

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

JUGE

Devroey, Xavier; Gambi, Alessio; Galeotti, Juan Pablo; Just, René; Kifetew, Fitsum; Panichella, Annibale; Panichella, Sebastiano

Published in:
Software Testing, Verification and Reliability

DOI:
[10.1002/stvr.1838](https://doi.org/10.1002/stvr.1838)

Publication date:
2023

Document Version
Peer reviewed version

[Link to publication](#)

Citation for published version (HARVARD):
Devroey, X, Gambi, A, Galeotti, JP, Just, R, Kifetew, F, Panichella, A & Panichella, S 2023, 'JUGE: An infrastructure for benchmarking Java unit test generators', *Software Testing, Verification and Reliability*, vol. 33, no. 3, e1838, pp. e1838. <https://doi.org/10.1002/stvr.1838>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

JUGE: An Infrastructure for Benchmarking Java Unit Test Generators

Xavier Devroey^{16*}, Alessio Gambi²⁸, Juan Pablo Galeotti³, René Just⁴, Fitsum Kifetew⁵, Annibale Panichella⁶, and Sebastiano Panichella⁷

¹ NADI, University of Namur, rue de Bruxelles 61, 5000 Namur, Belgium

² University of Passau, Innstrasse 41, 94032 Passau, Germany

³ Universidad de Buenos Aires, Argentina

⁴ University of Washington, Seattle, USA

⁵ Fondazione Bruno Kessler, via Sommarive 18, 38123 Trento, Italy

⁶ Delft University of Technology, Postbus 5, 2600 AA Delft, The Netherlands

⁷ Zurich University of Applied Sciences, Steinberggasse 13, 8400 Winterthur, Switzerland

⁸ IMC University of Applied Sciences, Am Campus Krems, Trakt G, 3500 Krems an der Donau, Austria

SUMMARY

Researchers and practitioners have designed and implemented various automated test case generators to support effective software testing. Such generators exist for various languages (*e.g.*, Java, C#, or Python) and various platforms (*e.g.*, desktop, web, or mobile applications). The generators exhibit varying effectiveness and efficiency, depending on the testing goals they aim to satisfy (*e.g.*, unit-testing of libraries vs. system-testing of entire applications) and the underlying techniques they implement. In this context, practitioners need to be able to compare different generators to identify the most suited one for their requirements, while researchers seek to identify future research directions. This can be achieved by systematically executing large-scale evaluations of different generators. However, executing such empirical evaluations is not trivial and requires substantial effort to select appropriate benchmarks, setup the evaluation infrastructure, and collect and analyze the results. In this Software Note, we present our *JUnit Generation Benchmarking Infrastructure* (JUGE) supporting generators (search-based, random-based, symbolic execution, *etc.*) seeking to automate the production of unit tests for various purposes (validation, regression testing, fault localization, *etc.*). The primary goal is to reduce the overall benchmarking effort, ease the comparison of several generators, and enhance the knowledge transfer between academia and industry by standardizing the evaluation and comparison process. Since 2013, several editions of a unit testing tool competition, co-located with the Search-Based Software Testing Workshop, have taken place where JUGE was used and evolved. As a result, an increasing amount of tools (over ten) from academia and industry have been evaluated on JUGE, matured over the years, and allowed the identification of future research directions. Based on the experience gained from the competitions, we discuss the expected impact of JUGE in improving the knowledge transfer on tools and approaches for test generation between academia and industry. Indeed, the JUGE infrastructure demonstrated an implementation design that is flexible enough to enable the integration of additional unit test generation tools, which is practical for developers and allows researchers to experiment with new and advanced unit testing tools and approaches. Copyright © 2022 John Wiley & Sons, Ltd.

Received . . .

KEY WORDS: JUGE; unit test generation; evaluation infrastructure; benchmarking

*Correspondence to: Xavier Devroey, NADI, University of Namur, rue de Bruxelles 61, 5000 Namur, Belgium. E-mail: xavier.devroey@unamur.be

1. INTRODUCTION

Over the last decades, researchers have come up with various techniques to automate the generation of test cases. In particular, unit test generators seek to automate the production of tests for various purposes (validation, regression testing, fault localization, *etc.*) using different techniques, including random search (*e.g.*, [1, 2]), search-based software testing (*e.g.*, [3–5]), and symbolic (*e.g.*, [6, 7]) and concolic execution (*e.g.*, [8, 9]).

Juristo *et al.* [10] identified three essential features each empirical evaluation should contribute to the software testing empirical body of knowledge. First, the evaluation should be fully defined, and the data should be analyzed with appropriate techniques to interpret the results. Second, the programs used for the evaluation and the setup and variables considered should represent the reality of the practice. Third, an evaluation should be replicable and come with a replication package to confirm previous results and reach an acceptable level of confidence in the hypothesis.

Similarly, to bridge the gap with industry, automated test case generators must come with solid evidence that the approach can also be applied in practice. For instance, evidence-based software engineering [11] can help practitioners make informed decisions about the choice of a generator based on the current best evidence from research. Those pieces of evidence come from empirical evaluations identifying the strengths and weaknesses of various generators.

In the case of unit test generators, conducting an empirical evaluation is not trivial. It requires an extensive manual effort to collect benchmarks (*i.e.*, Java classes for which to generate test cases), setup the evaluation and the evaluation infrastructure, collect and analyze the produced unit tests, and compare the results against the state-of-the-art. The primary goal of our JUnit Generation Benchmarking Infrastructure (JUGE) [12] is to reduce the overall effort and ease the comparison of several generators by standardizing the evaluation process. This standardization allows researchers to meet the requirements, enabling an effective contribution to the empirical body of knowledge in software testing.

JUGE is suited for evaluating and comparing *fully automated* black, white, and gray-box unit test generators. For instance, in past editions of the tool competition, JUGE has been applied to evaluate various types of tools relying on a variety of approaches, including search-based [3, 13, 14], random-based [1, 2, 15], and symbolic execution [7, 16]. In a nutshell, the generator takes the source code or the binaries of a Java project as input and generates unit tests for a given class or set of classes. A time budget limits the generation, and the generated tests are compared *w.r.t.* their structural coverage and mutation score. The benchmarks, the tests, and the intermediate results can be saved and archived to be added to a replication package and enable future comparisons without re-executing all the generators, thanks to the standardized evaluation process implemented in JUGE.

JUGE has been initially developed in the context of the tool competition, co-located with the Search-Based Software Testing (SBST) workshop. It has been built to be *extensible* and configurable to be used with different test case generators and deployed in different environments. It relies on standardized processes, including an *adapter* mechanism to run a tool and process the generated tests. *Isolation* of the generator and test executions are handled through containerization (using Docker). Similarly, *scalability* of the evaluations relies on the parallelization of different containers, which allows relying on standard technologies (*e.g.*, Docker commands and dashboards) to handle the overall evaluation smoothly.

Since 2013, ten editions of the tool competition have taken place and used the JUGE infrastructure to evaluate and compare automated unit test generators [17–26]. Consequently, JUGE has been improved and evolved over the years to integrate the latest advances from academia to enhance the comparison and best practices from industry to achieve high automation. Several tools have entered the competition [27–54] and matured over the years by fixing bugs evidenced by the evaluations using the JUGE infrastructure, but also by confronting the various approaches to different benchmarks to discover areas for improvement and future research directions. The current implementation is openly available on GitHub[†] and on Zenodo for long-term storage [12].

[†]<https://github.com/JUnitContest/JUGE>

The remainder of this paper is structured as follows: Section 2 discusses the background and related work of empirical evaluation and comparison of test case generators. Section 3 presents the challenges and requirements for building JUGE, as well as the design and implementation choices made to address those challenges. Section 4 provides guidelines for setting up an evaluation with JUGE (examples of evaluations setups using JUGE are discussed in Appendix A). Section 5 presents the impact of JUGE so far, while Section 6 discusses the lessons learned from building JUGE and running evaluation with it, as well as the selection of suitable benchmarks (*i.e.*, classes under test), and the future work. Finally, Section 7 concludes the paper.

2. BACKGROUND AND RELATED WORK

When designing a new test case generation technique, conducting empirical evaluations is paramount to position this new technique in the current software testing body of knowledge [10,55]. When the technique gains in maturity, developers will also rely on those empirical evaluations to make informed decisions about choosing a tool relevant to their industrial context [4, 11]. For instance, Melo *et al.* [56] designed a recommender for concurrent software testing techniques based on the characteristics of the software under test and the current body of knowledge in concurrent software testing.

2.1. Empirical evaluation guidelines

Over the years, several guidelines, benchmarks, and infrastructures have been developed to design, execute, and assess test case generators. For instance, Arcuri and Briand [57] defined guidelines for using statistical tools when evaluating and comparing randomized algorithms, which is the case with many automated test case generators. In their systematic review of the empirical evaluation of search-based test case generation, Ali *et al.* [58] identify the elements that should be reported in study designs. They found that search-based software testing has been focused on structural coverage and unit testing and that empirical studies should adopt a more rigorous and standardized execution and reporting approach. In particular, studies should account for random variation in the results using appropriate statistical hypothesis testing. They should then compare the techniques with other baselines to conclude that it brings any advantage.

More recently, in a significant effort to improve the review process in software engineering, Ralph *et al.* [59] defined *Empirical Standards* listing specific attributes expected when conducting empirical evaluation following a given research methodology. The empirical evaluation of automated test cases generation is classified under the umbrella of *Optimization Studies in Software Engineering: i.e., research studies that focus on the formulation of software engineering problems as search problems and apply optimization techniques to solve such problems* [59]. Among the essential characteristics of such studies, the standards require comparing the approach to an appropriate baseline and the distribution of the dataset (*i.e.*, the benchmarks used for the evaluation, if possible, and the results).

JUGE contributes to the general effort of improving the quality and reproducibility of empirical evaluations for unit test generators by (i) standardizing the evaluation process, using appropriate data analysis techniques, and (ii) enabling easy distribution of the benchmarks (*i.e.*, classes under tests used for the evaluation) and results, including the test cases, coverage, and mutation analysis, and statistical analysis for future comparisons and reproductions. Section 4 discusses the guidelines to design, execute, and report the results of an empirical evaluation using our infrastructure.

Similar efforts have been pursued in other areas of automated testing. Recently, FUZZBENCH [60], an open fuzzer benchmarking platform-as-a-service, has received much attention after its deployment, thanks to the support of Google. Similarly to JUGE, FUZZBENCH provides a standard procedure to benchmark fuzzer, using coverage and fault coverage on representative program datasets. It has been designed to be scalable and ensure fair resource allocation and platform independence. JUGE achieves the same goals by relying on containerization via Docker. It allows running the different generators in isolation while relying on standard technologies and tools

(e.g., Docker dashboards) to manage scalability, resource allocations, and independence from the underlying platform (more details about the infrastructure will be provided in Section 3).

2.2. Comparison of test case generators

Besides structural coverage, like lines or branch coverage, empirical evaluations rely on mutation analysis to compare different test case generators [61]. Mutation analysis [62] applies mutation operators, e.g., replacing an arithmetic operator, to a program under test to produce faulty variants (i.e., mutants) and executes a test suite on those variants. If a test fails on a particular mutant, this mutant is considered as *killed*. The mutation score, i.e., the ratio of killed mutants to the total number of mutants, is used to measure the faults detection capabilities of the test suite [63].

For now, JUGE supports structural coverage and mutation analysis of the generated tests. Other kinds of automated analysis can be plugged into the infrastructure's extensible architecture. Additionally, all the generated test suites are saved using a unique identifier and can be collected for additional manual inspection.

2.3. Benchmarks for software testing

Empirical evaluations can be performed on various benchmarks (i.e., classes under test). For instance, Fraser and Arcuri built SF110 [64], a corpus of 23,886 classes from 110 open-source projects used to evaluate and compare unit test generators. Other benchmarks follow a different approach by using actual bugs extracted from Java software systems. For instance, DEFECTS4J [65] is a collection of reproducible bugs and a supporting infrastructure widely used for evaluating software testing and debugging approaches. In its latest version (v2.0.0), DEFECTS4J contains 835 bugs from 17 Java software systems [66]. Similarly, BUGSWARM [67] is a toolkit designed to mine reproducible failures and corresponding fixes to evaluate fault-detection, localization, and repair approaches.

JUGE supports the definition of customized benchmarks. For instance, previous editions of the tool competition have used classes from DEFECTS4J's projects and classes collected from open-source projects. Section 6.2 discusses guidelines to select classes under test for an empirical evaluation based on our experience in the tool competition and related work.

3. JUGE INFRASTRUCTURE

We describe hereafter the challenges and requirements for building an automated infrastructure like JUGE. Based on those requirements, we present the architecture and implementation of JUGE.

3.1. Challenges and requirements

Building an evaluation infrastructure like JUGE poses several challenges. Indeed, to be reusable and have an impact on automated unit test generation, such infrastructure has to meet several requirements. We present and discuss the most important ones in the following paragraphs.

Extensible and configurable (C1). JUGE has to be extensible to adapt to different research contexts. For instance, JUGE has been used for the unit testing competition (described in Section 5) for which the overall execution time is limited. It should also be practical for large-scale evaluations running over several weeks. For this, adding new generators to the infrastructure should be easy to compare them to the state-of-the-art. Similarly, it should be easy to configure JUGE to use a given set of classes under test (i.e., benchmarks) for an evaluation. Finally, it should be easy to extend JUGE to use different measures on the generated tests and add new ones to adapt to the different research questions and hypotheses driving the empirical evaluation.

Isolation (C2). The execution of automated test case generators may have side effects on the rest of a system [68]. For instance, a generator using all available resources by default can negatively

impact other generators running in parallel. Additionally, unexpected behavior, such as failures caused by faults in the generator, can cause damage to the system. Moreover, evaluating the generated tests requires executing them, which might also cause undesirable side effects. JUGE has to include mechanisms to isolate the execution of the generator and the generated test from the rest of the host system to avoid any undesirable side effects or damages.

Performance and scalability (C3). JUGE should be able to run on various platforms depending on the resources available. For instance, it should run on a standalone machine but also on a computationally intensive server or in a distributed setting on several servers in parallel. It should also scale to extensive empirical evaluations involving several tools and hundreds of classes under test.

Standardization (C4). Finally, evaluations relying on JUGE should be standardized to ease replication. This includes configuration of JUGE for a given evaluation, recording, analysis, and reporting of the data, and archiving to provide a companion artifact for the evaluation.

3.2. Implementation of JUGE

Given the identified requirements, we developed JUGE. JUGE is *extensible* and can be *configured* for evaluating and comparing *fully automated* black, white, and grey-box unit test generators (C1). The generator expects as input the source code or the binaries of a Java project and generates unit tests for a given class or set of classes. The generation is limited by a time budget provided as input to the generator. The overall execution is limited by a global timeout (*i.e.*, to twice the time budget) to take the pre and post-processing of the generator into account. For each benchmark (*i.e.*, class under test), JUGE runs the test case generator with the given time budget. JUGE can be configured to repeat the executions a given amount of times to balance *performance* and *scalability* (C3), depending on the available resources and the required statistical power of the results. Once the generation is completed, JUGE can measure structural coverage, perform mutation analysis of the generated tests, and compare the different generators using sound statistical analysis in a *standard* manner (C4). Moreover, the benchmarks, generated tests, and other data can be collected and stored in an artifact.

JUGE is open-source, available on GitHub[‡] and packaged as a Docker image to ensure *isolation* from the host system (C2). It contains scripts and tools supporting (i) the *generation of unit tests* for a given set of classes under test and time budget; (ii) the *coverage and mutation analysis* of the generated tests; and (iii) the *statistical analysis* and comparison of different unit test generators.

As illustrated in Figure 1, JUGE relies on an *adapter*, called `runtool`, to wrap specific calls to a unit test generator (MYTOOL in Figure 1). This adapter offers an interface to the `benchmarktool`, in charge of orchestrating the evaluation of the unit test generator. The communication between the host and the JUGE docker container (B in Figure 1) is done via a common folder (A in Figure 1), mounted in the file tree structure of the image. This folder contains the executable binaries of the unit test generator and its `runtool` adapter. The generated tests, the metrics, and the statistical analysis results are saved in a subfolder (`results/`) to be made available to the host. The classes under test and the corresponding configuration file are saved in the Docker container (`benchmarks/`). Hence, to evaluate multiple tools, one can reuse the same container and only has to mount different folders, each containing the unit test generator and its `runtool` adapter.

3.3. Unit test generation

One of the main challenges when building JUGE was to define a generic protocol for the generation of unit tests able to handle various unit test generators. We rely on a set of conventions and a generic communication protocol between the `benchmarktool` and the `runtool` adapter.

[‡]<https://github.com/JUnitContest/JUGE>

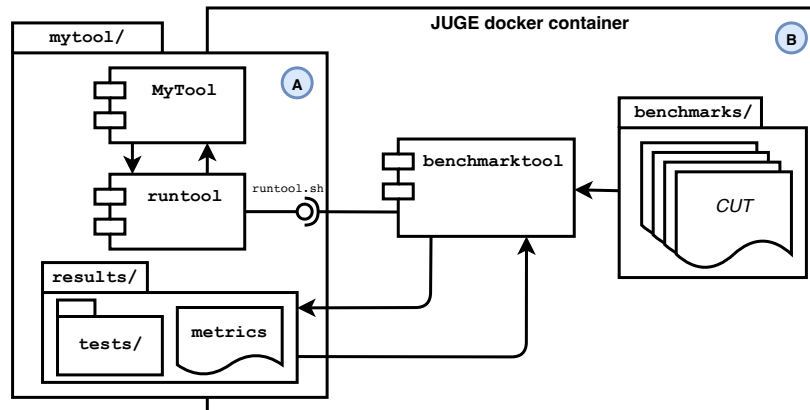


Figure 1. JUGE architecture overview

Listing 1: ITestingTool adapter interface

```

public interface ITestingTool {

    /**
     * List of additional class path entries required by a testing tool.
     * @return List of directories/jar files.
     */
    public List<File> getExtraClassPath();

    /**
     * Initialize the testing tool, with details about the code to be tested (
     * SUT).
     * Called only once.
     * @param src Directory containing source files of the SUT.
     * @param bin Directory containing class files of the SUT.
     * @param classPath List of directories/jar files (dependencies of the SUT)
     */
    public void initialize(File src, File bin, List<File> classPath);

    /**
     * Run the test tool, and let it generate test cases for a given class.
     * @param cName Name of the class for which unit tests should be
     * generated.
     * @param timeBudget How long the tool must run to test the class (in
     * milliseconds).
     */
    public void run(String cName, long timeBudget);

}

```

Conventions. By convention, the common folder (A in Figure 1) has to be named after the generator (mytool/ in our example) and mounted in the /home/ directory of the Docker container. For any generator, unit tests have to be generated in /home/mytool/temp/testcases/. Unit tests have to be stored as one or more Java test files containing JUnit tests for each class under test. Each Java test file has to declare a public class with a zero-argument public constructor, annotate test methods with @Test, and declare test methods public. Additional files may be saved to /home/mytool/temp/data/ for later offline analysis (e.g., for debugging of the generator).

Adapter. To make the liaison between the benchmarktool and the runtool adapter, the folder /home/mytool/ must contain a runtool executable script or binary

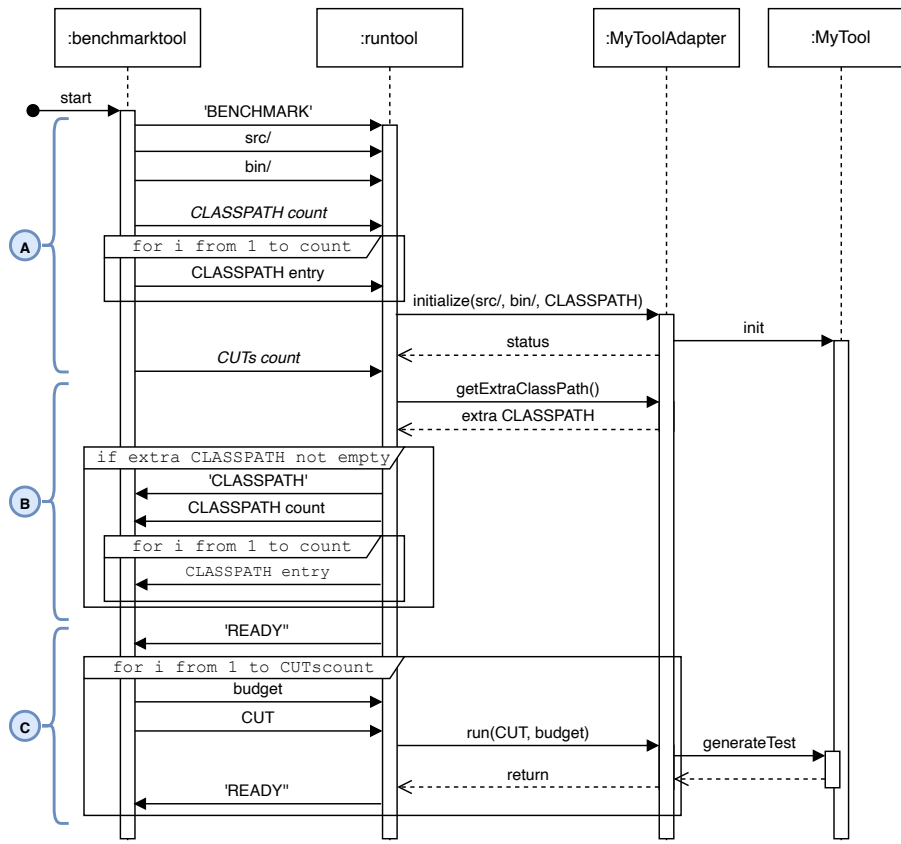


Figure 2. Communication protocol of the `runtool` adapter

(e.g., `runtool.sh` in Figure 1) that will be called by the `benchmarktool` to start the generation of unit tests. Typically, `runtool.sh` contains a single command launching the adapter. A customizable implementation of `runtool` is provided in the source code repository of our infrastructure.[§] To create an adapter for a new tool, one has to implement the different methods of the `ITestingTool.java` interface described in Listing 1.[¶] This effort should be minimal for any reasonably well-implemented tool. For instance, the implementation for `RANDOOOP` is 102 lines long, against around 150 lines of code for `EVOSUITE`-based implementations. The methods are then called, as specified by the protocol described in the following paragraph.

Communication protocol. The adapter has to support the protocol described in Figure 2. In the first part (A), the `benchmarktool` signals the start of a new evaluation by sending the 'BENCHMARK' message, followed by the paths to the source code and binaries of the software under test, the `CLASSPATH`, and the number of classes under test in the evaluation. Based on that information, the `runtool` adapter initializes the generator (in this case, `MYTOOL`).

After the initialization, the generator can signal that it will use additional `CLASSPATH` entries for its execution. The adapter notifies the `benchmarktool` of those additional entries (B in Figure 2). In the third part (C), the adapter notifies the `benchmarktool` that the generator is ready to start the evaluation by sending the 'READY' message. The `benchmarktool` then sends the time budget allocated for the generation and the class under test of the first `run` to the adapter that, in its turn, calls

[§]<https://github.com/JUnitContest/JUGE/tree/master/runtool>

[¶]<https://github.com/JUnitContest/JUGE/blob/master/runtool/src/main/java/sbst/runtool/ITestingTool.java>

the generator. After the generation, the adapter notifies the `benchmarktool` that the generator is ready for the next class under test.

3.4. Data collection

Once the test cases have been generated, JUGE can compute the different `metrics` for each test suite (*A* in [Figure 1](#)). Those metrics include: (i) the number of *flaky and non-compiling* tests, (ii) the *line and branch coverage*, and (iii) the *mutation score* of the generated tests. The JUGE infrastructure can be extended to support other kinds of metrics.

Flaky and non-compiling tests. First, if the test suite (one per class under test) does not compile, it is tagged and ignored in the subsequent analysis steps. Once compiled, the test suite is executed five times. Test methods (identified using the `@Test` annotation) producing different results between different executions are marked as flaky and ignored for the remainder of the analysis.

Line and branch coverage. JUGE relies on JACOCO [69] for statement and conditions coverage of the generated tests. Coverage information is furthermore used to reduce the subsequent mutation analysis time by restricting the execution of the tests against a given mutant to the tests effectively covering the lines modified by the mutant.

Mutation analysis. In the early versions of JUGE, we relied on PITEST [70] to generate and execute the mutants. However, it raised several issues for unit test generators relying on a dedicated test execution environment. For instance, test cases generated using EVOSUITE require executing with a dedicated runner to avoid flakiness, handle inputs and outputs, *etc.*, preventing JUGE from using the PITEST environment for test execution. To solve this issue, we refined the mutation analysis to use the default test execution environment, supporting ad-hoc test runners. We use PITEST to generate the various mutants and the results of the line coverage to reduce the analysis time by executing only tests reaching the mutated lines against each mutant. Additionally, we set a hard deadline (5 minutes by default) for the mutation analysis to avoid endless executions.

3.5. Data analysis

The generators can be compared based on quantitative analysis and the different measures collected during the analysis of the generated tests. For that, JUGE relies on Friedman's and post-hoc Conover's tests for multiple pairwise comparison [71] (available as an R script in JUGE). The former is a non-parametric test for significance, and it is widely used for multiple-problem analysis, where the problems correspond to the classes under test in our case. A significant *p*-value for this test indicates that the evaluated tools statistically differ *w.r.t.* to the overall performance score (*alternative hypothesis*). While Friedman's test does indicate whether the tools in the comparison are statistically different or not, it does not indicate for which pairs of tools such significance holds. Hence, the statistical analysis is complemented by using the post-hoc Conover's test for the pairwise comparison. Notice that the *p*-values produced by the post-hoc test are further adjusted with the Holm-Bonferroni procedure. This procedure corrects the statistical significance level (*p*-value=0.05) in case of multiple comparisons [21].

3.6. Internal architecture

The module `benchmarktool` (in [Figure 1](#)) is responsible for the orchestration and execution of the different steps of an evaluation. [Figure 3](#) presents the main classes of that module. From a configuration file, the `Main` class loads the list of the benchmarks (*i.e.*, the classes under test) to use for the evaluation in a `BenchmarkCollection`. Each `Benchmark` object contains information about the class under test, classpath, and binary and source directory to use.

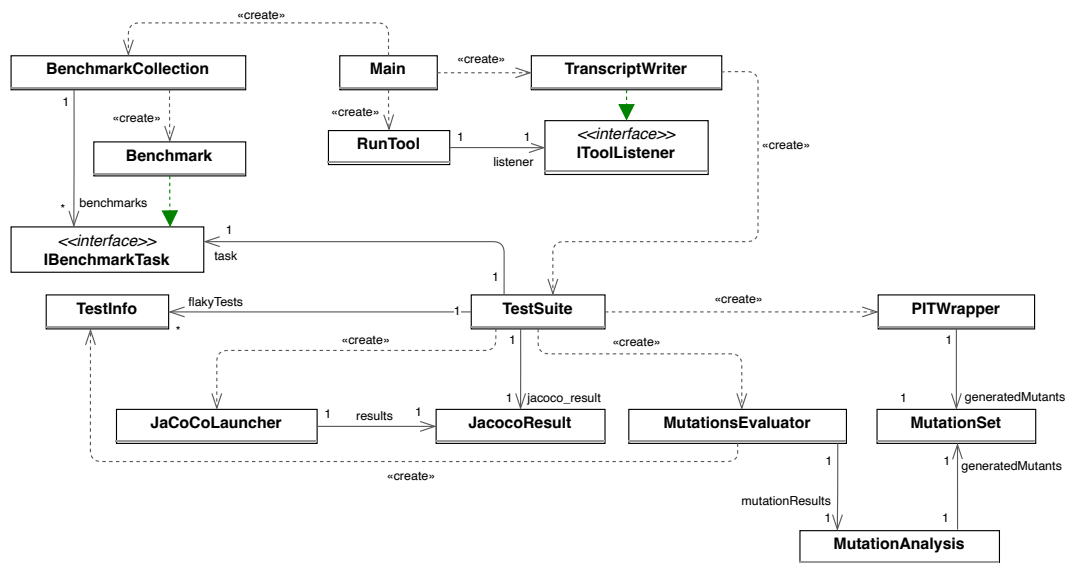


Figure 3. Internal architecture of the benchmarktool module

Once the benchmarks are loaded, the main class creates a `RunTool`, which is responsible for handling the execution and communication with the test case generator (the `sbst.benchmark.RunTool` class was renamed to `:benchmarktool` in Figure 2 to avoid confusion with the `:runtool` adapter module). This class opens an `SBSTChannel`, implemented as print and write streams on standard input and output, to communicate with the adapter (`:runtool` in Figure 2) of the generator following the protocol described in Section 3.3.

For each generated test suite, the `TranscriptWriter` class is responsible for collecting the different data (encapsulated in `TestInfo` objects) described in Section 3.4. The `JaCoCoLauncher` and `PITWrapper` encapsulate calls to (resp.) JACOCo, for code coverage information, and PITEST, for mutation analysis. As explained in 3.4, it was not possible to use PITEST to execute the full mutation analysis. The `PITWrapper` is therefore used to create the mutants, while a `MutationEvaluator` is in charge of executing the test suite on them and collecting the corresponding information.

4. SETTING UP AND RUNNING AN EVALUATION WITH JUGE

This section provides general guidelines for evaluating and comparing unit test generation tools with JUGE.

4.1. Evaluation setup

JUGE can be used to evaluate automated unit test generators that do not require human intervention during the generation process. It relies on the source code or binaries of a set of projects. JUGE is primarily designed to help carry out quantitative studies. It comes with support for structural coverage and mutation analysis and can be extended to support other measures.

Nevertheless, JUGE also allows qualitative analysis as all the tests generated during the evaluation are saved and can be inspected or reused. Additionally, as explained in Section 3.3, JUGE allows the unit test generators to save any additional data for later analysis. For instance, a search-based unit test generator can save intermediate fitness values to analyze the fitness landscape evolution.

Listing 2: Excerpt of a `benchmarks.list` configuration file

```

1 {
2   BCEL-1= {
3     src=/var/benchmarks/projects/bcel-6.0-src/src/main/java
4     bin=/var/benchmarks/projects/bcel-6.0-src/target/classes
5     classes=(org.apache.bcel.classfile.Utility)
6     classpath=(/var/benchmarks/projects/bcel-6.0-src/target/classes)
7   }
8   BCEL-2= {
9     [...]
10  }
11 }

```

Generator meta-parameters. Many unit test generators can be configured through meta-parameters (*e.g.*, mutation and crossover probabilities for search-based approaches). To ease the evaluation and processing of the results, we recommend considering each configuration as an individual generator with its adapter in a dedicated folder (*A* in Figure 1) and an explicit name reflecting the configuration. Configuring the generators with the right parameters to answer the research questions and reporting those configurations in the empirical study is of paramount importance to reduce the threats to validity and enable the replicability of the results.

JUGE meta-parameters. The infrastructure has two meta-parameters: the *time budget* and the *number of repetitions*. The time budget corresponds to the budget allocated to generate a set of test cases for a given benchmark (*i.e.*, a class under test). JUGE also uses the time budget to set a global timeout for each execution equal to twice the time budget. The time budget depends mainly on the type of approach used by the test case generator. For instance, previous research indicates that a time budget of three minutes is suited for a search-based generator like EVOSUITE [3, 13] but is not enough for symbolic execution approaches like TARDIS or SUSHI [44].

Similarly, the number of repetitions varies if the generator relies on an exact approach or uses randomness. For exact approaches, one execution is enough (unless one of the research questions considers the execution time, in which case, several repetitions are necessary). For randomized approaches (*e.g.*, search-based and random approaches), several repetitions are necessary to ensure the statistical power of the results. Arcuri and Briand [57] estimated that the number of repetitions is a compromise between the number of benchmarks used in the evaluation, the execution time of the generators, and the overall budget available to perform the evaluation. They concluded that each randomized generator should be executed 1,000 times and, if it is not possible, report the reasons and the total execution time of the entire evaluation. However, the number of repetitions (for a more significant number of benchmarks) should be at least 10.

4.2. Benchmarks

The selection of the benchmarks (*i.e.*, sets of classes under test) should follow a systematic approach and ensure that the benchmarks are diverse enough to reduce the threats to the validity of the research questions [72]. For instance, by considering projects from different application domains. Those projects (and classes under test) can come from existing benchmarks: *e.g.*, DEFECTS4J [65] or the previous editions of the tool competition relying on JUGE [20–24].

The benchmarks are described in a dedicated configuration file (`benchmarks.list`). Listing 2 provides an excerpt of `benchmarks.list` configuration file from the JUGE example benchmarks. Each benchmark has a unique identifier (line 2), the path to the root folder of the source files of the project (line 3), the path to the root folder containing the compiled classes (line 4), the list of classes under test (line 5), and the classpath with all the dependencies to use for the generation and coverage and mutation analysis (line 6). Once the benchmarks are defined, JUGE allows building a new Docker image (*B* in Figure 1) that can be instantiated multiple times in different containers to run the different tools.

4.3. Evaluation execution and results processing

Once the benchmarks and meta-parameters are defined, JUGE can start the evaluation by running different commands from the home directory of the tool in the Docker image (e.g., `/home/mytool` in the example of [Figure 1](#)). We summarize hereafter the main steps and commands to use during the evaluation.^{||} If the available hardware allows it, it is possible to run several Docker containers in parallel (instantiated from the same JUGE Docker image), each responsible for executing a different generator. One should, however, be cautious to avoid overloading the machine as it could impact the execution of the generators and provoke timeouts. Different Docker containers should be run on independent machines with the same hardware configuration. Practically, if this is not possible, we strongly recommend doing some initial tests to determine the adequate number of parallel Docker containers to avoid undesirable side effects.

Unit test generation. JUGE allows running multiple *rounds* of the tool's execution on the same benchmarks and with the same budget in one command. Each execution's results are placed in a folder named after the tool and the time budget (e.g., `/home/mytool/results_mytool_10` for a time budget of 10 seconds). For each benchmark and each round, JUGE creates a folder with the tests generated by the tool (e.g., `BCEL-1_1`, `BCEL-2_1`, etc. for the first round of executions). Those different folders also contain text files with the logs and additional data produced by the generator.

Data collection and analysis. After the generation of the unit tests, JUGE can perform a coverage and mutation analysis using JACOCo and PITest. The different results are stored in a dedicated CSV file (`transcript.csv`) for analysis. In our future work, we intend to extend JUGE to consider other types of measures, for instance, test case readability [73].

Once the different metrics have been computed for the different tools and budgets, the different results can be grouped in a single `results.csv` for a global quantitative analysis. JUGE can perform statistical analysis, as explained in Section 3.5, and produce a report with the results of the comparison.

4.4. Reporting, archiving, and reproducibility

One of the goals of the JUGE infrastructure is to enhance *repeatability* and *reproducibility* of both the results and statistical and qualitative analysis. For that, we strongly recommend submitting an *artifact* containing the following elements:

- the benchmarks, the generated tests, and additional data, if any,
- the files produced by the coverage and mutation, as well as any additional analysis,
- the results of the statistical analysis, together with any other data analysis scripts used for the evaluation.

Suppose some of the benchmarks are under a non-disclosure agreement. In that case, we strongly recommend adding benchmarks from open source systems to the evaluation and releasing those in the artifact. The design of such an artifact must be considered early in the study. We recommend, for instance, to fork the JUGE repository and update the benchmarks configuration and files to generate a Docker image used to perform the evaluation. The fork can then be easily saved in a data repository (like Zenodo,^{**} which has a GitHub integration) for long-term storage with a dedicated DOI.

In addition to the artifact, the reporting of the evaluation setup should mention the following elements:

^{||}Details on how to start the Docker container, and the different commands available in JUGE are available in the documentation at <https://github.com/JUnitContest/JUGE/blob/master/docs/>.
^{**}<https://zenodo.org>

- the randomized (or not) nature of the generators used in the evaluation;
- the meta-parameter configuration(s) of each generator;
- the meta-parameter configuration of JUGE (including the number of repetitions in the case of randomized generators) with a justification for those values;
- the total number of independent executions and the total execution time taken by the evaluation;
- the specifications of the hardware and the number of Docker containers running in parallel;
- the benchmarks selection procedure and the characteristics of the selected benchmarks relevant to the goals of the evaluation (*e.g.*, the number of lines of code of the projects and classes under test, the average McCabe's cyclomatic complexity of the benchmarks, *etc.*);
- any additional data collected and statistical analysis performed on the evaluation results with a proper justification (*e.g.*, see Arcuri and Briand [57] for a discussion on statistical analysis for randomized algorithms).

5. IMPACT OF JUGE

The JUGE infrastructure played a significant role in the *replication* of previous results regarding the structural coverage and mutation score achieved by automated unit test generators. The configurability of the infrastructure through the meta-parameters and the benchmarks considered for the various editions of the tool competition allowed us to assess the generated tests under various conditions. It independently confirmed that (i) search-based unit test generation (as implemented in EVOSUITE) achieves a better coverage and mutation score [20–22, 24–26]; and (ii) automatically generated tests can compete with manually written ones *w.r.t.* coverage and mutation score [22, 23].

5.1. Ten editions of the tool competition

The JUGE infrastructure and the tool competition also helped to push the boundaries of unit test generation by confronting industrial generators to academic ones and showcasing how research can contribute to the industrial practices [19, 25, 26]. Also, selecting various benchmarks from open source systems helped to improve the academic generators by confronting them with new classes under test, thereby increasing the generalisability of the underlying approaches. For instance, EVOSUITE has entered the competition multiple times with several algorithms (*whole suite approach* [74], MOSA [75], DYNAMOSA [13], *etc.*) and in 2019, the results of the competition lead to the fix of a major bug [43]. The results of EVOSUITE have also been recently independently confirmed using JUGE by Herlim *et al.* [76].

Table I describes the main characteristics of the different editions of the tool competition. Over the years, various tools have entered the competition and evolved. Among the different tools, RANDOOP is used as a baseline, and EVOSUITE has joined every year since the first edition. In 2015, 2021, and 2022, different industrial test case generation tools entered the competition.

The different editions have also tried different configurations *w.r.t.* to the execution of the tools and the time budget allocated for the generation. Before 2016, the time budget was left to the participants to decide (marked as - in Table I). Since 2016, the organizers have tried various time budgets to assess how the different tools react under a minimal budget: 10 seconds in 2017 and 2018 and 30 seconds in 2017 and 2021.

Similarly, the different editions have used classes under tests from various open-source projects to allow the distribution of the benchmarks after the competition. It allows one to replicate the results and the participants of the next edition to try their `runtool` adapter before submitting their tool to the competition. In 2016, the organizers used DEFECTS4J to generate regression tests and assess the tools' capability to expose real-world faults. Also, in 2019, 78 classes under test were initially

Table I. Editions of the tool competitions relying on the JUGE infrastructure with the generators, the time budgets (in seconds), the number of classes under test (#C), and the projects considered for the edition. Industrial tools are indicated by a start (*).

| Edition | Generators | Budgets (in sec.) | #C | Projects |
|-----------|---|--------------------------------|----|---|
| 2013 [17] | RANDOOP, EVOSUITE [27], T2 [28] | - | 77 | Apache Commons Lang, Apache Lucene, Barbecue, Joda Time, sqlsheet |
| 2014 [18] | RANDOOP, EVOSUITE [29], T3 [30] | - | 63 | Async Http Client, eclipse-cs, GData Java Client, Guava, Hibernate, JMLL, JWPL, Scribe, Twitter4j |
| 2015 [19] | RANDOOP, EVOSUITE (whole-suite) [31], EVOSUITE (MOSA) [34], GRT [32], JTEXPERT [33], T3 [35], undisclosed Commercial Tool (CT)* | - | 63 | Async Http Client, eclipse-cs, GData Java Client, Guava, Hibernate, JMLL, JWPL, Scribe, Twitter4j |
| 2016 [20] | RANDOOP, EVOSUITE (whole-suite) [36], JTEXPERT [37], T3 [38] | 60, 120, 240, 480 | 68 | DEFECTS4J [65] |
| 2017 [21] | RANDOOP, EVOSUITE (whole-suite) [39], JTEXPERT [40] | 10, 30, 60, 120, 240, 300, 480 | 69 | Apache Commons BCEL, Imaging, and Jxpath, Freehep, Gson, Re2J, LA4J, Okhttp |
| 2018 [22] | RANDOOP, EVOSUITE (whole-suite) [41], T3 [42] | 10, 60, 120, 240 | 59 | Dubbo, FastJason, JSoup, Okio, Redisson, Webmagic, Zxing |
| 2019 [23] | RANDOOP, EVOSUITE (DYNAMOSA) [43], SUSHI [44], TARDIS [44], T3 [45] | 10, 60, 120, 240 | 38 | Antlr4, AuthzForce, Dubbo, Fescar, FastJason, Imixs-Workflow, Okio, Spoon, Webmagic, Zxing |
| 2020 [24] | RANDOOP, EVOSUITE (DYNAMOSA) [46] | 60, 180 | 70 | Fescar/Seata, Guava, PdfBox, Spoon |
| 2021 [25] | RANDOOP, EVOSUITE [47], EVOSUITE (dynamic symbolic execution) [48], KEX [49]*, UTBOT [50]* | 30, 120 | 98 | Seata, Guava, FastJSON, Spoon, Weka, Okio |
| 2022 [26] | RANDOOP, EVOSUITE [52], BBC [51], KEX and KEX (reflection) [53]*, UTBOT and UTBOT (mocks) [54]* | 30, 120 | 65 | Seata, Guava, FastJSON, Spoon |

selected. However, due to issues with metrics computation (that have been fixed), the number of classes considered for the final ranking dropped to 38.

Running the tool competition every year is not trivial. One of the main challenges the different organizers face is the hardware infrastructure required due to the limited time between the submission of the different tools and the limit for providing the results (around two weeks). Both the generation of the tests and their evaluation using coverage and mutation analysis are time-consuming, requiring a powerful server or a cluster.

5.2. Overview of the results of the competition

As illustrated in Table I, several tools have entered the competition over the years. Figure 4 presents the averaged instruction and branch coverage, and averaged mutation score of the different tools per year, collected from the reports of the past editions [17–54]. As can be seen from the Figure, EVOSUITE has the best averaged structural coverage and mutation scores over the years. The evolution from one edition to another can be attributed to several factors, including bug correction and improvements of the underlying algorithms, selection of different time budgets and benchmarks. The full dataset providing averaged coverage, mutation score, and execution times per benchmark and time budget is available on Zenodo [77].^{††}

Among the different tools, RANDOOP, relying on feedback-directed random test generation [1], and EVOSUITE, relying on genetic algorithms [3], have entered the competition every year since

^{††}Due to incomplete data, the results of T2 from 2013 [28] are not included in Figure 4.

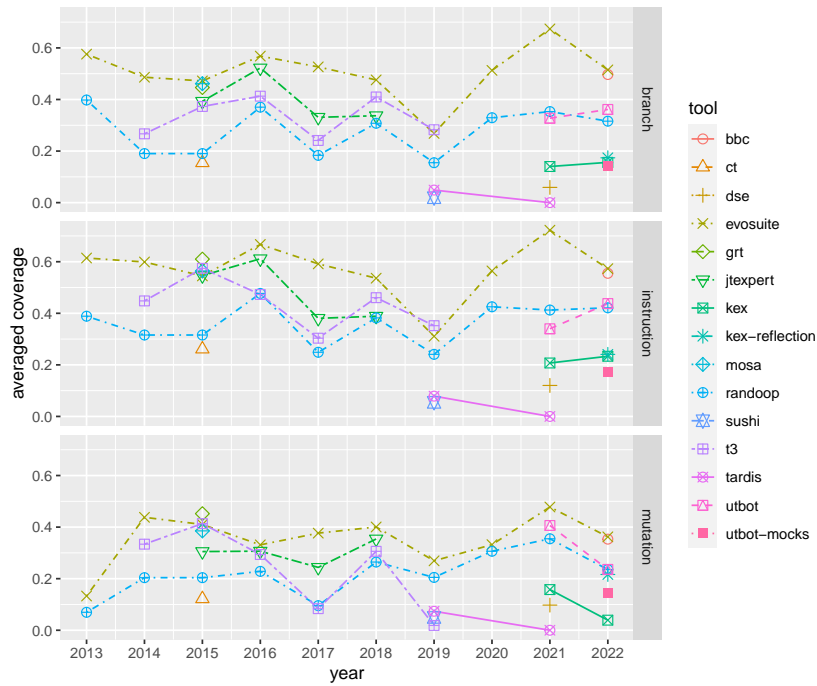


Figure 4. Averaged coverage evolution per tool over the years.

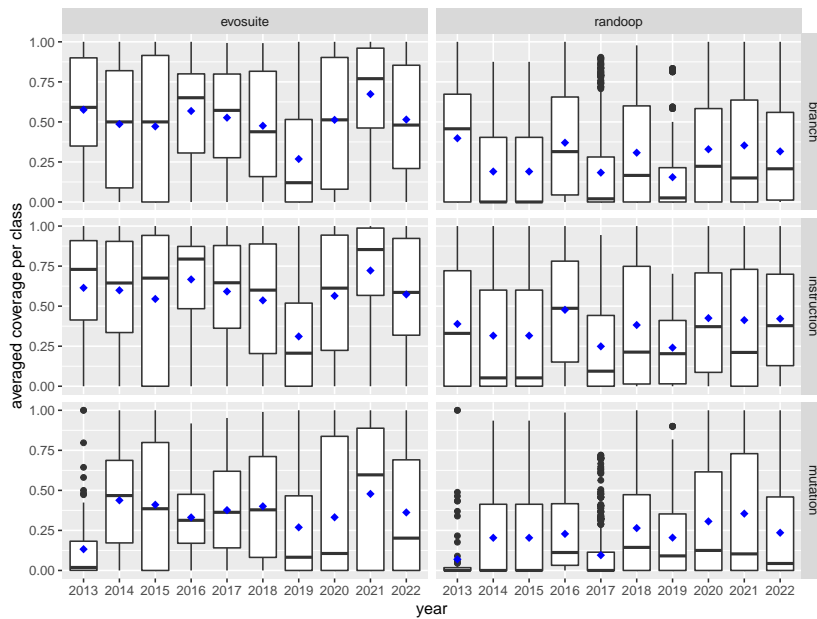


Figure 5. Averaged coverage evolution per benchmark for EvoSuite and Randoop. Diamond (◆) indicates the mean value.

2013. Figure 5 presents the distributions of the averaged coverage evolution per benchmark for each configuration of EVOSUITE and RANDOOP for the different editions of the competition (*i.e.*, each data point represents an averaged value of all the executions of the tool with one time budget on one benchmark). Generally, EVOSUITE performs better on individual benchmarks than RANDOOP

for unit test generation. This is confirmed when applying Friedman’s non-parametric with post-hoc Nemenyi test [78,79] on instructions coverage (RANDOOP ranking 1.77 and EVOSUITE 1.23), branch coverage (RANDOOP ranking 1.77 and EVOSUITE 1.23), and mutation score (RANDOOP ranking 1.73 and EVOSUITE 1.27). The rankings are significantly different as their corresponding average ranks differ by at least the given critical distance, here, 0.048, with a Friedman’s test p-value below 0.01 for the three cases.

6. DISCUSSION AND LESSONS LEARNED

Any empirical evaluation of automated unit test generation faces several technical and methodological challenges. JUGE seeks to alleviate those challenges by providing a standardized way of designing, running, and reporting such evaluations. Both the development of JUGE and the evaluation method reported in Section 4 took several years to develop. We discuss hereafter the main lessons learned and potential new applications of JUGE.

6.1. Building JUGE and evaluating generators

Adapt to different generators. The main technical challenges for such an infrastructure come from the diversity of the generators that can be considered for an empirical evaluation (*i.e.*, random-based, search-based, concolic/symbolic-based, *etc.*)(C1). It requires *isolating* (C2) the executions to avoid troubles in case of a bug in the generator (*e.g.*, erasing files from the host file system [68]) while still having a *standard communication interface* (C4). It is achieved through an adapter with a shared standard set of commands used by JUGE to interact with the generator. JUGE runs in a Docker container to guarantee the isolation of the generator from its environment during test case generation.

Performance, scalability, and statistical power. As for any empirical evaluation with a random-based generator, researchers have to balance the number of classes under test to consider reducing external validity with the number of executions to ensure enough statistical power, giving the external constraints on the overall execution time [57]. For instance, in the tool competition, the entire evaluation must be done in around two weeks. To cope with this limitation, organizers use sampling to select a subset of classes under test, limit the time budget (not more than 8 minutes), and the number of repetitions of the executions (between 6 and 10, depending on the year). As explained in Section 4.1, the time budget allocated to the generator depends on the type of approach and the goals of the evaluation.

JUGE supports parallel executions by running several Docker containers in parallel (the number of containers depends on the technical specifications of the host machine), allowing to *scale* while preserving the overall *performances* of an evaluation (C3). It allows relying on standard Docker technologies and tools (*e.g.*, dashboards) to handle the overall evaluation.

Configuration of the meta-parameters. In addition to the time budget and the number of repetitions, which can be configured for JUGE, the generators usually come with various meta-parameters that directly influence the generation process (C1). As explained in Section 4.1, such parameters should be carefully considered and reported to reduce the threats to the validity and enable the *replicability* of the results.

For instance, many test generators like EVOSUITE and RANDOOP include post-processing mechanisms that can be activated to minimize the generated tests [1, 3]. Such mechanisms are time-consuming and can be deactivated to reduce the overall execution time when evaluating properties such as coverage or the mutation score. However, deactivating test case minimization has a significant impact on other properties, such as structural properties, readability, the execution time of the tests, *etc.* Researchers should be aware of such impacts and carefully consider them when designing their studies.

Analysis of the generated tests. Automated test case generation is a challenging task and requires several mechanisms (*e.g.*, code instrumentation, handling I/O operations on the system under test, *etc.*) to be effective. Among the possible mechanisms is using a specific scaffolding for the generated tests: for instance, EVOSUITE controls elements that could be non-deterministic to avoid test flakiness. However, such mechanisms might cause undesired interactions with the infrastructure and, more specifically, with the tools used to analyze the generated tests (C1). It has been the case for EVOSUITE and the mutation analysis: the test runner (EvoRunner) used in the generated tests was not compatible with PITEST and required to use the mutated `.class` files directly instead of relying on the optimized PITEST infrastructure. In the latest version, JUGE includes options to parallelize the execution of the mutation analysis and reduce the overall execution time of the evaluation (C3).

Repeatability and reproducibility. JUGE considers each generator configuration (*e.g.*, EVOSUITE using a different generation algorithm) as a generator that will require its `runtool` adapter and corresponding shared folder, like, for instance, the different algorithms used with EVOSUITE in the competition. Those generators should be shared in a companion artifact to provide *standard* readily usable implementations to the research community (C4). As one of the main goals of JUGE is to provide a common platform where researchers and practitioners can plug their generators and compare them using various benchmarks, sharing generators and their corresponding adapters will greatly improve *repeatability* (*i.e.*, the same evaluation can be performed by the same team in the same conditions and produce the same results) and *reproducibility* (*i.e.*, the same evaluation can be performed by a different same team in similar conditions and produce the same results) of the results.

6.2. Benchmarks selection

As explained in Section 4.2, JUGE allows one to define their own set of benchmarks for a given evaluation. Section A.1 provides an example of benchmarks selection process, followed in the eighth edition of the tool competition [24]. As illustrated by the example, the competition mainly used an opportunistic approach by considering open-source projects built with Maven to ease the collection of projects' dependencies. The main selection criteria were availability and the complexity of the classes to test. Other criteria can also be taken into account.

In general, selecting representative benchmarks for an evaluation is not an easy task. The selection criteria depend on the goals of the evaluation but also the availability of the benchmarks. We provide some guidelines for benchmark selection based on our experience and related literature.

Representative. A good set of benchmarks should be representative of real-world software. This is a best practice adopted by the software engineering research community to ensure that the evaluation addresses a relevant problem [11, 13, 55, 60, 80]. The benchmarks can come from one or more systems depending on the experimental design. For instance, Almasi *et al.* [81] performed an evaluation of unit test generation on a closed-source industrial *case study*. Other evaluations studied unit test generation on various open-source (and openly accessible) systems [13, 74, 75, 82]. In this latter case, one should take care of selecting systems that are popular (*e.g.*, by looking at the number of stars on GitHub or the number of dependant projects in Maven central) and under active maintenance (*e.g.*, by looking at the commit frequency) to ensure that the results of the evaluation will be relevant to the software engineering industrial community.

Diverse. The benchmarks used for an evaluation should be diverse enough to mitigate threats to external validity. This diversity can, for instance, be achieved by considering benchmarks from multiple projects from diverse application domains. Another approach could be, instead of a random sampling like in Section A.1, selecting diverse benchmarks based on the coverage and mutation score a random generator achieves.

Adequate. The benchmarks should be selected adequately *w.r.t.* the goals of the evaluation. For instance, if one of the evaluation goals is to compare the coverage and mutation score of automatically generated and hand-written tests, one has to consider the coverage and mutation score of the existing test suites when selecting the benchmarks. *E.g.*, in 2019, the competition [23] introduced Spoon in the benchmarks as it has a hand-written test suite with high coverage. The results showed that none of the generators used in that edition could achieve a higher branch or line coverage, or mutation score.

Challenging. Depending on the kind of generation technique considered, some benchmarks might not be relevant or do not bring any interesting insights. For instance, Shamshiri *et al.* [83] showed that many classes in the SF110 corpus [64], a corpus of benchmarks widely used for unit test generation, can be covered using random search. Those benchmarks should be filtered out for evaluating search-based unit test generators as they will not provide interesting insights. For instance, the competition and other related work [13, 82] have considered McCabe’s cyclomatic complexity (*i.e.*, the number of independent paths in a control flow graph) of the benchmarks to filter out simple benchmarks, easily covered (*i.e.*, classes with a cyclomatic complexity lower than five).

Comparable. Generalizability is hard to achieve and requires several evaluations, preferably performed independently by different research teams. Selecting a subset of benchmarks that have already been used in other studies is interesting to allow comparison between different studies. It enables the meta-analysis of the results and cross-comparisons between different studies. For instance, Campos *et al.* [82], and Panichella *et al.* [13] reused the same benchmarks to perform large-scale empirical evaluations of search-based unit test generation algorithms, allowing a direct comparison and discussion of the results.

Documented. As stated in Section 4.4, both the selection procedure and the characteristics of the benchmarks should be reported together with the evaluation results. It is essential to the reviewing process to ensure that reviewers have enough information about the benchmarks to assess the paper describing the evaluation [59], but also to ensure that future research can compare and replicate the results on other benchmarks with similar characteristics.

Available. Finally, benchmarks coming from open-source projects should be made available to the research community. Building a good benchmark is not trivial and represents a substantial effort [65]. The community should share this effort by encouraging the best open-science practices and distribution of the benchmarks. JUGE provides a standard way of sharing and reusing such benchmarks (described in Section 4.2). ††

6.3. Future applications

The method described in Section 4 constitutes a standard that can be applied to unit test generation for other kinds of languages using an infrastructure similar to JUGE. For instance, Lukasczyk *et al.* [84] recently defined an approach to generate unit tests for Python. Of course, dynamically typed languages such as Python face additional challenges than those discussed in Section 3.1. Such challenges must be considered in the design of the infrastructure (*e.g.*, running type inference engines during pre-processing) and the selection of the benchmark (*e.g.*, considering classes with type annotations only, *etc.*), and reported in the description of the empirical evaluation.

Besides comparing unit test generators, the JUGE infrastructure can be used to generate large amounts of tests for various kinds of classes using different tools and configurations. It enables the continuous creation of an openly available corpus of automatically generated unit tests. Such a

††Benchmarks of previous tool competitions are available at <https://github.com/JUnitContest/JUGE/tree/master/infrastructure>.

corpus would (i) directly contribute to the body of empirical evidence on which decision-makers can rely to assess the usage of a unit test generator in their industrial context [11]; and (ii) enables further empirical evaluations on automatically generated tests without configuring and running the generators, which require a certain level of expertise. For instance, in a recent study, Panichella *et al.* [85] revisited previous studies on the presence of test smells in automatically generated tests and found that previous results vastly overestimated their presence. Among the different problems, they pointed out a misconfiguration of EVOSUITE and its minimization process, resulting in more prominent test cases more likely to contain certain smells. Building openly available corpora using the appropriate configuration for the generators, with a description of the characteristics applicable evaluations, would avoid such issues. More in general, having the JUGE infrastructure and its associated standards available to industrial and academic research communities can open the road for more systematic testing for other languages (e.g., Python, etc.) as well as the definition of testing pipelines to identify bugs and imperfections of systems in other application domains [86–89]. The availability of such technologies can also impact computer science education, with available tools that can be used in practical courses. Finally, we also expect future investigation supported by JUGE in the context of test code quality and cost-effectiveness of both automatically and manually generated tests [90,91].

Finally, regarding the infrastructure itself, we plan to add other types of analysis in addition to coverage and mutation score. Such analyses include performances of the generated tests (e.g., execution time and memory consumption) and readability.

7. CONCLUSION

JUGE sets a standard for properly assessing automated test case generators. It provides an infrastructure and a method to design, set up, and execute an empirical evaluation, collect and analyze the results, and produce a replication package to meet the requirements, enabling an effective contribution to the software testing empirical body of knowledge. It includes recommendations for selecting benchmarks and the parametrization of the generator and the infrastructure, depending on the research questions. JUGE was initially introduced and developed in the context of the tool competition and has been used with several generators and dozens of classes under test from various projects.

Finally, the JUGE infrastructure availability opens several directions for practitioners, who can rely on a large body of empirical evidence to assess automated test case generation usage in their context, and researchers, who can benefit from corpora of automatically generated tests for further empirical evaluations. JUGE also provides guidelines for evaluating unit test generation in other programming languages and for defining similar infrastructures in other domains.

ACKNOWLEDGEMENT

We would like to thank (in alphabetical order) Arthur Baars, Sebastian Bauersfeld, Matteo Biagiola, Ignacio Lebrero, Urko Rueda Molina, and Fiorella Zampetti for their contribution to the implementation of the JUGE infrastructure. We would also like to thank (in alphabetical order) Azat Abdullin, Marat Akhin, Giuliano Antoniol, Andrea Arcuri, Cyrille Artho, Mikhail Belyaev, Pietro Braione, Nikolay Bukharev, José Campos, Nelly Condori, Christoph Csallner, Giovanni Denaro, Gordon Fraser, Yann-Gaël Guéhéneuc, Masami Hagiya, Mainul Islam, Dmitry Ivanov, Gunel Jahangirova, Kiran Lakhota, Ignacio Manuel Lebrero Rial, Lei Ma, Alexey Menshutin, Arsen Nagdalian, Gilles Pesant, Simon Poulding, Wishnu Prasetya, Vincenzo Riccio, José Miguel Rojas, Abdelilah Sakti, Hiroyuki Sato, Sebastian Schweikl, Gleb Stromov, Yoshinori Tanabe, Paolo Tonella, Artem Ustinov, Sebastian Vogl, Tanja Vos, Mitsuharu Yamamoto, Fiorella Zampetti, and Cheng Zhang for their participation in previous editions of the competition and the feedback they provided on the infrastructure. Xavier Devroey was partially funded by the EU Horizon 2020 ICT-10-2016-RIA “STAMP” project (No.731529), the Vici “TestShift” project (No. VI.C.182.032) from the Dutch Science Foundation NWO, and the CyberExcellence (No. 2110186) project, funded by the Public Service of Wallonia (SPW Recherche). Alessio Gambi’s work was partially supported by the DFG project STUNT (DFG Grant Agreement n. FR 2955/4-1). Sebastiano Panichella and Annibale Panichella gratefully acknowledge the Horizon 2020 (EU Commission) support for the project *COSMOS* (DevOps for Complex

Cyber-physical Systems), Project No. 957254-COSMOS. René Just's work is partially supported by the National Science Foundation under grant CNS-1823172.

REFERENCES

1. Pacheco C, Lahiri SK, Ernst MD, Ball T. Feedback-Directed Random Test Generation. *29th International Conference on Software Engineering (ICSE'07)*, IEEE, 2007; 75–84, doi:10.1109/ICSE.2007.37.
2. Ma L, Artho C, Zhang C, Sato H, Gmeiner J, Ramler R. GRT: Program-Analysis-Guided Random Testing (T). *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, IEEE, 2015; 212–223, doi:10.1109/ASE.2015.49.
3. Fraser G, Arcuri A. EvoSuite: Automatic Test Suite Generation for Object-Oriented Software. *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering - SIGSOFT/FSE '11, ESEC/FSE '11*, ACM Press: New York, New York, USA, 2011; 416, doi:10.1145/2025113.2025179.
4. Alshahwan N, Gao X, Harman M, Jia Y, Mao K, Mols A, Tei T, Zorin I. Deploying Search Based Software Engineering with Sapienz at Facebook. *Search-Based Software Engineering. SSBSE 2018., LNCS*, vol. 11036, Springer, 2018, doi:10.1007/978-3-319-99241-9_1.
5. Derakhshanfar P, Devroey X, Panichella A, Zaidman A, Van Deursen A. Botsing, a Search-based Crash Reproduction Framework for Java. *35th IEEE/ACM International Conference on Automated Software Engineering (ASE '20)*, ACM/IEEE, 2020; 1278–1282, doi:10.1145/3324884.3415299.
6. Papadakis M, Malevis N. Automatic Mutation Test Case Generation via Dynamic Symbolic Execution. *2010 IEEE 21st International Symposium on Software Reliability Engineering*, IEEE, 2010; 121–130, doi:10.1109/ISSRE.2010.38.
7. Braione P, Denaro G, Mattavelli A, Pezzè M. Combining symbolic execution and search-based testing for programs with complex heap inputs. *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis - ISSTA 2017*, ACM Press: New York, New York, USA, 2017; 90–101, doi:10.1145/3092703.3092715.
8. Yun I, Lee S, Xu M, Jang Y, Kim T. QSYM : A practical concolic execution engine tailored for hybrid fuzzing. *27th USENIX Security Symposium (USENIX Security 18)*, USENIX Association: Baltimore, MD, 2018; 745–761.
9. Sen K. Concolic testing. *Proceedings of the twenty-second IEEE/ACM international conference on Automated software engineering - ASE '07*, 1, ACM Press, 2007; 571–572, doi:10.1145/1321631.1321746.
10. Juristo N, Moreno AM, Vegas S. Towards building a solid empirical body of knowledge in testing techniques. *ACM SIGSOFT Software Engineering Notes* sep 2004; **29**(5):1–4, doi:10.1145/1022494.1022544.
11. Dyba T, Kitchenham B, Jorgensen M. Evidence-based software engineering for practitioners. *IEEE Software* jan 2005; **22**(1):58–65, doi:10.1109/MS.2005.6.
12. Devroey X, Gambi A, Galeotti JP, Just R, Kifetew F, Panichella A, Panichella S. JUGE: Junit generation benchmarking infrastructure Jun 2021, doi:10.5281/zenodo.4904393. URL <https://doi.org/10.5281/zenodo.4904393>.
13. Panichella A, Kifetew FM, Tonella P. Automated Test Case Generation as a Many-Objective Optimisation Problem with Dynamic Selection of the Targets. *IEEE Transactions on Software Engineering* 2018; **44**(2):122–158, doi:10.1109/TSE.2017.2663435.
14. Sakti A, Pesant G, Gueheneuc YG. Instance Generator and Problem Representation to Improve Object Oriented Code Coverage. *IEEE Transactions on Software Engineering* mar 2015; **41**(3):294–313, doi:10.1109/TSE.2014.2363479.
15. Prasetya ISWB. T3, a Combinator-Based Random Testing Tool for Java: Benchmarking. *Future Internet Testing. FITTEST 2013, LNCS*, vol. 8432, Vos T, Lakhota K, Bauersfeld S (eds.), Springer, 2014; 101–110, doi:10.1007/978-3-319-07785-7_7.
16. Braione P, Denaro G, Mattavelli A, Pezzè M. SUSHI: A Test Generator for Programs with Complex Structured Inputs. *Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings*, i, ACM Press: New York, NY, USA, 2018; 21–24, doi:10.1145/3183440.3183472.
17. Bauersfeld S, Vos TEJ, Lakhota K, Poulding SM, Condori-Fernández N. Unit testing tool competition. *Sixth IEEE International Conference on Software Testing, Verification and Validation, ICST 2013 Workshops Proceedings, Luxembourg, Luxembourg, March 18-22, 2013*, IEEE Computer Society, 2013; 414–420, doi:10.1109/ICSTW.2013.55.
18. Bauersfeld S, Vos TEJ, Lakhota K. Unit testing tool competitions - lessons learned. *Future Internet Testing - First International Workshop, FITTEST 2013, Istanbul, Turkey, November 12, 2013, Revised Selected Papers, Lecture Notes in Computer Science*, vol. 8432, Vos TEJ, Lakhota K, Bauersfeld S (eds.), Springer, 2013; 75–94, doi:10.1007/978-3-319-07785-7_5.
19. Rueda U, Vos TEJ, Prasetya ISWB. Unit testing tool competition - round three. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 19–24, doi:10.1109/SBST.2015.12.
20. Rueda U, Just R, Galeotti JP, Vos TEJ. Unit testing tool competition: round four. *Proceedings of the 9th International Workshop on Search-Based Software Testing, SBST@ICSE 2016, Austin, Texas, USA, May 14-22, 2016*, ACM, 2016; 19–28, doi:10.1145/2897010.2897018.
21. Panichella A, Molina UR. Java unit testing tool competition - fifth round. *10th IEEE/ACM International Workshop on Search-Based Software Testing, SBST@ICSE 2017, Buenos Aires, Argentina, May 22-23, 2017*, IEEE, 2017; 32–38, doi:10.1109/SBST.2017.7.
22. Molina UR, Kifetew FM, Panichella A. Java unit testing tool competition: sixth round. *Proceedings of the 11th International Workshop on Search-Based Software Testing, ICSE 2018, Gothenburg, Sweden, May 28-29, 2018*, Galeotti JP, Gorla A (eds.), ACM, 2018; 22–29, doi:10.1145/3194718.3194728.

23. Kifetew FM, Devroey X, Rueda U. Java unit testing tool competition: seventh round. *Proceedings of the 12th International Workshop on Search-Based Software Testing, SBST@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Gorla A, Rojas JM (eds.), IEEE / ACM, 2019; 15–20, doi:10.1109/SBST.2019.00014.
24. Devroey X, Panichella S, Gambi A. Java Unit Testing Tool Competition - Eighth Round. *2020 IEEE/ACM 13th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2020, doi:10.1145/3387940.3392265.
25. Panichella S, Gambi A, Zampetti F, Riccio V. Sbst tool competition 2021. *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2021, doi:10.1109/SBST52555.2021.00011.
26. Gambi A, Jahangirova G, Riccio V, Zampetti F. Sbst tool competition 2022. *2022 IEEE/ACM 15th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2022, doi:10.1145/3526072.3527538.
27. Fraser G, Arcuri A. Evosuite at the SBST 2013 tool competition. *Sixth IEEE International Conference on Software Testing, Verification and Validation, ICST 2013 Workshops Proceedings, Luxembourg, Luxembourg, March 18-22, 2013*, IEEE Computer Society, 2013; 406–409, doi:10.1109/ICSTW.2013.53.
28. Prasetya ISWB. Measuring T2 against SBST 2013 benchmark suite. *Sixth IEEE International Conference on Software Testing, Verification and Validation, ICST 2013 Workshops Proceedings, Luxembourg, Luxembourg, March 18-22, 2013*, IEEE Computer Society, 2013; 410–413, doi:10.1109/ICSTW.2013.54.
29. Fraser G, Arcuri A. Evosuite at the second unit testing tool competition. *Future Internet Testing - First International Workshop, FITTEST 2013, Istanbul, Turkey, November 12, 2013, Revised Selected Papers, Lecture Notes in Computer Science*, vol. 8432, Vos TEJ, Lakhotia K, Bauersfeld S (eds.), Springer, 2013; 95–100, doi:10.1007/978-3-319-07785-7_6.
30. Prasetya ISWB. T3, a combinator-based random testing tool for java: Benchmarking. *Future Internet Testing - First International Workshop, FITTEST 2013, Istanbul, Turkey, November 12, 2013, Revised Selected Papers, Lecture Notes in Computer Science*, vol. 8432, Vos TEJ, Lakhotia K, Bauersfeld S (eds.), Springer, 2013; 101–110, doi:10.1007/978-3-319-07785-7_7.
31. Fraser G, Arcuri A. Evosuite at the SBST 2015 tool competition. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 25–27, doi:10.1109/SBST.2015.13.
32. Ma L, Artho C, Zhang C, Sato H, Hagiya M, Tanabe Y, Yamamoto M. GRT at the SBST 2015 tool competition. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 48–51, doi:10.1109/SBST.2015.19.
33. Sakti A, Pesant G, Guéhéneuc Y. Jtexpert at the third unit testing tool competition. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 52–55, doi:10.1109/SBST.2015.20.
34. Panichella A, Kifetew FM, Tonella P. Results for evosuite - MOSA at the third unit testing tool competition. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 28–31, doi:10.1109/SBST.2015.14.
35. Prasetya ISWB. T3: benchmarking at third unit testing tool contest. *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015, Florence, Italy, May 18-19, 2015*, Gay G, Antoniol G (eds.), IEEE Computer Society, 2015; 44–47, doi:10.1109/SBST.2015.18.
36. Fraser G, Arcuri A. Evosuite at the SBST 2016 tool competition. *Proceedings of the 9th International Workshop on Search-Based Software Testing, SBST@ICSE 2016, Austin, Texas, USA, May 14-22, 2016*, ACM, 2016; 33–36, doi:10.1145/2897010.2897020.
37. Sakti A, Pesant G, Guéhéneuc Y. Jtexpert at the fourth unit testing tool competition. *Proceedings of the 9th International Workshop on Search-Based Software Testing, SBST@ICSE 2016, Austin, Texas, USA, May 14-22, 2016*, ACM, 2016; 37–40, doi:10.1145/2897010.2897021.
38. Prasetya ISWB. Budget-aware random testing with T3: benchmarking at the SBST2016 testing tool contest. *Proceedings of the 9th International Workshop on Search-Based Software Testing, SBST@ICSE 2016, Austin, Texas, USA, May 14-22, 2016*, ACM, 2016; 29–32, doi:10.1145/2897010.2897019.
39. Fraser G, Rojas JM, Campos J, Arcuri A. Evosuite at the SBST 2017 tool competition. *10th IEEE/ACM International Workshop on Search-Based Software Testing, SBST@ICSE 2017, Buenos Aires, Argentina, May 22-23, 2017*, IEEE, 2017; 39–42, doi:10.1109/SBST.2017.6.
40. Sakti A, Pesant G, Guéhéneuc Y. Jtexpert at the SBST 2017 tool competition. *10th IEEE/ACM International Workshop on Search-Based Software Testing, SBST@ICSE 2017, Buenos Aires, Argentina, May 22-23, 2017*, IEEE, 2017; 43–46, doi:10.1109/SBST.2017.5.
41. Fraser G, Rojas JM, Arcuri A. Evosuite at the SBST 2018 tool competition. *Proceedings of the 11th International Workshop on Search-Based Software Testing, ICSE 2018, Gothenburg, Sweden, May 28-29, 2018*, Galeotti JP, Gorla A (eds.), ACM, 2018; 34–37, doi:10.1145/3194718.3194729.
42. Prasetya ISWB. T3 @sbst2018 benchmark, and how much we can get from asemantical testing. *Proceedings of the 11th International Workshop on Search-Based Software Testing, ICSE 2018, Gothenburg, Sweden, May 28-29, 2018*, Galeotti JP, Gorla A (eds.), ACM, 2018; 30–33, doi:10.1145/3194718.3194727.
43. Campos J, Panichella A, Fraser G. Evosuite at the SBST 2019 tool competition. *Proceedings of the 12th International Workshop on Search-Based Software Testing, SBST@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Gorla A, Rojas JM (eds.), IEEE / ACM, 2019; 29–32, doi:10.1109/SBST.2019.00017.
44. Braione P, Denaro G. SUSHI and TARDIS at the SBST2019 tool competition. *Proceedings of the 12th International Workshop on Search-Based Software Testing, SBST@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Gorla A, Rojas JM (eds.), IEEE / ACM, 2019; 25–28, doi:10.1109/SBST.2019.00016.
45. Prasetya ISWB. Random testing with austere budgeting in T3: benchmarking at SBST2019 testing tool contest. *Proceedings of the 12th International Workshop on Search-Based Software Testing, SBST@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Gorla A, Rojas JM (eds.), IEEE / ACM, 2019; 21–24, doi:10.1109/SBST.2019.00015.
46. Panichella A, Campos J, Fraser G. EvoSuite at the SBST 2020 Tool Competition. *2020 IEEE/ACM 13th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2020, doi:10.1145/3387940.

- 3392266.
47. Vogl S, Schweikl S, Fraser G, Arcuri A, Campos J, Panichella A. EvoSuite at the SBST 2021 Tool Competition. *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2021, doi: 10.1109/SBST52555.2021.00012.
 48. Lebrero Rial IM, Galeotti JP. EvoSuiteDSE at the SBST 2021 Tool Competition. *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2021, doi:10.1109/SBST52555.2021.00013.
 49. Abdullin A, Akhin M, Belyaev M. Kex at the 2021 SBST Tool Competition. *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2021, doi:10.1109/SBST52555.2021.00014.
 50. Ivanov D, Bukharev N, Menshutin A, Nagdalian A, Stromov G, Ustinov A. UtBot at the SBST2021 Tool Competition. *2021 IEEE/ACM 14th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2021, doi:10.1109/SBST52555.2021.00015.
 51. Derakhshanfar P, Devroey X. Basic Block Coverage for Unit Test Generation at the SBST 2022 Tool Competition. *2022 IEEE/ACM 15th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2022, doi: 10.1145/3526072.3527528.
 52. Schweikl S, Fraser G, Arcuri A. EvoSuite at the SBST 2022 Tool Competition. *2022 IEEE/ACM 15th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2022, doi:10.1145/3526072.3527526.
 53. Abdullin A, Akhin M, Belyaev M. Kex at the 2022 SBST Tool Competition. *2022 IEEE/ACM 15th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2022, doi:10.1145/3526072.3527527.
 54. Ivanov D, Menshutin A, Fokin D, Kamenev Y, Pospelov S, Kulikov E, Stroganov N. UtBot at the SBST2022 Tool Competition. *2022 IEEE/ACM 15th International Workshop on Search-Based Software Testing (SBST)*, IEEE / ACM, 2022, doi:10.1145/3526072.3527529.
 55. Vos TE, Marin B, Escalona MJ, Marchetto A. A Methodological Framework for Evaluating Software Testing Techniques and Tools. *2012 12th International Conference on Quality Software*, IEEE, 2012; 230–239, doi: 10.1109/QSIC.2012.16.
 56. Melo SM, Moura FM, Souza PSL, Souza SRS. SeleCTT: An Infrastructure for Selection of Concurrent Software Testing Techniques. *Proceedings of the IV Brazilian Symposium on Systematic and Automated Software Testing - SAST 2019*, ACM, 2019; 62–71, doi:10.1145/3356317.3356324.
 57. Arcuri A, Briand L. A Hitchhiker’s guide to statistical tests for assessing randomized algorithms in software engineering. *Software Testing, Verification and Reliability* 2014; **24**(3):219–250, doi:10.1002/stvr.1486.
 58. Ali S, Briand LC, Hemmati H, Panesar-Walawege RK. A Systematic Review of the Application and Empirical Investigation of Search-Based Test Case Generation. *IEEE Transactions on Software Engineering* nov 2010; **36**(6):742–762, doi:10.1109/TSE.2009.52.
 59. Ralph P, bin Ali N, Baltas S, Bianculli D, Diaz J, Dittrich Y, Ernst N, Felderer M, Feldt R, Filieri A, *et al.*. Empirical standards for software engineering research 2021.
 60. Metzman J, Szekeres L, Simon L, Sprabery R, Arya A. FuzzBench: an open fuzzer benchmarking platform and service. *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ACM, 2021; 1393–1403, doi:10.1145/3468264.3473932.
 61. Andrews J, Briand L, Labiche Y, Namin A. Using Mutation Analysis for Assessing and Comparing Testing Coverage Criteria. *IEEE Transactions on Software Engineering* aug 2006; **32**(8):608–624, doi:10.1109/TSE.2006.83.
 62. DeMillo R, Lipton R, Sayward F. Hints on Test Data Selection: Help for the Practicing Programmer. *Computer* apr 1978; **11**(4):34–41, doi:10.1109/C-M.1978.218136.
 63. Just R, Jalali D, Inozemtseva L, Ernst MD, Holmes R, Fraser G. Are mutants a valid substitute for real faults in software testing? *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering - FSE 2014*, FSE 2014, ACM Press, 2014; 654–665, doi:10.1145/2635868.2635929.
 64. Fraser G, Arcuri A. A Large-Scale Evaluation of Automated Unit Test Generation Using EvoSuite. *ACM Transactions on Software Engineering and Methodology* dec 2014; **24**(2):1–42, doi:10.1145/2685612.
 65. Just R, Jalali D, Ernst MD. Defects4J: a database of existing faults to enable controlled testing studies for Java programs. *Proceedings of the 2014 International Symposium on Software Testing and Analysis - ISSA 2014*, ACM Press, 2014; 437–440, doi:10.1145/2610384.2628055.
 66. Gay G, Just R. Defects4J as a Challenge Case for the Search-Based Software Engineering Community. *Symposium on Search-Based Software Engineering. SSBSE 2020.*, LNCS, vol. 12420, Springer, 2020; 255–261, doi:10.1007/978-3-030-59762-7_19.
 67. Tomassi DA, Dmeiri N, Wang Y, Bhowmick A, Liu YC, Devanbu PT, Vasilescu B, Rubio-Gonzalez C. BugSwarm: Mining and Continuously Growing a Dataset of Reproducible Failures and Fixes. *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, IEEE, 2019; 339–349, doi:10.1109/ICSE.2019.00048.
 68. Fraser G, Arcuri A. EvoSuite: On the Challenges of Test Case Generation in the Real World. *2013 IEEE Sixth International Conference on Software Testing, Verification and Validation*, IEEE, 2013; 362–369, doi: 10.1109/ICST.2013.51.
 69. Hoffmann MR, Mandrikov E, *et al.*. Jacoco java code coverage library 2014. URL <https://github.com/jacoco/jacoco>.
 70. Coles H, Laurent T, Henard C, Papadakis M, Ventresque A. PIT: a practical mutation testing tool for Java. *Proceedings of the 25th International Symposium on Software Testing and Analysis - ISSA 2016*, ACM Press: New York, New York, USA, 2016; 449–452, doi:10.1145/2931037.2948707.
 71. Conover WJ, Iman RL. Rank Transformations as a Bridge Between Parametric and Nonparametric Statistics. *The American Statistician* aug 1981; **35**(3):124, doi:10.2307/2683975.
 72. Nagappan M, Zimmermann T, Bird C. Diversity in software engineering research. *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2013*, ACM Press: New York, New York, USA, 2013; 466, doi:10.1145/2491411.2491415.

73. Daka E, Campos J, Fraser G, Dorn J, Weimer W. Modeling readability to improve unit tests. *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering - ESEC/FSE 2015*, 2015; 107–118, doi: 10.1145/2786805.2786838.
74. Fraser G, Arcuri A. Whole Test Suite Generation. *IEEE Transactions on Software Engineering* feb 2013; **39**(2):276–291, doi:10.1109/TSE.2012.14.
75. Panichella A, Kifetew FM, Tonella P. Reformulating Branch Coverage as a Many-Objective Optimization Problem. *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, IEEE, 2015; 1–10, doi:10.1109/ICST.2015.7102604.
76. Herlim RS, Hong S, Kim Y, Kim M. Empirical Study of Effectiveness of EvoSuite on the SBST 2020 Tool Competition Benchmark. *Search-Based Software Engineering. SSBSE 2021., LNCS*, vol. 12914, O’Reilly UM, Devroey X (eds.), Springer, 2021; 121–135, doi:10.1007/978-3-030-88106-1_9.
77. Devroey X. JUnitContest/junit-tool-competition-results: Tools Competition Results from 2013 to 2022 Dataset Oct 2022, doi:10.5281/zenodo.7194656. URL <https://doi.org/10.5281/zenodo.7194656>.
78. Japkowicz N, Shah M. *Evaluating Learning Algorithms: A Classification Perspective*. Cambridge University Press, 2011.
79. Panichella A. A Systematic Comparison of search-Based approaches for LDA hyperparameter tuning. *Information and Software Technology* feb 2021; **130**:106 411, doi:10.1016/j.infsof.2020.106411.
80. Juristo N, Moreno AM, Vegas S. Reviewing 25 Years of Testing Technique Experiments. *Empirical Software Engineering* Mar 2004; **9**(1/2):7–44, doi:10.1023/B:EMSE.0000013513.48963.1b.
81. Almasi MM, Hemmati H, Fraser G, Arcuri A, Benefelds J. An Industrial Evaluation of Unit Test Generation: Finding Real Faults in a Financial Application. *2017 IEEE/ACM 39th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP)*, IEEE, 2017; 263–272, doi:10.1109/ICSE-SEIP.2017.27.
82. Campos J, Ge Y, Albulian N, Fraser G, Eler M, Arcuri A. An empirical evaluation of evolutionary algorithms for unit test suite generation. *Information and Software Technology* 2018; **104**(August):207–235, doi:10.1016/j.infsof.2018.08.010.
83. Shamshiri S, Rojas JM, Fraser G, McMinn P. Random or Genetic Algorithm Search for Object-Oriented Test Suite Generation? *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference - GECCO '15*, vol. 28, ACM Press, 2015; 1367–1374, doi:10.1145/2739480.2754696.
84. Lukaczyk S, Kroiß F, Fraser G. Automated Unit Test Generation for Python. *Search-Based Software Engineering. SSBSE 2020., LNCS*, vol. 12420, Aleti A, Panichella A (eds.), Springer, 2020; 9–24, doi:10.1007/978-3-030-59762-7_2.
85. Panichella A, Panichella S, Fraser G, Sawant AA, Hellendoorn VJ. Revisiting Test Smells in Automatically Generated Tests: Limitations, Pitfalls, and Opportunities. *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, IEEE, 2020; 523–533, doi:10.1109/ICSME46990.2020.00056.
86. Di Sorbo A, Zampetti F, Visaggio CA, Penta DM, Panichella S. Automated identification and qualitative characterization of safety concerns reported in uav software platforms. *ACM Transactions on Software Engineering and Methodology* sep 2022; doi:10.1145/3564821.
87. Zampetti F, Kapur R, Di Penta M, Panichella S. An empirical characterization of software bugs in open-source cyber-physical systems. *Journal of Systems and Software* 2022; **192**:111425, doi:https://doi.org/10.1016/j.jss.2022.111425. URL <https://www.sciencedirect.com/science/article/pii/S0164121222001315>.
88. Birchler C, Khatiri S, Derakhshanfar P, Panichella S, Panichella A. Single and multi-objective test cases prioritization for self-driving cars in virtual environments. *ACM Trans. Softw. Eng. Methodol.* apr 2022; doi: 10.1145/3533818. URL <https://doi.org/10.1145/3533818>.
89. Birchler C, Ganz N, Khatiri S, Gambi A, Panichella S. Cost-effective simulationbased test selection in self-driving cars software with sdc-scissor. *2022 IEEE 29th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2022, doi:toappear.
90. Panichella A, Panichella S, Fraser G, Sawant A, Hellendoorn V. Test Smells 20 Years Later: Detectability, Validity, and Reliability. *Empirical Software Engineering (EMSE)* 2022; .
91. Grano G, Laaber C, Panichella A, Panichella S. Testing with fewer resources: An adaptive approach to performance-aware test case generation. *IEEE Transactions on Software Engineering* 2021; **47**(11):2332–2347, doi:10.1109/TSE.2019.2946773.

A. EXAMPLE OF EVALUATION USING JUGE

This appendix illustrates how JUGE can be used in practice. We use as an example the eighth edition of the Java unit testing tool competition [24] in which three of the authors were involved. The competition has occurred since 2013 and is co-located with the Search-Based Software Testing workshop. Participants willing to enter the competition have to provide (i) an executable version of their tool (potentially obfuscated, which enables the participation of industrial tools); and (ii) an implementation of the `runTool` adapter (described in Section 3.3) for their tool. In practice, this implementation consists in cloning the `runTool` project available in the JUGE repository and defining the methods of the `ITestingTool` interface described in Listing 1. The effort is minimal and consists of writing between 100 and 200 lines of code for a reasonably well-implemented tool (*i.e.*, tools relying on configuration parameters, for instance).

The competition aims to compare different tools on a diverse set of classes under test (*i.e.*, benchmarks) for different time budgets. As the organizers run the evaluation, the time budgets and number of executions

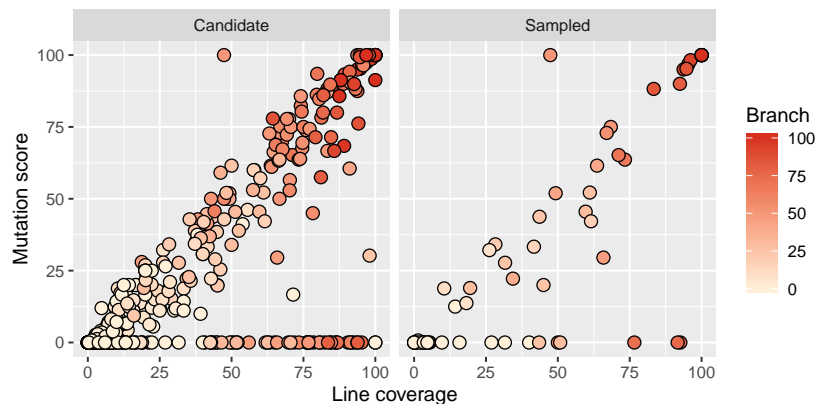


Figure 6. Example of reporting of the line coverage, branch coverage and mutation score for the candidate (382 classes) and selected benchmarks (60 classes, 20 per project) from the 2020 tool competition [24].

depend on the computational resources available. For the eighth edition, the tools were executed ten times for 60 and 180 seconds on each benchmark.

A.1. Benchmarks selection

Several ways exist to define a new set of benchmarks for a given set of projects. The projects selected for that edition were a mix of projects from the previous iterations of the competition to allow comparing results across years and popular GitHub projects built using Maven to ease dependency collection. The benchmarks were selected following a two-step procedure for the list of selected projects. In the first step, (i) identifying the packages in the project that contain classes relevant for the evaluation (*e.g.*, packages containing classes with the business logic); (ii) computing the McCabe's cyclomatic complexity for the different classes of those packages and remove classes with a complexity lower than five. This reduces the risk to sample classes with few branches, easily covered by randomly generated tests [13].

In the second step, a random generator (here, RANDOOP [1]) was executed with a low time budget (*e.g.*, ten seconds) on the remaining classes to filter out classes for which the generator could not generate any tests. This reduces the chances of facing technical difficulties while evaluating the different tools. Since the remaining classes were still too high, a subset of classes was randomly sampled for each project.

In addition to the previous steps, one can also use JUGE to perform a coverage and mutation analysis of the tests produced by the random generator and report the results for the candidate and sampled classes. Figure 6 shows the line and branch coverage and mutation score of the candidate and sampled classes of the eighth edition of the tool competition.

A.2. Execution

The eighth edition of the competition received only one submission (*i.e.*, EVOSUITE with DYNAMOSA) that was compared against RANDOOP, used as a baseline. The implementation of the `runtool` adapter is available in JUGE's GitHub repository. EVOSUITE and RANDOOP were executed ten times each against each benchmark. The executions were run in parallel (using Docker) on two servers: one with 40 CPU cores (2.30GHz) with 128 GB memory and one with 8 CPU cores (2.49GHz) with 160 GB memory. The coverage and mutation analysis were performed on the same machines. The total execution time for test generation and analysis took around four days.

A.3. Data analysis and ranking of the contestants

The competition combines the different metrics to ease the comparison of different generators. For that, it relies on a *scoring formula* [21]. This formula has been developed and refined during the different editions of the competition and takes into account the line and branch coverage, the mutation score, and the time budget used by the generator, and applies a penalty for flaky and non-compiling tests. In 2020, EVOSUITE ranked first. As a further example, Table II provides the ranking obtained through Friedman's test for the fifth edition of the tool competition [21]. EVOSUITE is ranked first with an average score of 1457, followed by JTEXPERT, T3, and RANDOOP. Table III gives the post-hoc Conover's test results for the same edition of the competition and indicates that the various comparisons are statistically significant, except for T3 and RANDOOP for which the *p*-value is above the confidence level of 0.05.

Table II. Example of scores and ranking obtained through Friedman’s test for the 5th edition of the tool competition [21].

| Tool | Score | Std.dev | Ranking |
|-------------|--------------|----------------|----------------|
| EVOsuite | 1457 | 192.72 | 1.55 |
| JTExpert | 849 | 102.03 | 2.71 |
| T3 | 526 | 82.43 | 2.81 |
| RANDOOOP | 448 | 34.75 | 2.92 |

Table III. Example of results of the post-hoc Conover’s test for the 5th edition of the tool competition [21].

| | EVOsuite | JTExpert | T3 | RANDOOOP |
|----------|----------|----------|------|----------|
| EVOsuite | - | - | - | - |
| JTExpert | < 0.01 | - | - | - |
| T3 | < 0.01 | 0.01 | - | - |
| RANDOOOP | < 0.01 | < 0.01 | 0.06 | - |

The scoring formula used in the competition provides an aggregated measure to rank the different tools based on their coverage and mutation analysis performances. Other aspects could be considered, like, for instance, the readability of the generated tests. Considering such aspects and including them in the scoring formula is part of the future work identified in the tenth edition of the competition [26].

Once the coverage and mutation analysis have been performed, the full results for a tool (*i.e.*, benchmarks, generated tests, and collected data) are sent to the tool’s authors for further analysis. Authors can further analyze the results for each round of execution of their tool on each benchmark.