issues that could be exposed through unit testing.

Research Question 2 (goal coverage): Targeting code coverage and a goal-based objective simultaneously results in no, or only a limited (negligible–small), drop in goal-based fitness compared with targeting a goal-based objective alone.

5.3 | Impact on Fault Detection (Research Question 3)

The third hypothesis that we raised was the following:

Hypothesis 3. *Targeting both coverage and a goal-based fitness function will not have an impact on the fault detection of generated test suites, as compared with targeting coverage or a goal-based fitness function alone.*

Regardless of the impact on the code coverage or attainment of non-coverage testing goals, targeting multiple objectives could change the specific inputs applied to the CUT. As a result, there

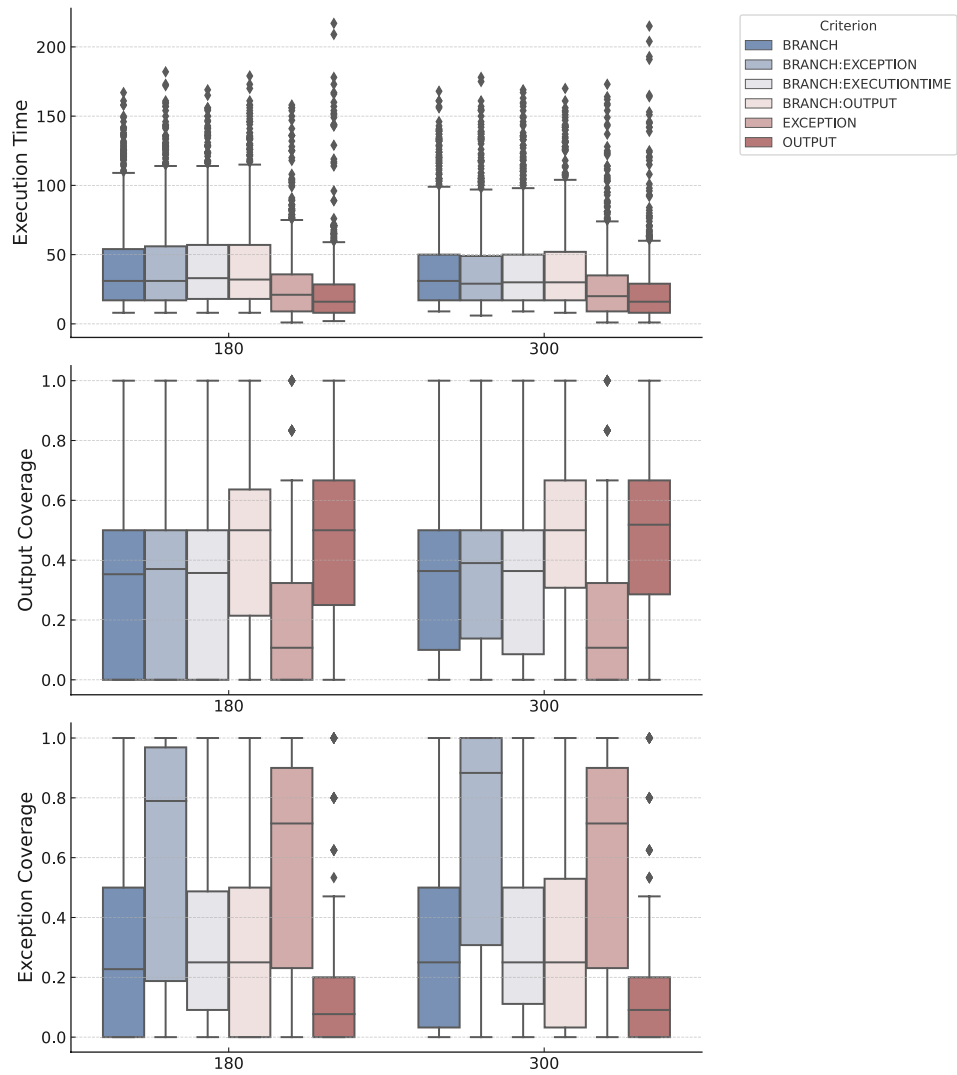


FIGURE 10 | Boxplots for Execution TTime, Exception Count, and Output Coverage divided by budget and configuration.

TABLE 9 | Averages and median values for Execution Time, Output Coverage and Exception Count, by configuration and budget.

Criterion	Budget	Execution Time		Output Coverage		Exception Count	
		Avg	Median	Avg	Median	Avg	Median
Branch	180	41.73	31.00	0.31	0.35	0.29	0.22
Branch and Exception	180	42.69	31.00	0.32	0.37	0.62	0.79
Branch and Execution Time	180	42.90	33.00	0.31	0.36	0.31	0.25
Branch and Output	180	42.89	31.00	0.43	0.49	0.30	0.25
Exception	180	27.48	21.00	0.19	0.11	0.59	0.71
Output	180	23.13	16.00	0.45	0.50	0.13	0.08
Branch	300	39.83	30.00	0.33	0.36	0.30	0.25
Branch and Exception	300	39.11	28.00	0.35	0.39	0.68	0.88
Branch and Execution Time	300	40.37	30.00	0.33	0.36	0.32	0.25
Branch and Output	300	40.38	30.00	0.46	0.50	0.32	0.25
Exception	300	26.87	20.00	0.19	0.11	0.59	0.71
Output	300	23.57	16.00	0.47	0.52	0.14	0.09

TABLE 10 | Significance tests and effect size for comparisons of **Exception Count** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Cliff's δ	Category
Branch vs. Branch and Exception	180	5.28×10^{-72}	Yes	−0.48	Large
Branch vs. Exception	180	2.68×10^{-84}	Yes	−0.56	Large
Branch and Exception vs. Exception	180	0.176	No	0.01	Negligible
Branch vs. Branch and Exception	300	3.06×10^{-80}	Yes	−0.51	Large
Branch vs. Exception	300	1.49×10^{-79}	Yes	−0.56	Large
Branch and Exception vs. Exception	300	0.152	No	0.10	Negligible

TABLE 11 | Significance tests and effect size for comparisons of **Output Coverage** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Cliff's δ	Category
Branch vs. Branch and Output	180	9.72×10^{-29}	Yes	−0.32	Small
Branch vs. Output	180	7.3×10^{-47}	Yes	−0.47	Medium
Branch and Output vs. Output	180	2.53×10^{-9}	Yes	−0.16	Small
Branch vs. Branch and Output	300	1.82×10^{-29}	Yes	−0.36	Medium
Branch vs. Output	300	1.53×10^{-44}	Yes	−0.49	Large
Branch and Output vs. Output	300	2.1×10^{-4}	Yes	−0.11	Negligible

TABLE 12 | Significance tests and effect size for comparisons of **Execution Time** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Effect size	Category
Branch vs. Branch and Execution Time	180	0.350	No	−0.01	Negligible
Branch vs. Branch and Execution Time	300	0.882	No	−0.01	Negligible

could be a change in the fault-revealing power of those test suites—it could either increase due to a change in the versatility of the test suite [10, 11, 41] or potentially decrease.

To assess this hypothesis, we consider two aspects of fault detection. First, the *likelihood of fault detection*—for each fault, the proportion of suites (for a particular configuration) that detect the fault to those generated for that configuration. Second, we consider the *number of failing tests*—how many test cases detect the fault when it is detected. We consider both so that we can examine both how likely a fault is to be detected and how much information exists to understand and debug the fault. If two configurations have the same likelihood of detection, one may offer more failing tests to use in the debugging process.

Figure 11 illustrates the likelihood of fault detection, and Table 13 offers descriptive statistics for each configuration and search budget. Table 13 lists, for each budget and fitness function configuration, the lowest, highest, average, median, 25th quartile and 75th quartile likelihood of fault detection observed for

the assessed faults, as well as the total number of faults detected over the set of 93 considered in this experiment. Immediately, we can see that the multi-objective configurations detect *more* faults—with Branch and Output detecting the most at the 180-s budget and Branch and Execution Time detecting the most at the 300-s budget. The multi-objective configurations are followed by Branch Coverage, then Output Coverage, and finally the Exception Count.

Research Question 3 (fault detection): Suites targeting multi-objective configurations detected more faults than single-objective configurations. Further, suites targeting Branch Coverage alone detected more faults than goal-based suites alone.

However, as the majority of faults are never detected, the median likelihood of fault detection is also zero for all configurations. The average, skewed by cases where faults are detected, is somewhat more informative. Targeting Branch and Output

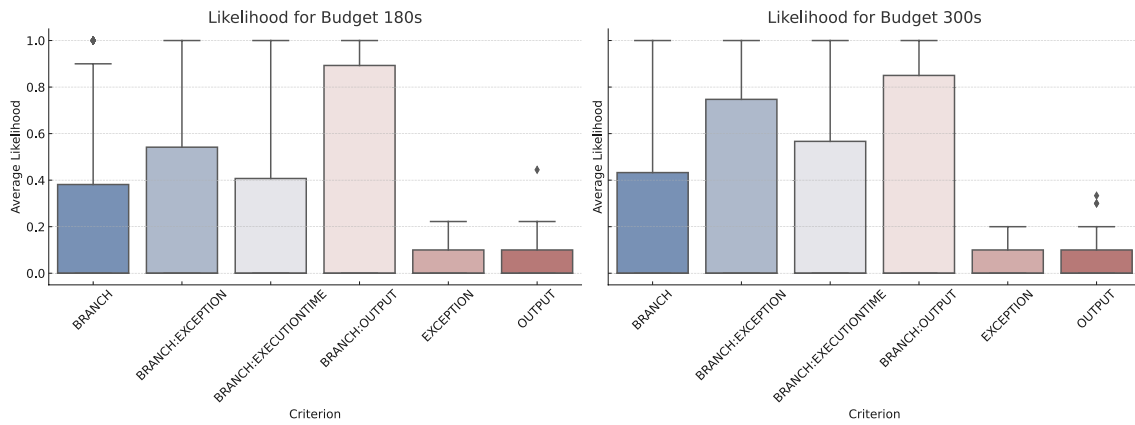


FIGURE 11 | Boxplots of the **likelihood of fault detection**, divided by budget and configuration.

TABLE 13 | Descriptive statistics on the **likelihood of fault detection**, split by budget and configuration.

Criterion	Budget	Min	25%	50%	75%	Max	Avg	# of bugs detected
Branch	180	0.00	0.00	0.00	0.38	1.00	0.24	28
Branch and Exception	180	0.00	0.00	0.00	0.54	1.00	0.27	30
Branch and Execution Time	180	0.00	0.00	0.00	0.41	1.00	0.25	34
Branch and Output	180	0.00	0.00	0.00	0.89	1.00	0.28	35
Exception	180	0.00	0.00	0.00	0.10	0.22	0.04	23
Output	180	0.00	0.00	0.00	0.10	0.44	0.09	27
Branch	300	0.00	0.00	0.00	0.43	1.00	0.26	32
Branch and Exception	300	0.00	0.00	0.00	0.75	1.00	0.30	35
Branch and Execution Time	300	0.00	0.00	0.00	0.57	1.00	0.28	39
Branch and Output	300	0.00	0.00	0.00	0.85	1.00	0.31	33
Exception	300	0.00	0.00	0.00	0.10	0.20	0.06	28
Output	300	0.00	0.00	0.00	0.10	0.33	0.13	31

Coverage yields the highest average likelihood of fault detection at both search budgets (28% and 31%), followed at both budgets by Branch and Exception Count (27% and 30%) and Branch and Execution Time (25% and 28%). This same ordering can be seen in the 75th percentile in Figure 11. Again, targeting the goal-based functions alone yields the lowest average likelihood of fault detection—with the worst performance from targeting the Exception Count alone.

We note that these results replicate the general trends observed in our previous work [10, 11]. While the exact results may differ due to changes made to EvoSuite and the stochastic nature of search-based test generation, we previously observed that targeting Branch Coverage alone yielded a higher likelihood of fault detection than targeting Exception Count or Output Coverage alone, and that multi-objective combinations also had an even higher average likelihood of fault detection.

Table 14 includes significance tests and effect sizes for the likelihood of fault detection. At both budgets, the multi-objective configurations do *not* yield significantly different results from targeting Branch Coverage alone in the likelihood of fault

detection. However, targeting Branch Coverage yields better results (with a large effect size) than targeting a goal-based fitness function at both budgets. Targeting Branch and Exception Count simultaneously also yields better results than targeting Exception Count alone, with a large effect size, at both budgets. Finally, targeting Branch Coverage and Output Coverage yields better results than targeting Output Coverage alone, with large effect sizes at both budgets.

Research Question 3 (fault detection): Suites targeting Branch Coverage alone or a multi-objective configuration outperform suites targeting a goal-based objective in the likelihood of fault detection with medium–large effect size.

Research Question 3 (fault detection): Suites targeting a multi-objective configuration fail to outperform suites targeting Branch Coverage alone with significance in the likelihood of fault detection. However, targeting a multi-objective configuration does increase the average and 75th percentile performance.

TABLE 14 | Significance tests and effect size for comparisons of **likelihood of fault detection** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Cliff's δ	Category
Branch vs. Branch and Exception	180	0.657	No	−0.035	Negligible
Branch vs. Branch and Execution Time	180	0.895	No	−0.014	Negligible
Branch vs. Branch and Output	180	0.457	No	−0.052	Negligible
Branch vs. Exception	180	6.39×10^{-5}	Yes	0.668	Large
Branch vs. Output	180	1.72×10^{-4}	Yes	0.625	Large
Branch and Exception vs. Exception	180	8.32×10^{-5}	Yes	0.792	Large
Branch and Output vs. Output	180	5.60×10^{-5}	Yes	0.810	Large
Branch vs. Branch and Exception	300	0.384	No	−0.061	Negligible
Branch vs. Branch and Execution Time	300	0.612	No	−0.046	Negligible
Branch vs. Branch and Output	300	0.461	No	−0.058	Negligible
Branch vs. Exception	300	1.41×10^{-5}	Yes	0.729	Large
Branch vs. Output	300	6.07×10^{-5}	Yes	0.669	Large
Branch and Exception vs. Exception	300	6.28×10^{-6}	Yes	0.937	Large
Branch and Output vs. Output	300	2.92×10^{-5}	Yes	0.838	Large

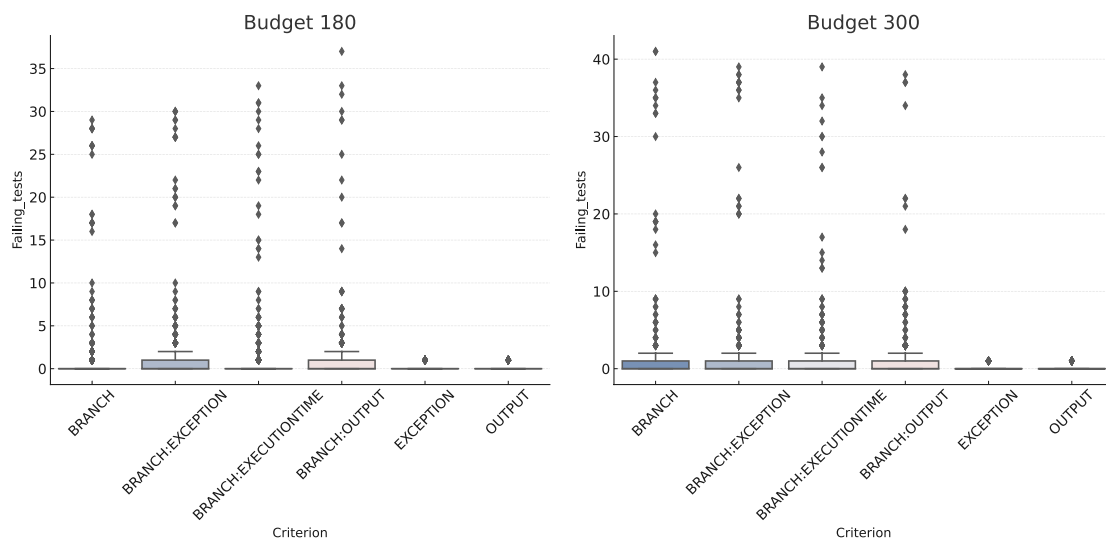


FIGURE 12 | Boxplots of the **number of failing tests**, divided by budget and configuration.

Figure 12 illustrates the number of failing tests, and Table 15 offers descriptive statistics for each configuration and search budget. Table 16 includes significance tests and effect sizes, when significance is found.

Here, we see largely similar trends to the likelihood of fault detection, with the median number of failing tests being 0 for all configurations. Targeting Branch and Exception Count yields the largest average number of failing tests at both budgets (1.13 and 1.28), followed by Branch and Execution Time (1.00) at the 180-s budget and Branch alone (1.08) at the 300-s budget. However, these results are in a relatively narrow range, and no multi-objective configuration is an outlier in terms of the number of tests that fail when a fault is detected. As shown in Table 16, targeting Branch Coverage alone or a multi-objective configuration

yields a larger number of failing tests than targeting a goal-based objective alone, with medium effect size.

Research Question 3 (fault detection): Suites targeting Branch Coverage alone or a multi-objective configuration outperform suites targeting a goal-based objective, in the number of failing tests, with medium effect size.

Research Question 3 (fault detection): Suites targeting a multi-objective configuration fail to outperform suites targeting Branch Coverage alone with significance in the number of failing tests.

TABLE 15 | Descriptive statistics on **number of failing tests**, split by budget and configuration.

Criterion	Budget	Avg	Min	25%	50%	75%	Max
Branch	180	0.89	0.00	0.00	0.00	0.00	29.00
Branch and Exception	180	1.13	0.00	0.00	0.00	1.00	30.00
Branch and Execution Time	180	1.00	0.00	0.00	0.00	0.00	33.00
Branch and Output	180	0.88	0.00	0.00	0.00	1.00	37.00
Exception	180	0.04	0.00	0.00	0.00	0.00	1.00
Output	180	0.09	0.00	0.00	0.00	0.00	1.00
Branch	300	1.08	0.00	0.00	0.00	1.00	41.00
Branch and Exception	300	1.28	0.00	0.00	0.00	1.00	39.00
Branch and Execution Time	300	1.04	0.00	0.00	0.00	1.00	39.00
Branch and Output	300	0.98	0.00	0.00	0.00	1.00	38.00
Exception	300	0.06	0.00	0.00	0.00	0.00	1.00
Output	300	0.13	0.00	0.00	0.00	0.00	1.00

TABLE 16 | Significance tests and effect size for comparisons of the **number of failing tests** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Cliff's δ	Category
Branch vs. Branch and Exception	180	0.043	No	−0.043	Negligible
Branch vs. Branch and Execution Time	180	0.518	No	−0.019	Negligible
Branch vs. Branch and Output	180	0.053	No	−0.049	Negligible
Branch vs. Exception	180	2.96×10^{-37}	Yes	0.35	Medium
Branch vs. Output	180	7.43×10^{-31}	Yes	0.35	Medium
Branch and Exception vs. Exception	180	1.90×10^{-36}	Yes	0.41	Medium
Branch and Output vs. Output	180	4.16×10^{-35}	Yes	0.36	Medium
Branch vs. Branch and Exception	300	0.014	No	−0.069	Negligible
Branch vs. Branch and Execution Time	300	0.124	No	−0.050	Negligible
Branch vs. Branch and Output	300	0.012	No	−0.059	Negligible
Branch vs. Exception	300	2.90×10^{-45}	Yes	0.33	Medium
Branch vs. Output	300	2.62×10^{-37}	Yes	0.33	Medium
Branch and Exception vs. Exception	300	2.29×10^{-38}	Yes	0.40	Medium
Branch and Output vs. Output	300	3.34×10^{-38}	Yes	0.41	Medium

5.4 | Impact on Test Suite Contents (Research Question 4)

Our fourth hypothesis was that:

Hypothesis 4. *Targeting both coverage and a goal-based fitness function will not have an impact on the size of the test suite and the average test length, as compared with targeting coverage or a goal-based fitness function alone.*

Targeting multiple objectives could increase the suite size or average test case length. Each unit test case contains one or more

interactions with the CUT. Each targeted objective imposes a set of obligations that must be covered in those interactions, and only particular input will ensure those obligations are met. Multi-objective optimization imposes a larger set of obligations than single-objective optimization.

Each test case can cover obligations of multiple criteria, meaning that one should not expect a linear increase in suite size or test case length during multi-objective optimization compared with single-objective optimization. However, some obligations require highly specific test input or setup, potentially necessitating additional specialized test cases or an increased number of

programme interactions. Therefore, some increase in suite size, test length or both could occur.

Figure 13 shows boxplots for the test suite size and average test case length, with Table 17 reporting median and average values for both. Tables 18 and 19 report the results of

significance tests and effect sizes (when significance is found) for both measurements, comparing single- and multi-objective configurations.

From Figure 13 and Table 18, we can immediately see that the distributions of test suite sizes vary significantly between

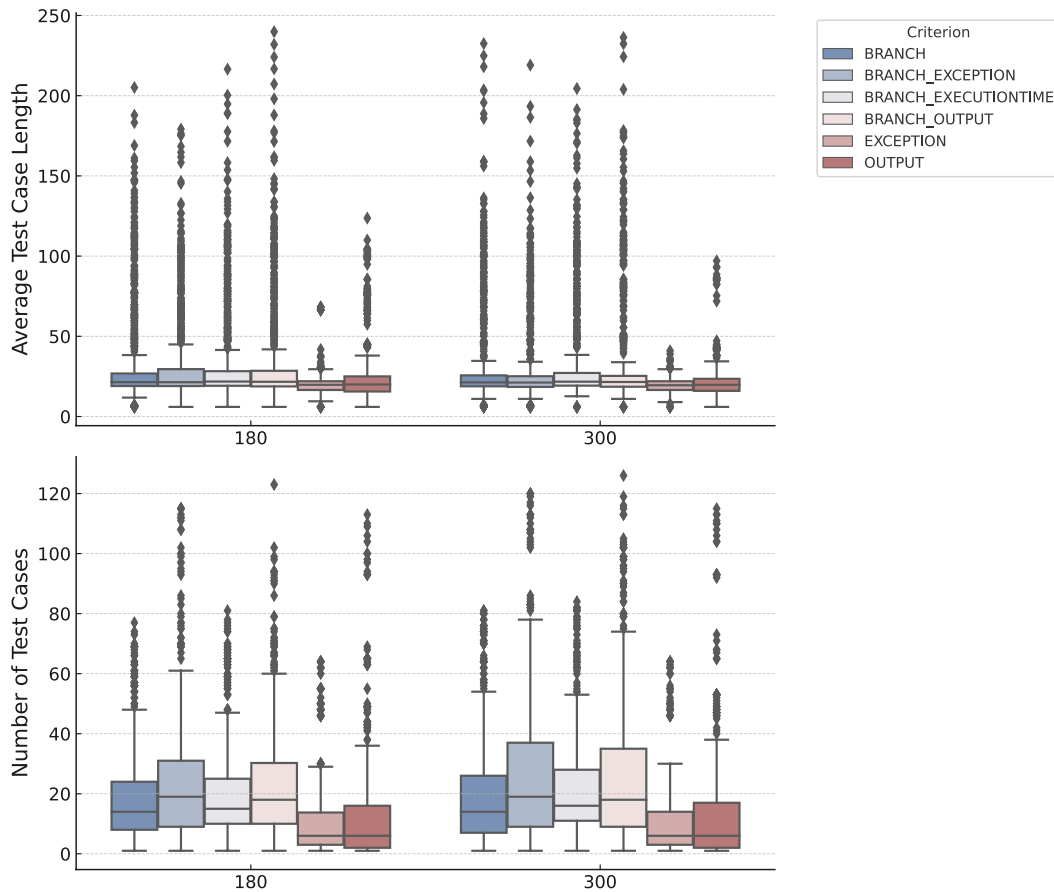


FIGURE 13 | Boxplots for the test suite size and the average test case length, divided by budget and configuration.

TABLE 17 | Average and median test suite size and average test case length, divided by budget and configuration.

Criterion	Budget	Suite size		Test case length	
		Avg	Median	Avg	Median
Branch	180	18.36	14.00	34.51	21.43
Branch and Exception	180	23.45	19.00	33.48	21.38
Branch and Execution Time	180	19.76	15.00	35.12	21.82
Branch and Output	180	22.89	18.00	35.90	21.67
Exception	180	9.92	6.00	20.35	20.00
Output	180	10.23	5.00	21.76	19.75
Branch	300	20.23	14.00	32.35	21.29
Branch and Exception	300	25.64	19.00	31.06	21.19
Branch and Execution Time	300	22.41	16.00	34.88	21.73
Branch and Output	300	25.96	18.00	34.76	21.54
Exception	300	9.98	6.00	19.85	20.00
Output	300	10.65	6.00	20.45	19.75

TABLE 18 | Significance comparisons and effect sizes for **test suite size** between single- and multi-objective configurations across different budgets. All *p*-values are significant.

Comparison	Budget	<i>p</i> -value	Cliff's δ	Effect size
Branch vs. Branch and Exception	180	1.41×10^{-7}	-0.29	Small
Branch vs. Branch and Output	180	1.95×10^{-6}	-0.26	Small
Branch vs. Branch and Execution Time	180	0.011	-0.12	Negligible
Branch and Exception vs. Exception	180	4.97×10^{-69}	0.84	Large
Branch and Output vs. Output	180	2.38×10^{-54}	0.68	Medium
Branch vs. Branch and Exception	300	4.35×10^{-7}	-0.27	Small
Branch vs. Branch and Output	300	1.80×10^{-6}	-0.27	Small
Branch vs. Branch and Execution Time	300	0.001	-0.13	Negligible
Branch and Exception vs. Exception	300	9.21×10^{-67}	0.92	Large
Branch and Output vs. Output	300	1.96×10^{-51}	0.73	Medium

TABLE 19 | Significance comparisons and effect sizes for **average test case length** between single- and multi-objective configurations across different budgets.

Comparison	Budget	<i>p</i> -value	Significant	Cliff's δ	Effect size
Branch vs. Branch and Exception	180	0.891	No	-0.001	Negligible
Branch vs. Branch and Output	180	0.374	No	-0.021	Negligible
Branch vs. Branch and Execution Time	180	0.490	No	-0.009	Negligible
Branch and Exception vs. Exception	180	4.97×10^{-30}	Yes	0.63	Medium
Branch and Output vs. Output	180	2.18×10^{-24}	Yes	0.45	Small
Branch vs. Branch and Exception	300	0.691	No	-0.003	Negligible
Branch vs. Branch and Output	300	0.300	No	-0.023	Negligible
Branch vs. Branch and Execution Time	300	0.002	Yes	-0.09	Negligible
Branch and Exception vs. Exception	300	4.30×10^{-26}	Yes	0.60	Medium
Branch and Output vs. Output	300	6.61×10^{-21}	Yes	0.52	Medium

configurations. Targeting code coverage and a goal-based fitness function simultaneously results in larger test suites than targeting either alone, with medium-large effect size compared with targeting a goal-based objective and negligible small effect size compared with targeting Branch Coverage alone.

Figure 13 and Table 19 also show that the average test case length tends to increase with multi-objective optimization compared with targeting a goal-based objective alone, with a small-medium effect size. However, the test length does not increase compared with targeting Branch Coverage alone—with only a negligible increase when targeting Branch and Execution Time.

Branch Coverage tends to have more obligations to cover than the Exception Count or Output Coverage, as a programme will generally have more branches in control flow than output

partitions or thrown exceptions. Covering the obligations of Branch Coverage requires more interactions with the CUT and requires that a larger number of specialized scenarios be set up and executed compared with a goal-based objective alone, increasing both suite size and test length.

There is a larger increase between Branch and Exception Count and Exception Count alone in both suite size and test length than between Branch and Output Coverage and Output Coverage alone. This is because the Exception Count depends on the number of exceptions discovered, which—in almost all cases—will be fewer than the number of required output partitions for the methods of the CUT. Further, tests that trigger exceptions may not achieve high coverage, as the execution path will end when the exception is triggered. If an exception is triggered early in the execution of a particular method, few coverage obligations will be achieved.

Research Question 4 (test suite contents): Both the test suite size and average test length increase with multi-objective optimization compared with when a goal-based criterion is targeted alone. Branch Coverage tends to impose more obligations than goal-based objectives, leading to the increase.

There is only a small increase in test suite size—and no increase in average test length—between Exception Count and Branch Coverage and Output Coverage and Branch Coverage versus Branch Coverage alone. As discussed above, the number of obligations for the goal-based objectives is small compared with the number for Branch Coverage, so only a small increase in suite size would be expected.

Research Question 4 (test suite contents): Targeting Exception or Output Coverage in addition to Branch Coverage leads to a small increase in test suite size compared with targeting Branch Coverage alone. However, there is no significant increase in test case length.

We see no or negligible increase in suite size and test length between Branch Coverage and Branch and Execution Time. The Execution Time fitness function differed from the others in that it had no ‘obligations’. Rather, the goal was simply to find the maximum execution time for a test suite during the generation process. Therefore, one would not expect a significant impact on the test suite size. Some impact on test case length would be reasonable, however, as increasing the number of interactions will increase the execution time. That said, the fitness function imposed a high penalty on test case length to prevent the generation of bloated test cases. Further, as shown in Table 12, actual attainment of the Execution Time goal was limited.

Research Question 4 (test suite contents): Targeting Execution Time in addition to Branch Coverage leads to no or negligible change in suite size or test case length compared with targeting Branch Coverage alone.

5.5 | Impact of Search Budget (Research Question 5)

Our final hypothesis was the following:

Hypothesis 5. *An increase in the search budget will not lead to increased attainment of each objective.*

An increased search budget could potentially increase the resulting attainment of each targeted objective. However, we must examine whether this is the case. In any case, we also hypothesize that the relative relationships between single- and multi-objective optimization will not fundamentally differ.

We observe that attainment generally did increase. We also observe that the observed trends generally held true. With regard to attained Branch Coverage, the increased search budget led to

higher coverage attainment (Table 2) and more suites reaching particular coverage thresholds (Table 8). However, this increase is approximately consistent across configurations, regardless of the targeted fitness functions. The same trends between configurations generally held at both search budgets with regard to total attained coverage, the particular obligations covered, and the rate of coverage attainment.

With regard to coverage of goal-based objectives, an increased search budget led to slightly higher median attainment of Exception Count and Output Coverage (Table 9). However, again, the same general trends were witnessed in comparisons of multi-objective and single-objective generation at both budgets for the most part. Two exceptions emerged (Tables 10 and 11). First, at a higher budget, targeting Branch and Exception yielded significantly better Exception Count than targeting Branch alone (when there was no significant difference at the lower budget). Second, a negligible difference at the lower search budget between targeting Branch and Output and targeting Output Coverage alone in terms of the achieved Output Coverage disappeared at a higher budget.

With regard to fault detection, the number of faults detected increased with the search budget. In addition, we see that the average and 75th percentile likelihood of fault detection also increased with the search budget (Table 13)—as well as the average number of failing tests (Table 15). The general relationships between single- and multi-objective configurations held, except that some effect sizes increased at the higher budget (Table 14).

Finally, an increased search budget generally led to little-to-no change in the median test suite size or test case length—however, there was a minor increase in the average suite size (Table 17). The observations with regard to single versus multi-objective generation held across budgets, with small amplifications at the larger search budget (e.g., an increased effect size for Branch and Output versus Output alone at a 300-s budget for the average test case length).

Research Question 5 (search budget): An increased search budget leads to increased Branch Coverage, goal attainment and fault detection, but does not substantially affect test suite size and average test length.

Research Question 5 (search budget): An increased search budget generally does not fundamentally change—but may increase the effect size of—relationships between single- and multi-objective optimization.

6 | Discussion

6.1 | Assessment of Hypotheses

Our study assessed five hypotheses about the relationships between coverage-directed test generation, goal-directed test generation, and multi-objective optimization targeting both

coverage and testing goals. Here, we summarize our findings with regard to these hypotheses.

Hypothesis 1. *The inclusion of goal-based fitness functions as additional generation targets will not have an impact on the attainment of code coverage, as compared with targeting coverage alone.*

Ultimately, our observations **fail to refute this hypothesis**. We found that adding a second goal-based fitness function does not have a significant impact on the final Branch Coverage attained by test suites. Further, in the majority of cases, the addition of a goal-based fitness function does not change the likelihood of covering particular test obligations. Almost all differences in the average coverage of individual obligations are within 10% of when Branch Coverage is targeted alone. Finally, there are almost no significant differences between configurations with regard to the rate of attainment of Branch Coverage.

Hypothesis 2. *Targeting both coverage and a goal-based fitness function will not have an impact on the attainment of goal-based fitness functions, as compared with targeting coverage or a goal-based fitness function alone.*

Our observations **partially refute this hypothesis**. First, we observed that targeting code coverage *alone* leads to worse goal attainment than directly targeting a goal-based objective, adding evidence to our previous observations [11] that coverage is a prerequisite for goal attainment but does not guarantee attainment.

We observed that targeting Branch Coverage and Exception Count did not yield significant differences from targeting Exception Count alone. However, the multi-objective configuration did have a higher median performance. As observed in our previous work [10, 11], the addition of Branch Coverage can offer feedback that leads to the discovery of more exceptions. However, such cases are rare. Targeting Branch and Output Coverage did yield significant differences from targeting Output Coverage alone, with worse results from the multi-objective configuration. However, the effect size is small at the 180-s budget and negligible at the 300-s budget.

In short, we cannot reject this hypothesis in all situations. When we do, the results (i.e., for Output Coverage) indicate that there is a potential loss in performance, but only a small one.

Hypothesis 3. *Targeting both coverage and a goal-based fitness function will not have an impact on the fault detection of generated test suites, as compared with targeting coverage or a goal-based fitness function alone.*

Our observations **partially refute this hypothesis**. Suites targeting multi-objective configurations detected more faults than single-objective configurations. Suites targeting Branch Coverage alone detected fewer faults than multi-objective configurations, but they did detect more faults than suites targeting goal-based objectives.

In addition, suites targeting Branch Coverage alone or a multi-objective configuration outperform suites targeting a

goal-based objective in the likelihood of fault detection with medium–large effect size. However, suites targeting a multi-objective configuration fail to outperform suites targeting Branch Coverage alone with significance in the likelihood of fault detection. That said, targeting a multi-objective configuration does increase the average and 75th percentile performance.

These findings reinforce our previous observations [10, 11] that coverage is needed to discover faults but does not guarantee the selection of the specific inputs needed to trigger a failure. Targeting coverage yields more failures than only targeting testing goals. Targeting a goal *in addition* to coverage resulted in the discovery of more faults than coverage alone by biasing the test input used in the generated suite. However, many faults still remain undetected. Future research should consider additional goal-based fitness functions and aim to discover which functions can best shape coverage towards an increased likelihood of fault detection.

Hypothesis 4. *Targeting both coverage and a goal-based fitness function will not have an impact on the size of the test suite and the average test length, as compared with targeting coverage or a goal-based fitness function alone.*

Our observations **partially refute this hypothesis**. Both the test suite size and average test length significantly increase with multi-objective optimization compared with when a goal-based criterion is targeted alone. Branch Coverage tends to impose more obligations than goal-based objectives, leading to the increase.

Targeting Exception or Output Coverage in addition to Branch Coverage also leads to a negligible–small increase in test suite size compared with targeting Branch Coverage alone. However, there is no significant increase in test case length.

Targeting Execution Time in addition to Branch Coverage leads to no or negligible changes in suite size or test length compared with targeting Branch Coverage alone, as the Execution Time fitness function does not have distinct test obligations.

Hypothesis 5. *An increase in the search budget will not lead to increased attainment of each objective.*

Our observations **partially refute this hypothesis**. An increased search budget leads to increased Branch Coverage, goal attainment, and fault detection, but does not substantially affect suite size and test length. Additionally, an increased search budget does not fundamentally change—but may amplify—relationships between single- and multi-objective optimization.

6.2 | Threats to Validity

6.2.1 | External Validity

For this study, we focused on case examples of real faults from the Defects4J dataset. The use of this dataset introduces certain threats to external validity. First, the faults used in the study

represent only 14 Java projects. This is a relatively small number of projects. Nevertheless, we believe that Defects4J offers enough case examples that our results are generalizable to, at minimum, other small to medium-sized Java projects. Further, as Defects4J is used extensively in search-based test generation research [59], the use of Defects4J examples enables comparisons of our results with other research and eases replication.

The set of specific faults used from the Defects4J dataset may also introduce selection bias, as certain types of faults or certain projects may be overrepresented or underrepresented. While we lacked experimental resources to consider all faults, we worked to ensure that we drew a proportional sample. We initially selected 206 faults, then retained 93 faults in the final experiment. This set remains large enough to offer a broad range of case examples, and we do not believe that any single project is overrepresented.

We have based our research on a single test generation framework, EvoSuite. There are many search-based methods of generating tests, and these methods may yield different results. Unfortunately, no other generation framework offers the same variety of fitness functions, particularly goal-based fitness functions. Therefore, a more thorough comparison of tools cannot be made at this time. In addition, by focusing on a single generation framework, we ensure that all test suites are compared in a controlled and fair manner.

Within EvoSuite, we also only employed one multi-objective algorithm, whole suite generation. Other algorithms may yield different results, as search objectives may be targeted through different mechanisms. We chose this algorithm to enable comparison to past research, and chose to focus on a single algorithm to perform a focused and detailed analysis of the data collected. We believe that the general trends observed would hold regardless of the algorithm, even if specific results varied. In future work, we will consider the influence of the algorithm more closely.

6.2.2 | Internal Validity

Evolutionary algorithms inherently introduce randomness, affecting result consistency. To mitigate this, we conducted multiple trials, aiming to average out randomness and stabilize outcomes. To control experiment cost, we only generated 10 test suites for each class, budget, and fitness function configuration. A larger number of repetitions may yield different results. However, given the consistency of our results, we believe that this is a sufficient number of trials to draw stable conclusions from.

6.2.3 | Conclusion Validity

Conclusion validity depends on our choice of statistical tests and the assumptions underlying those tests. Data segmentation allowed for targeted analysis of different budget and fitness function configurations, with descriptive statistics and box plots providing an initial overview that could be used to validate the results of statistical analyses. We have favoured non-parametric methods, as distribution characteristics were not known a priori, and normality cannot be assumed.

7 | Conclusion

Past research has suggested the potential benefit of blending code coverage and goal-based fitness functions. While multi-objective generation has been previously studied, how these objectives interact—and, in particular, the interaction between coverage and goal-based fitness functions—has not been studied in depth. Therefore, in this study, we assessed and explored five hypotheses about this interaction and its effects on code coverage, goal attainment, fault detection, the size of the test suite, the length of test cases and the impact of the search budget.

Ultimately, our observations suggest that there are more benefits than drawbacks in targeting multiple objectives over a single objective. Targeting multiple objectives does not reduce code coverage, and goal attainment is either not reduced or only minorly reduced. At the same time, targeting multiple objectives can lead to the detection of more faults and a higher average likelihood of fault detection. Multi-objective optimization does lead to larger test suites but imposes only a small increase over suites targeting code coverage alone, and test case length is not significantly increased.

The benefits of multi-objective optimization are often more limited than hypothesized in past research, but the improvements in fault detection are still sufficient enough to recommend multi-objective optimization over targeting coverage or testing goals alone. Our study offers insight into how coverage and goal-based objectives interact during multi-objective test generation, offering guidance to researchers and testers and a starting point for future research on multi-objective test generation.

In future work, we would like to continue to explore these—and other—hypotheses with an expanded scope and and consideration of additional experimental variables. We will target a wider variety of projects and faults and will also vary the metaheuristic algorithms used to perform multi-objective generation (e.g., contrasting whole suite generation and MOSA). In addition, we will consider combinations of more than two fitness functions. Our past research found that EvoSuite's default combination of eight fitness functions performed worse at fault detection than simply targeting Branch Coverage under the same budget, as competing objectives and the overhead of calculating fitness limited test suite evolution. However, a subset of more than two and less than eight functions may yield highly effective results. Finally, we will also explore situations where fitness functions can be given different weights during optimization—e.g., targeting both a goal-based and coverage-based function and giving heavier weight to the goal-based function.

Acknowledgements

This research was supported by Vetenskapsrådet grants 2019-05275 and 2020-05272. Computing resources were provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by Vetenskapsrådet grant agreement 2022-06725.

Data Availability Statement

The data that support the findings of this study are openly available in Exploring the Interaction of Code Coverage and Non-Coverage at <https://doi.org/10.5281/zenodo.11047567>.

References

1. M. Aniche, *Effective Software Testing: A Developer's Guide* (Simon and Schuster, 2022).
2. Y. T. Chen, R. Gopinath, A. Tadakamalla, et al., "Revisiting the Relationship Between Fault Detection, Test Adequacy Criteria, and Test Set Size," in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering* (2020), 237–249.
3. G. Gay, M. Staats, M. Whalen, and M. Heimdahl, "The Risks of Coverage-Directed Test Case Generation," *IEEE Transactions on Software Engineering* 41 (2015): 803–819, <https://doi.org/10.1109/TSE.2015.2421011>.
4. T. T. Chekam, M. Papadakis, Y. Le Traon, and M. Harman, "An Empirical Study on Mutation, Statement and Branch Coverage Fault Revelation That Avoids the Unreliable Clean Program Assumption," in *2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)*, (IEEE, 2017), 597–608.
5. A. S. Burke, E. K. Chen, T. Y. Clark, et al., "An Orchestrated Survey of Methodologies for Automated Software Test Case Generation," *Journal of Systems and Software* 86, no. 8 (2013): 1978–2001.
6. P. McMinn, "Search-Based Software Test Data Generation: A Survey," *Software Testing, Verification and Reliability* 14 (2004): 105–156.
7. A. Arcuri, "It Really Does Matter How You Normalize the Branch Distance in Search-Based Software Testing," *Software Testing, Verification and Reliability* 23, no. 2 (2013): 119–147.
8. H. Hemmati, "How Effective Are Code Coverage Criteria?," in *2015 IEEE International Conference on Software Quality, Reliability and Security* (2015), 151–156.
9. L. Inozemtseva and R. Holmes, "Coverage Is Not Strongly Correlated With Test Suite Effectiveness," in *Proceedings of the 36th International Conference on Software Engineering*, ACM, New York, NY, USA, ICSE 2014 (2014): 435–445, <https://doi.org/10.1145/2568225.2568271>.
10. G. Gay, "Generating Effective Test Suites by Combining Coverage Criteria," in *Proceedings of the Symposium on Search-Based Software Engineering*, vol. 2017 (Springer Verlag, SSBSE, 2017).
11. A. Salahirad, H. Almulla, and G. Gay, "Choosing the Fitness Function for the Job: Automated Generation of Test Suites That Detect Real Faults," *Software Testing, Verification and Reliability* 29, no. 4–5 (2019): e1701, <https://doi.org/10.1002/stvr.1701>.
12. A. Groce, M. A. Alipour, and R. Gopinath, "Coverage and Its Discontents," in *Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software*, ACM, New York, NY, USA, Onward'14 (2014), 255–268, <https://doi.org/10.1145/2661136.2661157>.
13. Y. Meng, G. Gay, and M. Whalen, "Ensuring the Observability of Structural Test Obligations," *IEEE Transactions on Software Engineering* 46, no. 7 (2018): 748–772, <https://doi.org/10.1109/TSE.2018.2869146>.
14. D. Istanbulu, M. Zimmer, and G. Gay, "How Do Different Types of Testing Goals Affect Test Case Design?," in *IFIP International Conference on Testing Software and Systems* (Springer, 2023), 97–114.
15. R. Feldt, S. Poulding, D. Clark, and S. Yoo, "Test Set Diameter: Quantifying the Diversity of Sets of Test Cases," in *2016 IEEE International Conference on Software Testing, Verification and Validation (ICST)* (IEEE, 2016), 223–233.
16. R. Feldt, R. Torkar, T. Gorschek, and W. Afzal, "Searching for Cognitively Diverse Tests: Towards Universal Test Diversity Metrics," in *2008 IEEE International Conference on Software Testing Verification and Validation Workshop*, (IEEE, 2008), 178–186.
17. F. Palomba, A. Panichella, A. Zaidman, R. Oliveto, and A. De Lucia, "Automatic Test Case Generation: What If Test Code Quality Matters?," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*, (2016), 130–141.
18. S. Shamshiri, R. Just, J. M. Rojas, G. Fraser, P. McMinn, and A. Arcuri, "Do Automatically Generated Unit Tests Find Real Faults? An Empirical Study of Effectiveness and Challenges," in *Proceedings of the 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, ACM, New York, NY, USA, ASE 2015 (2015).
19. M. H. Moghadam, M. Saadatmand, M. Borg, M. Bohlin, and B. Lisper, "An Autonomous Performance Testing Framework Using Self-Adaptive Fuzzy Reinforcement Learning," *Software Quality Journal* 30, no. 1 (2022): 127–159.
20. N. Alshahwan and M. Harman, "Coverage and Fault Detection of the Output-Uniqueness Test Selection Criteria," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis* (ACM, ISSTA, 2014), 181–192, <https://doi.org/10.1145/2610384.2610413>.
21. M. Pezze and M. Young, *Software Test and Analysis: Process, Principles, and Techniques* (John Wiley and Sons, 2006).
22. E. Barr, M. Harman, P. McMinn, M. Shahbaz, and S. Yoo, "The Oracle Problem in Software Testing: A Survey," *IEEE Transactions on Software Engineering* 41, no. 5 (2015): 507–525, <https://doi.org/10.1109/TSE.2014.2372785>.
23. M. Staats, M. W. Whalen, A. Rajan, and M. P. Heimdahl, "Coverage Metrics for Requirements-Based Testing: Evaluation of Effectiveness," in *Proceedings of the Second NASA Formal Methods Symposium* (NASA, 2010).
24. O. I. de Normalización, *ISO 26262: Road Vehicles: Functional Safety* (ISO, 2011), <https://books.google.se/books?id=3gcAjwEACAAJ>.
25. Y. Moy, E. Ledniet, H. Delseny, V. Wiels, and B. Monate, "Testing or Formal Verification: DO-178C Alternatives and Industrial Experience," *IEEE Software* 30, no. 3 (2013): 50–57, <https://doi.org/10.1109/MS.2013.43>.
26. A. Parsai and S. Demeyer, "Comparing Mutation Coverage Against Branch Coverage in an Industrial Setting," *International Journal on Software Tools for Technology Transfer* 22, no. 4 (2020): 365–388.
27. J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis With Applications to Biology, Control, and Artificial Intelligence* (MIT Press, 1992).
28. S. Ali, L. C. Briand, H. Hemmati, and R. K. Panesar-Walawege, "A Systematic Review of the Application and Empirical Investigation of Search-Based Test Case Generation," *IEEE Transactions on Software Engineering* 36, no. 6 (2010): 742–762.
29. J. Malburg and G. Fraser, "Combining Search-Based and Constraint-Based Testing," in *Proceedings of the 2011 26th IEEE/ACM International Conference on Automated Software Engineering*, IEEE Computer Society, Washington, DC, USA, ASE'11 (2011), 436–439, <https://doi.org/10.1109/ASE.2011.6100092>.
30. R. Feldt and S. Poulding, "Broadening the Search in Search-Based Software Testing: It Need Not Be Evolutionary," in *2015 IEEE/ACM 8th International Workshop on Search-Based Software Testing (SBST)* (2015), 1–7.
31. G. Gay, "To Call, or Not to Call: Contrasting Direct and Indirect Branch Coverage in Test Generation," in *Proceedings of the 11th International Workshop on Search-Based Software Testing* (ACM, SBST, 2018).
32. A. Kanapala and G. Gay, "Mapping Class Dependencies for Fun and Profit," in *Proceedings of the Symposium on Search-Based Software Engineering* (Springer Verlag, SSBSE, 2018).
33. A. Ramírez, J. R. Romero, and S. Ventura, "A Survey of Many-Objective Optimisation in Search-Based Software Engineering," *Journal of Systems and Software* 149 (2019): 382–395.
34. A. Panichella, F. M. Kifetew, and P. Tonella, "Automated Test Case Generation as a Many-Objective Optimization Problem With Dynamic Selection of the Targets," *IEEE Transactions on Software Engineering* 44, no. 2 (2017): 122–158.

35. A. Arcuri, "Test Suite Generation With the Many Independent Objective (MIO) Algorithm," *Information and Software Technology* 104 (2018): 195–206, <https://doi.org/10.1016/j.infsof.2018.05.003>.
36. J. Campos, Y. Ge, N. Albulian, G. Fraser, M. Eler, and A. Arcuri, "An Empirical Evaluation of Evolutionary Algorithms for Unit Test Suite Generation," *Information and Software Technology* 104 (2018): 207–235, <https://doi.org/10.1016/j.infsof.2018.08.010>.
37. S. Lukaszczuk, F. Kroiß, and G. Fraser, "An Empirical Study of Automated Unit Test Generation for Python," *Empirical Software Engineering* 28, no. 2 (2023): 36, <https://doi.org/10.1007/s10664-022-10248-w>.
38. A. Panichella, F. M. Kifetew, and P. Tonella, "A Large Scale Empirical Comparison of State-of-the-Art Search-Based Test Case Generators," *Information and Software Technology* 104 (2018): 236–256.
39. N. Alshahwan, X. Gao, M. Harman, et al., "Deploying Search Based Software Engineering With Sapienz at Facebook," in *Search-Based Software Engineering* (Springer International Publishing, 2018), 3–45.
40. M. Kechagia, X. Devroey, A. Panichella, G. Gousios, and A. van Deursen, "Effective and Efficient API Misuse Detection via Exception Propagation and Search-Based Testing," in *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, Association for Computing Machinery, New York, NY, USA, ISSTA 2019* (2019), 192–203, <https://doi.org/10.1145/3293882.3330552>.
41. J. M. Rojas, J. Campos, M. Vivanti, G. Fraser, and A. Arcuri, "Combining Multiple Coverage Criteria in Search-Based Unit Test Generation," in *Search-Based Software Engineering*, eds. M. Barros and Y. Labiche (Springer International Publishing, Volume 9275 of Lecture Notes in Computer Science, 2015), 93–108, https://doi.org/10.1007/978-3-319-22183-0_7.
42. K. Lakhotia, M. Harman, and P. McMinn, "A Multi-Objective Approach to Search-Based Test Data Generation," in *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*, ACM, New York, NY, USA, GECCO'07 (2007), 1098–1105, <https://doi.org/10.1145/1276958.1277175>.
43. S. Yoo and M. Harman, "Using Hybrid Algorithm for Pareto Efficient Multi-Objective Test Suite Minimisation," *Journal of Systems and Software* 83, no. 4 (2010): 689–701.
44. M. Weiglhofer, G. Fraser, and F. Wotawa, "Using Coverage to Automate and Improve Test Purpose Based Testing," *Information and Software Technology* 51, no. 11 (2009): 1601–1617.
45. Z. Zhou, Y. Zhou, C. Fang, Z. Chen, and Y. Tang, "Selectively Combining Multiple Coverage Goals in Search-Based Unit Test Generation," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering* (Association for Computing Machinery, ASE'22, 2023), <https://doi.org/10.1145/3551349.3556902>.
46. P. McMinn, M. Harman, G. Fraser, and G. M. Kapfhammer, "Automated Search for Good Coverage Criteria: Moving From Code Coverage to Fault Coverage Through Search-Based Software Engineering," in *Proceedings of the 9th International Workshop on Search-Based Software Testing*, ACM, New York, NY, USA, SBST'16 (2016), 43–44, <https://doi.org/10.1145/2897010.2897013>.
47. H. Almulla and G. Gay, "Learning How to Search: Generating Effective Test Cases Through Adaptive Fitness Function Selection," *Empirical Software Engineering* 27, no. 2 (2022): 38, <https://doi.org/10.1007/s10664-021-10048-8>.
48. R. Just, D. Jalali, and M. D. Ernst, "Defects4J: A Database of Existing Faults to Enable Controlled Testing Studies for Java Programs," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis* (ACM, ISSTA, 2014), 437–440, <https://doi.org/10.1145/2610384.2628055>.
49. J. M. Rojas, M. Vivanti, A. Arcuri, and G. Fraser, "A Detailed Investigation of the Effectiveness of Whole Test Suite Generation," *Empirical Software Engineering* 22, no. 2 (2017): 852–893, <https://doi.org/10.1007/s10664-015-9424-2>.
50. S. Vogl, S. Schweikl, G. Fraser, A. Arcuri, J. Campos, and A. Panichella, "Evosuite at the SBST 2021 Tool Competition," in *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)* (IEEE, 2021), 28–29.
51. A. Panichella, F. M. Kifetew, and P. Tonella, "Automated Test Case Generation as a Many-Objective Optimisation Problem With Dynamic Selection of the Targets," *IEEE Transactions on Software Engineering* 44, no. 2 (2017): 122–158.
52. K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation* 6, no. 2 (2002): 182–197.
53. V. J. Manès, H. Han, C. Han, et al., "The Art, Science, and Engineering of Fuzzing: A Survey," *IEEE Transactions on Software Engineering* 47, no. 11 (2021): 2312–2331, <https://doi.org/10.1109/TSE.2019.2946563>.
54. J. W. Tukey, *Exploratory Data Analysis*. Addison-Wesley Series in Behavioral Science (Addison-Wesley Pub. Co., 1977).
55. R. McGill, J. W. Tukey, and W. A. Larsen, "Variations of Box Plots," *American Statistician* 32, no. 1 (1978): 12–16.
56. F. Wilcoxon, "Individual Comparisons by Ranking Methods," *Biometrics Bulletin* 1, no. 6 (1945): 80–83, <https://doi.org/10.2307/3001968>.
57. H. B. Mann and D. R. Whitney, "On a Test of Whether One of Two Random Variables Is Stochastically Larger Than the Other. The Annals of Mathematical Statistics," (1947), 50–60.
58. N. Cliff, "Dominance Statistics: Ordinal Analyses to Answer Ordinal Questions," *Psychological Bulletin* 114, no. 3 (1993): 494–509.
59. G. Gay and R. Just, "Defects4J as a Challenge Case for the Search-Based Software Engineering Community," in *Search-Based Software Engineering*, eds. A. Aleti and A. Panichella (Springer International Publishing, 2020), 255–261.