

A Systematic Literature Review of How Mutation Testing Supports Test Activities

Qianqian Zhu, Annibale Panichella, Andy Zaidman

Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628CD Delft, The Netherlands.

Corresponding author:

Qianqian Zhu

Email address: qianqian.zhu@email.com

ABSTRACT

Mutation testing has been very actively investigated by researchers since the 1970s and remarkable advances have been achieved in its concepts, theory, technology and empirical evidence. While the latest realisations have been summarised by existing literature review, we lack insight into how mutation testing is *actually* applied. Our goal is to identify and classify the main applications of mutation testing and analyse the level of replicability of empirical studies related to mutation testing. To this aim, this paper provides a systematic literature review on the *application perspective* of mutation testing based on a collection of 159 papers published between 1981 and 2015. In particular, we analysed in which testing activities mutation testing is used, which mutation tools and which mutation operators are employed. Additionally, we also investigated how the core inherent problems of mutation testing, i.e. the equivalent mutant problem and the high computational cost, are addressed during the actual usage. The results show that most studies use mutation testing as an assessment tool targeting unit tests, and many of the supporting techniques for making mutation testing applicable in practice are still underdeveloped. Based on our observations, we made nine recommendations for the future work, including an important suggestion on how to report mutation testing in testing experiments in an appropriate manner.

1 INTRODUCTION

Mutation testing is defined by Jia and Harman [1] as a fault-based testing technique which provides a testing criterion called the *mutation adequacy score*. This score can be used to measure the effectiveness of a test set in terms of its ability to detect faults [1]. The principle of mutation testing is to introduce syntactic changes into the original program to generate faulty versions (called *mutants*) according to well-defined rules (mutation operators) [2]. Mutation testing originated in the 1970s with works from Lipton [3], DeMillo et al. [4] and Hamlet [5] and has been a very active research field over the last few decades. The activeness of the field is in part evidenced by the extensive survey of more than 390 papers on mutation testing that Jia and Harman published in 2011 [1]. Jia and Harman's survey highlights the research achievements that have been made over the years, including the development of tools for a variety of languages and empirical studies performed [1]. Additionally, they highlight some of the actual and inherent problems of mutation testing, amongst others: (1) the high computational cost caused by generating and executing the numerous mutants and (2) the tremendous time-consuming human investigation required by test oracle problem and equivalent mutants detection.

While Jia and Harman's survey provides us with a great overview of the latest realisations in research, we lack insight into how mutation testing is *actually applied*. Specifically, we are interested in analysing in which testing activities mutation testing is used, which mutation tools are employed and which mutation operators are used. Additionally, we want to investigate how the aforementioned problems of the high computational cost and the considerable human effort required are dealt with when applying mutation testing. In order to steer our research, we aim to fulfil the following objectives:

- to identify and classify the applications of mutation testing in testing activities;
- to analyse how the main problems are coped with when applying mutation testing;

- to provide guidelines for applying mutation testing in testing experiments;
- to identify gaps in current research and to provide recommendations for future work.

As systematic literature reviews have been shown to be good tools to summarise existing evidence concerning a technology and identify gaps in current research [6], we will follow this approach for reaching our objectives. We only consider the articles which provide sufficient details on how mutation testing is used in their studies, which require at least brief specification about the adopted mutation tool, mutation operators or mutation score. Moreover, we selected only papers that use Mutation Testing as a tool for evaluating or improving other testing activities rather than focusing on the development of mutation tools, operators or challenges and open issues for mutation analysis. This resulted in a collection containing 159 papers published from 1981 to 2015. We analysed this collection in order to answer the following two research questions:

RQ1: *How mutation testing is used in testing activities?*

This research question aims to identify and classify the main software testing tasks where mutation testing is applied. In particular, we are interested in the following key aspects: (1) in which circumstances mutation testing is used (e.g. assessment tool), (2) which testing activities are involved (e.g. test data generation, test case prioritisation), (3) which test level it targets (e.g. unit level) and (4) which testing strategies it supports (e.g. structural testing). The above four detailed aspects are defined to characterise the essential features related to usage of mutation testing and the testing activities involved. With these elements in place, we can provide an in-depth analysis of the applications of mutation testing.

RQ2: *How are empirical studies related to the mutation testing designed and reported?*

The objective of this question is to synthesise empirical evidence related to mutation testing. The case studies or experiments play an inevitable role in a research study. The design and demonstration of the evaluation methods should ensure the replicability. The replicability means that the subject, the basic methodology, as well as the result, should be clearly pointed out in the article. In particular, we are interested in how the articles report the following information related to mutation testing: (1) mutation tools, (2) mutation operators, (3) mutant equivalence problem, (4) techniques for reduction of computational cost and (5) subject programs used in the case studies. After gathering this information, we can draw conclusions from the distributions of related techniques adopted under the above five facets and thereby provide guidelines for applying mutation testing.

The remainder of this review is organised as follows: Section 2 provides an overview on background notions on Mutation testing. Section 3 details the main procedures we followed to conduct the systematic literature review and describes our inclusion and exclusion criteria. Section 4 presents the discussion of the our findings, particularly Section 4.3 summarises the answers to the research questions while Section 4.4 provides the recommendation for the future research. Section 5 discusses the threats to validity, and the Section 6 concludes the paper.

2 BACKGROUND

In order to level the playing field, we first provide the basic concepts related to mutation testing, i.e., its fundamental hypothesis and generic process, including the *Competent Programmer Hypothesis*, the *Coupling Effect*, *mutation operators* and the *mutation score*. Subsequently, we discuss the benefits and limitations of mutation testing. After that, we present a historical overview of mutation testing where we mainly address the studies that concern the application of mutation testing.

2.1 Basic Concepts

2.1.1 Fundamental Hypothesis

Mutation testing starts with the assumption of the *Competent Programmer Hypothesis* (introduced by Demillo et al. [4] in 1978): “The competent programmers create programs that are close to being correct.” This hypothesis implies that the potential faults in the programs delivered by the competent programmers are just very simple mistakes; these defects can be corrected by a few simple syntactical changes. Inspired by the above hypothesis, mutation testing typically applies small syntactical changes to original programs, thus implying that the faults that are seeded resemble faults made by “competent programmers”.

At first glance, it seems that the programs with complex errors cannot be explicitly generated by mutation testing. However, the *Coupling Effect*, which was coined by Demillo et al. [4] states that “Test data that distinguishes all programs differing from a correct one by only simple errors is so sensitive that

it also implicitly distinguishes more complex errors”. This means complex faults are *coupled* to simple faults. This hypothesis was later supported by Offutt [7, 8] through empirical investigations over the domain of mutation testing. In his experiments, he used first order mutants, which are created by applying the mutation operator to the original program once, to represent simple faults. Conversely, higher-order mutants, which are created by applying the mutation operator to the original program more than once, stand for complex faults. The results showed that the test data generated for 1-order mutants killed a higher percentage of mutants when applied to higher-order mutants thus yielding positive empirical evidence about the *Coupling Effect*. Besides, there has been a considerable effort in validating the coupling effect hypothesis, amongst others the theoretical studies of Wah [9–11] and Kapoor [12].

2.1.2 The Generic Mutation Testing Process

After introducing the fundamental hypotheses of mutation testing, we are going to give a detailed description of the *generic process* of mutation testing:

Given a program P and a test suite T , a mutation engine makes syntactic changes (defined as mutation operators) to the program P , thereby generating a set of mutant programs \mathbb{M} . After that, each mutant $P_m \in \mathbb{M}$ is executed against T to verify whether tests fail or not.

Here is an example of a mutation operator, i.e. Arithmetic Operator Replacement (AOR), on a statement “ $X = a + b$ ”. The produced mutants include “ $X = a - b$ ”, “ $X = a \times b$ ” and “ $X = a \div b$ ”.

The execution results of T on $P_m \in \mathbb{M}$ are compared with P : (1) if the output of P_m is different from P , then P_m is *killed* by T ; (2) otherwise, i.e. the output of P_m is same as P , this leads to either (2.1) P_m is *equivalent* to P , which means that they are syntactically different but functionally equivalent; or (2.2) T is not adequate to detect the mutants, which requires test case augmentation.

The result of mutation testing can be represented as the *mutation score* (also referred as mutation coverage or mutation adequacy), which is defined as:

$$\text{mutation score} = \frac{\# \text{killed mutants}}{\# \text{nonequivalent mutants}} \quad (1)$$

From the above equation, we can see that mutant equivalence detection is done before calculating the mutation score, as the denominator explicitly mentions nonequivalent mutants. Budd and Angluin [13] have theoretically proven that deciding the equivalence of two programs is not measurable. Meanwhile, in their systematic literature survey Madeyski et al. [14] have also indicated that the equivalent mutant problem takes an enormous amount of time in practice.

A mutation testing system in large part can be regarded as a language system [15] since the programs under test must be parsed, modified and executed. The main components of mutation testing consist of the mutant creation engine, the equivalent mutant detector and the test execution runner. The first prototype of a mutation testing system for Fortran was proposed by Budd and Sayward [16] in 1977. Since then, numerous mutation tools have been developed for different languages, such as Mothra [17] for Fortran, Proteum [18] for C, Mujava [19] for Java, and SQLMutation [20] for SQL.

2.1.3 Benefits & Limitations

Mutation testing is widely considered as a “high end” test criterion [15]. This is in part due to the fact that mutation testing is extremely hard to satisfy because of the massive number of mutants. However, many empirical studies found that it is much stronger than other test adequacy criteria in terms of fault exposing capability, e.g. Mathur and Wong [21], Frankl et al. [22] and Li et al. [23]. In addition to comparing mutation testing with other test criteria, there have also been empirical studies comparing real faults and mutants. The most well-known research work on such a topic is by Andrews et al. [24]: they suggest that when using carefully selected mutation operators and after removing equivalent mutants, mutants can provide a good indication of the fault detection ability of a test suite. As a result, we consider the benefits of mutation testing to be:

- better fault exposing capability compare to other test coverage criteria, e.g. all-use
- a good alternative to real faults which can provide a good indication of the fault detection ability of a test suite

The limitations of mutation testing are inherent. Firstly, both the generation and execution of a vast number of mutants are computationally expensive. Secondly, the equivalent mutants detection is also an inevitable stage of mutation testing which is a prominent undeterminable problem, thereby requiring human effort to investigate. We thus consider the major limitations of mutation testing to be:

- the high computational cost caused by the tremendous amount of mutants
- the undecidable Equivalent Mutant Problem resulting in the difficulty of fully automating the equivalent mutant analysis

In order to deal with the above two limitations, many efforts have been made to reduce the computational cost and propose heuristic methods to detect equivalent mutants. As for the high computational cost, Offutt and Untch [25] performed a literature review in which they summarised the approaches to reduce computational cost into three strategies: *do fewer*, *do smarter* and *do faster*. These three types were later classified into two classes by Jia and Harman [1]: reduction of the generated mutants and reduction of the execution cost. Mutant sampling (e.g. [26, 27]), mutant clustering (e.g. [28, 29]) and selective mutation (e.g. [30–32]) are the most well-known techniques for reducing the number of mutants while maintaining efficacy of mutation testing to an acceptable degree. For reduction of the execution expense, researchers have paid much attention to weak mutation (e.g. [33–35]) and mutant schemata (e.g. [36, 37]).

To overcome the Equivalent Mutant Problem, there are mainly three categories classified by Madeyski et al. [14]: (1) detecting equivalent mutants, such as Baldwin and Sayward [38] (using compiler optimisations), Hierons et al. [39] (using program slicing), Martin and Xie [40] (through change-impact analysis), Ellims et al. [41] (using running profile), and du Bousquet and Delaunay [42] (using model checker); (2) avoiding equivalent mutant generation, such as Mresa and Bottaci [31] (through selective mutation), Harman et al. [43] (using program dependence analysis), and Adamopoulos et al. [44] (using co-evolutionary search algorithm); (3) suggesting equivalent mutants, such as bayesian learning [45], dynamic invariants analysis [46], and coverage change examination (e.g. [47]).

2.2 Historical Overview

In this part, we are going to present a chronological overview of important research in the area of mutation testing. As the focus of our review is the application perspective of mutation testing, we mainly address the studies that concern the application of mutation testing. In the following paragraphs, we will first give a brief summary of the development of mutation testing, and — due to the sheer size of the research body — we will then highlight some notable studies on applying mutation testing.

Mutation testing was initially introduced as a fault-based testing method which was regarded significantly better in detecting errors than the *covering measure* approach [48]. Since then mutation testing has been actively investigated and studied thereby resulting in remarkable advances in its concepts, theory, technology and empirical evidence. The main interests on mutation testing includes (1) defining mutation operator [49], (2) developing mutation testing systems [17, 19, 33], (3) reducing the cost of mutation testing [30, 36], (4) overcoming the equivalent mutant detection [14], and (5) empirical studies with mutation testing [24]. For more literature on mutation testing, we refer to the existing surveys of DeMillo [50], Offutt and Untch [25], Jia and Harman [1] and Offutt [2].

In the meanwhile, mutation testing has also been applied to support other testing activities, such as test data generation and test strategy evaluation. The early application of mutation testing can be traced back to the 1980s [51–54]). Ntafos is one of the very first researchers to use mutation analysis as a measure of test set effectiveness. Ntafos applied mutation operators (e.g. constant replacement) to the source code of 14 Fortran programs [52]. The generated test suites were based on three test strategies, i.e. random testing, branch testing and data-flow testing, and were evaluated regarding mutation score.

DeMillo and Offutt [35] are the first to automate test data generation guided by fault-based testing criteria. Their method is called Constraint-based testing (CBT). They transformed the conditions under which mutants will be killed (necessity and sufficiency condition) to the corresponding algebraic constraints (using constraint template table). The test data then was automatically generated by solving the constraint satisfaction problem using heuristics. Their proposed constraint-based automatic test data generator is limited and was only validated on five laboratory-level Fortran programs. Other remarkable approaches of the automatic test data generation includes a paper by Zhang et al. [55], who adopted Dynamic Symbolic Execution, and a framework by Papadakis and Malevris [56] in which three techniques,

i.e. Symbolic Execution, Concolic testing and Search-based testing, were used to support the automatic test data generation.

Apart from test data generation, mutation testing is widely adopted to assess the cost-effectiveness of different test strategies. The work above by Ntafos [52] is one of the early studies on applying mutation testing. Recently, there has been a considerable effort in the empirical investigation of structural coverage and fault-finding effectiveness, including Namin and Andrews [57] and Inozemtseva et al. [58]. Also of interest is assertion coverage proposed by Zhang and Mesbah [59] and observable modified condition/decision coverage (OMC/DC) presented by Whalen et al. [60]; these novel test criteria were also evaluated via mutation testing.

Test case prioritisation is one of the practical approaches to reducing the expense of regression testing by rescheduling test cases to expose the faults more quickly. Mutation testing has also been applied in supporting test case prioritisation. Among these studies are influential papers by Rothermel et al. [61] and Elbaum et al. [62] who proposed a new test case prioritisation method based on the rate of mutants killing. Moreover, Do and Rothermel [63, 64] measured the effectiveness of different test case prioritisation strategies via mutation faults, since Andrews et al. [24]'s empirical study suggested that mutation faults can be representative of real faults.

The test-suite reduction is another test activity we identified which is supported by mutation testing. The research work of Offutt et al. [65] is the first to target test-suite reduction strategies especially for mutation testing. They proposed *Ping-Pong* reduction heuristics to select test cases based on their mutation scores. Another notable work is Zhang et al. [66] that investigated test-suite reduction techniques on Java programs with real-world JUnit test suites via mutation testing.

Another portion of the application of mutation testing is debugging, such as fault localisation. Influential examples include an article by Zhang et al. [67] in which mutation analysis is adopted to investigate the effect of coincidental correctness upon coverage-based fault localiser, and a novel fault localisation method by Papadakis et al. [68], [69] who used mutants to identify the faulty program statements.

3 RESEARCH METHOD

In this section, we describe the main procedures we took to conduct this review. We adopted the methodology of the systematic literature review. A systematic literature review [6] is a means of aggregating and evaluating all the related primary studies under a research scope in an unbiased, thorough and trustworthy way. Unlike the general literature review, systematic literature review aims to eliminate bias and incompleteness through a systematic mechanism [70]. Kitchenham [6] presented comprehensive and reliable guidelines for applying the systematic literature review to the field of software engineering. The guidelines cover three main phases: (i) planning the review, (ii) conducting the review, and (iii) reporting the review. Each step is well-defined and well-structured. By following these guidelines, we can reduce the likelihood of generating biased conclusions and sum all the existing evidence in a manner that is fair and seen to be fair.

The principle of the systematic literature review [71] is to convert the information collection into a systematic research study; this research study first defines several specific research questions and then searches for the best answers accordingly. These research questions and search mechanisms (consisting of study selection criteria and data extraction strategy) are included in a review protocol, a detailed plan to perform the systematic review. After developing the review protocol, the researchers need to validate this protocol for further resolving the potential ambiguity.

Following the main stages of the systematic review, we will introduce our review procedure in four parts: we will first specify the research questions, and then present the study selection strategy and data extraction framework. In the fourth step, we will show the validation results of the review protocol. The overview of our systematic review process is shown in Figure 1.

3.1 Research Questions

The research questions are the most critical part of the review protocol. The research questions determine study selection strategy and data extraction strategy. In this review, our objective is to examine the primary applications of mutation testing and identify the problems and gaps, therefore, we can provide guidelines for applying mutation testing and recommendations for future work. To achieve these goals and starting with our most vital interest, the application perspective of mutation testing, we naturally further divide it

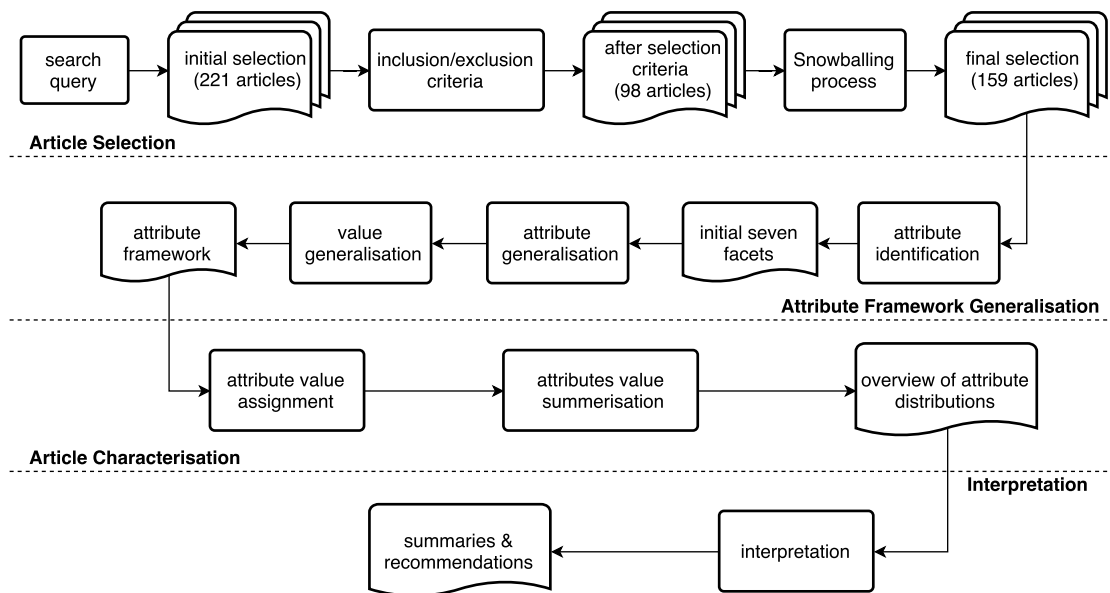


Figure 1. Overview of Systematic Review Process [72]

into two aspects: (1) how mutation testing is used and (2) how the relevant empirical studies are reported. For the first aspect, we aim to identify and classify the main applications of mutation testing:

RQ1: How is mutation testing used in testing activities?

In order to understand how mutation testing is used, we should first determine in which circumstances it is used. The usages might range from using mutation testing as a way to assess how other testing approaches perform or mutation testing might be a building block of an approach altogether. This leads to RQ1.1:

RQ1.1: *Which role does mutation testing play in testing activities?*

There is a broad range of specific testing activities in which mutation testing can be of help, e.g. fault localisation, test data generation, etc. RQ1.2 seeks to uncover these activities.

RQ1.2: *Which testing activities does mutation testing usually support?*

In Jia and Harman's survey [1] of mutation testing, they found that most approaches work at the unit testing level. In RQ1.3 we will investigate whether the application of mutation testing is also mostly done at the unit testing level, or whether other levels of testing are also considered important.

RQ1.3: *Which test level does mutation testing usually target?*

Jia and Harman [1] have also indicated that mutation testing is most often used in a white box testing context. In RQ1.4 we explore what other strategies can also benefit from the application of mutation testing.

RQ1.4: *Which testing strategies does mutation testing frequently support?*

For the second aspect, we are going to synthesise empirical evidence related to mutation testing:

RQ2: How are empirical studies related to the mutation testing designed and reported?

A plethora of mutation testing tools exist and have been surveyed by Jia and Harman [1]. Little is known which ones are most applied and why these are more popular. RQ2.1 tries to fill this knowledge gap by providing insight into which tools are used more frequently in a particular context.

RQ2.1: *Which mutation tools have been frequently used?*

The mutation tools that we surveyed implement different mutation operators. Also, the various mutation approaches give different names to virtually the same mutation operators. RQ2.2 explores what mutation operators each method or tool has to offer and how mutation operators can be compared.

RQ2.2: *Which mutation operators have been used more frequently?*

The equivalent mutant problem, i.e. the situation where a mutation leads to change that is not observable in behaviour, is one of the most significant open issues in mutation testing. Both Jia and Harman [1] and Madeyski et al. [14] highlighted some of the most remarkable achievements in the area, but we have a lack of knowledge when it comes to how the equivalent mutant problem is coped with during the actual application of mutation testing. RQ2.3 aims to seek answers for exactly this question.

RQ2.3: *Which approaches are used to overcome the equivalent mutant problem more often when applying mutation testing?*

As mutation testing is computationally expensive, techniques to reduce costs are important. Selective Mutation and Weak Mutation are the most widely studied cost reduction techniques [1], but it is unclear which reduction techniques are actually used when applying mutation testing, which is the exact topic of RQ2.4.

RQ2.4: *Which techniques are used to reduce the computational cost more frequently when applying mutation testing?*

In order to better understand in which context mutation testing is applied, we want to look into the programming languages that have been used in the experiments. But also the size of the case study systems is of interest, as it can be an indication of the maturity of certain tools. Finally, we are also explicitly looking at whether the case study systems are available for replication purposes (in addition to the check for availability of the mutation testing tool in RQ2.1).

RQ2.5: *What are the most common subjects used in the experiments (in terms of programming language, size and data availability)?*

3.2 Study Selection Strategy

Initial Study Selection:

We started with searching queries in online platforms, including Google Scholar, Scopus, ACM Portal, IEEE explore as well as Springer, Wiley, Elsevier Online libraries, to collect papers containing the keywords “mutation testing” or “mutation analysis” in their titles, abstracts and keywords. Meanwhile, to ensure the high quality of the selected papers, we only considered the articles published in seven top journals and ten top conferences (as listed in Table 1) dating from 1971 as data sources. The above 17 venues are chosen because we regard them to be the leaders in Software Engineering and the major venues that report a high proportion of research on software testing, debugging, software quality and validation. Moreover, we excluded article summaries, interviews, reviews, workshops, panels and poster sessions from the search. If the paper’s language is not English or its full-text is not available, we also excluded such a paper. After this step, 221 papers were initially selected.

Inclusion/Exclusion Criteria:

Since we are interested in how mutation testing is *applied* in practice, we need selection criteria to include the papers that use mutation testing as a tool for evaluating or improving other testing activities and exclude the papers focusing on the development of mutation tools and operators, or challenges and open issues for mutation analysis. Moreover, the selected articles should also provide sufficient evidence for answering the research questions. Therefore, we define two inclusion/exclusion criteria for study selection. The inclusion/exclusion criteria are as follows:

1. The article must focus on the supporting role of mutation testing in testing activities. This criterion excludes the research *solely* on mutation testing itself, such as defining mutation operators, developing mutation systems, investigating ways to solve open issues related to mutation analysis and comparisons between mutation testing and other testing techniques.
2. The article exhibits sufficient evidence that mutation testing is used to support testing related activities. The sufficient evidence means that the article must clearly describe how the mutation

| Type | Venue Name | No. of papers After Applying Search queries | No. of papers After Applying In./Ex. Criteria | No. of papers After Snowballing procedure | |
|------------|--|---|---|---|-----|
| Journal | Transaction on Software Engineering (TSE) | 19 | 9 | 19 | |
| | Journal of Empirical Software Engineering (EMSE) | 4 | 3 | 5 | |
| | Journal on Software Testing, Verification and Reliability (STVR) | 33 | 16 | 19 | |
| | Journal Software Maintenance and Evolution (JSME) | 0 | 0 | 0 | |
| | Transaction on Reliability (TR) | 1 | 1 | 1 | |
| | Transaction on Software Engineering and Methodology (TOSEM) | 3 | 2 | 3 | |
| | Journal of Systems and Software (JSS) | 17 | 8 | 8 | |
| | Information and Software Technology (IST) | 0 | 0 | 2 | |
| | Software Quality Journal (JSQ) | 0 | 0 | 2 | |
| Conference | International Conference on Software Engineering (ICSE) | 29 | 9 | 16 | |
| | European Software Engineering Conference / International Symposium on the Foundations of Software Engineering (ESEC/FSE) | 6 | 1 | 8 | |
| | International Conference on Software Testing, Verification, Validation (ICST) | 46 | 24 | 21 | |
| | International Symposium on Software Testing and Analysis (ISSTA) | 14 | 3 | 9 | |
| | International Conference on Automated Software Engineering (ASE) | 7 | 3 | 6 | |
| | International Conference on Software Maintenance and Evolution (ICSME/ICSM) | 6 | 3 | 9 | |
| | International Symposium on Empirical Software Engineering and Measurement (ESEM/ISESE) | 2 | 1 | 3 | |
| | Proceedings International Symposium on Search Based Software Engineering (SSBSE) | 0 | 0 | 0 | |
| | Proceedings International Conference on Quality Software (QSIC) | 8 | 5 | 6 | |
| | International Symposium on Software Reliability Engineering (ISRE) | 26 | 10 | 20 | |
| | Proceedings Asia Pacific Software Engineering Conference (APSEC) | 0 | 0 | 1 | |
| | Proceedings of the International Conference on Testing Computer Software (TCS) | 0 | 0 | 1 | |
| | Total | | 221 | 98 | 159 |

Note: the venues marked in **bold** font are not initially selected, but were added after the Snowballing procedure.

Table 1. Venues Involved in study selection

testing is involved in the testing activities. The author(s) must state at least one of the following details about the mutation testing in the article: mutation tool, mutation operators, mutation score. This criterion also excludes theoretical studies on mutation testing.

The first author then carefully read the titles and abstracts to check whether the papers in the initial collection belong to our set of selected papers based on the inclusion/exclusion criteria. If it is unclear from the titles and abstracts whether mutation testing was applied, the entire article especially the experiment part was read as well. After we have applied the inclusion/exclusion criteria, 98 papers remained.

Snowballing Procedure:

After selecting 98 papers from digital databases and applying our selection criteria, there is still a high potential to miss articles of interest. According to Brereton et al. [71]'s work, they pointed out that most online platforms do not provide adequate support for systematic identification of relevant papers. To overcome this shortfall of online databases, we then adopted a both backward and forward snowballing strategies [73] to find missing papers. Snowballing refers to using the list of references in a paper or the citations to the paper to identify additional papers [73]. Using the references and the citations respectively is referred to as backward and forward snowballing [73].

We used the 98 papers as the start set and performed a backward and forward snowballing procedure recursively until no further papers could be added to our set. During the snowballing procedure, we extended the initially selected venues to minimise the chance of missing related papers. The snowballing process resulted in another 61 articles (and four additional venues).

3.3 Data Extraction Strategy

Data extracted from the papers are used to answer the research questions we proposed. Based on our research questions, we draw seven facets of interest that are highly relevant to the information we need to answer the questions. The seven facets are: (1) the roles of mutation testing in testing activities;

(2) the testing activities; (3) the mutation tools used in experiments; (4) the mutation operators used in experiments; (5) the description of equivalent mutant problem; (6) the description of cost reduction techniques for mutation testing; (7) the subjects involved in experiments. An overview of these facets is given in Table 2.

For each facet, we first just read the corresponding details in each paper and extracted the exact text from the papers. During the reading procedure, we started to identify and classify more specific attributes of interest under each facet and assigned values to each attribute. The values of each attribute were generalised and modified during the reading process: we combined some values together or divided one into several smaller groups. In this way, we generated an attribute framework, and then we used the framework to characterise each paper. Therefore, we can show quantitative results for each attribute to support our answers. And the attribute framework can also be further used for validation and replication of the review work.

(1) roles of mutation testing in testing activities:

The first facet concerns the role of mutation testing in testing activities drawn from RQ1.1. We identified two classes for the function of mutation testing: assessment and guide. When mutation testing is used as a measure of test effectiveness concerning fault-finding capability, we classify this role as “assessment”. While for the “guide” role, mutation testing is adopted to improve the testing effectiveness as guidance, i.e., it is an inherent part of an approach.

To identify and classify the role of mutation testing, we mainly read the description of mutation testing in *experiment part* of each paper. If we find the phrases which have the same meanings as “evaluate fault-finding ability” or “assess the testing effectiveness” in a paper, we then classify the paper into the class of “assessment”. In particular, when used as a measure of testing effectiveness, mutation testing is usually conducted at the end of the experiment; this means mutation testing is not involved in the generation or execution of test suites. Unlike the “assessment” role, if mutation testing is adopted to help to generate test suites or run test cases, we then classify these paper into the “guide” set. In this case, mutation testing is not used in the final step of the experiment.

(2) testing activities:

The second facet focuses on testing activities. Three attributes are relevant to testing activities: the categories of testing activities (RQ1.2), test levels (RQ1.3) and testing strategies (RQ1.4). To identify the categories of testing activities, we group similar testing activities based on information in title and abstract. The testing activities we identified so far consist of 11 classes: test data generation, test-suite reduction/selection, test strategy evaluation, test case minimisation, test case prioritisation, test oracle, fault localisation, programming repairing, development scheme evaluation, model clone detection and model review. We classify the papers by reading the description appeared in *title and abstract*.

For test level, the values are based on the concept of test level and the authors’ specification. More precisely, we are considering five test levels: unit, integration, system, others and n/a. To characterise the test level, we search the exact words “unit”, “integration”, “system” in the article, as these four test levels are regular terms and cannot be replaced by other synonyms. If there is no relevant result after searching in a paper, we then classify the paper’s test level into “n/a”, i.e. no specification regarding the test level. In addition, for the paper which is difficult for us to categorise into any of the four phases, such as testing of the grammar of a programming language and spreadsheet testing, we mark this situation as “others”.

For testing strategies, coarse-grained classification is adequate to gain an overview of the distribution of testing strategies. We identified five classes according to the test design techniques: structural testing, specification-based testing, similarity-based testing, hybrid testing and others [74, 75]. Among the structural testing and specification-based testing classes, we further divided into traditional version and enhanced one based on whether the regular testing is improved by other methods.

To be classified into the “structural testing” class, the paper should either contain the keywords of “structure-based”, “code coverage-based” or “white box”, or use structural test design techniques, such as statement testing, branch testing and condition testing. For the “specification-based testing” class, the articles should either contain the keywords of “black box”, “requirement-based” or “specification-based”, or use specification-based test design techniques, such as equivalence partitioning, boundary value analysis, decision tables and state transition testing. The similarity-based method aims to maximise the diversity of test cases to improve the test effectiveness; this technique is mainly based on test case relationship rather than software artefacts. Therefore, similarity-based testing does not belong to either structural testing or specification-based testing. The hybrid testing combines structural testing and specification testing

together. Besides, several cases are using static analysis, code review or other techniques which cannot fit the above classes; in such situations, we mark the value as “others”.

Furthermore, to classify “enhanced” version of structural and specification-based testing we rely on whether other testing methods were adopted to improve the traditional testing. For instance, Whalen et al. [60] combined the MC/DC coverage metric with a notion of *observability* to ensure the fault propagation conditions. Papadakis and Malevris [76] proposed an automatic mutation test case generation via dynamic symbolic execution. To distinguish such instances from the traditional structural and specification-based testing, we marked as “enhanced”.

(3) mutation tools used in experiment:

For the mutation tools (derived from RQ2.1), we are interested in their types, but also in their availability. Our emphasis on tool availability is instigated to address possible replication of the studies. The values of “Yes” or “No” for the tool availability depends on whether the mutation tool is open to the public. The tool type intends to provide further analysis of mutation tool, which is based on the whether the tool is self-developed and whether the tool itself is a complete mutation testing system. We identified five types of mutation tools: existing, partially-based, self-written, hand-seeded and n/a. The “existing” tool must be a complete mutation testing system, while “partially-based” means these tools are used as a base or a framework for mutation testing. The example for “partially-based” tools are EvoSuite, jFuzz, TrpAutoRepair and GenProg. The self-written tool category represents those tools that have been developed by the authors of the study. The “hand-seeded” value means the mutants were generated manually in the studies. Besides, we defined “n/a” value in addition to the “tool types” attribute; the value of “n/a” marks the situation where lacks of a description of mutation tools including tool names/citations and whether hand-seeded or not.

(4) mutation operators used in experiment:

As for the mutation operators (related to RQ2.2), we focus on two attributes: description level and generic classification. The former is again designed to assess the repeatability issue related to mutation testing. The description degree depends on the way that the authors presented the mutation operators used in their studies, consisting of three values: “well-defined”, “not sufficient” and “n/a”. If the paper showed that the complete list of mutation operators is available, then we classify such paper into “well-defined”. The available full list includes two main situations: (1) the authors listed each name(s) of mutation operators and/or specified how the mutation operators make changes to programs in the articles; (2) the studies adopted existing tools and mentioned the used mutation operator (including the option were all or the default set of mutation operators provided by that tool were used). The well-defined category thus enables the traceability of the complete list of mutation operators. For the remaining set, the incomplete list, if there is some information about the mutation operators in the article but not enough for replication of the whole list of mutation operators, then we classify the paper into “not sufficient”. The typical example is that the author used such words as “etc.”, “such as” or “e.g.” in the specification of the mutation operators; this indicates that only some mutation operators are explicitly listed in the paper, but not all. The last value, “n/a”, means no description of the mutation operators was given in the paper at all.

In order to compare the mutation operators from different tools to analyse the popularity of involved mutation operators amongst the papers, we collected the information about mutation operators mentioned in the articles. Notably, we only consider the articles which are classified as “well-defined”. We excluded the papers with “not sufficient” label as their lists of mutation operators are not complete as this might result in biased conclusions based on incomplete information. Moreover, during the reading process, we found that different mutation testing tools use slightly different names for their mutation operators. For example, in MuJava [19], the mutation operator which replaces relational operators with other relational operators is called “Relational Operator Replacement”, while that is named “Conditionals Boundary Mutator” in PIT [77]. Therefore, we saw a need to compose a generic classification of mutation operators, which enables us to more easily compare mutation operators from different tools or definitions.

The method we adopted here to generate the generic classification is to group the similar mutation operators together among all the existing mutation operators in the literature based on how they mutate the programs. Firstly, mutation testing can be applied to both program source code and program specification. Thus, we classified the mutation operators into two top-level groups: program mutation and specification mutation operators. In particular, we are more interested in the program mutation, so we further divided program mutation testing into three sub-categories: expression-level, statement-level and others. The expression-level mutation operators focus on the inner components of the statements, such as operators

and operands, while the statement-level ones mutate the at least one single statement. For the “others” class, it includes mutation operators related to the programming language’s unique features, e.g. Objected-Oriented specific mutation operators. Our generic classification of mutation operators is as follows:

1. Specification mutation

2. Program mutation

- (a) Expression-level

- i. **arithmetic operator:** it mutates the arithmetic operators (including addition “+”, subtraction “-”, multiplication “*”, division “/”, modulus “%”, unary operators “+”, “-”, and short-cut operators “++”, “--”) ¹ by replacement, insertion or deletion.
 - ii. **relational operator:** it mutates the relational operators (including “>”, “>=”, “<”, “<=”, “==”, “!=”) by replacement.
 - iii. **conditional operator:** it mutates the conditional operators (including and “&”, or “|”, exclusive or “^”, short-circuit operator “&&”, “||”, and negation “!”) by replacement, insertion or deletion.
 - iv. **shift operator:** it mutates the shift operators (including “>>”, “<<” and “>>>”) by replacement.
 - v. **bitwise operator:** it mutates the bitwise operators (including bitwise and “&”, bitwise or “|”, bitwise exclusive or “^” and bitwise negation “~”) by replacement, insertion or deletion.
 - vi. **assignment operator:** it mutates the assignment operators (including the plain operator “=” and short-cut operators “+=”, “-=”, “*=”, “/=”, “%=”, “&=”, “|=”, “^=”, “<<=”, “>>=”, “>>>=”) by replacement. Besides, the plain operator “=” is also changed to “==” in some cases.
 - vii. **absolute value:** it mutates the arithmetic expression by preceding unary operators including ABS (computing the absolute value), NEGABS (compute the negative of the absolute value) and ZPUSH (testing whether the expression is zero. If the expression is zero, then the mutant is killed; otherwise execution continues and the value of the expression is unchanged) ².
 - viii. **constant:** it changes the literal value including increasing/decreasing the numeric values, replacing the numeric values by zero or swapping the boolean literal (true/false).
 - ix. **variable:** it substitutes a variable with another already declared variable of the same type and/or of the compatible type.
 - x. **type:** it replaces a type with the other compatible types including type casting. ³
 - xi. **conditional expression:** it replaces the conditional expression by true/false so that the statements following the conditional always execute or skip.
 - xii. **parenthesis:** it changes the precedence of the operation by deleting, adding or removing the parentheses.

- (b) Statement-level

- i. **return statement:** it mutates return statement in the method calls including return value replacement or return statement swapping.

¹The syntax of these operators might vary slightly in different languages. Here we just used the operators in Java as an example. So as the same in (ii) - (vi) operators.

²The definition of this operator is from the Mothra [17] system. In some cases, this operator only applies the absolute value replacement.

³The types of the variables varies in different programming languages.

³The changes between the objects of the parent and the child are excluded which belongs to “OO-specific”

- ii. **switch statement:** it mutates `switch` statements by making different combinations of the `switch` labels (`case/default`) or the corresponding block statement.
- iii. **if statement:** it mutates `if` statements including removing additional semicolons after conditional expressions, adding an `else` branch or replacing last `else if` symbol to `else`.
- iv. **statement deletion:** it deletes statements including removing the method calls or removing each statement⁴.
- v. **statement swap:** it swaps the sequences of statements including rotating the order of the expressions under the use of the `comma` operator, swapping the contained statements in `if-then-else` statements and swapping two statements in the same scope.
- vi. **brace:** it moves the closing brace up or down by one statement.
- vii. **goto label:** it changes the destination of the `goto` label.
- viii. **loop trap:** it introduces a guard (trap after n^{th} loop iteration) in front of the loop body. The mutant is killed if the guard is evaluated the n^{th} time through the loop.
- ix. **bomb statement:** it replaces each statement by a special `Bomb()` function. The mutant is killed if the `Bomb()` function is executed which ensures each statement is reached.
- x. **control-flow disruption (break/continue):** it disrupts the normal control flow by adding, removing, moving or replacing `continue/break` labels.
- xi. **exception handler:** it mutates the exception handlers including changing the `throws`, `catch` or `finally` clauses.
- xii. **method call:** it changes the number or position of the parameters/arguments in a method call, or replace a method name with other method names that have the same or compatible parameters and result type.
- xiii. **do statement:** it replaces `do` statements with `while` statements.
- xiv. **while statement:** it replaces `while` statements with `do` statements.

(c) Others

- i. **OO-specific:** the mutation operators related to O(bject)-O(riented) Programming features [78], such as Encapsulation, Inheritance and Polymorphism, e.g. `super` keyword insertion.
- ii. **SQL-specific:** the mutation operators related to SQL-specific features [20], e.g. replacing `SELECT` to `SELECT DISTINCT`.
- iii. **Java-specific**⁵: the mutation operators related to Java-specific features [78] (the operators in Java-Specific Features), e.g. `this` keyword insertion.
- iv. **JavaScript-specific:** the mutation operators related to JavaScript-specific features [79] (including DOM, JQUERY, and XMLHttpRequest operators), e.g. `var` keyword deletion.
- v. **SpreadSheet-specific:** the mutation operators related to SpreadSheet-specific features [80], e.g. changing range of cell areas.
- vi. **AOP-specific:** the mutation operators related to A(spect)-O(riented)-P(rogramming) features [81, 82], e.g. removing pointcut.
- vii. **concurrent mutation:** the mutation operators related to concurrent programming features [83, 84], e.g. replacing `notifyAll()` with `notify()`.
- viii. **Interface mutation:** the mutation operators related to Interface-specific features [85, 86], suitable for use during integration testing.

⁴To maintain the syntactical validity of the mutants, semicolons or other symbols, such as `continue` in Fortran, are retained.

⁵This set of mutation operators originated from Java features but not limited to Java language, since other languages can share certain features, e.g., `this` keyword is also available in C++ and C#, and `static` modifier is supported by C and C++ as well.

(5) description of the equivalent mutant problem & (6) description of cost reduction techniques for mutation testing:

The fifth and sixth facets aim to show how the most significant problems are coped with when applying mutation testing (related to RQ2.3 and RQ2.4 respectively). We composed the list of techniques based on both our prior knowledge and the descriptions given in the papers. We identified seven methods for dealing with the equivalent mutant problem and five for reducing computational cost except for “n/a” set (more details are given in Table 2).

For the equivalent mutant problem, we started by searching the keywords “equivalen*” and “equal” in each paper to target the context of equivalent mutants issue. Then we extracted the corresponding text from the articles. If there are no relevant findings in a paper, we mark this article as “n/a” which means the authors did not mention how they overcame the equivalent mutant problem. Here it should be noted that we only considered the description related to the equivalent mutant problem given by the authors; this means we excluded the internal heuristic mechanisms adopted by the existing tools if the author did not point out such internal approaches. For example, the tool of JAVALANCHE [87] ranks mutations by impact to help users detect the equivalent mutants. But if the authors who used JAVALANCHE did not specify that internal feature, we will not label the paper into the class that used the approach of “ranking the mutations”.

For the cost reduction techniques, we read the experiment part carefully to extract the reduction mechanism from the papers. In addition, we excluded the runtime optimisation and selective mutation. Because the former one, runtime optimisation, is an internal optimisation adopted during the tool implementation, thereby such information is more likely to be reported in the tool documentation. We did not consider the runtime optimisation to avoid incomplete statistics. As for the second one, selective mutation, we assume it is adopted by all papers since it is nearly impossible to implement and use all the operators in practice. If a paper does not contain any description of the reduction methods in the experiment part, we mark this article as “n/a”.

(7) subjects involved in the experiment:

For the subject programs in the evaluation part, we are interested in three aspects: programming language, size and data availability. From the programming language, we can obtain an overall idea of how established mutation testing is in each programming language domain and what the current gap is. From the subject size, we can see the scalability issue related to mutation testing. From the data availability situation, we can assess the replicability of the studies.

For the programming language, we extracted the programming language of the subjects involved in the experiment in these articles, such as Java, C, SQL, etc. If the programming language of the subject programs is not clearly pointed out, we mark it as “n/a”. Note, more than one languages might be involved in a single experiment.

For the subject size, we defined four categories according to the lines of code (LOC): preliminary, small, medium and large. If the subject size is less than 100 LOC, then we classify it into the “preliminary” category. If the size is between 100 to 10K LOC, we consider it “small”, while between 10K and 1M LOC we appraised it as “medium”. If the size is greater than 1M LOC, we consider it as “large”. Since our size scale is based on LOC, if the LOC of the subject is not given, or other metrics are used, we mark it as “n/a”. To assign the value to a paper, we always take the biggest subjects used in the papers.

For the data available, we defined two classes: Yes and No. “Yes” means *all* subjects in the experiments can be openly accessible; this can be identified either from the keywords “open source”, SIR [88], GitHub⁶, SF100 [89] or SourceForge⁷, or from the open link provided by the authors. It is worth noting that if one of the subjects used in a study is not available, we classify the paper into “No”.

The above facets of interest and corresponding attributes and detailed specification of values are listed in Table 2.

| Facet | Attribute | Value | Description |
|--------------------|----------------|--|--|
| Roles | classification | assessment guide | assessing the fault-finding effectiveness improving other testing activities as guidance |
| Testing Activities | category | test data generation test-suite reduction/selection | creating test input data reducing the test suite size while maintaining its fault detection ability |

⁶<https://github.com/>

⁷<https://sourceforge.net/>

| | | | |
|---------------------|------------------------|---|---|
| | | test strategy evaluation | evaluating test strategies by carrying out the corresponding whole testing procedure, including test pool creation, test case selection and/or augmentation and testing results analysis. |
| | | test-case minimisation | simplifying the test case by shortening the sequence and removing irrelevant statements |
| | | test case prioritisation | reordering the execution sequence of test cases |
| | | test oracle | generating or selecting test oracle data |
| | | fault localisation | identifying the detective part of a program given the test execution information |
| | | program repairing | generating patches to correct detective part of a program |
| | | development scheme evaluation | evaluating the practice of software development process via observational studies or controlled experiments, such as Test driven development (TDD) |
| | | model clone detection | identifying similar model fragments within a given context |
| | | model review | determining the quality of the model at specification level using static analysis techniques |
| | test level | unit | testing activities focus on unit level. Typical example of unit testing includes: using unit testing tools, such as Junit and Nunit, intra-method testing, intra-class testing. |
| | | integration | testing activities focus on integration level. Typical example of integration testing includes: caller/callee and inter-class testing |
| | | system | testing activities focus on system level. Typical examples of system testing includes: high-level model-based testing techniques and high-level specification abstraction methods |
| | | others | testing activities are not related to source code. Typical example includes: grammar. |
| | | n/a | no specification about the testing level in the article. |
| | testing strategy | structural | white-box testing, uses the internal structure of the software to derive test cases, such as statement testing, branch testing and condition testing |
| | | enhanced structural | adopting other methods to improve the traditional structural testing, mutation-based techniques, information retrieval knowledge, observation notations and assertion coverage |
| | | specification-based | viewing software as a black box with input and output, such as equivalence partitioning, boundary value analysis, decision tables and state transition testing |
| | | enhanced specification-based | adopting other methods to improve the traditional specification-based testing, such as mutation testing. |
| | | similarity-based | maximising the diversity of test cases to improve the test effectiveness |
| | | hybrid | combining structural testing and specification testing together |
| | | others | using static analysis, or focusing on other testing techniques which cannot fit in above six classes |
| Mutation Tools | availability | Yes/No | Yes: open to the public; No: no valid open access |
| | type | existing tool partially-based self-written | a complete mutation testing system used as a base or framework for mutation testing developed by the authors and the open link of the tool is also accessible |
| | | hand-seeded | generating mutants manually based on the mutation operators |
| | | n/a | no description of the adopted mutation testing tool |
| Mutation Operators | description Level | well-defined not sufficient | the complete list of mutation operators is available the article provides some information about mutation operators but the information is not enough for replication |
| | | n/a | no description of the mutation operators |
| | generic classification | refer to Section 3.3 (4) | refer to Section 3.3 (4) |
| Equivalence Solver | methods | not killed as equivalent not killed as nonequivalent no investigation | treating mutants not killed as equivalent treating mutants not killed as nonequivalent no distinguishing between equivalent mutants and nonequivalent ones |
| | | manual model checker | manual investigation using model checker to remove functionally equivalent mutants |
| | | reduce likelihood | generating mutants that are less likely to be equivalent, such as using behaviour-affecting variables, carefully-designed mutation operators and constraints binding |
| | | deterministic model | adopting the deterministic model to make the equivalence problem decidable |
| | | n/a | no description of mutant equivalence detector |
| Reduction Technique | methods | mutant sample | randomly select a subset of mutants for testing execution based on fixed selection ratio |
| | | fixed number | select a subset of mutants based on fixed number |

| | | | |
|--------------|----------------|---|--|
| | | weak mutation | compare internal state of the mutant and the original program immediately after the mutated statement(s) |
| | | higher-order | reduce the number of mutants by selecting higher-order mutants which contains more than one faults |
| | | selection strategy | generate less mutants by selecting where to mutate based on random algorithm or other techniques |
| | | n/a | no description of reduction techniques (except for runtime optimisation and selective mutation) |
| Subject | language | Java, C, C#, etc. | various programming languages |
| | size (maximum) | preliminary | <100 LOC |
| | | small | 100 ~ 10K LOC |
| medium | | 10K ~ 1M LOC | |
| large | | > 1M LOC | |
| | n/a | no description of program size regarding LOC | |
| availability | Yes/No | Yes: open to the public; No: no valid open access | |

Table 2. Attribute Framework

3.4 Review Protocol Validation

The review protocol is a critical element of a systematic literature review and researchers need to specify and carry out procedures for its validation [71]. The validation procedure aims to eliminate the potential ambiguity and unclear points in the review protocol specification. In this review, we conduct the review protocol validation among the three authors. We also used the results to improve our review protocol. The validation focuses on two things: selection criteria and attribute framework, including the execution of two pilot runs of study selection procedure and data extraction process.

3.4.1 Selection Criteria Validation

We performed a pilot run of the study selection process, for which we randomly generated ten candidate papers from selected venues (including articles out of our selection scope) and carried out the paper selection among three authors independently based on the inclusion/exclusion criteria. After that, the three authors compared and discussed the selection results. The results show that for 9 out of 10 papers, the authors had an immediate agreement. The three authors discussed the one paper that showed disagreement, leading to a revision of the first inclusion/exclusion criterion. In the first exclusion criterion, we added “solely” to the end of the sentence “...This criterion excludes the research on mutation testing itself...”. By adding “solely” to the first criterion, we include articles whose main focus is mutation testing, but also cover the application of mutation testing.

3.4.2 Attribute Framework Validation

To execute the pilot run of the data extraction process, we randomly select ten candidate papers from our selected collection. These 10 papers are classified by all three authors independently using the attribute framework that we defined earlier. The discussion that follows from this process leads to revisions of our attribute framework. Firstly, we clarified that the information extracted from the papers must have exactly the same meaning as described by the authors; this mainly means that we cannot further interpret the information. If the article does not provide any clear clue for a certain attribute, we use the phrase “not specified” (“n/a”) to mark this situation. By doing so, we can minimise the potential misinterpretation of the articles.

Secondly, we ensured that the values of the attribute framework are as complete as possible, so that for each attribute we can always select a value. For instance, when extracting testing activities information from the papers, we can simply choose one or several options of the 11 categories provided by the predefined attribute framework. The purpose of providing all possible values to each attribute is to assist data extraction in an unambiguous and trustworthy manner. Through looking at the definitions of all potential values for each attribute, we can easily target unclear or ambiguous points in data extraction strategy. If there are missing values for certain attributes, we can only add the additional data definition to extend the framework. The attribute framework can also be of clear guideline for future replication. Furthermore, we can then present quantitative distributions for each attribute in later discussion to support our answers to research questions.

To achieve the above two goals, we made revisions to several attributes as well as values. The specified modifications are listed as follows:

| Testing Activity | Assessment | Guide | Total |
|--------------------------------|------------|-------|-------|
| test data generation | 31 | 28 | 59 |
| test strategy evaluation | 55 | 3 | 58 |
| test case prioritisation | 10 | 6 | 16 |
| test oracle | 9 | 4 | 13 |
| test-suite selection/reduction | 9 | 4 | 13 |
| fault localisation | 7 | 4 | 11 |
| program repairing | 1 | 1 | 2 |
| test case minimisation | 1 | 1 | 2 |
| development scheme evaluation | 0 | 1 | 1 |
| model clone detection | 1 | 0 | 1 |
| model review | 1 | 0 | 1 |
| Total | 112 | 49 | 159 |

Table 3. Testing Activities Summary

Mutation Tools: Previously, we combined tool availability and tool types together by defining three values: self-written, existing and not available; this is not clear to distinguish available tools from unavailable ones. Therefore, we further defined two attributes, i.e., tool availability and tool types.

Mutation Operators: We added “description level” to address the interest of how mutation operators are specified in the articles; this also helps in the generalisation of mutation operator classification.

Reduction Techniques: We added the “fixed number” value to this attribute.

Subjects: We changed the values of “data availability” from “open source”, “industrial” or “self-defined” to “Yes” or “No”. Since the previous definitions can not distinguish between available dataset and unavailable ones.

4 REVIEW RESULTS

After developing the review protocol, we conducted the task of article characterisation accordingly. Given the attribute assignment under each facet, we are now at the stage of interpreting the observations and reporting our results. In the following section, we discuss our review results following the sequence of our research questions. While Section 4.1 deals with the observations related to how mutation testing is applied (RQ1), Section 4.2 will present the RQ2-related discussion. For each sub-research question, we will first show the distribution of the relevant attributes and our interpretation of the results (marked as **Observation**). Each answer to a sub-research question is also summarised at the end. All the metadata of all the surveyed papers regarding their characterisations are shown in Table 19-21 in Appendix 6.

4.1 RQ1: How is the mutation testing used in testing activities?

4.1.1 RQ1.1 & RQ1.2: Which role does mutation testing play in each testing activity?

Observation.

We opted to discuss the two research questions RQ1.1 and RQ1.2 together, because it gives us the opportunity to analyse per testing activity (e.g., test data generation) whether mutation testing is used as a way to guide the technique, or whether mutation testing is used as a technique to assess some (new) approach. Consider Table 3, in which we report the role mutation testing plays onto the two columns “Assessment” and “Guide” (see Table 2 for the explanation about our attribute framework), while the testing activities are projected onto the rows. The table is then populated with our survey results, with the additional note that some papers belong to multiple categories.

As Table 3 shows, test data generation and test strategy evaluation occupies the majority of testing activities (accounting for 73.6%, 117 instances). The remaining are test case prioritisation (10.1%), test oracle generation/selection (8.2%) and test suite reduction/selection (8.2%). Only two instances studied test-case minimisation; this shows mutation testing has not been widely used to simplify test cases by shortening the test sequence and removing irrelevant statements.

As the two roles (assessment and guide) are used quite differently depending on the testing activities, we will discuss them separately. Also, for the “guiding” role, for which we see an increasing number of

applications in recent decades, we find a number of hints and insights for future researchers to consider, which explains why we will analyse this part in a more detailed way when compared to the description of mutation testing as a tool for assessment.

(1) *Assessment.*

We observed that mutation testing mainly serves as an assessment tool to evaluate the fault-finding ability of the corresponding test techniques (70.4%) as it is widely considered as a “high end” test criterion [15]. In order to do so, mutation testing typically generates a large number of mutants of a program, which are sometimes also combined with natural defects or hand-seeded ones. The results of the assessment are usually quantified as metrics of fault finding capability: mutation score (or mutation coverage, mutation adequacy) and killed mutants are the most common metrics in mutation testing. Besides, in test-case prioritisation, the Average Percentage of Faults Detected (APFD) [90], which measures the rate of fault detection per percentage of test suite execution, is also popular.

Amongst the papers in our set, we also found 18 studies that performed *mutant analysis*, which means that the researchers are trying to get a better understanding about mutation faults, e.g. which faults are more valuable in a particular context. A good example of this mutant analysis is the hard mutant problem investigated by Czemerinski et al. [91] where they analysed the failure rate for the hard-to-kill mutants (killed by less than 20% of test cases) using the domain partition approach.

(2) *Guide.*

To provide insight into how mutation testing acts as guidance to improve testing methods per test activity, we will highlight the most significant research efforts to demonstrate why mutation testing can be of benefit as a building block to guide other testing activities. In doing so, we hope the researchers in this field can learn from the current achievements so as to explore other interesting applications of mutation testing in the future.

Firstly, let us start with test data generation, which attracts most interest when mutation testing is adopted as a building block (28 instances). The main idea of mutation-based test data generation is to generate test data that can effectively kill mutants. For automatic test data generation, killing mutants serves as a condition to be satisfied by test data generation mechanisms, such as constraint-based techniques and search-based algorithms; in this way, mutation-based test data generation can be transformed into the structural testing problem. The mutation killable condition can be classified into three steps as suggested by Offutt and Untch [25]: reachability, necessity and sufficiency. When observing the experiments contained in the papers that we surveyed (except the model-based testing), we see that with regard to the killable mutant condition most papers (78.5%) are satisfied with a weak mutation condition (necessity), while a strong mutation condition (sufficiency) appears less (28.6%). The same is true when comparing first-order mutants (92.9%) to higher-order mutants (7.1%). Except for the entirely automatic test data generation, Baudry et al. [92–94] focused on the automation of the test case enhancement phase: they optimised the test cases regarding mutation score via genetic and bacteriological algorithms, starting from an initial test suite. von Mayrhauser et al. [95] and Smith and Williams [96] augmented test input data using the requirement of killing as many mutants as possible.

The second and third most-frequent use cases when applying mutation testing to guide the testing efforts come from test case prioritisation (6 instances) and the test strategy evaluation (6 instances). For test case prioritisation, the goal is to detect faults as early as possible in the regression testing process. The incorporation of measures of fault proneness into prioritisation techniques is one of the directions to overcome the limitation of the conventional coverage-based prioritisation methods. As relevant substitutes of real faults, mutants are used to approximate the fault-proneness capability to reschedule the testing sequences. Qu et al. [97] ordered test cases according to prior fault detection rate using both hand-seeded and mutation faults. Kwon et al. [98] proposed a linear regression model to reschedule test cases based on Information Retrieval and coverage information, where the coefficients in the model are determined by mutation testing. Moreover, Rothermel et al. [61, 90] and Elbaum et al. [62] compared different approaches of test-case prioritisation, among which included the prioritisation in order of the probability of exposing faults estimated by the killed mutants information. In Qi et al. [99]’s study, they adopted a similar test-case prioritisation method to improve patch validation during program repairing.

Thirdly, as for the fault localisation (4 instances), the locations of mutants are used to assist the localisation of “unknown” faults (the faults which have been detected by at least one test case, but that have still to be located [68]). The motivation of this approach is based on the following observation:

“Mutants located on the same program statements frequently exhibit a similar behaviour” [68]. Thus, the identification of an “unknown” fault could be obtained thanks to a mutant at the same (or close) location. Taking advantage of the implicit link between the behaviour of “unknown” faults with some mutants, Murtaza et al. [100] used the traces of mutants and prior faults to train a decision tree to identify the faulty functions. Also, Papadakis et al. [68, 69] and Moon et al. [101] ranked the suspiciousness of “faulty” statements based on their passing and failing test executions of the generated mutants.

When it comes to the test oracle problem (4 instances), mutation testing can also be of benefit for driving the generation of assertions, as the prerequisite for killing the mutant is to distinguish the mutant from the original program. In Fraser and Zeller [102]’s study, they illustrated how they used mutation testing to generate test oracles: assertions, as commonly used oracles, are generated based on the trace information of both the unchanged program and the mutants recorded during the executions. First, for each difference between the runs on the original program and its mutants, the corresponding assertion is added. After that, these assertions are minimised to find a sufficient subset to detect all the mutants per test case; this becomes a minimum set covering problem. Besides, Staats et al. [103] and Gay et al. [104] selected the most “effective” oracle data by ranking variables (trace data) based on their killed mutants information.

Mutation-based test-suite reduction (4 instances) relies on the number of killed mutants as a heuristic to perform test-suite reduction, instead of the more frequently used traditional coverage criteria, e.g., statement coverage. The intuition behind this idea is that the reduction based on the mutation faults can produce a better-reduced test suite with less or no loss in fault-detection capability. The notable examples include an empirical study carried out by Shi et al. [105] who compared the trade-offs among various test-suite reduction techniques based on statement coverage and killed mutants.

Zooming in on the test strategy evaluation (3 instances), we observe, on the one hand, the idea of incorporating an estimation of fault-exposure probability into test data adequacy criteria intrigued some researchers. Among them are Chen et al. [106]: in their influential work they examined the fault-exposing potential (FEP) coverage adequacy which is estimated by mutation analysis. Their findings show quite small, but statistically significant increases in the fault-detection ability of FEP-based test suites compared to statement-based ones. On the other hand, a mutation-based test strategy was also investigated and evaluated under the whole testing procedure including test case generation, augmentation and evaluation. A remarkable example includes an observational user study conducted by Smith and Williams [107] with four software testers to explore the cost-effectiveness of mutation testing for manually augmenting test cases. Their results indicate that mutation testing was regarded as an effective but relatively expensive technique for writing new test cases.

From the above guide roles of the testing activities, we can see that mutation testing is mainly used as an indication of the potential defects: either (1) to be killed in test data generation, test case prioritisation, test-suite reduction and test strategy evaluation, or (2) to be suspected in the fault localisation. In most cases, mutation testing serves as *where-to-check* constraints, i.e., introducing a modification in a certain statement or block. In contrast, only four studies applied mutation testing to solving the test oracle problem, which targets the *what-to-check* issue. The *what-to-check* problem is not a problem unique to mutation testing, but rather an inherent challenge of test data generation. As mentioned above, mutation testing can not only help in precisely targeting at *where* to check, but also suggesting *what* to check for [102] (see the first recommendation labeled as **R1** in Section 4.4). In this way, mutation testing could be of benefit to improve the test code quality.

After we had analysed how mutation testing is applied to guide various testing activities, we are now curious to better understand how these mutation-based testing methods were evaluated, especially because mutation testing is commonly used as an assessment tool. Therefore, we summed up the evaluation fault types among the articles labelled as “guide” in Table 4. We can see 37 cases (75.5%), which is the addition of the first and the third rows in Table 4 (33 + 4), still adopted mutation faults to assess the effectiveness. Among these studies, four instances [103, 104, 108, 109] realised the potentially biased results caused by the same set of mutants being used in both guidance and assessment. They partitioned the mutants into different groups and used one for evaluation set. Besides, one study [55] used a different mutation tool while the other [85] adopted different mutation operators to generate mutants intending to eliminate bias. These findings signal an open issue: how to find an adequate fault set instead of mutation faults to effectively evaluate mutation-based testing methods? (see the second recommendation labeled as **R2** in Section 4.4) Although hand-seeded faults and real bugs could be an option, searching for such an

| Evaluation Fault Type | Total |
|-------------------------------|-------|
| mutation faults | 33 |
| hand-seeded faults | 7 |
| hand-seeded + mutation faults | 4 |
| no evaluation | 3 |
| real faults | 2 |

Table 4. Guide Role Summary

| Test Level | Total |
|-------------|-------|
| n/a | 73 |
| unit | 64 |
| integration | 9 |
| other | 9 |
| system | 8 |

Table 5. Test Level Summary

adequate fault set increases the difficulty when applying mutation testing as guidance.

Summary.

Test data generation and test strategy evaluation occupy the majority of testing activities when applying mutation testing (73.6%). Mutation testing mainly serves as an assessment tool to evaluate the fault-finding ability of various testing techniques (70.4%). While as guidance, mutation testing is primarily used in test data generation (28 instances) and test-case prioritisation (6 instances). From the above observations, we draw one open issue and one recommendation for the “guide” role of mutation testing. The open issue is how to find an adequate fault set instead of mutation faults to effectively evaluate mutation-based testing methods. The recommendation is mutation testing can suggest not only *where* to check but also *what* to check. *Where* to check widely used to generate killable mutant constraints in different testing activities; while *what* to check is seldom adopted to improve the test data quality.

4.1.2 RQ1.3: Which test level does mutation testing usually target at?

Observations.

Table 5 presents the summary of the test level distribution across the articles. We observe that the authors of 73 papers do not provide a clear clue about the test level they target (the class marked as “n/a”). This is a clear and open invitation for future investigations in the area to be more clear about the essential elements of testing activities such as the test level. For the remainder of our analysis of RQ1.3, we excluded the papers labelled as “n/a” when calculating percentages, i.e., our working set is 86 (159 – 73) papers.

Looking at Table 5, mutation testing mainly targets the unit testing level (74.4%), an observation which is in accordance with the results in Jia and Harman’s survey [1]. One of the underlying causes for the popularity of the unit level could be the origin of mutation testing. The principle of mutation testing is to introduce small syntactic changes to the original program; this means the mutation operators only focus on small parts of the program, such as arithmetical operators and *return* statements. Thus, such small changes mostly reflect the abnormal behaviour of unit-level functionality.

While unit testing is by far the most observed test level category in our set of papers, higher-level testing, such as integration testing, system testing, can also benefit from the application of mutation testing. Here we highlight several research works as examples: Hao et al. [110] and Do and Rothermel [64] used the programs with system test cases as their subjects in case studies. Hou et al. [85] studied interface-contract mutation in support of integration testing under the context of component-based software. Li et al. [111] proposed a two-tier testing method (one for integration level, the other for system level) for graphical user interface (GUI) software testing. Rutherford et al. [112] defined and evaluated adequacy criteria under system-level testing for distributed systems. In Denaro et al. [113]’s study, they proposed a test data generation approach using data flow information for inter-procedural testing of object-oriented

programs.

The important point we discovered here is that all the aforementioned studies did not restrict mutation operators to model integration errors or system ones. In other words, the traditional program mutations can be applied to higher-level testing. Amongst these articles, the mutation operators adopted are mostly at the unit level, e.g. Arithmetic Mutation Replacement, Relational Mutation Replacement. The mutation operators designed for higher-level testing, e.g. [114, 115], are seldom used in these studies. This reveals a potential direction for future research: the cross-comparison of different levels of mutation operators and testing activities at different test levels (see the third recommendation labeled as **R3** in Section 4.4). The investigation of different level of mutations can explore the effectiveness of mutation faults at different test levels, such as the doubts whether integration-level mutation is better than unit-level mutation when assessing testing techniques at the integration-level. In the same vein, an analysis of whether mutants are a good alternative to real/hand-seeded ones (proposed by Andrews et al. [24]) at higher levels of testing also seems like an important avenue to check out.

In addition, we created a class “others” in which we list 9 papers that we found hard to classify in any of the other four test phases. These works can be divided into three groups: grammar-based testing [116–118], spreadsheet-related testing [80, 119, 120] and SQL-related testing [121–123]. The application of mutation testing on the “other” set indicates that the definition of mutation testing is actually quite broad, thus potentially leading to even more intriguing possibilities [2]: What else can we mutate?

Summary.

The application of mutation testing is mostly done at the unit-level testing (74.4%), while we did still observe several investigations targeting higher level testing efforts. What is worth noticing, is that those studies focusing on higher-level testing still apply traditional (unit-level) mutation operators instead of specific ones designed to model higher-level defects. This reveals a clear gap in current research: the application of mutation testing at higher test levels. Hence, we recommend future research to explore the relationship between different level mutations and different test levels. For example, empirical evaluations on whether mutations at the integration level can represent real faults at the integration testing level. In a different context, but still very important to highlight: 45.9% of papers did not clearly specify their target test level(s). For reasons of clarity, understandability and certainly replicability, it is very important to understand exactly at what level the testing activities take place. It is thus a clear call to arms to researchers to better describe these essential testing activity features.

4.1.3 RQ1.4: Which testing strategies does mutation testing support more frequently?

Observations.

In Table 6 we summarised the distribution of testing strategies based on our coarse-grained classification (e.g. structural testing, specification-based testing) as mentioned in Table 2. Looking at Table 6, structure-based testing comes first among all the testing strategies (59.1%, 94 instances). The underlying cause could be that structural testing is still the main focus of testing strategies in the software testing context. The other testing strategies have also been supported by mutation testing: (1) specification-based testing accounts for 54 cases; (2) hybrid testing for five instances (combination of structural and structure-based testing); (3) three cases applying mutation testing in similarity-based testing; (4) 15 instances in others, e.g. static analysis.

One interesting finding is that the enhanced structural testing ranks second among the seven groups of testing strategies. The enhanced structural testing uses other approaches to overcome its inherent fault-revealing disadvantages (i.e. “coverage is not strongly correlated with test effectiveness” [58]), including mutation-based techniques, information retrieval knowledge, observation notations and assertion coverage. The popularity of enhanced structural testing reveals the awareness of the shortage of conventional coverage-based testing strategies has increased.

Compared to enhanced structural testing, enhanced specification-based testing did not attract much interest. The 11 instances mainly adopted mutation testing (e.g. Qu et al. [97] and Papadakis et al. [124]) to improve the testing strategies.

Summary.

Mutation testing has been widely applied in support of different testing strategies. From the observation, the testing strategies other than white box testing can also benefit from the application of mutation testing, such as specification-based testing, hybrid testing and similarity-based testing. But structural testing is more popular than the others (59.1%). Moreover, techniques like mutation-based techniques and

| Testing Strategies | Total |
|--|-------|
| structural testing | 48 |
| structural testing (enhanced) | 46 |
| specification-based testing | 43 |
| others | 15 |
| specification-based testing (enhanced) | 11 |
| hybrid testing | 5 |
| similarity-based testing | 3 |

Table 6. Testing Strategies summary

| Availability | Types | Total |
|--------------|------------------------------------|-------|
| Yes | existing | 78 |
| | partially-based | 7 |
| | self-written | 1 |
| No | n/a (no information) | 40 |
| | existing (given the name/citation) | 17 |
| | self-written | 11 |
| | hand-seeded | 9 |

Table 7. Mutation Tool Summary

information retrieval knowledge are also being adopted to improve the traditional structural-based testing, which typically only relies on the coverage information of software artefacts. This serves an indication of the increasing realisation of the limitations of coverage-based testing strategies.

4.2 R2: How are empirical studies related to the mutation testing designed and reported?

4.2.1 RQ2.1: Which mutation tools have been used more frequently?

Observations.

We are interested in getting insight into the types (as defined in Table 2) of mutation testing tools that are being used and into their availability. Therefore, we tabulated the different types of mutation testing tools and their availability in Table 7. As shown in Table 7, 49.1% of the studies adopted existing tools which are open to the public; this matches our expectation: as mutation testing is not the main focus of the studies, if there exists a well-performing and open-source mutation testing tool, the researchers are likely willing to use these tools. However, we also encountered 12 cases of using self-implemented tools and nine studies that manually applied mutation operators. A likely reason for implementing a new mutation testing tool or for applying mutation operators manually is that existing tools do not satisfy a particular need of the researchers. In addition, most existing mutation testing tool are typically targeting one specific language and a specific set of mutation operators [2] and they are not always easy to extend, e.g., when wanting to add a newly-defined mutation operator. As such, providing more flexible mechanisms for creating new mutation operators in mutation testing tools is an important potential direction for the future research [2] (see the fourth recommendation labeled as **R4** in Section 4.4).

Unfortunately, there are also still 77 studies (47%) that do not provide access to the tools, in particular, 40 papers did not provide any information about the tools at all, a situation that we marked as “n/a” in Table 7. This unclarity should serve as a notice to researchers: the mutation testing tool is one of the basic elements in mutation testing and lack of information on it seriously hinders replicability of the experiments.

Having discussed the tool availability and types, we are wondering which existing open-source mutation testing tools are most popular. The popularity of the tools can not only reveal their level of maturity, but also provide a reference for researchers entering the field to help them choose a tool. To this end, we summarised the names of mutation tools for different programming languages in Table 8. Table 8 shows that we encountered 18 mutation tools in total. Most tools target at one programming language

| Language | Tool | Total |
|-------------------|------------------------------|-------|
| Java | MuJava/ μ -java/Muclipse | 31 |
| | PIT/PiTest | 7 |
| | JAVALANCHE | 7 |
| | MAJOR | 5 |
| | Jumble | 2 |
| | Sofya | 1 |
| | Jester | 1 |
| C | Proteum | 11 |
| | MiLu | 2 |
| | SMT-C | 1 |
| Fortran | Mothra | 4 |
| SQL | SQLMutation/JDAMA | 2 |
| | SchemaAnalyst | 1 |
| C# | GenMutants | 1 |
| | PexMutator | 1 |
| JavaScript | MUTANDIS | 2 |
| AspectJ | AjMutator | 1 |
| UML specification | MoMuT::UML | 1 |

Table 8. Existing Mutation Tool Summary

(except for Mothra [17] which can support C as well). We encountered eight mutation tools for Java, with the top 3 most-used being MuJava [125], PIT [126]x and JAVALANCHE [87]. We found that three mutation tools for C are used, where Proteum [18] is most-frequently applied.

In Jia and Harman [1]’s literature review, they summarised 36 mutation tools developed between 1977 and 2009. When comparing their findings (36 tools) to ours (18 tools), we find that there are 12 tools in common. The potential reason for us not covering the other 24 is that we only consider peer-reviewed conference papers and journals; this will likely filter some papers which applied the other 24 mutation tools. Also important to stress, is that the goal of Jia and Harman’s survey is different to ours: while we focus on the application of mutation tools, their study surveys articles that introduce mutation testing tools. In doing so, we still managed to discover 8 mutation tools which are not covered by Jia and Harman: (1) two tools are for Java: PIT and Sofya; (2) one for C: SMT-C; (3) one for SQL: SchemaAnalyst; (4) one for UML: MoMuT::UML; (5) two for C#: GenMutants and PexMutator; (6) one for JavaScript: MUTANDIS. Most of these tools were released after 2009, which makes them too new to be included in the review of Jia and Harman. Moreover, we can also witness the trend of the development of the mutation testing for programming languages other than Java and C means compared to Jia and Harman [1]’s data.

Summary.

Around 50% of the articles that we have surveyed adopt existing (open source) tools, while in a few cases (21 in total) the authors implemented their own tools or seeded the mutants by hand. This calls for a more flexible mutation generation engine that allows to easily add mutation operators or certain types of constraints. Furthermore, we found 77 papers that did not provide any information about the mutation tools they used in their experiments; this should be a clear call to arms to the research community to be more precise when reporting on mutation testing experiments. We have also gained insight into the most popular tools for various programming languages, e.g., MuJava for Java and Proteum for C. We hope this list of tools can be a useful reference for new researchers who want to apply mutation analysis.

4.2.2 RQ2.2: Which mutation operators have been used more frequently?

Observations

For the mutation operators, we first present the distribution of the three description levels (as mentioned in Table 2) in Table 9. As Table 9 shows, 61.6% (98 instances) of the studies that we surveyed specify

| Description Level | Total |
|-------------------|-------|
| well-defined | 98 |
| n/a | 39 |
| not sufficient | 22 |

Table 9. Description Level of Mutation Operators Summary

the mutation operators that they use, while more than one-third of the articles do not provide enough information about the mutation operators to replicate the studies. These 61 instances that are labelled as “n/a” and “not sufficient” indicate the necessity of raising the awareness of the use of mutation in the testing experiment as claimed by Namin and Kakarla [127]. They highlighted that mutation testing is highly sensitive to external threats caused by the influential factors including mutation operators. Therefore, one suggestion here is to report the complete list of the mutation operators when applying mutation testing in testing experiments: the authors can provide either a self-contained list in the paper or the relevant citations which can be used to trace the complete list.

After that, based on our generic classification of the mutation operators (as defined in Section 3.3 (4)), we characterised the 98 papers labelled as “well-defined”. In addition to the overall distribution of the mutation operators regardless of the programming languages, we are also interested in the differences of the mutation operators for different languages as the differences could indicate potential gaps in the existing mutation operator sets for certain programming languages. In Table 10 we project the different languages onto seven columns and our predefined mutation operator categories onto the rows, thus presenting the distribution of the mutation operators used in the literature under our research scope.

Overall, we can see that program mutation is more popular than specification mutation from the Table 10. Among the program mutation operators, the arithmetic, relational and conditional operators are the top 3 mutation operators. These three operators often used together in most cases as their total numbers of applications are similar. The underlying cause of the popularity of these three operators could be that the three operators are among Offutt et al. [30]’s 5 sufficient mutation operators. Moreover, the expression-level operators are more popular than the statement-level ones. As for the statement-level mutation operators, statement deletion, method call and `return` statement are the top 3 mutation operators.

When we compare the mutation operators used in different languages to our mutation operator categories, we see that there exist differences between different programming languages, just like we assumed. Table 11 leads to several interesting findings that reveal potential gaps in various languages:

1. for Java, seven mutation operators at the expression and statement level (except `goto` label which is not supported in Java) are not covered: `type`, `bomb` statement, `do` statement, `brace`, `loop trap`, `while` statement and `if` statement
2. for C, only one potential mutation operator, method call, was not covered. The C programming language does not provide direct support for exception handling
3. for C++, 3 expression-level, 10 statement-level and the OO-specific operators are not used in literature
4. for C#, only a basic set of mutation operators are covered
5. for Fortran, the earliest programming language mutation testing was applied to, the relevant studies covered a basic set
6. for SQL, since the syntax of SQL is quite different from the imperative programming languages, only six operators at the expression level and SQL-specific ones are used
7. for JavaScript, only three mutation operators other than JavaScript-specific ones are adopted in existing studies.

⁸The Java-specific operator here refers to the `static` modifier change (including insertion and deletion).

| Level | Operator | Java | C | C++ | C# | Fortran | SQL | JavaScript | Total |
|-------------------------------|------------------------|--------------------|---|----------------|----|---------|-----|------------|-------|
| Specification Mutation | | 2 | 2 | 1 | - | - | - | - | 22 |
| Program Mutation | | 43 | 9 | 3 | 2 | 2 | 4 | 1 | 77 |
| Expression-level | arithmetic operator | 42 | 8 | 3 | 2 | 2 | 2 | - | 66 |
| | relational operator | 38 | 7 | 3 | 2 | 2 | 2 | - | 62 |
| | conditional operator | 39 | 6 | 3 | 2 | 2 | 2 | - | 61 |
| | assignment operator | 24 | 3 | 2 | - | - | - | - | 29 |
| | bitwise operator | 26 | 3 | 2 | - | - | - | - | 31 |
| | shift operator | 26 | 1 | 2 | - | - | - | - | 29 |
| | constant | 12 | 3 | 2 | 1 | 2 | 2 | - | 27 |
| | variable | 8 | 3 | 2 | 1 | 2 | 2 | 1 | 23 |
| | absolute value | 5 | 3 | 1 | 2 | 1 | 2 | - | 15 |
| | conditional expression | 2 | 2 | - | - | - | - | - | 5 |
| | parenthesis | 1 | 1 | - | - | - | - | - | 2 |
| | type | - | 2 | - | - | - | - | - | 2 |
| | Statement-level | statement deletion | 6 | 3 | - | 2 | 2 | - | - |
| method call | | 7 | - | 2 | 1 | - | - | 1 | 13 |
| return statement | | 6 | 2 | - | - | 2 | - | - | 11 |
| exception handler | | 1 | - | 1 | 1 | - | - | - | 5 |
| goto label | | - | 2 | - | - | 2 | - | - | 5 |
| control-flow disruption | | 2 | 1 | 2 | - | - | - | - | 4 |
| statement swap | | 2 | 1 | 2 | - | - | - | - | 4 |
| bomb statement | | - | 1 | - | - | 2 | - | - | 4 |
| switch statement | | 2 | 2 | - | - | - | - | - | 4 |
| do statement | | - | 1 | - | - | 2 | - | - | 3 |
| brace | | - | 2 | - | - | - | - | - | 2 |
| loop trap | | - | 2 | - | - | - | - | - | 2 |
| while statement | | - | 2 | - | - | - | - | - | 2 |
| if statement | | - | - | - | - | - | - | - | - |
| Others | OO-specific | 17 | - | - | 1 | - | - | - | 21 |
| | Java-specific | 12 | - | 1 ⁸ | - | - | - | - | 12 |
| | SQL-specific | - | - | - | - | - | 4 | - | 4 |
| | Concurrent mutation | 3 | - | - | - | - | - | - | 3 |
| | AOP-specific | - | - | - | - | - | - | - | 2 |
| | Interface mutation | 1 | - | - | - | - | - | - | 2 |
| | Spreadsheet-specific | - | - | - | - | - | - | - | 2 |
| JavaScript-specific | - | - | - | - | - | - | 1 | 1 | |

Table 10. Mutation Operators Used In the Literature

| | MuJava/ μ -java/Muclipse | PIT/PiTest | JAVALANCHE | MAJOR | Jumble | Sofya | Jester | Proteum | MiLu | SMT-C | Mothra | SQLMutation/JDAMA | SchemaAnalyst | GenMutants | PexMutator | MUTANDIS | AjMutator | MoMut::UML |
|-------------------------|------------------------------|------------|------------|-------|--------|-------|--------|---------|------|-------|--------|-------------------|---------------|------------|------------|----------|-----------|------------|
| Specification Mutation | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · |
| arithmetic operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| relational operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| conditional operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| assignment operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| bitwise operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| shift operator | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < | < |
| constant | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · |
| variable | · | < | · | · | < | · | · | < | · | < | · | < | · | · | · | < | · | · |
| absolute value | · | · | · | < | · | · | · | < | · | · | < | < | · | < | < | < | · | · |
| conditional expression | · | < | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| parenthesis | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| type | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| statement deletion | · | < | < | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| method call | · | · | · | · | · | < | · | < | · | · | < | < | · | · | · | < | · | · |
| return statement | · | < | · | · | < | · | · | < | · | · | < | < | · | · | · | < | · | · |
| if statement | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| exception handler | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| goto label | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| control-flow disruption | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| statement swap | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| bomb statement | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| switch statement | · | < | · | · | < | · | · | < | · | < | · | < | · | · | · | < | · | · |
| do statement | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| brace | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| loop trap | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| while statement | · | · | · | · | · | · | · | < | · | · | < | < | · | · | · | < | · | · |
| OO-specific | < | < | · | · | · | < | · | < | · | · | · | · | · | · | · | · | · | · |
| Java-specific | < | < | · | · | · | < | · | < | · | · | · | · | · | · | · | · | · | · |
| SQL-specific | · | · | · | · | · | · | · | · | · | · | · | < | < | · | · | · | · | · |
| JavaScript-specific | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | < | · | · |
| AOP-specific | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | · | < | · |

Table 11. Comparison of Mutation Operators in Existing Mutation Tools

The above findings are just the initial results in which we neither did further analysis to chart the syntax differences of these languages nor investigate the possibility of the equivalent mutants caused by our classification. Moreover, for some languages, e.g. JavaScript, the relevant studies are too few to draw any definitive conclusions. Here, we can only say that for different languages the existing studies did not cover all the mutation operators that we listed in Table 10: some are caused by the differences in the syntax, while the others could point to potential gaps. The distribution of the mutation operators of different languages summarised in Table 10 can be a reference for further investigations into mutation operator differences in various programming languages.

Furthermore, our generic classification of the existing mutation operators can be of benefit to compare mutation tools in the future. Thereby, we compared the existing mutation testing tools (as listed in Table 8) to our mutation operator categories in Table 11. The result shows that none of the existing mutation testing tools we analysed can cover all types of operators we classified. For seven mutation testing tools for Java, they mainly focus on the expression-level mutations and only four kinds of statement-level mutators are covered. Furthermore, MuJava, PIT and Sofya provide some OO-specific operators, whereas PIT only supports one type, the Constructor Calls Mutator. For the three mutation testing tools for C that we have considered, Proteum covers the most mutation operators. SMT-C is an exceptional case of the traditional mutation testing which targets at semantic mutation testing. For the tools designed for C#, OO-specific operators are not present. Another interesting finding when we compared Table 10 and Table 11, is that the `if` statement mutator is not used in literature but it is supported by SMT-C. This observation indicates that not all the operators provided by the tools are used in the studies when applying mutation testing, which reinforces our message of the need for “well-defined” mutation operators when reporting mutation testing studies.

| Equivalence Detector | Total |
|-----------------------------|-------|
| n/a | 91 |
| manual investigation | 29 |
| not killed as equivalent | 16 |
| no investigation | 11 |
| reduce likelihood | 7 |
| model checker | 6 |
| deterministic model | 1 |
| not killed as nonequivalent | 1 |

Table 12. Methods for Overcoming E(quivalent) M(utant) P(robblem) Summary

Summary.

For the mutation operators, we focused on two attributes: their description level and a generic classification across tools and programming languages. When investigating the description level of the mutation operators that are used in studies, we found that 61.6% (98 instances) explicitly defined the mutation operators used. This leads us to strongly recommend improving the description level for the sake of replicability. Furthermore, the distribution of mutation operators based on our predefined categories shows the lacking of certain mutation operators in some programming languages among the existing (and surveyed) mutation testing tools. A possible avenue for future work is to see which of the *missing* mutation operators can be implemented for the programming languages lacking these operators.

4.2.3 RQ2.3: Which approaches are used to overcome the equivalent mutant problem more often when applying mutation testing?

Observations.

In Table 12 we summarised our findings of how the studies that we surveyed deal with the equivalent mutant problem. More specifically, Table 12 presents how many times we encountered each of the approaches for solving the equivalent mutant problem. When looking at the results, we firstly observe that in 57.2% of the cases we assigned “n/a”; one possible reason for this high number of papers not providing information is that we excluded the internal methods of overcoming the equivalent mutant problem, meaning that mutation tools have it built in already. Moreover, as a non-decidable problem [13], the equivalent mutant problem is still under-developed [14].

As shown in Table 12, there are only 14 instances (6 “model checker” + 7 “reduce likelihood”+1 “deterministic model”) actually adopting equivalent mutant detectors by using automatic mechanisms. In the remainder of the papers, the problem of equivalent mutants is solved by (1) manual analysis, (2) making assumptions (treating mutants not killed as either equivalent or nonequivalent), and (3) no investigation. In particular, the manual investigation (29 instances) and the method of treating mutants not killed as equivalent (16 instances) are more commonly used than other methods.

We can only speculate as to the reasons behind the above situation: Firstly, most studies use mutation testing as an evaluation mechanism or guiding heuristic, rather than their main research topic. So the authors are maybe not willing to spare too much effort in dealing with problems associated with mutation testing. Moreover, looking at the internal features of existing tools used in literature (in Table 13), we found that only four tools adopt certain techniques to address the equivalent mutant problem. Most of the tools did not assist in dealing with the equivalent mutant problem. Thereby, the aforementioned three solutions, (1) manual analysis, (2) making assumptions or (3) no investigation, are considered by the authors. If there exists a well-developed auxiliary tool that can be seamlessly connected to the existing mutation systems for helping the authors detect equivalent mutants, this tool might be more than welcomed. We recommend that future research on the equivalent mutant problem can further implement their algorithms in such an auxiliary tool and make it open to the public (see the fifth recommendation labeled as **R5** in Section 4.4).

Secondly, mutation score is mainly used as a relative comparison for estimating the effectiveness of different techniques. Sometimes, mutation testing is only used to generate likely faults; equivalent mutants have no impact on the other measures such as the Average Percentage of Fault Detection rate (APFD) [90]. Furthermore, the definition of the mutation score is also modified by some authors: they

| Language | Tool | Equivalent mutants | Mu- | Cost Reduction |
|-------------------|------------------------------|----------------------------------|---------|---|
| Java | MuJava/ μ -java/Muclipse | n/a | | MSG, bytecode translation (BCEL) [125] |
| | PIT/PiTest | n/a | | Bytecode translation (ASM), coverage-based test selection [126] |
| | JAVALANCHE | Ranking mutations by impact [87] | | MSG, bytecode translation (ASM), coverage-based test selection, parallel execution [87] |
| | MAJOR | n/a | | Compiler-integrated, coverage-based test selection [128] |
| | Jumble | n/a | | Bytecode translation (BCEL), conventional test selection [129] |
| | Sofya | n/a | | Bytecode translation (BCEL) [130] |
| | Jester | n/a | | n/a |
| C | Proteum | n/a | | Mutant sample [18] |
| | MiLu | n/a | | Higher-order mutants, test harness [131] |
| | SMT-C | n/a | | Interpreter-based, weak mutation [132] |
| Fortran | Mothra | n/a | | Interpreter-based [1] |
| SQL | SQLMutation/JDAMA | Constraint binding [20] | binding | n/a |
| | SchemaAnalyst | n/a | | n/a |
| C# | GenMutants | n/a | | n/a |
| | PexMutator | n/a | | Compiler-based [55] |
| JavaScript | MUTANDIS | Reduce likelihood [79] | likeli- | Selection strategy [79] |
| AspectJ | AjMutator | Static analysis [133] | | Compiler-based [133] |
| UML specification | MoMuT::UML | n/a | | Compiler-based [134] |

Note: "n/a" in the table means we did not find any relevant information recorded in literature or websites, and some tools might adopt certain techniques but did not report such information in the sources we can trace.

Table 13. Inner features of Existing Mutation Tool

used the total number of mutants as the denominator instead of number of nonequivalent mutants. The equivalent mutant problem seems to not pose a significant threat to the validation of the testing techniques involved in these studies.

However, we should not underestimate the impact of the equivalent mutant problem on the accuracy of the mutation score. Previous empirical results indicated that 10 to 40 percent of mutants are equivalent [135, 136]. What's more, in Schuler and Zeller's study [47], they further claimed that around 45% of all undetected mutants turned out to be equivalent; this observation leads to the assumption that by simply treating mutants not killed as equivalent mutations, we could be overestimating the mutation score. Therefore, we recommend performing more large-scale investigations on whether the equivalent mutant problem has a strong impact on the accuracy of the mutation score.

Summary.

The techniques for equivalent mutant detection are not commonly used when applying mutation testing. The main approaches that are used are the manual investigation and treating mutants not killed as equivalent. Based on the results, we recommend further research on the equivalent mutant problem can develop a mature and useful auxiliary tool which can easily link to the existing mutation system. Such an extra tool assists people to solve the equivalent mutant problem when applying mutation testing more efficiently. Moreover, research on whether the equivalent mutant problem has a high impact on the accuracy of the mutation score is still needed, as the majority did not consider the equivalent mutant problem as a significant threat to validation of the testing activities. Also, 57.2% of the studies are lacking an explanation as to how they are dealing with overcoming the equivalent mutants problem; this again calls for more attention on reporting mutation testing *appropriately*.

| Cost Reduction Technique | Total |
|--------------------------|-------|
| n/a | 108 |
| fixed number | 25 |
| weak mutation | 15 |
| mutant sample | 9 |
| selection strategy | 4 |
| higher-order | 1 |

Table 14. Cost Reduction Summary

4.2.4 RQ2.4: Which techniques are used to reduce the computational cost more frequently when applying mutation testing?

Observations.

Since mutation testing requires high computational demands, cost reduction is necessary for applying mutation testing, especially in the industrial environment. We summed up the uses of such computational cost reduction techniques when using mutation testing in Table 14. Please note that we excluded the runtime optimisation and selective mutation techniques. We opted to exclude this because the runtime optimisation is related to tool implementation which is not very likely to appear in the papers under our research scope, while the second one, selective mutation, is adopted by all the papers.

First of all, we noticed that 108 articles (67.9%) did not mention any reduction techniques. If we take into account those papers that used the runtime optimisation and selective mutation, one plausible explanation for the numerous “n/a” instances is a lack of awareness of *properly* reporting mutation testing, as we mentioned earlier. Secondly, random selection of the mutants based on a fixed number comes next (25 instances), followed by weak mutation (15 instances) and mutant sampling (9 cases).

But why is the technique of using a “fixed number” of mutants more popular than the others? We speculate that this could be because of the fact that choosing a certain number of mutants is more realistic in real software development: the total number of mutants generated by mutation tools is huge; while, realistically, only a few faults are made by the programmer during implementation. By fixing the number of mutants, it becomes easier to control the mutation testing process. Instead, relying on the weak mutation condition would require additional implementation efforts to modify the tools. Also of importance to note is the difference between the “fixed number” and “mutant sample” choice: while the first one implies a fixed number of mutants, the second one relies on fixed sampling rate. Compared to using a fixed number, mutant sampling sometimes cannot achieve the goal of reducing the mutants number efficiently. In particular, it is difficult to set one sample ratio if the size of the subjects varies greatly. For example, consider the following situation: one subject has 100,000 mutants while the other has 100 mutants. When the sample ratio is set to 1%, the first subject still has 1000 mutants left, while the number of mutants for the second one is reduced to one.

Moreover, we performed a further analysis of the mutation tools in Table 13. We find that most tools adopted certain types of cost reduction techniques to overcome the high computational expense problem. For mutation testing tools for Java, bytecode translation is frequently adopted while Mutant Schemata Generation (MSG) is used in two tools, MuJava and JAVALANCHE. Another thing to highlight is that MiLu used a special test harness to reduce runtime [131]. This test harness is created containing all the test cases and settings for running each mutant. Therefore, only the test harness need to be executed while each mutant runs as an internal function call during the testing process.

Selective mutation is also widely applied in nearly all the existing mutation testing tools (as shown in Table 11). This brings us to another issue, namely whether the selected subset of mutation operators is sufficient to represent the whole mutation operator set? When adopting selective mutation, some configurations are based on prior empirical evidence, e.g., Offutt et al. [30]’s five sufficient Fortran mutation operators and Siami et al. [32]’s 28 sufficient C mutation operators. However, most cases are not supported by the empirical or theoretical studies that show a certain subset of mutation operators can represent the whole mutation operator set. Thereby, we recommend more empirical studies on selective mutation in programming languages other than Mothra and C (see the sixth recommendation labeled as R6 in Section 4.4).

Based on the above discussion, we infer that mutation testing’s high computational problem can

be adequately controlled using the state-of-art reduction techniques, especially selective mutation and runtime optimisation.

Summary.

Selective mutation is the most widely used method for reducing the high computational cost of mutation testing. However, in most cases, there are no existing studies to support the prerequisite that selecting a particular subset of mutation operators is sufficient to represent the whole mutation operator set. Therefore, one recommendation is to conduct more empirical studies on selective mutation in various programming languages. Without regard to runtime optimisation and selective mutation, random selection of the mutants based on a fixed number (25 papers) is the most popular technique used to reduce the computational cost. The following ones are weak mutation and mutant sampling. Besides, a high percentage of the papers (67.9%) did not report any reduction techniques used to cope with computational cost when applying mutation testing; this again should serve as a reminder for our research community to pay more attention to *properly* reporting mutation testing in testing experiments.

4.2.5 RQ2.5: What are the most common subjects used in the experiments?

Observations.

To analyse the most common subjects used in the experiments, we focus on three attributes of the subject programs, i.e., programming language, size and data availability. We will discuss these three attributes one by one in the following paragraphs.

Table 15 shows the distribution of the programming languages. We can see that Java and C dominate the application domain (66.0%, 105 instances). While JavaScript is an extensively used language in the web application domain, we only found two cases of it being described in the studies we surveyed. The potential reasons for this uneven distribution are unbalanced data availability and the complex nature of building a mutation testing system. The first cause, uneven data availability, is likely instigated through the fact that existing, well-defined software repositories such as SIR [88], SF100 [89] are based on C and Java. We have not encountered such repositories for JavaScript, C# or SQL. Furthermore, it is easier to design a mutation system targeting one programming language. This stands in contrast to many web applications, that are often composed out of a combination of JavaScript, HTML, CSS, etc. This thus increases the difficulty of developing a mutation system for these combinations of programming languages. It is also worth noticing that we have not found any research on a purely functional programming language in our research scope.

When considering the maximum size of the subject programs, studies involving preliminary (<100 LOC), small (100~10K LOC) subjects or studies with no information about programs size (“n/a” instances in Table 16) occupy the 79% (126 instances) of papers in our collection. This high percentage of preliminary, small and “n/a” subjects indicates that mutation testing is rarely applied to programs whose size is higher than 10K LOC. We did find that only 31 studies use medium size subjects, which corresponds to 19% of cases. Finally, we observed only two cases where mutation testing is applied in more large-scale projects: (i) Qi et al. [99] adopted mutation testing to prioritise test cases to speed up patch validation during program repairing and (ii) von Mayrhauser et al. [95] proposed a new test data generation technique based on an Artificial Intelligence (AI) planner and evaluated their method on an Automated Cartridge System (ACS). These two instances show the full potential of mutation testing to be employed as a practical testing tool for large industrial systems.

With regard to the concern of data availability, we observe the following: 49.1% of the studies provide open access to their experiments. Together with 5 instances of “n/a” in Table 15 and 59 in Table 16 (including subjects which cannot be measured as LOC, e.g. Spreadsheet applications), it is worth noticing that subject programs used in the experiment should be clearly specified. In addition, basic information on programming language, size and subject should also be clearly specified in the articles to ensure replicability.

Summary.

For the subject systems used in the experiments in our survey, we discussed three aspects: programming language, size of subject programs and data availability. For programming languages, Java and C are the most common programming languages used in the experiments when applying mutation testing. There is a clear challenge in creating more mutation testing tools for other programming languages, especially in

⁸Here, we did not consider the number of “n/a” value when calculating the percentage (159 – 59 = 100).

| Language | Total |
|----------------------------------|-------|
| Java | 74 |
| C | 31 |
| C# | 6 |
| Fortran | 6 |
| n/a | 5 |
| C++ | 4 |
| Lustre/Simulink | 8 |
| SQL | 4 |
| Eiffel | 3 |
| Spreadsheet | 3 |
| AspectJ | 3 |
| C/C++ | 2 |
| JavaScript | 2 |
| Ada | 1 |
| Kermeta | 1 |
| Delphi | 1 |
| Enterprise JavaBeans application | 1 |
| ISO C++ grammar | 1 |
| PLC | 1 |
| Sulu | 1 |
| XACML | 1 |
| XML | 1 |
| other specification languages | 10 |

Table 15. Programming Language Summary

| Subject Size | Total |
|------------------------|-------|
| n/a | 59 |
| small (100~10K LOC) | 58 |
| medium (10K~1M LOC) | 31 |
| preliminary (<100 LOC) | 9 |
| large (>1M LOC) | 2 |

Table 16. Subject Size Summary

the area of web applications and functional programming (see the seventh recommendation labeled as **R7** in Section 4.4).

As for the maximum size of subject programs, small to medium scale projects (100~1M) are widely used when applying mutation testing. Together with two large-scale cases, we can see the full potential of mutation testing as a practical testing tool for large industrial systems. We recommend more research on large-scale systems to further explore scalability issues (see the eighth recommendation labeled as **R8** in Section 4.4).

The third aspect we consider is data availability. Only 49.1% of the studies that we surveyed provide access to the subjects used. This again calls for more attention on reporting test experiments *appropriately*: the authors should explicitly specify the subject programs used in the experiment, covering at least the details of programming language, size and source.

4.3 Summary of Research Questions

We will now revisit the research questions and answer them in the light of our observations.

RQ1: How mutation testing is used in testing activities?

Mutation testing is mainly used as a fault-based evaluation method (70.4%) in different testing activities. It assesses the fault-finding ability of various testing techniques through the mutation score or the number

| Data Availability | Total |
|-------------------|-------|
| no | 81 |
| yes | 78 |

Table 17. Data Availability Summary

of killed mutants. Adopting mutation testing to improve other testing activities as a guide was first proposed by DeMilli and Offutt [35] in 1991 when they used it to generate test data. As a “high end” test criterion, mutation testing started to gain popularity as a building block in different testing activities, like test data generation (28 instances), test case prioritisation (6 cases) and test strategy evaluation (3 instances). However, using mutation testing as part of new test approaches raises a challenge in itself, namely how to efficiently evaluate mutation-based testing? Besides, we found one limitation related to the “guide” role of mutation testing: mutation testing usually serves as a *where*-to-check constraint rather than a *what*-to-check improvement. Another finding of the application of mutation testing is that it often targets unit-level testing (74.4%), while only a small amount of studies featuring higher-level testing show the benefit of mutation testing. As a result, we conclude that the current state of application of mutation testing is still rather limited.

RQ2: How are empirical studies related to the mutation testing designed and reported?

First of all, for the mutation testing tools and mutation operators used in literature, we found that 49.1% of the articles adopted existing (open-source) mutation testing tools, such as MuJava for Java and Proteum for C. In contrast, we did encounter a few cases (21 in total) where the authors implemented their own tools or seeded mutants by hand. Furthermore, to investigate the distribution of mutation operators in the studies, we created a generic classification of mutation operators as shown in Section 3.3 (4). The results indicate that certain programming languages lack specific mutation operators, at least as far as the mutation tools that we have surveyed concern.

Moreover, when looking at the two most significant problems related to mutation testing, the main approaches to dealing with the equivalent mutant problem are (1) treating mutants not killed as equivalent and (2) not investigating the equivalent mutants at all. In terms of cost reduction techniques, we observed that the “fixed number of mutants” is the most popular technique used to reduce computational cost, although we should mention that we did not focus on built-in reduction techniques.

The above findings suggest that the existing techniques designed to support the application of mutation testing are largely still under development: a mutation testing tool with a more complete set of mutation operators or a flexible mutation generation engine to which mutation operators can be added, is still needed. In the same vein, a more mature and efficient auxiliary tool for assisting in overcoming the equivalent mutant problem is needed. Furthermore, we have observed that we lack insight into the impact of selective mutation on mutation testing; this suggests a deeper understanding of mutation testing is required. For example, if we know what particular kinds of faults mutation is good at finding or how useful a particular type of mutant is when testing real software, we can then design the mutation operators accordingly.

Based on the distribution of subject programs used in testing experiments or case studies, Java and C are the most common programming languages used in the experiments. Also, small to medium scale projects (100~1M LOC) are the most common subjects employed in the literature.

Besides, from the statistics of the collection, we found a considerable amount of papers did not provide a sufficiently clear or thorough specification when reporting mutation testing in their empirical studies. We summarised the poorly-specified aspects of mutation testing in Table 18. As a result, we call for more attention on reporting mutation testing appropriately. The authors should provide at least the following details in the articles: the mutation tool (preferably with a link to its source code), mutation operators used in experiments, how to deal with the equivalent mutant problem, how to cope with high computational cost and details on the subject program (see the ninth recommendation labeled as **R9** in Section 4.4).

4.4 Recommendation of Future Research

In this section, we will summarise the recommendations for the future research based on the discussion in two research questions above (in Section 4.1.1-4.3). We propose nine recommendations for the future research as follows:

| The poorly-specified aspects in reporting mutation testing | Number of papers |
|--|------------------|
| test level | 73 |
| mutation tool source | 77 |
| mutation operators | 61 |
| equivalent mutant problem | 91 |
| reduction problem | 108 |
| subject program source | 81 |

Table 18. Poorly-specified aspects in empirical studies

- **R1: Mutation testing can not only be used as *where-to-check* constraints but also to suggest *what to check to improve test code quality*.**

As shown in Table 3 in Section 4.1.1, when mutation testing serves as a “guide”, mutants generated by the mutation testing system are mainly used to suggest the location to be checked, i.e., *where-to-check* constraints. For example, the location of mutants is used to assist the localisation of “unknown” faults in fault localisation. The mutation-based test data generation also used the position information to generate killable mutant conditions. However, mutation testing is not widely considered to be a benefit to improve test code quality by suggesting *what to check*, especially in test oracle problem. The *what-to-check* direction can be one opportunity for future research in mutation testing as a “guide” role.

- **R2: For testing approaches that are guided by a mutation approach, more focus can be given to finding an appropriate way to evaluate mutation-based testing in an efficient manner.**

When looking at the evaluation types in Table 4 in Section 4.1.1, we observe that 75.5% of the mutation-based testing techniques still adopt mutation faults to assess their effectiveness. This raises the question of whether the conclusions might be biased. As such, we open the issue of finding an appropriate way to evaluate mutation-based testing in an efficient manner.

- **R3: Study the higher-level application of mutation testing.**

In Section 4.1.2 we made the observation that mutation testing seems to mainly target the unit-level testing, accounting for 74.4% of the studies we surveyed. This reveals a potential gap in how mutation testing is currently applied. It is thus our recommendation that researchers pay more attention to higher-level testing, such as integration testing and system testing. The research community should not only investigate potential differences in applying mutation testing at the unit-level or at a higher level of testing, but also explore whether the conclusions based on unit-level mutation testing could still apply to higher-level mutation testing. A pertinent question in this area could for example be whether an integration mutation fault can be considered as an alternative to a real bug at the integration level.

- **R4: The design of a more flexible mutation generation engine that allows for the easy addition of new mutation operators.**

As shown in Table 7 in Section 4.2.1, 49.1% of the articles adopted the existing tools which are open-source, while we also found 21 instances of researchers implementing their own tool or seeding the mutants by hand. Furthermore, in Table 10 and Table 11, we can see certain existing mutation testing tools lack certain mutation operators. These findings imply that existing mutation testing tools can not always satisfy all kinds of needs, and new types of mutation operators are also potentially needed. Since most existing mutation testing tools have been initialized for one particular language and a specific set of mutation operators, we see a clear need for a more flexible mutation generation engine to which new mutation operators can be added easily [2].

- **R5: A mature and efficient auxiliary tool to detect equivalent mutants that can be easily integrated with existing mutation tools.**

In Section 4.2.3, the problem of equivalent mutants is mainly solved by manual analysis, assumptions (treating mutants not killed as either equivalent or nonequivalent) or no investigation at all during application. This observation leads to doubt about the efficacy of the state-of-art equivalent mutant detection. In the meanwhile, if there is a mature and efficient auxiliary tool which can easily link to the existing mutation system, the auxiliary tool can be a practical solution for the equivalent mutant problem when applying mutation testing. As a result, we call for a well-developed and easy to integrate auxiliary tool for the equivalent mutant problem.

- **R6: More empirical studies on the selective mutation method can pay attention to programming languages other than Mothra and C.**

As mentioned in Section 4.2.4, selective mutation is used by all the studies in our research scope. However, the selection of a subset of mutation operators in most papers is not well supported by existing empirical studies, except for Mothra [30] and C [32]. Selective mutation requires more empirical studies to explore whether a certain subset of mutation operator can be applied in different programming languages.

- **R7: More attention should be given to other programming languages, especially web applications and functional programming projects.**

As discussed in RQ2.5 in Section 4.2.5, Java and C are the most common programming languages that we surveyed. While JavaScript and functional programming languages are seldom applied. JavaScript, as one of the most popular languages in developing web applications, calls for more attention from the researchers. In the meanwhile, functional programming languages, such as Lisp and Haskell, are still playing an inevitable role in the implementation of programs; thus also require more focus in future studies.

- **R8: Application of mutation testing in large-scale systems to explore scalability issues.**

From Table 16 in Section 4.2.5 we learn that the application of mutation testing to large-scale programs whose size is greater than 1M LOC rarely happens (only two cases). In order to effectively apply mutation testing in industry, the scalability issue of mutation testing requires more attention. We recommend future research to use mutation testing in more large-scale systems to explore scalability issues.

- **R9: Authors should provide at least the following details in the articles: mutation tool source, mutation operators used in experiments, how to deal with the equivalent mutant problem, how to cope with high computational cost and subject program source.**

From Table 18 in Section 4.4 we remember that a considerable amount of papers reported mutation testing in a poor manner. To help progress research in the area more quickly and to allow for replication studies, we all need to take care to be careful in how we report mutation testing in empirical studies. We consider the above five elements to be essential when reporting mutation testing.

5 THREATS TO THE VALIDITY OF THIS REVIEW

We have presented our methodology for performing this systematic literature review and its findings in the previous sections. As conducting a literature review largely relies on manual work, there is the concern that different researchers might end up with slightly different results and conclusions. To eliminate this potential risk caused by researcher bias as much as possible, we follow the guidelines for performing systematic literature reviews by Kitchenman [6], Wohlin [73], Brereton et al. [71] whenever possible. In particular, we keep a detailed record of procedures made throughout the review process by documenting all the metadata from article selection to characterisation see (Table 19-21 in Appendix).

In this section, we are going to describe the main threats to validity of this review and discuss how we attempted to mitigate the risks regarding four aspects: the article selection, the attribute framework, the article characterisation and the result interpretation.

5.1 Article Selection

Mutation testing is an active research field, and a plethora of realisations have been achieved as shown in Jia and Harman's thorough survey [1]. To address the main interest of our review, i.e. actual application of mutation testing, we need to define inclusion/exclusion criteria to include papers of interest and exclude irrelevant ones. But this also introduces a potential threat to the validity of our study: unclear article selection criteria. To minimise the ambiguity caused by the selection strategies, we carried out a pilot run of the study selection process to validate our selection criteria among the three authors. This selection criteria validation led to a tiny revision. Besides, if there is any doubt about whether a paper belongs in our selected set, we had an internal discussion to see whether the paper should be included or not.

The venues in Table 1 were selected because we considered them to be the key venues in software engineering and most relevant to software testing, debugging, software quality and validation. This presumption might result in an incomplete paper collection. In order to mitigate this threat, we also adopted snowballing to extend our set of papers from pre-selected venues to reduce the possibility of missing influential papers.

Although we made efforts to minimise the risks with regard to article selection, we cannot make definitive claims about the completeness of this review. We have one major limitation related to the article selection: we only considered top conference or journal papers to ensure the high quality while we excluded article summaries, interviews, reviews, workshops, panels and poster sessions. Vice versa, sometimes we were also confronted with a vague use of the "mutation testing" terminology, in particular, some papers used the term "mutation testing", while they are doing fault seeding, e.g. Lyu et al. [137]. The main difference between mutation testing and error seeding is the way how to introduce defects in the program [138]: mutation testing follows certain rules while error seeding adds the faults directly without any particular techniques.

5.2 Attribute Framework

We consider the attribute framework to be the most subjective step in our approach: the generalisation of the attribute framework could be influenced by the researcher's experience as well as the reading sequence of the papers. To generate a useful and reasonable attribute framework, we followed a two-step approach: (1) we first wrote down the facets of interest according to our research questions and then (2) derived corresponding attributes of interest. Moreover, for each attribute, we need to ensure all possible values of each attribute are available, as well as a precise definition of each value. In this manner, we can target and modify the unclear points in our framework quickly. In particular, we conducted a pilot run for specifically for validating our attribute framework. The results led to several improvements to the attribute framework and demonstrated the applicability of the framework.

5.3 Article Characterisation

Thanks to the complete definitions of values for each attribute, we can assign the value(s) to articles in a systematic manner. However, applying the attribute framework to the research body is still a subjective process. To eliminate subtle differences caused by our interpretation, we do no further interpretation of the information extracted from the papers in the second pilot run of validation. In particular, if there is no specified detail in a paper, we then mark as "n/a". Furthermore, we listed our data extraction strategies about how to identify and classify the values of each attribute in Section 3.3.

5.4 Result Interpretation

Researcher bias could cause a potential threat to validity when it comes to the result interpretation, i.e., the author might seek what he expected for in the review. We reduce the bias by (1) selecting all possible papers in a manner that is fair and seen to be fair and (2) discuss our findings based on statistical data we collected from the article characterisation. Also, our results are discussed among all the authors to reach an agreement.

6 CONCLUSION

In this paper we have reported on a systematic literature review on the *application perspective* of mutation testing, clearly contrasting previous systematic reviews that surveyed the whole field of mutation testing and that did not specifically go into how mutation testing is applied (e.g. [1]). We have characterised the studies that we have found on the basis of seven facets: (1) the role that mutation testing has in testing

activities; (2) the testing activities (including categories, test level and testing strategies); (3) the mutation tools used in the experiments; (4) the mutation operators used in the experiments; (5) the description of the equivalent mutant problem; (6) the description of cost reduction techniques for mutation testing; and (7) the subjects software systems involved in the experiments (in terms of programming language, size and data availability). These seven facets pertain to our two main research questions: **RQ1** *How is mutation testing used in testing activities?* and **RQ2** *How are empirical studies related to mutation testing designed and reported?*

Our main procedure to conduct this systematic literature review is shown in Figure 1. To collect all the relevant papers under our research scope, we started with search queries in online libraries considering 17 venues. We selected the literature that focuses on the supporting role of mutation testing in testing activities with sufficient evidence to suggest that mutation testing is used. After that, we performed a snowballing procedure to collect missing articles, thus resulting in a final selection of 159 papers in 21 venues. Through a detailed reading of this research body, we derived an attribute framework that was consequently used to characterise the studies in a structured manner. The resulting systematic literature review can be of benefit for researchers in the area of mutation testing. Specifically, we provide (1) guidelines on how to apply and subsequently report on mutation testing in testing experiments and (2) recommendations for future work.

The derived attribute framework is shown in Table 2. This attribute framework generalises and details the essential elements related to the *actual* application of mutation testing, such as in which circumstances mutation testing is used and which mutation testing tool is selected. In particular, a generic classification of mutation operators is constructed to study and compare the mutation operators used in the experiments described. This attribute framework can be used as a reference for researchers when describing mutation operators. We then presented the characterisation data of all the surveyed papers in Table 19-21 in Appendix. Based on our analysis of the results (in Section 4), four points are key to remember:

1. Most studies use mutation testing as an assessment tool; they target the unit level. Not only should we pay more attention to higher-level and specification mutation, but we should also study how mutation testing can be employed to improve the test code quality. Furthermore, we also encourage researchers to investigate and explore more interesting applications for mutation analysis in the future by asking such questions as: what else can we mutate? (Section 4.1.1-4.1.2)
2. Many of the supporting techniques for making mutation testing truly applicable are still under-developed. Also, existing mutation tools are not complete with regard to the mutation operators they offer. The two key problems, namely the equivalent mutant detection problem and the high computation cost of mutation testing issues, are not well-solved in the context of our research body (Section 4.2.1-4.2.4).
3. A deeper understanding of mutation testing is required, such as what particular kinds of faults mutation testing is good at finding. This would help the community to develop new mutation operators as well as overcome some of the inherent challenges (Section 4.3).
4. The awareness of *appropriately* reporting mutation testing in testing experiments should be raised among the researchers (Section 4.3).

In summary, the work described in this paper makes following contributions:

1. A systematic literature review of 159 studies that apply mutation testing in scientific experiments, which includes an in-depth analysis of how mutation testing is applied and reported on.
2. A detailed attribute framework that generalises and details the essential elements related to *actual* use of mutation testing
3. A generic classification of mutation operators that can be used to compare different mutation testing tools.
4. An actual characterisation of all the selected papers based on the attribute framework.
5. A series of recommendations for future work including important suggestions on how to report mutation testing in testing experiments in an appropriate manner.

REFERENCES

- [1] Y. Jia and M. Harman, “An analysis and survey of the development of mutation testing,” *Software Engineering, IEEE Transactions on*, vol. 37, no. 5, pp. 649–678, 2011.
- [2] J. Offutt, “A mutation carol: Past, present and future,” *Information and Software Technology*, vol. 53, no. 10, pp. 1098–1107, 2011.
- [3] R. Lipton, “Fault diagnosis of computer programs,” *Student Report, Carnegie Mellon University*, 1971.
- [4] R. A. DeMillo, R. J. Lipton, and F. G. Sayward, “Hints on test data selection: Help for the practicing programmer,” *Computer*, no. 4, pp. 34–41, 1978.
- [5] R. G. Hamlet, “Testing programs with the aid of a compiler,” *Software Engineering, IEEE Transactions on*, no. 4, pp. 279–290, 1977.
- [6] B. Kitchenham, “Guidelines for performing systematic literature reviews in software engineering,” in *Technical Report EBSE-2007-01*, 2007.
- [7] A. Offutt, “The coupling effect: fact or fiction,” in *ACM SIGSOFT Software Engineering Notes*, vol. 14, pp. 131–140, ACM, 1989.
- [8] A. J. Offutt, “Investigations of the software testing coupling effect,” *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 1, no. 1, pp. 5–20, 1992.
- [9] K. Wah, “Fault coupling in finite bijective functions,” *Software Testing, Verification and Reliability*, vol. 5, no. 1, pp. 3–47, 1995.
- [10] K. Wah, “A theoretical study of fault coupling,” *Software testing, verification and reliability*, vol. 10, no. 1, pp. 3–45, 2000.
- [11] K. H. T. Wah, “An analysis of the coupling effect i: single test data,” *Science of Computer Programming*, vol. 48, no. 2, pp. 119–161, 2003.
- [12] K. Kapoor, “Formal analysis of coupling hypothesis for logical faults,” *Innovations in Systems and Software Engineering*, vol. 2, no. 2, pp. 80–87, 2006.
- [13] T. A. Budd and D. Angluin, “Two notions of correctness and their relation to testing,” *Acta Informatica*, vol. 18, no. 1, pp. 31–45, 1982.
- [14] L. Madeyski, W. Orzeszyna, R. Torkar, and M. Józala, “Overcoming the equivalent mutant problem: A systematic literature review and a comparative experiment of second order mutation,” *Software Engineering, IEEE Transactions on*, vol. 40, no. 1, pp. 23–42, 2014.
- [15] P. Ammann and J. Offutt, *Introduction to software testing*. Cambridge University Press, 2008.
- [16] T. Budd and F. Sayward, “Users guide to the pilot mutation system,” *Yale University, New Haven, Connecticut, Technique Report*, vol. 114, 1977.
- [17] K. N. King and A. J. Offutt, “A fortran language system for mutation-based software testing,” *Software: Practice and Experience*, vol. 21, no. 7, pp. 685–718, 1991.
- [18] M. E. Delamaro, J. C. Maldonado, and A. Mathur, “Proteum—a tool for the assessment of test adequacy for c programs users guide,” in *PCS*, vol. 96, pp. 79–95, 1996.
- [19] Y.-S. Ma, J. Offutt, and Y.-R. Kwon, “Mujava: a mutation system for java,” in *Proceedings of the 28th international conference on Software engineering*, pp. 827–830, ACM, 2006.
- [20] J. Tuya, M. J. Suárez-Cabal, and C. De La Riva, “Mutating database queries,” *Information and Software Technology*, vol. 49, no. 4, pp. 398–417, 2007.
- [21] A. P. Mathur and W. E. Wong, “An empirical comparison of data flow and mutation-based test adequacy criteria,” *Software Testing, Verification and Reliability*, vol. 4, no. 1, pp. 9–31, 1994.
- [22] P. G. Frankl, S. N. Weiss, and C. Hu, “All-uses vs mutation testing: an experimental comparison of effectiveness,” *Journal of Systems and Software*, vol. 38, no. 3, pp. 235–253, 1997.
- [23] N. Li, U. Praphamontripong, and J. Offutt, “An experimental comparison of four unit test criteria: Mutation, edge-pair, all-uses and prime path coverage,” in *Software Testing, Verification and Validation Workshops, 2009. ICSTW’09. International Conference on*, pp. 220–229, IEEE, 2009.

- [24] J. H. Andrews, L. C. Briand, and Y. Labiche, "Is mutation an appropriate tool for testing experiments?[software testing]," in *Software Engineering, 2005. ICSE 2005. Proceedings. 27th International Conference on*, pp. 402–411, IEEE, 2005.
- [25] A. J. Offutt and R. H. Untch, "Mutation 2000: Uniting the orthogonal," in *Mutation testing for the new century*, pp. 34–44, Springer, 2001.
- [26] A. T. Acree Jr, "On mutation.," tech. rep., DTIC Document, 1980.
- [27] W. E. Wong and A. P. Mathur, "Reducing the cost of mutation testing: An empirical study," *Journal of Systems and Software*, vol. 31, no. 3, pp. 185–196, 1995.
- [28] S. Hussain, "Mutation clustering," *Ms. Th., King's College London, Strand, London*, 2008.
- [29] C. Ji, Z. Chen, B. Xu, and Z. Zhao, "A novel method of mutation clustering based on domain analysis.," in *SEKE*, vol. 9, pp. 422–425, 2009.
- [30] A. J. Offutt, A. Lee, G. Rothermel, R. H. Untch, and C. Zapf, "An experimental determination of sufficient mutant operators," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 5, no. 2, pp. 99–118, 1996.
- [31] E. S. Mresa and L. Bottaci, "Efficiency of mutation operators and selective mutation strategies: An empirical study," *Software Testing Verification and Reliability*, vol. 9, no. 4, pp. 205–232, 1999.
- [32] A. Siami Namin, J. H. Andrews, and D. J. Murdoch, "Sufficient mutation operators for measuring test effectiveness," in *Proceedings of the 30th international conference on Software engineering*, pp. 351–360, ACM, 2008.
- [33] W. E. Howden, "Weak mutation testing and completeness of test sets," *IEEE Transactions on Software Engineering*, no. 4, pp. 371–379, 1982.
- [34] A. J. Offutt and S. D. Lee, "An empirical evaluation of weak mutation," *IEEE Transactions on Software Engineering*, vol. 20, no. 5, pp. 337–344, 1994.
- [35] R. DeMillo and A. J. Offutt, "Constraint-based automatic test data generation," *Software Engineering, IEEE Transactions on*, vol. 17, no. 9, pp. 900–910, 1991.
- [36] R. H. Untch, "Mutation-based software testing using program schemata," in *Proceedings of the 30th annual Southeast regional conference*, pp. 285–291, ACM, 1992.
- [37] R. Untch, A. J. Offutt, and M. J. Harrold, "Mutation testing using mutant schemata," in *Proc. International Symposium on Software Testing and Analysis (ISSTA)*, pp. 139–148, 1993.
- [38] D. Baldwin and F. Sayward, "Heuristics for determining equivalence of program mutations.," tech. rep., DTIC Document, 1979.
- [39] R. Hierons, M. Harman, and S. Danicic, "Using program slicing to assist in the detection of equivalent mutants," *Software Testing, Verification and Reliability*, vol. 9, no. 4, pp. 233–262, 1999.
- [40] E. Martin and T. Xie, "A fault model and mutation testing of access control policies," in *Proceedings of the 16th international conference on World Wide Web*, pp. 667–676, ACM, 2007.
- [41] M. Ellims, D. Ince, and M. Petre, "The csaw c mutation tool: Initial results," in *Testing: Academic and Industrial Conference Practice and Research Techniques-MUTATION, 2007. TAICPART-MUTATION 2007*, pp. 185–192, IEEE, 2007.
- [42] L. du Bousquet and M. Delaunay, "Towards mutation analysis for lustre programs," *Electronic Notes in Theoretical Computer Science*, vol. 203, no. 4, pp. 35–48, 2008.
- [43] M. Harman, R. Hierons, and S. Danicic, "The relationship between program dependence and mutation analysis," in *Mutation testing for the new century*, pp. 5–13, Springer, 2001.
- [44] K. Adamopoulos, M. Harman, and R. M. Hierons, "How to overcome the equivalent mutant problem and achieve tailored selective mutation using co-evolution," in *Genetic and evolutionary computation conference*, pp. 1338–1349, Springer, 2004.
- [45] A. M. R. Vincenzi, E. Y. Nakagawa, J. C. Maldonado, M. E. Delamaro, and R. A. F. Romero, "Bayesian-learning based guidelines to determine equivalent mutants," *International Journal of Software Engineering and Knowledge Engineering*, vol. 12, no. 06, pp. 675–689, 2002.

- [46] D. Schuler, V. Dallmeier, and A. Zeller, "Efficient mutation testing by checking invariant violations," in *Proceedings of the eighteenth international symposium on Software testing and analysis*, pp. 69–80, ACM, 2009.
- [47] D. Schuler and A. Zeller, "(un-) covering equivalent mutants," in *2010 Third International Conference on Software Testing, Verification and Validation*, pp. 45–54, IEEE, 2010.
- [48] T. A. Budd, R. J. Lipton, R. A. DeMillo, and F. G. Sayward, *Mutation analysis*. Yale University, Department of Computer Science, 1979.
- [49] H. Agrawal, R. DeMillo, R. Hathaway, W. Hsu, W. Hsu, E. Krauser, R. J. Martin, A. Mathur, and E. Spafford, "Design of mutant operators for the c programming language," tech. rep., Technical Report SERC-TR-41-P, Software Engineering Research Center, Department of Computer Science, Purdue University, Indiana, 1989.
- [50] R. A. DeMillo, "Test adequacy and program mutation," in *Software Engineering, 1989. 11th International Conference on*, pp. 355–356, May 1989.
- [51] S. C. Ntafos, "On testing with required elements," in *Proceedings of COMPSAC*, vol. 81, pp. 132–139, 1981.
- [52] S. C. Ntafos, "An evaluation of required element testing strategies," in *Proceedings of the 7th international conference on Software engineering*, pp. 250–256, IEEE Press, 1984.
- [53] S. C. Ntafos, "On required element testing," *Software Engineering, IEEE Transactions on*, no. 6, pp. 795–803, 1984.
- [54] J. W. Duran and S. C. Ntafos, "An evaluation of random testing," *Software Engineering, IEEE Transactions on*, no. 4, pp. 438–444, 1984.
- [55] L. Zhang, T. Xie, L. Zhang, N. Tillmann, J. De Halleux, and H. Mei, "Test generation via dynamic symbolic execution for mutation testing," in *Software Maintenance (ICSM), 2010 IEEE International Conference on*, pp. 1–10, IEEE, 2010.
- [56] M. Papadakis and N. Malevris, "Automatically performing weak mutation with the aid of symbolic execution, concolic testing and search-based testing," *Software Quality Journal*, vol. 19, no. 4, pp. 691–723, 2011.
- [57] A. S. Namin and J. H. Andrews, "The influence of size and coverage on test suite effectiveness," in *Proceedings of the eighteenth international symposium on Software testing and analysis*, pp. 57–68, ACM, 2009.
- [58] L. Inozemtseva and R. Holmes, "Coverage is not strongly correlated with test suite effectiveness," in *Proceedings of the 36th International Conference on Software Engineering*, pp. 435–445, ACM, 2014.
- [59] Y. Zhang and A. Mesbah, "Assertions are strongly correlated with test suite effectiveness," in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE*, pp. 214–224, 2015.
- [60] M. Whalen, G. Gay, D. You, M. P. Heimdahl, and M. Staats, "Observable modified condition/decision coverage," in *Proceedings of the 2013 International Conference on Software Engineering*, pp. 102–111, IEEE Press, 2013.
- [61] G. Rothermel, R. H. Untch, C. Chu, and M. J. Harrold, "Prioritizing test cases for regression testing," *Software Engineering, IEEE Transactions on*, vol. 27, no. 10, pp. 929–948, 2001.
- [62] S. Elbaum, A. G. Malishevsky, and G. Rothermel, "Test case prioritization: A family of empirical studies," *Software Engineering, IEEE Transactions on*, vol. 28, no. 2, pp. 159–182, 2002.
- [63] H. Do and G. Rothermel, "A controlled experiment assessing test case prioritization techniques via mutation faults," in *Software Maintenance, 2005. ICSM'05. Proceedings of the 21st IEEE International Conference on*, pp. 411–420, IEEE, 2005.
- [64] H. Do and G. Rothermel, "On the use of mutation faults in empirical assessments of test case prioritization techniques," *Software Engineering, IEEE Transactions on*, vol. 32, no. 9, pp. 733–752, 2006.

- [65] A. J. Offutt, J. Pan, and J. M. Voas, "Procedures for reducing the size of coverage-based test sets," in *Proceedings of the Twelfth International Conference on Testing Computer Software*, Citeseer, 1995.
- [66] L. Zhang, D. Marinov, L. Zhang, and S. Khurshid, "An empirical study of junit test-suite reduction," in *Software Reliability Engineering (ISSRE), 2011 IEEE 22nd International Symposium on*, pp. 170–179, IEEE, 2011.
- [67] X. Wang, S.-C. Cheung, W. K. Chan, and Z. Zhang, "Taming coincidental correctness: Coverage refinement with context patterns to improve fault localization," in *Proceedings of the 31st International Conference on Software Engineering*, pp. 45–55, IEEE Computer Society, 2009.
- [68] M. Papadakis and Y. Le Traon, "Using mutants to locate "unknown" faults," in *Software Testing, Verification and Validation (ICST), 2012 IEEE Fifth International Conference on*, pp. 691–700, IEEE, 2012.
- [69] M. Papadakis and Y. Le Traon, "Metallaxis-fl: mutation-based fault localization," *Software Testing, Verification and Reliability*, vol. 25, no. 5-7, pp. 605–628, 2015.
- [70] L. Kysh, "Difference between a systematic review and a literature review." <https://dx.doi.org/10.6084/m9.figshare.766364.v1>, 2013. [Online; accessed 4-August-2016].
- [71] P. Brereton, B. A. Kitchenham, D. Budgen, M. Turner, and M. Khalil, "Lessons from applying the systematic literature review process within the software engineering domain," *Journal of systems and software*, vol. 80, no. 4, pp. 571–583, 2007.
- [72] B. Cornelissen, A. Zaidman, A. Van Deursen, L. Moonen, and R. Koschke, "A systematic survey of program comprehension through dynamic analysis," *IEEE Transactions on Software Engineering*, vol. 35, no. 5, pp. 684–702, 2009.
- [73] C. Wohlin, "Guidelines for snowballing in systematic literature studies and a replication in software engineering," in *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering*, p. 38, ACM, 2014.
- [74] D. Graham, E. Van Veenendaal, and I. Evans, *Foundations of software testing: ISTQB certification*. Cengage Learning EMEA, 2008.
- [75] A. van Deursen, "Software Testing in 2048." <https://speakerdeck.com/avandeursen/software-testing-in-2048>, 1 2016. [Online; accessed 13-Sep-2016].
- [76] M. Papadakis and N. Malevris, "Automatic mutation test case generation via dynamic symbolic execution," in *Software reliability engineering (ISSRE), 2010 IEEE 21st international symposium on*, pp. 121–130, IEEE, 2010.
- [77] "Available mutation operations (PIT)." <http://pitest.org/quickstart/mutators/>. [Online; accessed 10-August-2016].
- [78] Y.-S. Ma and J. Offutt, "Description of class mutation operators for java," *Electronics and Telecommunications Research Institute, Korea*, 2005.
- [79] S. Mirshokraie, A. Mesbah, and K. Pattabiraman, "Efficient javascript mutation testing," in *Software Testing, Verification and Validation (ICST), 2013 IEEE Sixth International Conference on*, pp. 74–83, IEEE, 2013.
- [80] S. Außerlechner, S. Fruhmann, W. Wieser, B. Hofer, R. Spörk, C. Mühlbacher, and F. Wotawa, "The right choice matters! smt solving substantially improves model-based debugging of spreadsheets," in *2013 13th International Conference on Quality Software*, pp. 139–148, IEEE, 2013.
- [81] R. Delamare, B. Baudry, S. Ghosh, S. Gupta, and Y. Le Traon, "An approach for testing pointcut descriptors in aspectj," *Software Testing, Verification and Reliability*, vol. 21, no. 3, pp. 215–239, 2011.
- [82] D. Xu and J. Ding, "Prioritizing state-based aspect tests," in *2010 Third International Conference on Software Testing, Verification and Validation*, pp. 265–274, IEEE, 2010.
- [83] S. Hong, M. Staats, J. Ahn, M. Kim, and G. Rothermel, "The impact of concurrent coverage metrics on testing effectiveness," in *Software Testing, Verification and Validation (ICST), 2013 IEEE Sixth International Conference on*, pp. 232–241, IEEE, 2013.

- [84] S. Hong, M. Staats, J. Ahn, M. Kim, and G. Rothermel, "Are concurrency coverage metrics effective for testing: a comprehensive empirical investigation," *Software Testing, Verification and Reliability*, vol. 25, no. 4, pp. 334–370, 2015.
- [85] S.-S. Hou, L. Zhang, T. Xie, H. Mei, and J.-S. Sun, "Applying interface-contract mutation in regression testing of component-based software," in *Software Maintenance, 2007. ICSM 2007. IEEE International Conference on*, pp. 174–183, IEEE, 2007.
- [86] H. Yoon and B. Choi, "Effective test case selection for component customization and its application to enterprise javabeans," *Software Testing, Verification and Reliability*, vol. 14, no. 1, pp. 45–70, 2004.
- [87] D. Schuler and A. Zeller, "Javalanche: efficient mutation testing for java," in *Proceedings of the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering*, pp. 297–298, ACM, 2009.
- [88] H. Do, S. Elbaum, and G. Rothermel, "Supporting controlled experimentation with testing techniques: An infrastructure and its potential impact," *Empirical Software Engineering*, vol. 10, no. 4, pp. 405–435, 2005.
- [89] G. Fraser and A. Arcuri, "Sound empirical evidence in software testing," in *Proceedings of the 34th International Conference on Software Engineering*, pp. 178–188, IEEE Press, 2012.
- [90] G. Rothermel, R. H. Untch, C. Chu, and M. J. Harrold, "Test case prioritization: An empirical study," in *Software Maintenance, 1999.(ICSM'99) Proceedings. IEEE International Conference on*, pp. 179–188, IEEE, 1999.
- [91] H. Czemerinski, V. Braberman, and S. Uchitel, "Behaviour abstraction adequacy criteria for api call protocol testing," *Software Testing, Verification and Reliability*, 2015.
- [92] B. Baudry, V. Le Hanh, J.-M. Jézéquel, and Y. Le Traon, "Building trust into oo components using a genetic analogy," in *Software Reliability Engineering, 2000. ISSRE 2000. Proceedings. 11th International Symposium on*, pp. 4–14, IEEE, 2000.
- [93] B. Baudry, F. Fleurey, J.-M. Jézéquel, and Y. Le Traon, "Genes and bacteria for automatic test cases optimization in the .net environment," in *Software Reliability Engineering, 2002. ISSRE 2003. Proceedings. 13th International Symposium on*, pp. 195–206, IEEE, 2002.
- [94] B. Baudry, F. Fleurey, J.-M. Jézéquel, and Y. Le Traon, "From genetic to bacteriological algorithms for mutation-based testing," *Software Testing, Verification and Reliability*, vol. 15, no. 2, pp. 73–96, 2005.
- [95] A. von Mayrhauser, M. Scheetz, E. Dahlman, and A. E. Howe, "Planner based error recovery testing," in *Software Reliability Engineering, 2000. ISSRE 2000. Proceedings. 11th International Symposium on*, pp. 186–195, IEEE, 2000.
- [96] B. H. Smith and L. Williams, "On guiding the augmentation of an automated test suite via mutation analysis," *Empirical Software Engineering*, vol. 14, no. 3, pp. 341–369, 2009.
- [97] X. Qu, M. B. Cohen, and G. Rothermel, "Configuration-aware regression testing: an empirical study of sampling and prioritization," in *Proceedings of the 2008 international symposium on Software testing and analysis*, pp. 75–86, ACM, 2008.
- [98] J.-H. Kwon, I.-Y. Ko, G. Rothermel, and M. Staats, "Test case prioritization based on information retrieval concepts," in *Software Engineering Conference (APSEC), 2014 21st Asia-Pacific*, vol. 1, pp. 19–26, IEEE, 2014.
- [99] Y. Qi, X. Mao, and Y. Lei, "Efficient automated program repair through fault-recorded testing prioritization," in *2013 IEEE International Conference on Software Maintenance*, pp. 180–189, IEEE, 2013.
- [100] S. S. Murtaza, N. Madhavji, M. Gittens, and Z. Li, "Diagnosing new faults using mutants and prior faults (nier track)," in *Software Engineering (ICSE), 2011 33rd International Conference on*, pp. 960–963, IEEE, 2011.
- [101] S. Moon, Y. Kim, M. Kim, and S. Yoo, "Ask the mutants: Mutating faulty programs for fault localization," in *2014 IEEE Seventh International Conference on Software Testing, Verification and Validation*, pp. 153–162, IEEE, 2014.

- [102] G. Fraser and A. Zeller, "Mutation-driven generation of unit tests and oracles," *Software Engineering, IEEE Transactions on*, vol. 38, no. 2, pp. 278–292, 2012.
- [103] M. Staats, G. Gay, and M. P. Heimdahl, "Automated oracle creation support, or: how i learned to stop worrying about fault propagation and love mutation testing," in *Proceedings of the 34th International Conference on Software Engineering*, pp. 870–880, IEEE Press, 2012.
- [104] G. Gay, M. Staats, M. Whalen, and M. P. Heimdahl, "Automated oracle data selection support," *IEEE Transactions on Software Engineering*, vol. 41, no. 11, pp. 1119–1137, 2015.
- [105] A. Shi, A. Gyori, M. Gligoric, A. Zaytsev, and D. Marinov, "Balancing trade-offs in test-suite reduction," in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pp. 246–256, ACM, 2014.
- [106] W. Chen, R. H. Untch, G. Rothermel, S. Elbaum, and J. Von Ronne, "Can fault-exposure-potential estimates improve the fault detection abilities of test suites?," *Software Testing, Verification and Reliability*, vol. 12, no. 4, pp. 197–218, 2002.
- [107] B. H. Smith and L. Williams, "Should software testers use mutation analysis to augment a test set?," *Journal of Systems and Software*, vol. 82, no. 11, pp. 1819–1832, 2009.
- [108] Y. Lou, D. Hao, and L. Zhang, "Mutation-based test-case prioritization in software evolution," in *Software Reliability Engineering (ISSRE), 2015 IEEE 26th International Symposium on*, pp. 46–57, IEEE, 2015.
- [109] D. Hao, L. Zhang, X. Wu, H. Mei, and G. Rothermel, "On-demand test suite reduction," in *Proceedings of the 34th International Conference on Software Engineering*, pp. 738–748, IEEE Press, 2012.
- [110] D. Hao, L. Zhang, L. Zhang, G. Rothermel, and H. Mei, "A unified test case prioritization approach," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 24, no. 2, p. 10, 2014.
- [111] P. Li, T. Huynh, M. Reformat, and J. Miller, "A practical approach to testing gui systems," *Empirical Software Engineering*, vol. 12, no. 4, pp. 331–357, 2007.
- [112] M. J. Rutherford, A. Carzaniga, and A. L. Wolf, "Evaluating test suites and adequacy criteria using simulation-based models of distributed systems," *Software Engineering, IEEE Transactions on*, vol. 34, no. 4, pp. 452–470, 2008.
- [113] G. Denaro, A. Margara, M. Pezze, and M. Vivanti, "Dynamic data flow testing of object oriented systems," in *Proceedings of the 37th International Conference on Software Engineering-Volume 1*, pp. 947–958, IEEE Press, 2015.
- [114] M. E. Delamaro, J. Maidonado, and A. P. Mathur, "Interface mutation: An approach for integration testing," *IEEE Transactions on Software Engineering*, vol. 27, no. 3, pp. 228–247, 2001.
- [115] P. R. Mateo, M. P. Usaola, and J. Offutt, "Mutation at the multi-class and system levels," *Science of Computer Programming*, vol. 78, no. 4, pp. 364–387, 2013.
- [116] M. Hennessy and J. F. Power, "Analysing the effectiveness of rule-coverage as a reduction criterion for test suites of grammar-based software," *Empirical Software Engineering*, vol. 13, no. 4, pp. 343–368, 2008.
- [117] H. S. Chae, G. Woo, T. Y. Kim, J. H. Bae, and W.-Y. Kim, "An automated approach to reducing test suites for testing retargeted c compilers for embedded systems," *Journal of Systems and Software*, vol. 84, no. 12, pp. 2053–2064, 2011.
- [118] F. Belli and M. Beyazit, "Exploiting model morphology for event-based testing," *IEEE Transactions on Software Engineering*, vol. 41, no. 2, pp. 113–134, 2015.
- [119] B. Hofer and F. Wotawa, "Why does my spreadsheet compute wrong values?," in *Software Reliability Engineering (ISSRE), 2014 IEEE 25th International Symposium on*, pp. 112–121, IEEE, 2014.
- [120] B. Hofer, A. Perez, R. Abreu, and F. Wotawa, "On the empirical evaluation of similarity coefficients for spreadsheets fault localization," *Automated Software Engineering*, vol. 22, no. 1, pp. 47–74, 2015.
- [121] J. Tuya, M. J. Suárez-Cabal, and C. De La Riva, "Full predicate coverage for testing sql database queries," *Software Testing, Verification and Reliability*, vol. 20, no. 3, pp. 237–288, 2010.

- [122] G. M. Kapfhammer, P. McMinn, and C. J. Wright, "Search-based testing of relational schema integrity constraints across multiple database management systems," in *2013 IEEE Sixth International Conference on Software Testing, Verification and Validation*, pp. 31–40, IEEE, 2013.
- [123] P. McMinn, C. J. Wright, and G. M. Kapfhammer, "The effectiveness of test coverage criteria for relational database schema integrity constraints," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 25, no. 1, p. 8, 2015.
- [124] M. Papadakis, C. Henard, and Y. Le Traon, "Sampling program inputs with mutation analysis: Going beyond combinatorial interaction testing," in *Software Testing, Verification and Validation (ICST), 2014 IEEE Seventh International Conference on*, pp. 1–10, IEEE, 2014.
- [125] Y.-S. Ma, J. Offutt, and Y. R. Kwon, "Mujava: an automated class mutation system," *Software Testing, Verification and Reliability*, vol. 15, no. 2, pp. 97–133, 2005.
- [126] "Mutation testing systems for Java compared." http://pitest.org/java_mutation_testing_systems/. [Online; accessed 25-August-2016].
- [127] A. S. Namin and S. Kakarla, "The use of mutation in testing experiments and its sensitivity to external threats," in *Proceedings of the 2011 International Symposium on Software Testing and Analysis*, pp. 342–352, ACM, 2011.
- [128] R. Just, F. Schweiggert, and G. M. Kapfhammer, "Major: An efficient and extensible tool for mutation analysis in a java compiler," in *Proceedings of the 2011 26th IEEE/ACM International Conference on Automated Software Engineering*, pp. 612–615, IEEE Computer Society, 2011.
- [129] S. A. Irvine, T. Pavlinic, L. Trigg, J. G. Cleary, S. Inglis, and M. Utting, "Jumble java byte code to measure the effectiveness of unit tests," in *Testing: Academic and industrial conference practice and research techniques-MUTATION, 2007. TAICPART-MUTATION 2007*, pp. 169–175, IEEE, 2007.
- [130] W. Motycka, "Installation Instructions." <http://sofya.unl.edu/doc/manual/installation.html>, 7 2013. [Online; accessed 25-August-2016].
- [131] Y. Jia and M. Harman, "Milu: A customizable, runtime-optimized higher order mutation testing tool for the full c language," in *Practice and Research Techniques, 2008. TAIC PART'08. Testing: Academic & Industrial Conference*, pp. 94–98, IEEE, 2008.
- [132] H. Dan and R. M. Hierons, "Smt-c: A semantic mutation testing tools for c," in *2012 IEEE Fifth International Conference on Software Testing, Verification and Validation*, pp. 654–663, IEEE, 2012.
- [133] R. Delamare, B. Baudry, and Y. Le Traon, "Ajmutator: a tool for the mutation analysis of aspectj pointcut descriptors," in *Software Testing, Verification and Validation Workshops, 2009. ICSTW'09. International Conference on*, pp. 200–204, IEEE, 2009.
- [134] W. Krenn, R. Schlick, S. Tiran, B. Aichernig, E. Jobstl, and H. Brandl, "Momut:: Uml model-based mutation testing for uml," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, pp. 1–8, IEEE, 2015.
- [135] A. J. Offutt and W. M. Craft, "Using compiler optimization techniques to detect equivalent mutants," *Software Testing, Verification and Reliability*, vol. 4, no. 3, pp. 131–154, 1994.
- [136] A. J. Offutt and J. Pan, "Automatically detecting equivalent mutants and infeasible paths," *Software testing, verification and reliability*, vol. 7, no. 3, pp. 165–192, 1997.
- [137] M. R. Lyu, Z. Huang, S. K. Sze, and X. Cai, "An empirical study on testing and fault tolerance for software reliability engineering," in *Software Reliability Engineering, 2003. ISSRE 2003. 14th International Symposium on*, pp. 119–130, IEEE, 2003.
- [138] "Mutation Testing and Error Seeding-White Box Testing Techniques." <http://www.softwaretestinggenius.com/mutation-testing-and-error-seeding-white-box-testing-techniques>. [Online; accessed 28-July-2016].
- [139] M. S. AbouTrab, M. Brockway, S. Counsell, and R. M. Hierons, "Testing real-time embedded systems using timed automata based approaches," *Journal of Systems and Software*, vol. 86, no. 5, pp. 1209–1223, 2013.

- [140] B. K. Aichernig, H. Brandl, E. Jöbstl, and W. Krenn, "Efficient mutation killers in action," in *2011 Fourth IEEE International Conference on Software Testing, Verification and Validation*, pp. 120–129, IEEE, 2011.
- [141] B. K. Aichernig, H. Brandl, E. Jöbstl, W. Krenn, R. Schlick, and S. Tiran, "Killing strategies for model-based mutation testing," *Software Testing, Verification and Reliability*, vol. 25, no. 8, pp. 716–748, 2015.
- [142] S. Ali, L. C. Briand, M. J.-u. Rehman, H. Asghar, M. Z. Z. Iqbal, and A. Nadeem, "A state-based approach to integration testing based on uml models," *Information and Software Technology*, vol. 49, no. 11, pp. 1087–1106, 2007.
- [143] J. H. Andrews and Y. Zhang, "General test result checking with log file analysis," *Software Engineering, IEEE Transactions on*, vol. 29, no. 7, pp. 634–648, 2003.
- [144] J. H. Andrews, L. C. Briand, Y. Labiche, and A. S. Namin, "Using mutation analysis for assessing and comparing testing coverage criteria," *Software Engineering, IEEE Transactions on*, vol. 32, no. 8, pp. 608–624, 2006.
- [145] K. Androutsopoulos, D. Clark, H. Dan, R. M. Hierons, and M. Harman, "An analysis of the relationship between conditional entropy and failed error propagation in software testing," in *Proceedings of the 36th International Conference on Software Engineering*, pp. 573–583, ACM, 2014.
- [146] G. Antoniol, L. C. Briand, M. D. Penta, and Y. Labiche, "A case study using the round-trip strategy for state-based class testing," in *Software Reliability Engineering, 2002. ISSRE 2003. Proceedings. 13th International Symposium on*, pp. 269–279, IEEE, 2002.
- [147] P. Arcaini, A. Gargantini, and E. Riccobene, "Using mutation to assess fault detection capability of model review," *Software Testing, Verification and Reliability*, vol. 25, no. 5-7, pp. 629–652, 2015.
- [148] A. Arcuri and L. Briand, "Adaptive random testing: An illusion of effectiveness?," in *Proceedings of the 2011 International Symposium on Software Testing and Analysis*, pp. 265–275, ACM, 2011.
- [149] N. Asoudeh and Y. Labiche, "Multi-objective construction of an entire adequate test suite for an efsm," in *Software Reliability Engineering (ISSRE), 2014 IEEE 25th International Symposium on*, pp. 288–299, IEEE, 2014.
- [150] R. Baker and I. Habli, "An empirical evaluation of mutation testing for improving the test quality of safety-critical software," *IEEE Transactions on Software Engineering*, vol. 39, no. 6, pp. 787–805, 2013.
- [151] A. Bandyopadhyay and S. Ghosh, "Test input generation using uml sequence and state machines models," in *2009 International Conference on Software Testing Verification and Validation*, pp. 121–130, IEEE, 2009.
- [152] S. Bardin, M. Delahaye, R. David, N. Kosmatov, M. Papadakis, Y. Le Traon, and J.-Y. Marion, "Sound and quasi-complete detection of infeasible test requirements," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, pp. 1–10, IEEE, 2015.
- [153] F. Belli, M. Beyazit, T. Takagi, and Z. Furukawa, "Mutation testing of "go-back" functions based on pushdown automata," in *2011 Fourth IEEE International Conference on Software Testing, Verification and Validation*, pp. 249–258, IEEE, 2011.
- [154] A. Bertolino, S. Daoudagh, F. Lonetti, and E. Marchetti, "Automatic xacml requests generation for policy testing," in *2012 IEEE Fifth International Conference on Software Testing, Verification and Validation*, pp. 842–849, IEEE, 2012.
- [155] L. C. Briand, M. Di Penta, and Y. Labiche, "Assessing and improving state-based class testing: A series of experiments," *Software Engineering, IEEE Transactions on*, vol. 30, no. 11, pp. 770–783, 2004.
- [156] L. C. Briand, Y. Labiche, and Y. Wang, "Using simulation to empirically investigate test coverage criteria based on statechart," in *Proceedings of the 26th International Conference on Software Engineering*, pp. 86–95, IEEE Computer Society, 2004.

- [157] L. C. Briand, Y. Labiche, and Q. Lin, "Improving statechart testing criteria using data flow information," in *Software Reliability Engineering, 2005. ISSRE 2005. 16th IEEE International Symposium on*, pp. 10–pp, IEEE, 2005.
- [158] Y. Cheon, "Abstraction in assertion-based test oracles," in *Seventh International Conference on Quality Software (QSIC 2007)*, pp. 410–414, IEEE, 2007.
- [159] P. Chevalley and P. Thévenod-Fosse, "An empirical evaluation of statistical testing designed from uml state diagrams: the flight guidance system case study," in *Software Reliability Engineering, 2001. ISSRE 2001. Proceedings. 12th International Symposium on*, pp. 254–263, IEEE, 2001.
- [160] H. Czemerinski, V. Braberman, and S. Uchitel, "Behaviour abstraction coverage as black-box adequacy criteria," in *Software Testing, Verification and Validation (ICST), 2013 IEEE Sixth International Conference on*, pp. 222–231, IEEE, 2013.
- [161] R. A. DeMillo and A. J. Offutt, "Experimental results from an automatic test case generator," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 2, no. 2, pp. 109–127, 1993.
- [162] D. Di Nardo, F. Pastore, and L. Briand, "Generating complex and faulty test data through model-based mutation analysis," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, pp. 1–10, IEEE, 2015.
- [163] H. Do, S. Mirarab, L. Tahvildari, and G. Rothermel, "An empirical study of the effect of time constraints on the cost-benefits of regression testing," in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of software engineering*, pp. 71–82, ACM, 2008.
- [164] S. H. Edwards and Z. Shams, "Comparing test quality measures for assessing student-written tests," in *Companion Proceedings of the 36th International Conference on Software Engineering*, pp. 354–363, ACM, 2014.
- [165] S. H. Edwards, "Black-box testing using flowgraphs: an experimental assessment of effectiveness and automation potential," *Softw. Test., Verif. Reliab.*, vol. 10, no. 4, pp. 249–262, 2000.
- [166] C. Fang, Z. Chen, K. Wu, and Z. Zhao, "Similarity-based test case prioritization using ordered sequences of program entities," *Software Quality Journal*, vol. 22, no. 2, pp. 335–361, 2014.
- [167] G. Fraser and N. Walkinshaw, "Assessing and generating test sets in terms of behavioural adequacy," *Software Testing, Verification and Reliability*, vol. 25, no. 8, pp. 749–780, 2015.
- [168] G. Fraser, M. Staats, P. McMinn, A. Arcuri, and F. Padberg, "Does automated white-box test generation really help software testers?," in *Proceedings of the 2013 International Symposium on Software Testing and Analysis*, pp. 291–301, ACM, 2013.
- [169] G. Fraser and A. Arcuri, "Achieving scalable mutation-based generation of whole test suites," *Empirical Software Engineering*, vol. 20, no. 3, pp. 783–812, 2015.
- [170] G. Gay, M. Staats, M. Whalen, and M. P. Heimdahl, "The risks of coverage-directed test case generation," *Software Engineering, IEEE Transactions on*, vol. 41, no. 8, pp. 803–819, 2015.
- [171] M. Gligoric, A. Groce, C. Zhang, R. Sharma, M. A. Alipour, and D. Marinov, "Comparing non-adequate test suites using coverage criteria," in *Proceedings of the 2013 International Symposium on Software Testing and Analysis*, pp. 302–313, ACM, 2013.
- [172] A. Gonzalez-Sanchez, R. Abreu, H.-G. Gross, and A. J. Van Gemund, "Prioritizing tests for fault localization through ambiguity group reduction," in *Automated Software Engineering (ASE), 2011 26th IEEE/ACM International Conference on*, pp. 83–92, IEEE, 2011.
- [173] R. Gopinath, C. Jensen, and A. Groce, "Code coverage for suite evaluation by developers," in *Proceedings of the 36th International Conference on Software Engineering*, pp. 72–82, ACM, 2014.
- [174] A. Gupta and P. Jalote, "Test inspected unit or inspect unit tested code?," in *Empirical Software Engineering and Measurement, 2007. ESEM 2007. First International Symposium on*, pp. 51–60, IEEE, 2007.
- [175] M. Harman, Y. Jia, and W. B. Langdon, "Strong higher order mutation-based test data generation," in *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*, pp. 212–222, ACM, 2011.

- [176] N. E. Holt, R. Torkar, L. Briand, and K. Hansen, "State-based testing: Industrial evaluation of the cost-effectiveness of round-trip path and sneak-path strategies," in *Software Reliability Engineering (ISSRE), 2012 IEEE 23rd International Symposium on*, pp. 321–330, IEEE, 2012.
- [177] K. Jamrozik, G. Fraser, N. Tillmann, and J. De Halleux, "Augmented dynamic symbolic execution," in *Automated Software Engineering (ASE), 2012 Proceedings of the 27th IEEE/ACM International Conference on*, pp. 254–257, IEEE, 2012.
- [178] E. Jee, D. Shin, S. Cha, J.-S. Lee, and D.-H. Bae, "Automated test case generation for fbd programs implementing reactor protection system software," *Software Testing, Verification and Reliability*, vol. 24, no. 8, pp. 608–628, 2014.
- [179] S. A. Jolly, V. Garousi, and M. M. Eskandar, "Automated unit testing of a scada control software: an industrial case study based on action research," in *Software Testing, Verification and Validation (ICST), 2012 IEEE Fifth International Conference on*, pp. 400–409, IEEE, 2012.
- [180] U. Kanewala and J. M. Bieman, "Using machine learning techniques to detect metamorphic relations for programs without test oracles," in *2013 IEEE 24th International Symposium on Software Reliability Engineering (ISSRE)*, pp. 1–10, IEEE, 2013.
- [181] M. Khalil and Y. Labiche, "On the round trip path testing strategy," in *Software Reliability Engineering (ISSRE), 2010 IEEE 21st International Symposium on*, pp. 388–397, IEEE, 2010.
- [182] S.-W. Kim, J. A. Clark, and J. A. McDermid, "Investigating the effectiveness of object-oriented testing strategies using the mutation method," *Software Testing, Verification and Reliability*, vol. 11, no. 4, pp. 207–225, 2001.
- [183] K. Koster and D. Kao, "State coverage: a structural test adequacy criterion for behavior checking," in *The 6th Joint Meeting on European software engineering conference and the ACM SIGSOFT symposium on the foundations of software engineering: companion papers*, pp. 541–544, ACM, 2007.
- [184] J. S. Kracht, J. Z. Petrovic, and K. R. Walcott-Justice, "Empirically evaluating the quality of automatically generated and manually written test suites," in *Quality Software (QSIC), 2014 14th International Conference on*, pp. 256–265, IEEE, 2014.
- [185] Y. Le Traon, B. Baudry, and J.-M. Jézéquel, "Design by contract to improve software vigilance," *IEEE Transactions on Software Engineering*, vol. 32, no. 8, pp. 571–586, 2006.
- [186] Y. Le Traon, T. Mouelhi, and B. Baudry, "Testing security policies: going beyond functional testing," in *The 18th IEEE International Symposium on Software Reliability (ISSRE'07)*, pp. 93–102, IEEE, 2007.
- [187] S. C. Lee and J. Offutt, "Generating test cases for xml-based web component interactions using mutation analysis," in *Software Reliability Engineering, 2001. ISSRE 2001. Proceedings. 12th International Symposium on*, pp. 200–209, IEEE, 2001.
- [188] Y. Lei and J. H. Andrews, "Minimization of randomized unit test cases," in *Software Reliability Engineering, 2005. ISSRE 2005. 16th IEEE International Symposium on*, pp. 10–pp, IEEE, 2005.
- [189] M.-H. Liu, Y.-F. Gao, J.-H. Shan, J.-H. Liu, L. Zhang, and J.-S. Sun, "An approach to test data generation for killing multiple mutants," in *Software Maintenance, 2006. ICSM 2006. IEEE International Conference on*, pp. 113–122, IEEE, 2006.
- [190] F. Lorber, "Model-based mutation testing of synchronous and asynchronous real-time systems," in *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*, pp. 1–2, IEEE, 2015.
- [191] P. Loyola, M. Staats, I.-Y. Ko, and G. Rothermel, "Dodona: automated oracle data set selection," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, pp. 193–203, ACM, 2014.
- [192] J. Mayer and R. Guderlei, "On random testing of image processing applications," in *2006 Sixth International Conference on Quality Software (QSIC'06)*, pp. 85–92, IEEE, 2006.
- [193] J. Mayer and C. Schneckenburger, "An empirical analysis and comparison of random testing techniques," in *Proceedings of the 2006 ACM/IEEE international symposium on Empirical software engineering*, pp. 105–114, ACM, 2006.

- [194] A. Milani Fard, M. Mirzaaghaei, and A. Mesbah, "Leveraging existing tests in automated test generation for web applications," in *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*, pp. 67–78, ACM, 2014.
- [195] T. Miller and P. Strooper, "A case study in model-based testing of specifications and implementations," *Software Testing, Verification and Reliability*, vol. 22, no. 1, pp. 33–63, 2012.
- [196] S. Mirshokraie, "Effective test generation and adequacy assessment for javascript-based web applications," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, pp. 453–456, ACM, 2014.
- [197] J.-M. Mottu, S. Sen, M. Tisi, and J. Cabot, "Static analysis of model transformations for effective test generation," in *2012 IEEE 23rd International Symposium on Software Reliability Engineering*, pp. 291–300, IEEE, 2012.
- [198] S. Mouchawrab, L. C. Briand, and Y. Labiche, "Assessing, comparing, and combining statechart-based testing and structural testing: An experiment," in *Empirical Software Engineering and Measurement, 2007. ESEM 2007. First International Symposium on*, pp. 41–50, IEEE, 2007.
- [199] S. Mouchawrab, L. C. Briand, Y. Labiche, and M. Di Penta, "Assessing, comparing, and combining state machine-based testing and structural testing: a series of experiments," *Software Engineering, IEEE Transactions on*, vol. 37, no. 2, pp. 161–187, 2011.
- [200] A. J. Offutt and S. Liu, "Generating test data from soft specifications," *Journal of Systems and Software*, vol. 49, no. 1, pp. 49–62, 1999.
- [201] M. Papadakis and N. Malevris, "Mutation based test case generation via a path selection strategy," *Information and Software Technology*, vol. 54, no. 9, pp. 915–932, 2012.
- [202] M. Patrick, R. Alexander, M. Oriol, and J. A. Clark, "Subdomain-based test data generation," *Journal of Systems and Software*, vol. 103, pp. 328–342, 2015.
- [203] F. Pinte, N. Oster, and F. Saglietti, "Techniques and tools for the automatic generation of optimal test data at code, model and interface level," in *Companion of the 30th international conference on Software engineering*, pp. 927–928, ACM, 2008.
- [204] A. Pretschner, T. Mouelhi, and Y. Le Traon, "Model-based tests for access control policies," in *2008 1st International Conference on Software Testing, Verification, and Validation*, pp. 338–347, IEEE, 2008.
- [205] X. Qu, M. B. Cohen, and K. M. Woolf, "Combinatorial interaction regression testing: A study of test case generation and prioritization," in *Software Maintenance, 2007. ICSM 2007. IEEE International Conference on*, pp. 255–264, IEEE, 2007.
- [206] X. Qu, M. Acharya, and B. Robinson, "Configuration selection using code change impact analysis for regression testing," in *Software Maintenance (ICSM), 2012 28th IEEE International Conference on*, pp. 129–138, IEEE, 2012.
- [207] I. Rubab, S. Ali, L. Briand, and Y. Le Traon, "Model-based testing of obligations," in *2014 14th International Conference on Quality Software*, pp. 1–10, IEEE, 2014.
- [208] D. Schuler and A. Zeller, "Assessing oracle quality with checked coverage," in *2011 Fourth IEEE International Conference on Software Testing, Verification and Validation*, pp. 90–99, IEEE, 2011.
- [209] D. Schuler and A. Zeller, "Checked coverage: an indicator for oracle quality," *Software Testing, Verification and Reliability*, vol. 23, no. 7, pp. 531–551, 2013.
- [210] S. Segura, R. M. Hierons, D. Benavides, and A. Ruiz-Cortés, "Automated test data generation on the analyses of feature models: A metamorphic testing approach," in *2010 Third International Conference on Software Testing, Verification and Validation*, pp. 35–44, IEEE, 2010.
- [211] S. R. Shahamiri, W. M. Wan-Kadir, S. Ibrahim, and S. Z. M. Hashim, "Artificial neural networks as multi-networks automated test oracle," *Automated Software Engineering*, vol. 19, no. 3, pp. 303–334, 2012.
- [212] W. Shelton, N. Li, P. Ammann, and J. Offutt, "Adding criteria-based tests to test driven development," in *2012 IEEE Fifth International Conference on Software Testing, Verification and Validation*, pp. 878–886, IEEE, 2012.

- [213] Q. Shi, Z. Chen, C. Fang, Y. Feng, and B. Xu, "Measuring the diversity of a test set with distance entropy," *IEEE Transactions on Reliability*, vol. 65, no. 1, pp. 19–27, 2015.
- [214] A. Shi, T. Yung, A. Gyori, and D. Marinov, "Comparing and combining test-suite reduction and regression test selection," in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, pp. 237–247, ACM, 2015.
- [215] K. Shrestha and M. J. Rutherford, "An empirical evaluation of assertions as oracles," in *2011 Fourth IEEE International Conference on Software Testing, Verification and Validation*, pp. 110–119, IEEE, 2011.
- [216] M. Staats, M. W. Whalen, and M. P. Heimdahl, "Better testing through oracle selection (nier track)," in *Proceedings of the 33rd International Conference on Software Engineering*, pp. 892–895, ACM, 2011.
- [217] M. Staats, P. Loyola, and G. Rothermel, "Oracle-centric test case prioritization," in *Software Reliability Engineering (ISSRE), 2012 IEEE 23rd International Symposium on*, pp. 311–320, IEEE, 2012.
- [218] M. Stephan and J. R. Cordy, "Model clone detector evaluation using mutation analysis.," in *ICSME*, pp. 633–638, 2014.
- [219] R. P. Tan and S. Edwards, "Evaluating automated unit testing in sulu," in *2008 1st International Conference on Software Testing, Verification, and Validation*, pp. 62–71, IEEE, 2008.
- [220] S. Tasiran, M. E. Keremoğlu, and K. Muşlu, "Location pairs: a test coverage metric for shared-memory concurrent programs," *Empirical Software Engineering*, vol. 17, no. 3, pp. 129–165, 2012.
- [221] S. R. Vergilio, J. C. Maldonado, M. Jino, and I. W. Soares, "Constraint based structural testing criteria," *Journal of Systems and Software*, vol. 79, no. 6, pp. 756–771, 2006.
- [222] M. Vivanti, A. Mis, A. Gorla, and G. Fraser, "Search-based data-flow test generation," in *Software Reliability Engineering (ISSRE), 2013 IEEE 24th International Symposium on*, pp. 370–379, IEEE, 2013.
- [223] H. Wang, K. Zhai, and T. Tse, "Correlating context-awareness and mutation analysis for pervasive computing systems," in *Quality Software (QSIC), 2010 10th International Conference on*, pp. 151–160, IEEE, 2010.
- [224] X. Wang, L. Zhang, and P. Tanofsky, "Experience report: how is dynamic symbolic execution different from manual testing? a study on klee," in *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, pp. 199–210, ACM, 2015.
- [225] A. Watanabe and K. Sakamura, "A specification-based adaptive test case generation strategy for open operating system standards," in *Proceedings of the 18th international conference on Software engineering*, pp. 81–89, IEEE Computer Society, 1996.
- [226] E. Weyuker, T. Goradia, and A. Singh, "Automatically generating test data from a boolean specification," *IEEE Transactions on Software Engineering*, vol. 20, no. 5, p. 353, 1994.
- [227] X. Xie, J. W. Ho, C. Murphy, G. Kaiser, B. Xu, and T. Y. Chen, "Testing and validating machine learning classifiers by metamorphic testing," *Journal of Systems and Software*, vol. 84, no. 4, pp. 544–558, 2011.
- [228] D. Xu, O. El-Ariss, W. Xu, and L. Wang, "Testing aspect-oriented programs with finite state machines," *Software Testing, Verification and Reliability*, vol. 22, no. 4, pp. 267–293, 2012.
- [229] J. Xuan and M. Monperrus, "Test case purification for improving fault localization," in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pp. 52–63, ACM, 2014.
- [230] S. Yoo and M. Harman, "Test data regeneration: generating new test data from existing test data," *Software Testing, Verification and Reliability*, vol. 22, no. 3, pp. 171–201, 2012.
- [231] D. You, S. Rayadurgam, M. Whalen, M. P. Heimdahl, and G. Gay, "Efficient observability-based test generation by dynamic symbolic execution," in *Software Reliability Engineering (ISSRE), 2015 IEEE 26th International Symposium on*, 2015.

- [232] Y. Zhan and J. A. Clark, “A search-based framework for automatic testing of matlab/simulink models,” *Journal of Systems and Software*, vol. 81, no. 2, pp. 262–285, 2008.
- [233] L. Zhang, D. Hao, L. Zhang, G. Rothermel, and H. Mei, “Bridging the gap between the total and additional test-case prioritization strategies,” in *Software Engineering (ICSE), 2013 35th International Conference on*, pp. 192–201, IEEE, 2013.
- [234] L. Zhang, S.-S. Hou, C. Guo, T. Xie, and H. Mei, “Time-aware test-case prioritization using integer linear programming,” in *Proceedings of the eighteenth international symposium on Software testing and analysis*, pp. 213–224, ACM, 2009.
- [235] J. Zhang, J. Chen, D. Hao, Y. Xiong, B. Xie, L. Zhang, and H. Mei, “Search-based inference of polynomial metamorphic relations,” in *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*, pp. 701–712, ACM, 2014.
- [236] C. Zhou and P. Frankl, “Jdama: Java database application mutation analyser,” *Software Testing, Verification and Reliability*, vol. 21, no. 3, pp. 241–263, 2011.

A ARTICLE CHARACTERISATION RESULTS

| Paper | Role | Testing Act. | | | Mutation Tool | | | MOP | EMP | RT | Subject | | | Mutant Analysis |
|-------------------------------|------|--|------------|------------|---------------|-------------|--------------|------------|--------------------------|-----------------------------|--------------|------|--------|-----------------|
| | | Category | Test Level | Strategy | Avail. | Type | Name | Des. Level | Method | Method | Lan. | Size | Avail. | |
| AbouTrab et al. [139] | A | test strategy evaluation | n/a | spec. | N | hand. | | well | n/a | n/a | C | n/a | N | |
| Aichernig et al. [140] | G | test data generation | int. | spec. (E) | Y | exi. | μ java | not suff. | n/a | n/a | java | n/a | N | |
| Aichernig et al. [141] | G | test data generation | n/a | spec. (E) | Y | exi. | μ java | well | manual invest. | n/a | UML | S | N | |
| Ali et al. [142] | A | test strategy evaluation | int. | spec. | N | n/a | | well | n/a | n/a | java | P | Y | |
| Andrews et al. [143] | A | test oracle | unit | spec. | N | self. | Self-written | well | n/a | mutant sample | C | S | N | |
| Andrews et al. [144] | A | test strategy evaluation | n/a | struc. | N | exi. | [143] | well | not killed as equivalent | mutant sample | C | S | Y | * |
| Androutsopoulos et al. [145] | A | test strategy evaluation | n/a | others | Y | exi. | SMT-C | well | n/a | fixed number, weak mutation | C | M | N | |
| Antoniol et al. [146] | A | test strategy evaluation | n/a | spec. | N | n/a | | well | n/a | n/a | C++ | S | N | * |
| Arcaini et al. [147] | A | model review | n/a | others | N | n/a | | well | model checker | n/a | spec. | n/a | Y | * |
| Arcuri and Briand [148] | A | test strategy evaluation | n/a | spec. | Y | exi. | MuJava | well | not killed as equivalent | mutant sample | java | S | N | |
| Asoudeh and Labiche [149] | A | test data generation | n/a | spec. (E) | Y | exi. | MAJOR | well | n/a | n/a | java | n/a | N | |
| Außerlechner et al. [80] | A | fault localisation | other | others | N | n/a | | well | n/a | n/a | spread-sheet | n/a | Y | |
| Baker and Habli [150] | A | test strategy evaluation | n/a | struc. | Y, N | exi., hand. | MILU | well | manual invest. | n/a | C, Ada | P | N | * |
| Bandyopadhyay and Ghosh [151] | A | test data generation | unit | spec. | Y | exi. | MuJava | well | n/a | n/a | java | S | N | * |
| Bardin et al. [152] | G | test data generation | unit | hybrid | Y | exi. | MuJava | well | n/a | weak mutation | java | S | Y | |
| Baudry et al. [92] | G | test data generation | unit | struc. (E) | N | exi. | μ Slayer | well | manual invest. | n/a | Eiffel | n/a | Y | |
| Baudry et al. [93] | G | test data generation | sys. | struc. (E) | N | self. | | well | manual invest. | n/a | C# | n/a | N | |
| Baudry et al. [94] | G | test data generation | unit, sys. | struc. (E) | N | n/a | | well | reduce likelihood | n/a | Eiffel, C# | n/a | N | |
| Belli and Beyazit [118] | G | test data generation | other | spec. (E) | N | n/a | | well | reduce likelihood | fixed number | n/a | n/a | N | |
| Belli et al. [153] | A | test data generation, test strategy evaluation | n/a | spec. | N | n/a | | well | deterministic model | n/a | spec. | n/a | N | * |
| Bertolino et al. [154] | A | test data generation, test strategy evaluation | n/a | spec. | N | exi. | [40] | well | n/a | n/a | XACML | n/a | N | |
| Briand et al. [155] | A | test strategy evaluation | unit | spec. | N | n/a | | well | manual invest. | n/a | java, c++ | S | Y | |
| Briand et al. [156] | A | test strategy evaluation | n/a | spec. | N | n/a | | n/a | not killed as equivalent | n/a | Java, C++ | n/a | N | |
| Briand et al. [157] | A | test strategy evaluation | n/a | hybrid | N | n/a | | well | n/a | n/a | java | S | Y | |
| Cai et al. and Lyu [?] | A | test strategy evaluation | n/a | struc. | N | n/a | | n/a | n/a | n/a | C | S | N | |
| Chae et al. [117] | A | test data generation, test-suite reduction | other | struc. | N | n/a | | well | n/a | n/a | C | M | N | |
| Chen et al. [106] | G | test strategy evaluation | n/a | struc. (E) | Y | exi. | Proteum | n/a | no invest. | n/a | C | S | Y | |

| | | | | | | | | | | | | | | |
|------------------------------------|-----|---|------------|--------------------|---|-------------|---------------------------|-----------|--------------------------|------------------------------|--------------|-----|---|---|
| Cheon [158] | A | test oracle | unit | spec. | N | hand. | manual | not suff. | manual invest. | fixed number | java | P | N | |
| Chevalley and Thévenod-Fosse [159] | A | test data generation | n/a | spec. | N | self. | Self-written | well | manual invest. | n/a | java | S | N | |
| Czemerinski et al. [160] | A | test strategy evaluation | unit | spec. | Y | exi. | MuJava | well | no invest. | mutant sample | java | S | Y | |
| Czemerinski et al. [91] | A | test strategy evaluation | unit | spec. | Y | exi. | μ-JAVA | well | not killed as equivalent | n/a | Java | S | Y | * |
| Delamare et al. [81] | A | test oracle | unit | struc. | Y | exi. | AjMutator | well | manual invest. | n/a | AspectJ | S | N | |
| DeMilli and Offutt [?] | G | test data generation | unit | struc. (E) | Y | exi. | Mothra | n/a | manual invest. | weak mutation | Fortran | P | N | |
| DeMillo and Offutt [161] | G | test data generation | unit | struc. (E) | Y | exi. | Mothra | n/a | manual invest. | weak mutation | Fortran | n/a | N | |
| Denaro et al. [113] | A | test data generation | int. | struc. | Y | exi. | PIT | n/a | n/a | n/a | Java | S | Y | * |
| Di Nardo et al. [162] | G | test data generation | sys. | spec. (E) | N | n/a | | well | n/a | selection strategy | spec. | n/a | N | |
| Do and Rothermel [63] | A | test case prioritisation | unit | struc. | N | self. | Self-written | well | n/a | n/a | Java | n/a | Y | |
| Do and Rothermel [64] | A | test case prioritisation | unit, sys. | struc., spec. | N | self. | Self-written | well | n/a | n/a | Java | n/a | Y | |
| Do et al. [163] | A | test case prioritisation | n/a | struc., struc. (E) | N | self. | Self-written | well | n/a | fixed number | java | M | Y | |
| Duran and Ntafos [54] | A | test strategy evaluation | n/a | hybrid | N | n/a | | n/a | n/a | n/a | Fortran | n/a | N | |
| Edwards and Shams [164] | A | test strategy evaluation | unit | struc. | Y | exi. | Javalanche | n/a | not killed as equivalent | n/a | java | n/a | N | |
| Edwards [165] | A | test strategy evaluation | n/a | spec. | N | n/a | | well | manual invest. | n/a | spec. | S | N | |
| Elbaum et al. [62] | G | test case prioritisation | n/a | struc. (E) | Y | exi. | Proteum | n/a | n/a | n/a | C | M | Y | |
| Fang et al. [166] | A | test case prioritisation | unit | sim. | Y | exi. | Jumble | n/a | not killed as equivalent | n/a | Java | M | Y | * |
| Fraser and Walkinshaw [167] | A,G | test data generation, test strategy evaluation | unit | struc. (E) | Y | exi., part. | Javalanche, EvoSuite | well | n/a | n/a | java | S | N | |
| Fraser and Zeller [102] | G | test data generation, test oracle, test-case minimisation | unit | struc. (E) | Y | exi. | μTEST, Javalanche | not suff. | no invest. | weak mutation | Java | n/a | Y | |
| Fraser et al. [168] | G | test data generation | unit | struc. | Y | exi. | MAJOR | n/a | n/a | fixed number | java | S | Y | |
| Fraser et al. [169] | G | test data generation | unit | struc. (E) | Y | part. | EvoSuite | well | no invest. | mutant sample, weak mutation | Java | n/a | Y | |
| Gay et al. [104] | G | test oracle | unit | struc. (E) | N | n/a | | well | model checker | fixed number | simulink | n/a | N | |
| Gay et al. [170] | A | test data generation, test suite reduction | unit | struc., struc. (E) | N | exi. | [?] | not suff. | model checker | fixed number, weak mutation | Lustre | S | N | |
| Gligoric et al. [171] | A | test strategy evaluation | n/a | struc. | N | exi., exi. | Javalanche, [143] | not suff. | no invest. | n/a | Java, C | M | Y | |
| Gonzalez-Sanchez et al. [172] | A | fault localisation | n/a | struc., struc. (E) | Y | part. | Zoltar | n/a | n/a | n/a | C | M | Y | |
| Gopinath et al. [173] | A | test strategy evaluation | unit | struc. | Y | exi. | PIT | n/a | n/a | n/a | Java | M | Y | |
| Gupta and Jalote [174] | A | fault localisation, program repairing | unit | others | N | hand. | manual | well | n/a | n/a | java | S | N | |
| Hao et al. [109] | G | test-suite reduction | n/a | struc. (E) | Y | exi. | Proteum, Mu-Java | n/a | n/a | mutant sample | java, c | S | Y | |
| Hao et al. [110] | A | test case prioritisation | unit, sys. | struc. | Y | exi. | [143], Mujava, Javalanche | n/a | n/a | fixed number | Java, C | M | Y | |
| Harman et al. [175] | G | test data generation | n/a | struc. (E) | Y | exi. | MiLu | not suff. | not killed as equivalent | higher-order | C | S | N | |
| Hennessy and Power [116] | A | test-suite reduction | other | others | N | self. | Self-written | well | n/a | weak mutation | C++ Grammar | M | Y | |
| Hofer and Wotawa [119] | A | fault localisation | other | others | N | n/a | | well | n/a | n/a | spread-sheet | n/a | Y | |

| | | | | | | | | | | | | | | |
|---------------------------------|-----|---|------------|---------------|------|-------------|----------------|-----------|--------------------------|--------------------|--------------|-----|---|---|
| Hofer et al. [120] | A | fault localisation | other | others | N | n/a | | not suff. | n/a | n/a | spread-sheet | n/a | Y | |
| Holt et al. [176] | A | test strategy evaluation | n/a | spec. | N | hand. | Hand-seeded | well | n/a | n/a | C++ | S | N | |
| Hong et al. [83] | A | test strategy evaluation | n/a | struc. | N | n/a | | well | n/a | n/a | Java | S | N | |
| Hong et al. [84] | A | test strategy evaluation | n/a | struc. | N | n/a | | well | n/a | n/a | java | S | N | |
| Hou et al. [85] | G | test strategy evaluation | int. | spec. (E) | N | n/a | | well | n/a | weak mutation | Java | S | Y | |
| Inozemtseva et al. [58] | A | test strategy evaluation | unit | struc. | Y | exi. | PIT | n/a | not killed as equivalent | n/a | Java | M | Y | |
| Jamrozik et al. [177] | G | test data generation | unit | struc. (E) | N | n/a | | n/a | n/a | n/a | C# | n/a | N | |
| Jee et al. [178] | A | test data generation | unit | struc. | N | n/a | | well | n/a | n/a | PLC | n/a | N | |
| Jolly et al. [179] | A | test strategy evaluation | unit | spec. | N | exi. | Mutant Power | n/a | n/a | fixed number | C# | M | N | |
| Kanewala and Bieman [180] | A | test oracle | unit | others | Y | exi. | μ java | well | n/a | n/a | java | P | N | |
| Kapfhammer et al. [122] | A | test data generation | other | others | N | n/a | | well | n/a | n/a | SQL | n/a | N | |
| Khalil and Labiche [181] | A | test strategy evaluation | n/a | spec. | Y | exi. | MuJava | well | not killed as equivalent | n/a | java | S | Y | |
| Kim et al. [182] | A | test strategy evaluation | n/a | spec., struc. | N | self. | Self-written | not suff. | n/a | n/a | Java | S | N | * |
| Koster and Kao [183] | A | test strategy evaluation | unit | struc. (E) | Y | exi. | MuJava | well | n/a | n/a | java | S | Y | |
| Kracht et al. [184] | A | test data generation | unit | struc. | Y | exi. | MAJOR | well | n/a | n/a | Java | M | Y | |
| Kwon et al. [98] | G | test case prioritisation | unit | struc. (E) | Y | exi. | MAJOR | not suff. | n/a | fixed number | Java | S | Y | |
| Le Traon et al. [185] | A | test strategy evaluation | unit | spec. | N | exi. | μ Slayer | well | manual invest. | n/a | Eiffel | n/a | Y | |
| Le Traon et al. [186] | A | test strategy evaluation, test-suite selection | sys. | spec. | N | n/a | | well | reduce likelihood | n/a | spec. | M | N | |
| Lee and Offutt [187] | G | test data generation | int. | spec. (E) | N | n/a | | well | n/a | n/a | XML | n/a | Y | |
| Lei and Andrews [188] | A | test case minimisation | unit | struc. | N | n/a | | well | not killed as equivalent | n/a | C | S | N | |
| Li et al. [111] | A | test strategy evaluation | int., sys. | spec. | N | n/a | | well | manual invest. | n/a | java | S | N | |
| Liu et al. [189] | G | test data generation | unit | struc. (E) | N | self. | Self-written | well | manual invest. | weak mutation | Java | P | N | |
| Lorber [190] | G | test data generation | n/a | spec. (E) | Y | exi. | MoMuT::UML | n/a | n/a | n/a | spec. | n/a | N | |
| Lou et al. [108] | G | test case prioritisation | unit | struc. (E) | Y | exi. | Javalanche | well | n/a | fixed number | java | M | Y | |
| Loyala et al. [191] | A | test oracle | unit | others | Y | exi. | MAJOR | n/a | n/a | n/a | Java | S | Y | |
| Mayer and Guderlei [192] | A | test data generation, test strategy evaluation | n/a | spec. | Y | exi. | MuJava | well | n/a | n/a | java | n/a | N | |
| Mayer and Schneckenburger [193] | A | test strategy evaluation | n/a | spec. | Y | exi. | MuJava | well | n/a | n/a | java | n/a | N | |
| Mcminn et al. [123] | A | test strategy evaluation, test data generation | other | struc. | Y | exi. | SchemaAnalyst | well | reduce likelihood | n/a | SQL | n/a | N | |
| Milani et al. [194] | A | test data generation | unit | spec. | Y | exi. | MUTANDIS [79] | well | reduce likelihood | n/a | JavaScript | M | Y | |
| Miller and Strooper [195] | A | test strategy evaluation | n/a | spec. | Y, N | exi., hand. | MuJava, manual | well | manual invest. | mutant sample | java | n/a | N | |
| Mirshokraie [196] | A,G | test data generation, test oracle, test strategy evaluation | unit | struc. (E) | Y | exi. | MUTANDIS | n/a | n/a | n/a | JavaScript | n/a | N | |
| Moon et al. [101] | G | fault localisation | n/a | struc. (E) | Y | exi. | Proteum | not suff. | n/a | selection strategy | C | M | Y | |
| Mottu et al. [197] | A | test data generation | n/a | struc. (E) | N | n/a | | well | n/a | n/a | AspectJ | n/a | Y | |
| Mouchawrab et al. [198] | A | test data generation | n/a | spec., struc. | Y | exi. | Mujava | n/a | manual invest. | n/a | Java | S | Y | * |
| Mouchawrab et al. [199] | A | test data generation | n/a | spec., struc. | Y | exi. | Mujava | n/a | manual invest. | n/a | Java | S | Y | |
| Murtaza et al. [100] | G | fault localisation | n/a | struc. (E) | N | exi. | [143] | not suff. | n/a | n/a | C | M | Y | |

| | | | | | | | | | | | | | |
|-------------------------------|---|--|------|--------------------|---|--------------|------------------|-----------|--|---------------|----------------|-----|-----|
| Namin and Andrews [57] | A | test strategy evaluation | unit | struc. | Y | exi. | Proteum | well | not killed as equivalent | fixed number | C | S | Y |
| Ntafos [52] | A | test strategy evaluation | n/a | struc., spec. | N | exi. | [48] | well | n/a | n/a | Fortran | n/a | N |
| Ntafos [53] | A | test strategy evaluation | n/a | struc. | N | n/a | | n/a | n/a | n/a | Fortran | P | Y |
| Offutt and Liu [200] | A | test data generation | int. | spec. | Y | exi. | Mothra | well | no invest. | n/a | n/a | n/a | N |
| Offutt et al. [65] | G | test-suite reduction | unit | struc. (E) | Y | exi. | Mothra | well | n/a | n/a | Fortran | P | N |
| Papadakis and Maleveris [76] | G | test data generation | unit | struc. (E) | Y | part. | jFuzz | well | not killed as equivalent, manual invest. | n/a | Java | S | Y |
| Papadakis and Maleveris [56] | G | test data generation | unit | struc. (E) | Y | part. | jFuzz | well | manual invest. | weak mutation | Java | n/a | Y |
| Papadakis and Malveris [201] | G | test data generation | unit | struc. (E) | N | n/a | | well | not killed as equivalent | weak mutation | Delphi, C, n/a | P | N |
| Papadakis et al. [68] | G | fault localisation | unit | struc. (E) | Y | exi. | Proteum | n/a | no invest. | mutant sample | C | S | Y |
| Papadakis et al. [124] | A | test-suite selection | n/a | spec. (E) | N | n/a | | well | not killed as equivalent | n/a | C | S | Y |
| Papadakis et al. [69] | G | fault localisation | unit | struc. (E) | Y | exi. | Proteum | not suff. | no invest. | mutant sample | C | M | Y |
| Patrick et al. [202] | G | test data generation | n/a | hybrid | Y | exi. | MuJava | n/a | n/a | n/a | java | S | N |
| Pinte et al. [203] | A | test data generation | unit | struc. | Y | exi. | MuJava | n/a | n/a | n/a | java | n/a | N |
| Pretschner et al. [204] | A | test data generation, test strategy evaluation | n/a | spec. | N | n/a | | well | n/a | n/a | spec. | n/a | N |
| Qi et al. [99] | G | program repairing | n/a | struc. (E) | Y | part. | TrpAutoRepair | not suff. | n/a | n/a | C | L | Y |
| Qu et al. [97] | G | test case prioritisation | n/a | spec. (E) | N | exi. | [143] | n/a | n/a | fixed number | C | M | Y |
| Qu et al. [205] | A | test data generation, test case prioritisation | n/a | spec. | N | exi. | [143] | not suff. | n/a | fixed number | C | M | Y |
| Qu et al. [206] | A | test-suite selection | n/a | spec. | N | exi. | [143] | n/a | n/a | n/a | C/C++ | M | N |
| Rothermel et al. [90] | G | test case prioritisation | n/a | struc., struc. (E) | Y | exi. | Proteum | n/a | n/a | n/a | C | S | Y |
| Rothermel et al. [61] | G | test case prioritisation | n/a | struc. (E) | Y | exi. | Proteum | n/a | not killed as nonequivalent | n/a | C | S | Y |
| Rubab et al. [207] | A | test data generation | n/a | spec. | N | hand. | | well | n/a | n/a | spec. | n/a | N |
| Rutherford et al. [112] | A | test strategy evaluation | sys. | spec. | Y | exi. | MuJava | well | n/a | n/a | java | S | N |
| Schuler and Zeller [208] | A | test strategy evaluation | unit | struc. | Y | exi. | Javalanche | n/a | n/a | n/a | java | M | Y |
| Schuler and Zeller [209] | A | test strategy evaluation | unit | struc. | Y | exi. | Javalanche | n/a | n/a | n/a | java | M | Y |
| Segura et al. [210] | A | test data generation | n/a | spec. | Y | exi. | MuClipse | well | manual invest. | n/a | java | S | Y |
| Shahamiri et al. [211] | A | test oracle | n/a | spec. | N | n/a | | not suff. | n/a | n/a | C# | n/a | N |
| Shelton et al. [212] | G | development scheme evaluation | unit | others | Y | exi. | MuJava | n/a | manual invest. | n/a | java | n/a | N * |
| Shi et al. [213] | A | test strategy evaluation | n/a | sim. | Y | Self-written | mutate.py | not suff. | manual invest. | fixed number | C | M | Y |
| Shi et al. [105] | G | test-suite reduction | n/a | struc., struc. (E) | Y | exi. | PIT | well | n/a | n/a | Java | S | Y |
| Shi et al. [214] | A | test-suite reduction | unit | struc. | Y | exi. | PIT | n/a | n/a | n/a | java | M | Y |
| Shrestha and Rutherford [215] | A | test oracle | unit | struc. | Y | exi. | MuJava | well | manual invest. | n/a | java | S | Y |
| Smith and Williams [96] | G | test data generation | unit | struc. (E) | Y | exi. | MuClipse | well | no invest. | n/a | Java | S | Y |
| Smith and Williams [107] | G | test strategy evaluation | unit | struc. (E) | Y | exi. | Jumble, MuClipse | well | n/a | n/a | java | S | N * |
| Staats et al. [216] | A | test oracle | n/a | others | N | n/a | | n/a | model checker | fixed number | Lustre | n/a | N |
| Staats et al. [103] | G | test oracle | unit | struc. (E) | N | exi. | [?] | not suff. | model checker | fixed number | Lustre | S | N |
| Staats et al. [217] | A | test case prioritisation | unit | struc. | Y | exi. | Sofya | well | n/a | fixed number | java | S | N |

| | | | | | | | | | | | | | | |
|-----------------------------|---|--|------------------|--------------------|---|-------|------------------------------|-----------|-----------------------------------|--------------------|-------------|-----|---|---|
| Stephan and Cordy [218] | A | model clone detection | n/a | others | N | n/a | | well | n/a | random mutation | simulink | n/a | N | |
| Tan and Edwards [219] | A | test data generation, test strategy evaluation | unit | struc. | N | exi. | sulu tool | well | n/a | n/a | sulu | n/a | N | |
| Tasiran et al. [220] | A | test strategy evaluation | n/a | struc. | N | n/a | | well | n/a | n/a | java | n/a | N | |
| Tuya et al. [121] | A | test strategy evaluation | other | struc. | Y | exi. | SQLMutation | well | reduce likelihood, manual invest. | n/a | SQL | n/a | Y | |
| Vergilio et al. [221] | A | test strategy evaluation | n/a | struc. (E) | Y | exi. | Proteum | well | manual invest. | n/a | C | n/a | N | |
| Vivanti et al. [222] | G | test data generation | unit | struc., struc. (E) | Y | part. | EvoSuite | n/a | n/a | weak mutation | Java | n/a | Y | |
| von Mayrhauser et al. [95] | G | test data generation | n/a | spec. (E) | N | n/a | | well | n/a | n/a | n/a | L | N | |
| Wang et al. [223] | A | test strategy evaluation | n/a | sim. | Y | exi. | MuClipse | well | not killed as equivalent | n/a | Java | S | N | * |
| Wang et al. [224] | A | test data generation | n/a | struc. | N | exi. | [143] | n/a | n/a | fixed number | C | M | Y | |
| Watanabe and Sakamura [225] | A | test data generation | int. | spec. | N | n/a | | well | n/a | n/a | spec. | n/a | N | |
| Weyuker et al. [226] | A | test data generation | n/a | spec. | N | n/a | | well | n/a | n/a | n/a | n/a | N | * |
| Whalen et al. [60] | A | test strategy evaluation | n/a | struc. (E) | N | exi. | [?] | not suff. | model checker | fixed number | Simulink | n/a | N | |
| Xie et al. [227] | A | test strategy evaluation | n/a | others | Y | exi. | MuJava | well | manual invest. | fixed number | java | n/a | Y | |
| Xu and Ding [82] | A | test case prioritisation | n/a | spec. | N | n/a | | well | n/a | n/a | AspectJ | S | N | |
| Xu et al. [228] | A | test strategy evaluation | n/a | spec. | N | hand. | | well | n/a | n/a | AspectJ | S | N | |
| Xuan and Monperrus [229] | A | fault localisation | unit | others | Y | exi. | PIT | well | n/a | fixed number | Java | M | Y | * |
| Yoo and Harman [230] | A | test data generation | unit | struc. | Y | exi. | MuJava | well | no invest. | n/a | Java | n/a | Y | |
| Yoon and Choi [86] | A | test-suite selection | int. | hybrid | N | self. | TECC | well | n/a | selection strategy | EJB | n/a | N | |
| You et al. [231] | A | test data generation | n/a | struc. (E) | N | exi. | [?] | not suff. | no invest. | fixed number | Lustre | n/a | N | |
| Zhan and Clark [232] | G | test data generation, test-suite reduction | n/a (code level) | struc., struc. (E) | N | self. | | not suff. | manual invest. | n/a | simulink | n/a | N | |
| Zhang and Mesbah [59] | A | test strategy evaluation | unit | struc. (E) | Y | exi. | PIT | well | not killed as equivalent | n/a | Java | M | Y | |
| Zhang et al. [233] | A | test case prioritisation | unit | struc. | Y | exi. | MuJava | n/a | n/a | n/a | Java | M | Y | |
| Zhang et al. [67] | A | fault localisation | n/a | struc. (E) | Y | exi. | Proteum | not suff. | n/a | n/a | C | S | Y | * |
| Zhang et al. [234] | A | test case prioritisation | n/a | struc. | Y | exi. | Jester | well | n/a | fixed number | java | S | Y | |
| Zhang et al. [55] | G | test data generation | n/a | struc. (E) | Y | exi. | PexMutator, GenMutants | well | manual invest. | weak mutation | C# | S | Y | |
| Zhang et al. [66] | A | test-suite reduction | unit | struc. | Y | exi. | MuJava | well | n/a | fixed number | Java | M | Y | |
| Zhang et al. [235] | A | test oracle | n/a | spec. | Y | exi. | MuClipse | not suff. | n/a | n/a | java, C/C++ | S | Y | |
| Zhou and Frankl [236] | A | test data generation, test strategy evaluation | n/a | struc. | Y | exi. | JDAMA (based on SQLMutation) | well | reduce likelihood, manual invest. | weak mutation | Java, sql | n/a | Y | |

Table 19. Article Characterisation Results

B ARTICLE CHARACTERISATION RESULTS REGARDING THE“GUIDE” ROLE

| Paper | Testing Act. | Usage | Fault Type | Notes |
|------------------------------|---|--|------------------------|--------------------------------|
| Aichernig et al. [140] | test data generation (automatic) | generate test data to kill mutants (model-based) | mutation | traditional |
| Aichernig et al. [141] | test data generation (automatic) | generate test data to kill mutants (model-based) | mutation | traditional |
| Bardin et al. [152] | test data generation (automatic) | generate weak mutant killable conditions | no evaluation | |
| Baudry et al. [92] | test data generation (optimisation) | use mutation score to optimise test data based on an initial test set | mutation | |
| Baudry et al. [93] | test data generation (optimisation) | use mutation score to optimise test data based on an initial test set | mutation | |
| Baudry et al. [94] | test data generation (optimisation) | use mutation score to optimise test data based on an initial test set | mutation | |
| Belli and Beyazit [118] | test data generation (automatic) | generate test data to kill mutants (model-based) | mutation | |
| Chen et al. [106] | test strategy evaluation | use mutants to estimate fault-exposing potential | Hand-seeded | |
| DeMilli and Offutt [?] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |
| DeMillo and Offutt [161] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |
| Di Nardo et al. [162] | test data generation (automatic) | generate test data to kill mutants (model-based) | mutation | |
| Elbaum et al. [62] | test case prioritisation | use mutants to estimate fault-exposing potential | Hand-seeded | |
| Fraser and Walkinshaw [167] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |
| Fraser and Zeller [102] | test data generation (automatic), test oracle, test-case minimisation | generate strong mutant killable conditions (mutant impact) | mutation | |
| Fraser et al. [168] | test data generation (manual vs automatic) | use mutation analysis to generate assertion based on the behavior of the program | hand-seeded + mutation | |
| Fraser et al. [169] | test data generation (automatic) | generate weak/strong mutant killable conditions for test data generation | mutation | |
| Gay et al. [104] | test oracle | rank variables (for test oracle) according to killed mutants | mutation | evaluation set |
| Hao et al. [109] | test-suite reduction | use mutants to collect statistics on loss in fault-detection capability at the level of individual statements for various levels of confidence, and use these to construct a fault-detection- loss table | Hand-seeded + mutation | later version + evaluation set |
| Harman et al. [175] | test data generation (automatic) | generate strong-higher-order mutant killable conditions | mutation | |
| Hou et al. [85] | test strategy evaluation | propose Interface-Contract Mutation coverage as a test adequacy criterion | Hand-seeded + mutation | different operators |
| Jamrozik et al. [177] | test data generation (automatic) | augment test input data using mutation testing | mutation | |
| Kwon et al. [98] | test case prioritisation | use mutants to determine coefficients of the linear regression model with IR and coverage information | mutation | |
| Lee and Offutt [187] | test data generation (automatic) | generate test data to kill mutants (model-based) | mutation | |
| Liu et al. [189] | test data generation (automatic) | generate weak multiple-mutant-killable conditions | mutation | |
| Lorber [190] | test data generation (automatic) | generate test data to kill mutants (model-based) | no evaluation | |
| Lou et al. [108] | test case prioritisation | use mutants (in prior version) to estimate fault-detection capability | mutation | later version, evaluation set |
| Mirshokraie [196] | test oracle | use mutation analysis to generate assertion | mutation | |
| Moon et al. [101] | fault localisation | assign suspicious value to mutants according to test execution information | Hand-seeded | |
| Murtaza et al. [100] | fault localisation | use traces of mutants and prior faults to train the predication model (decision tree) | real | |
| Offutt et al. [65] | test-suite reduction | use traces of mutants and prior faults to train the predication model (decision tree) | real | |
| Papadakis and Maleveris [76] | test data generation (automatic) | generate strong mutant killable conditions | mutation | |
| Papadakis and Maleveris [56] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |
| Papadakis and Malveris [201] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |

| | | | | |
|----------------------------|--|--|------------------------|----------------|
| Papadakis et al. [68] | fault localisation | assign suspicious value to mutants according to test execution information | Hand-seeded | |
| Papadakis et al. [69] | fault localisation | assign suspicious value to mutants according to test execution information | Hand-seeded | |
| Patrick et al. [202] | test data generation (optimisation) | optimise subdomain to kill mutants | mutation | |
| Qi et al. [99] | program repairing | prioritise test cases in patch validation according to fault-exposing potential | mutation | |
| Qu et al. [97] | test case prioritisation | order test cases according to prior fault detection information using both hand-seeded and mutation faults | Hand-seeded + mutation | |
| Rothermel et al. [90] | test case prioritisation | use mutants to estimate fault-exposing potential | Hand-seeded | evaluation set |
| Rothermel et al. [61] | test case prioritisation | use mutants to estimate fault-exposing potential | Hand-seeded | |
| Shelton et al. [212] | development scheme evaluation | complementary to Test-Driven Development to add extra test cases | mutation | |
| Shi et al. [105] | test-suite reduction | to kill same mutants as the requirement | mutation | |
| Smith and Williams [96] | test data generation (augmentation) | to kill mutants as the requirement | mutation | |
| Smith and Williams [107] | test strategy evaluation | use mutation coverage to generate test data | mutation | |
| Staats et al. [103] | test oracle | rank variables (for test oracle) according to killed mutants | mutation | evaluation set |
| Vivanti et al. [222] | test data generation (automatic) | generate weak mutant killable conditions | mutation | |
| von Mayrhauser et al. [95] | test data generation (augmentation) | augment test input data to kill mutants | no evaluation | |
| Zhan and Clark [232] | test data generation (automatic), test-suite reduction | generate strong mutant killable conditions (model-based) | mutation | |
| Zhang et al. [55] | test data generation (automatic) | generate weak mutant killable conditions | mutation | different tool |

Table 20. Article Characterisation results regarding the “guide” Role

C MUTATION OPERATOR CHARACTERISATION RESULTS REGARDING THE “WELL-DEFINED” ENTITY

| Paper | Mutation Operators |
|------------------------------------|--|
| AbouTrab et al. [139] | specification |
| Aichernig et al. [141] | specification, arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Ali et al. [142] | specification |
| Andrews et al. [143] | absolute value, arithmetic op, conditional op, relational op, program |
| Andrews et al. [144] | arithmetic op, relational op, constant, statement deletion, program |
| Androutopoulos et al. [145] | arithmetic op, program |
| Antoniol et al. [146] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, method call, control-flow disruption, constant, variable, statement swap, program |
| Arcaini et al. [147] | specification |
| Arcuri and Briand [148] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Asoudeh and Labiche [149] | arithmetic op, conditional op, relational op, shift op, absolute value, program |
| Außerlechner et al. [80] | spreadsheet-specific, program |
| Baker and Habli [150] | arithmetic op, relational op, conditional op, bitwise op, assignment op, statement deletion, constant, program |
| Bandyopadhyay and Ghosh [151] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific |
| Bardin et al. [152] | arithmetic op, conditional op, relational op, program |
| Baudry et al. [92] | exception handler, arithmetic op, conditional op, relational op, statement deletion, constant, variable, method call, OO-specific, program |
| Baudry et al. [93] | exception handler, arithmetic op, conditional op, relational op, statement deletion, constant, variable, method call, OO-specific, program |
| Baudry et al. [94] | exception handler, arithmetic op, conditional op, relational op, statement deletion, constant, variable, method call, OO-specific, program |
| Belli and Beyazit [118] | specification |
| Belli et al. [153] | specification |
| Bertolino et al. [154] | specification |
| Briand et al. [155] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, method call, control-flow disruption, constant, variable, statement swap, exception handler, Java-specific, program |
| Briand et al. [157] | arithmetic op, constant, method call, relational op, return statement, statement deletion, program |
| Chae et al. [117] | variable, program |
| Chevalley and Thévenod-Fosse [159] | variable, arithmetic op, relational op, constant, parentheses, program |
| Czemerinski et al. [160] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Czemerinski et al. [91] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Delamare et al. [81] | AOP-specific, program |
| Di Nardo et al. [162] | specification |
| Do and Rothermel [63] | arithmetic op, relational op, conditional op, method call, OO-specific, program |
| Do and Rothermel [64] | arithmetic op, relational op, conditional op, method call, OO-specific, program |
| Do et al. [163] | arithmetic op, relational op, conditional op, method call, OO-specific, program |
| Edwards [165] | arithmetic op, relational op, conditional op, absolute value, program |
| Fraser and Walkinshaw [167] | constant, conditional op, arithmetic op, statement deletion, program |
| Fraser et al. [169] | arithmetic op, relational op, constant, variable, conditional op, statement deletion, OO-specific, bitwise op, program |
| Gay et al. [104] | arithmetic op, relational op, conditional op, constant, variable, program |
| Gupta and Jalote [174] | OO-specific, constant, arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, control-flow disruption, method call, statement swap, variable, conditional exp, program |
| Hennessy and Power [116] | absolute value, arithmetic op, conditional op, relational op, program |
| Hofer and Wotawa [119] | spreadsheet-specific, program |
| Holt et al. [176] | specification |
| Hong et al. [83] | concurrent mutation, program |
| Hong et al. [84] | concurrent mutation, program |
| Hou et al. [85] | arithmetic op, conditional op, relational op, constant, variable, Interface mutation, program |
| Jee et al. [178] | specification |
| Kanewala and Bieman [180] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Kapfhammer et al. [122] | SQL-specific, program |
| Khalil and Labiche [181] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Koster and Kao [183] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Kracht et al. [184] | arithmetic op, conditional op, relational op, shift op, absolute value, program |
| Le Traon et al. [185] | exception handler, arithmetic op, conditional op, relational op, statement deletion, constant, variable, OO-specific, program |
| Le Traon et al. [186] | specification |
| Lee and Offutt [187] | specification |
| Lei and Andrews [188] | absolute value, arithmetic op, conditional op, relational op, program |
| Li et al. [111] | arithmetic op, conditional op, bitwise op, constant, variable, Java-specific, OO-specific, return statement, program |
| Liu et al. [189] | arithmetic op, relational op, conditional op, constant, variable, program |
| Lou et al. [108] | constant, conditional op, arithmetic op, statement deletion, program |
| Mayer and Guderlei [192] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Mayer and Schneckenburger [193] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Mcminn et al. [123] | SQL-specific, program |
| Milani et al. [194] | method call, variable, JS-specific, program |

| | |
|-------------------------------|---|
| Miller and Strooper [195] | specification, arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Mottu et al. [197] | specification |
| Namin and Andrews [57] | arithmetic op, bitwise op, return statement, conditional op, assignment op, type, relational op, goto label, loop trap, brace, switch statement, conditional exp, while statement, variable, program |
| Ntafos [52] | constant, variable, arithmetic op, relational op, conditional op, statement deletion, return statement, goto label, do statement, bomb statement, program |
| Offutt and Liu [200] | constant, variable, arithmetic op, relational op, conditional op, statement deletion, return statement, goto label, do statement, bomb statement, absolute value, program |
| Offutt et al. [65] | constant, variable, arithmetic op, relational op, conditional op, statement deletion, return statement, goto label, do statement, bomb statement, absolute value, program |
| Papadakis and Maleveris [76] | absolute value, arithmetic op, relational op |
| Papadakis and Maleveris [56] | absolute value, arithmetic op, conditional op, relational op, program |
| Papadakis and Malveris [201] | absolute value, arithmetic op, conditional op, relational op, program |
| Papadakis et al. [124] | specification |
| Pretschner et al. [204] | specification |
| Rubab et al. [207] | specification |
| Rutherford et al. [112] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Segura et al. [210] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Shi et al. [105] | arithmetic op, shift op, bitwise op, relational op, conditional op, assignment op, conditional exp, return statement, OO-specific, constant, statement deletion, variable, switch statement, program |
| Shrestha and Rutherford [215] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Smith and Williams [96] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Smith and Williams [107] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, constant, return statement, switch statement, OO-specific, Java-specific, program |
| Staats et al. [217] | arithmetic op, conditional op, relational op, method call, program |
| Stephan and Cordy [218] | specification |
| Tan and Edwards [219] | arithmetic op, relational op, conditional op, conditional exp, statement deletion, program |
| Tasiran et al. [220] | concurrent mutation, program |
| Tuya et al. [121] | arithmetic op, conditional op, relational op, absolute value, constant, variable, SQL-specific, program |
| Vergilio et al. [221] | bomb statement, conditional exp, statement deletion, return statement, goto label, control-flow disruption, while statement, do-while statement, loop trap, statement swap, brace, switch statement, arithmetic op, bitwise op, conditional op, shift op, relational op, assignment o, parentheses, type, variable, constant, program |
| von Mayrhauser et al. [95] | specification |
| Wang et al. [223] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, OO-specific, Java-specific, program |
| Watanabe and Sakamura [225] | relational op, constant, method call, program |
| Weyuker et al. [226] | specification |
| Xie et al. [227] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Xu and Ding [82] | AOP-specific, program |
| Xu et al. [228] | specification |
| Xuan and Monperrus [229] | arithmetic op, shift op, bitwise op, return statement, conditional op, assignment op, program |
| Yoo and Harman [230] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Yoon and Choi [86] | interface mutation, program |
| Zhang and Mesbah [59] | relational op, arithmetic op, shift op, bitwise op, conditional op, assignment op, return statement, statement deletion, program |
| Zhang et al. [234] | absolute value, arithmetic op, conditional op, relational op, program |
| Zhang et al. [55] | absolute value, arithmetic op, conditional op, relational op, program |
| Zhang et al. [66] | arithmetic op, relational op, conditional op, bitwise op, assignment op, shift op, program |
| Zhou and Frankl [236] | arithmetic op, conditional op, relational op, absolute value, constant, variable, SQL-specific, program |

Table 21. Mutation Operator Characterisation results regarding the “well-defined” Entity