



A comprehensive study of code-removal patches in automated program repair

Davide Ginelli¹ · Matias Martinez² · Leonardo Mariani¹ · Martin Monperrus³

Accepted: 22 November 2021 / Published online: 5 May 2022
© The Author(s) 2022, corrected publication 2022

Abstract

Automatic Program Repair (APR) techniques can promisingly help reduce the cost of debugging. Many relevant APR techniques follow the *generate-and-validate* approach, that is, the faulty program is iteratively modified with different change operators and then validated with a test suite until a plausible patch is generated. In particular, Kali is a generate-and-validate technique developed to investigate the possibility of generating plausible patches by only removing code. Former studies show that indeed Kali successfully addressed several faults. This paper addresses the single and particular case of code-removal patches in automated program repair. We investigate the reasons and the scenarios that make their creation possible, and the relationship with patches implemented by developers. Our study reveals that code-removal patches are often insufficient to fix bugs, and proposes a comprehensive taxonomy of code-removal patches that provides evidence of the problems that may affect test suites, opening new opportunities for researchers in the field of automatic program repair.

Keywords Automatic program repair · Code-removal patches · Software testing · Debugging

Communicated by: Mark Harman

✉ Davide Ginelli
davide.ginelli@unimib.it

Matias Martinez
matias.martinez@uphf.fr

Leonardo Mariani
leonardo.mariani@unimib.it

Martin Monperrus
monperrus@kth.se

¹ University of Milano - Bicocca, Milan, Italy

² Université Polytechnique Hauts-de-France, Valenciennes, France

³ KTH Royal Institute of Technology, Stockholm, Sweden

1 Introduction

The cost of software failures is often significantly high. For instance, a report conducted by the Consortium for IT Software Quality (CISQ) indicates that the cost of poor-quality software in the US in 2018 has been approximately \$2.84 trillion, with 37.46% due to software failures, and 16.87% due to activities performed to find and fix bugs (Krasner 2018). Similarly, multiple studies show that developers can spend up to 75% of their time in debugging and fixing activities (Britton et al. 2013; Undo Software 2014; Coralogix 2018).

Automatic Program Repair (APR) techniques represent a possible solution to reduce the cost of dealing with bugs, thanks to their capability of automatically suggesting fixes (Gazzola et al. 2019; Meyvis and Yoon 2021). Among the many strategies, a number of APR techniques follow the *generate-and-validate* approach, that is, they modify a faulty program according to different procedures (e.g., using search-based (Le Goues et al. 2012) or enumerative procedures (Weimer et al. 2013)) and then validate the generated program against a test suite. Notably, GenProg (Le Goues et al. 2012; Colanzi and McMinn 2018) is one of the first APR systems that implements the generate-and-validate approach, using genetic programming to modify the faulty program based on three repair operators and running test cases to validate the modified program. The three operators used in GenProg can (1) add a statement copied from another location of the same program, (2) replace a statement of the source code with another one copied from another point of the same program, and (3) simply delete a code element (e.g., a statement).

Kali (Qi et al. 2015) is a repair technique that exclusively focuses on the latter: code removal. Kali achieves code-removal using three operators: a) an operator to remove a code element (e.g., a statement); b) an operator to replace a condition so as to force the execution of only a specific branch (e.g., the `if` branch); c) an operator to add a `return` statement that interrupts the execution of a procedure. Interestingly, recent empirical studies have shown that Kali is effective in finding test-suite adequate patches in large and complex systems (Qi et al. 2015; Long and Rinard 2016). Our intuition is that the very presence of code-removal patches has some meaning, and that code-removal patches may be useful in some ways. The open research question is to what extent the presence of a code-removal patch indicates the need of removing or changing code, or suggests that the program is not well tested. Our paper is the first that studies code-removal patches in depth: we investigate the reasons and the scenarios that make the generation of code-removal patches possible, and the relation between code-removal patches and correct human-written patches.

For example, if a developer modifies a program feature without updating the corresponding test cases, the program might fail during the execution of the test cases. The code-removal patch would simply remove the last changes of the programmer putting the feature in its original state, making the program pass the test cases again. This is an important hint for the developers who can exploit the code-removal patch to reveal the problem in the test cases. Our paper proposes a comprehensive taxonomy of code-removal patches, which both improve the level of understanding of program repair techniques and can be exploited in future research on program debugging. In particular, we provide evidence that code-removal patches may allow 1) to generate higher quality patches, avoiding the creation of patches that work only because, for instance, the faulty functionality is no longer executed, and 2) to help developers during the debugging phase, speeding up the analysis and the understanding of the bugs.

Given the width of possible code-removal patches, we qualify the problem domain as follows. We consider faults revealed by either one failing or one crashing test case only, since

it is a situation commonly encountered in practice. For instance, the Bears benchmark (Madeiral et al. 2019), which is a benchmark of 251 reproducible bugs from 72 different projects, has 71.32% of builds with a single failing (38.65% of the total) or crashing (32.67% of the total) test case. We take the bug data from the Repairator-Experiments repository (Urli et al. 2018), which is a repository that includes more than 4 thousands builds with a single failing or crashing test case. On those 4k+ faults, we perform a deep qualitative analysis, and identify major contributions for the understanding of code-removal patches.

In a nutshell, the contributions of this paper are:

- An analysis of nearly 2k failed builds from 674 real Java software projects and their corresponding code-removal patches, that make possible to have a clear view about the effectiveness of code-removal patches on a wide range of programs;
- The definition of a comprehensive taxonomy of code-removal patches that can be exploited to better understand the current limitations of program repair techniques, showing the potential issues that may cause the acceptance of wrong code-removal patches;
- A comparative analysis of human patches and code-removal patches on the same bug, showing the possibility to exploit automatic code-removal patches to give valuable information to developers about the cause of the problem, e.g., errors in test cases, or flaky tests;
- A freely accessible open science repository¹ containing all the artifacts and the outcome of the analysis, which can be used by other researchers in the program repair community.

The paper is organized as follows. Section 2 provides the background knowledge necessary to understand our study. Section 3 explains the experimental procedure we used to conduct our study, providing details about our choices, and explaining the reasons behind them. Section 4 presents the results related to our analysis and provides the answers to our research questions. Section 5 describes the threats to the validity of our results. Section 6 presents the related work. Section 7 provides final remarks.

2 Background

This section provides definitions and background information useful to understand our study. In particular, we describe the tool we used to generate code-removal patches, we present the repository of build failures that we analyzed in our study, and finally we provide the fundamental definitions useful in the domain of APR about test cases and patches.

2.1 Program Repair with Code-removal

Kali is an original automatic repair system developed by Qi et al. for C programs, that only removes functionality (Qi et al. 2015). The reason for Kali is that generate-and-validate approaches often end up repairing software systems by deleting functionality. For this reason,

¹ <https://github.com/repairator/open-science-repairator/>

Kali has been designed to investigate how a system that simply and only removes functionality is effective compared to other, more sophisticated, repair strategies.

Definition: Code-removal Patch A code-removal patch is a patch that simply removes functionality, by deleting, or skipping code. This latter case can be done through the replacement of a condition in order to force the execution of a specific branch, or through the addition of a `return` statement in a function body. Despite functionality removal, a code-removal patch may change a program making a full test suite to pass (Qi et al. 2015).

Listing 1 shows an example of a code-removal patch generated by Kali for the bug php-309892-309910, related to the PHP's standard library function `substr_compare`. In this case, the code-removal patch changes the `if` condition adding the instruction `&& !(1)`, thus the body of the `if` statement is skipped, because the condition is always evaluated as false. This patch is semantically-equivalent to the fix implemented by the developers, which entirely removes the `if` statement (Cambronero et al. 2019).

```

1 -   if (len > s1_len - offset) {
2 +   if (len > s1_len - offset && !(1)) {
3       len = s1_len - offset;
4   }
```

Listing 1: Example of code-removal patch generated by Kali for the bug php-309892-309910.

2.2 Choice of jKali Implementation

In addition to the original Kali (Qi et al. 2015) for C programs, we are aware of two other implementations of the approach, Astor's jKali (Martinez and Monperrus 2016) and Arja-Kali (Yuan and Banzhaf 2020), both for Java programs. To choose the implementation to use for our experiments, we analyzed the publicly available data.

In Table 1, we report the number of code-removal patches generated by Arja-Kali and Astor jKali in previous published research.

The first column (*Experiment*) indicates the reference to the experiment from which the data has been extracted from. The second column (*Benchmark*) indicates the benchmark and the number of total bugs on which the repair tools have been tested, The third column (*Arja-Kali*) and the fourth column (*Astor jKali*) report the number of code-removal patches generated by Arja-Kali and Astor jKali, respectively.

Based on these data, the results obtained by the different implementations are similar, with some exceptions. In the cases that involve the Defects4J benchmark, Arja-Kali seems to perform better than Astor jKali. According to the results reported by Yuan and Banzhaf (2020), Arja-Kali is able to generate 33 plausible patches for the 224 bugs of Defects4J (Just et al. 2014). However, in the experiment conducted by Durieux et al. (2019), we noticed that for 12 bugs that were considered patched by Yuan and Banzhaf (2020), Arja-Kali was not able to generate any patch. This mostly happens for the bugs related to the project Apache

Commons Lang, where 7 out of 9 bugs were not patched in the new experiment conducted by Durieux et al. (2019).

This aspect is confirmed also by the results of the experiment conducted by Liu et al. (2020) that compares some program repair tools on the Defects4J benchmark to evaluate the efficiency of test-based automated program repair. In this experiment, it is reported that Arja-Kali is not able to generate patches for bugs related to Apache Commons Lang project. Moreover, the experiment reports that the correctness rate of patches generated by Arja-Kali is 4.6%, while the correctness rate of patches generated by Astor jKali is 16%.

We manually tested the patches generated by Arja-Kali and Astor jKali for the bugs of Defects4J reported in the experiment conducted by Durieux et al. (2019). In that experiment, it is reported that Arja-Kali is able to generate 72 patches, while Astor jKali 27. However, applying the patches to the faulty programs, we observed that only 23 out of 72 patches generated by Arja-Kali work. For other 3 cases (Closure-133, Math-8, and Math-85), we did not find the patches in the repository linked with the paper².

We noticed that most of the patches (46) generated by Arja-Kali makes the programs pass the failing test cases, while introducing new bugs that make other tests to fail. Thus, the patches were not likely tested with the whole test suite, but only with a subset of the tests. This hypothesis is confirmed by the analysis of the log files associated with the execution of Arja-Kali (e.g. the one associated with Closure-3³), that report a number of positive test cases used to evaluate the patches that is smaller than the number of positive test cases used for fault localization.

These results suggest that Arja-Kali may not have a reliable behavior. On the other hand, all the patches generated by Astor jKali that we tested actually work. We thus decided to select Astor jKali for our study, that is, the most reliable Kali tool for Java based on the empirical evidence reported so far.

2.3 Repairator and Repairator-experiments

Repairator is a software engineering bot that monitors program bugs discovered during Continuous Integration, and tries to fix them automatically (Monperrus et al. 2019). Repairator implements multiple repair strategies, including the generation of code-removal patches using jKali.

Repairator-Experiments is an open science repository⁴ that contains the metadata information of the Travis CI builds that Repairator tried to repair (Urli et al. 2018). It hosts 14,132 failed builds (collected in the period February 2017 - September 2018) for 1,609 Java open-source projects hosted on GitHub. It provides detailed information about the builds, such as the event that triggered the build, the number of failing and crashing test cases, and the type of failures. Moreover, for every build, Repairator-Experiments stores the source code associated with the failing commit.

We used the builds hosted in Repairator in our study.

² https://github.com/program-repair/RepairThemAll_experiment

³ https://github.com/program-repair/RepairThemAll_experiment/blob/d469678068870bee75943f8b23331d31d0320bed/results/Defects4J/Closure/3/Kali/0/repair.log

⁴ <https://github.com/repairator/repairator-experiments>

Table 1 Number of code-removal patches generated in published experiments (Durieux et al. 2019; Liu et al. 2020; Yuan and Banzhaf 2020)

Experiment	Benchmark	Arja-Kali	Astor jKali
(Durieux et al. 2019)	Bears (251)	15	10
(Durieux et al. 2019)	Bugs.jar (1,158)	24	8
(Durieux et al. 2019)	Defects4J (395)	72	27
(Durieux et al. 2019)	IntroClassJava (297)	5	5
(Durieux et al. 2019)	QuixBugs (40)	2	2
(Liu et al. 2020)	Defects4J (395)	65	25
(Yuan and Banzhaf 2020)	Defects4J (224)	33	22

2.4 Definition of Failing and Crashing Test Cases

The projects in the Repairator-Experiments repository are equipped with automated unit test cases. In our study we used these tests to validate the patches generated by jKali.

When a test fails, we distinguish two main cases: failing test cases and crashing test cases.

Definition: Failing Test Case A failing test case is a test case that fails due to a violated assertion. Typically, it means that the actual values generated by the program under test violate one of the assertions present in the tests. Per the terminology of JUnit, this is a test failure (Langr et al. 2015). Thus, a failing test case reports an invalid test result.

Definition: Crashing Test Case A crashing test case is a test case for which the program under test generates an exception during its execution. The exception can be caught or uncaught. In the former case, the exception is caught using a `try-catch` statement in the test case that, for example, uses the instruction `Assertion.fail()`⁵ in the body of the `catch` to make the test fail. In the latter case, the execution terminates with an uncaught exception, per terminology of JUnit, causing a test error (Langr et al. 2015). A frequent case is the case of the program raising a `NullPointerException`⁶, which occurs when a program attempts to use an object that has not a value.

2.5 Selection of the Builds Used in the Experiment

We focus on the builds that have only one failing test case or only one crashing test case and are hosted in the Repairator repository. The restriction to builds with a single failing or crashing test is motivated by the fact that multiple failing or crashing test cases suggest the presence of several bugs in the code, which may significantly complicate the creation of a patch that fixes all of them. This would imply discarding code-removal patches that make the program pass only a subset of the failing/crashing test cases, potentially biasing the results.

Moreover, with our study we want to investigate which are the most common types of failure and exceptions for which a code-removal patch can work and why it works. Considering mixed scenarios with builds characterized by many failing and crashing test cases would have made manual analysis harder and results not as clear cut.

⁵ [https://junit.org/junit4/javadoc/4.12/org/junit/Assert.html#fail\(\)](https://junit.org/junit4/javadoc/4.12/org/junit/Assert.html#fail())

⁶ <https://docs.oracle.com/en/java/javase/15/docs/api/java.base/java/lang/NullPointerException.html>

2.6 Test-suite Adequate and Correct Patches

Even when a generate-and-validate APR technique modifies a program until none of the available test cases fail or crash, the resulting program is not necessarily correct. In fact, test suites are typically inadequate to cover every expected behavior of a program, and thus programs changed using the patches generated by APR techniques may yet be incorrect (Yu et al. 2017).

Due to this aspect, we need to distinguish between two types of patches: test-suite adequate and correct patches (Martinez et al. 2016).

Definition: Test-suite-adequate Patch (Also Known as Plausible Patch) It is a patch that makes the program pass all the available test cases, but it might potentially be incorrect. This situation occurs when the test suite does not cover every relevant behaviors of the program. For example, a patch that simply removes the faulty instruction may make a program pass all test cases without producing a correct fix.

Definition: Correct Patch It is a patch that makes the program pass all the available test cases and satisfies the requirements of the application. In particular, in this study a patch is considered correct if it is 1) identical or 2) semantically-equivalent to the human-written patch.

3 Experimental Methodology

This section describes our experimental procedure and discusses the design choices relevant to our study.

3.1 Goals & Research Questions

The goal of this study is to understand the nature of test-suite-adequate code-removal patches and to investigate how these patches relate to human patches, with a thorough qualitative study. To our knowledge, this is novel in program repair research, nobody has ever studied this important point.

In particular, the study aims to answer to the following research questions:

- RQ1 What is the relation between assertion failures and the generation of test-suite-adequate code-removal patches?** This research question investigates the ability of jKali to generate code-removal patches for tests that fail due to the violation of an assertion
- RQ2 What is the relation between crashing tests and the generation of test-suite-adequate code-removal patches?** This research question investigates the ability of jKali to generate code-removal patches for the faults revealed by crashing test cases.
- RQ3 To what extent can code-removal patches, even if incorrect, give valuable information to developers?** The goal of this research question is to study if a code-removal patch can give valuable information about the cause of the problem, regardless of its correctness.

RQ4 How do developers fix the failed builds associated with a test-suite-adequate code-removal patch? The goal of this research question is to study if developers fix bugs according to some patterns when the bug can be also addressed with a code-removal patch. We also investigate the semantic and syntactic similarities between the fix produced by developers and the patch produced by jKali.

Our experimental methodology is organized in three main phases: the collection of build failures, presented in Section 3.2; the analysis of the collected build failures, to identify the ones amenable to code-removal patches, presented in Section 3.3; and the analysis of the code-removal patches and their comparison to developers' patches, presented in Section 3.4.

3.2 Data Collection

The first phase of our study consists of collecting the build failures, and the related artifacts, necessary to answer to RQs 1-4. As source of build failures, we consider the Repairator-Experiments repository, which includes thousands of past failed builds annotated with metadata as described in Section 2.3. For each failed build, we consider several artifacts:

- The *failure-related artifacts*.
- The *code-removal patches* generated for the failed builds.
- The *human patches* associated with the failed builds that also have a code-removal patch.

All these data have been saved on a GitHub repository⁷ for the sake of scientific reproducibility. We describe below how we obtain these artifacts.

Failure-related Artifacts The builds relevant to our study are all the builds stored in the Repairator-Experiments repository with either one failing or one crashing test case. We found 2,381 builds with only one failing test case, and 1,724 builds with only one crashing test case. Overall 4,105 out of the 14,132 builds (29.05%) satisfy our selection criterion. For each failed build, we collect from the Repairator-Experiments repository the source code version that made the build fail, the metadata associated with the failed build (e.g., the build id), and finally the information about the failure (e.g., the name of the class and the test method that fails or crashes).

Code-removal Patches For every selected build, we ran jKali to generate code-removal patches, if such patches exist. We did not use the code-removal patches already present in the Repairator-Experiments repository because they have been generated with a previous version of jKali, which was affected by several bugs now fixed. These new runs ensured that a project code associated with a failed build is still executable. In fact, sometime there are problems with deletion of branches or dependencies that now make the execution of jKali impossible.

We set jKali to find every possible patch in the search space, not stopping after finding the first patch. The search space ordering used by jKali is given by the suspiciousness score of the

⁷ <https://github.com/repairator/open-science-repairator/>

faulty lines computed with fault localization, implemented in GZoltar (Campos et al. 2012). GZoltar returns a list of lines associated with a suspiciousness score, ranked from the most suspicious to the least suspicious one, according to the Ochiai formula. jKali, starting from the most suspicious line, tries to generate a code-removal patch by changing the program in the considered line. The code-removal patch can be generated using any of the following three operators as follows: 1) remove a code element, 2) replace a condition in order to force the execution of a specific branch, and 3) add a return statement to interrupt the execution of a procedure. Thus, for every failed build, there could be 0, 1 or multiple code-removal patches.

Our experiment was performed using a CPU Intel Xeon E5-2690v4 (2x14 cores) with 128GB of Memory per node. All nodes are running Ubuntu Focal (20.04 LTS). jKali was run in parallel on different builds thanks to the PFS file system available on Swedish National Infrastructure for Computing (SNIC).

Given the computational nature of the task, we run jKali with a timeout of 3 hours for every build. This 3 hours period comprises all phases of jKali, from downloading the source code from the repository to the execution of the repair attempts. The core repair loop of jKali itself is set to run for a maximum of 100 minutes. This setup is based on previous research (Durieux et al. 2019), in which repair tools required 13.5 minutes on average to generate a patch. Thus, allocating 100 minutes is 7.4 times the average repair time, and it reduces the probability of stopping the repair process due to a timeout. Using this configuration, in our experiments jKali never reached the timeout.

This process produces in total 129 code-removal patches on the considered bug benchmark (65 patches for the 28 builds with a failing test case and 64 patches for the 24 builds with a crashing test case).

Human Patches The data collection step finally includes the identification of the human patches associated with the builds for which jKali created a code-removal patch. To determine the corresponding human patch, we combine automatic and manual analysis: we use the Travis CI API⁸ to automatically determine the commits that may include the patch and also we manually check the presence and appropriateness of the human patch.

The process to retrieve the human patch consists of three steps:

1. Find the information about the failed build using Travis CI's API⁹. Among the available data, the most relevant for the evaluation process were: build number, commit causing the failure, and the repository of the project;
2. Find the list of builds subsequent to the failed one using Travis CI API¹⁰ and the build number. Filter the list of builds related to the same branch/pull request of the failing commit and check the state in order to find the first "passed" one;
3. Extract the commit associated with the passed build and check the diff with the failing build, read the commit messages that explain the applied changes to better understand what developers did and why. If only one commit with a simple change is extracted, we confirm the patch by directly applying this transformation to the failed build and running the test suite. Unless hundreds of lines spanning multiple files are part of the diff, these

⁸ <https://docs.travis-ci.com/user/developer/>

⁹ <https://developer.travis-ci.com/resource/build#Build>

¹⁰ <https://developer.travis-ci.com/resource/builds#Builds>

changes are applied locally to the failing build to check if the build passes the test cases. If the answer is positive, these changes are considered as the ground truth patch derived from the human patch. When there are too many changes between the failing build and the passed one, we do not derive any ground truth human patch, since it is infeasible for us to precisely isolate the changes that actually contribute to fixing the bug.

A single person was involved in the three steps described above. We relied on the information associated with the build using the Travis CI API. Overall, this combination of automated and manual analysis has taken 3.5 hours per code-removal patch on average and 21 days in total.

3.3 Analysis of Failed Continuous Integration Builds

The second part of the study consists in the analysis of the collected builds, which is organized in the following two steps:

- we perform a sanity check of the builds,
- we retrieve essential information about the failure.

Sanity Check of Builds The goal of the first step is to check if the characteristics of the builds used in our experiment, e.g., the number of failing/crashing test cases, are the same as the ones reported in the previous Repairmator experiment (Urli et al. 2018). Indeed, even if a build failed in the past, there is no guarantee that the same build fails with the same error after several months: this is due to changing external dependencies, closed third-party services that the application under test uses, flakiness¹¹ of the tests, etc. In addition, our selection criterion requires builds that include only one failing or crashing test case. To sum up, the sanity check ensures the following two conditions:

- the build process terminates correctly,
- the execution of the tests terminates with either a failing or a crashing test case.

We thus downloaded the 4,105 selected builds and executed their test cases, discarding the builds that did not pass the sanity check. At the time of the experiment in March 2020, this step results in 2,187 builds discarded for the following reasons:

- 635 builds had errors during the build phase (e.g., because some dependencies are not resolvable);
- 393 builds passed all the test cases, suggesting the presence of flaky tests causing the failure on the first place;
- 75 builds generated a timeout during the execution of the test cases;
- 1,084 builds had multiple failing and crashing test cases.

¹¹ As observed in a study conducted by Durieux et al. (2020), 0.80% of the builds of their dataset, which contains 3,286,773 Travis CI builds, are flaky. In particular, 46.77% of the restarted builds (1.72% of the builds of their dataset) change their state from *failed* to *passed*. Thus, builds that fail due to flaky tests are not unusual.

After the sanity check, the number of relevant builds for our study amounts to 1,918. In particular, 949 builds have only one failing test case, and 969 builds have only one crashing test case.

Retrieval of Information about the Failure The second step consists of collecting qualitative and quantitative information about the failed builds. For the builds with one crashing test case only, we collect the type of the exception responsible for the failure (e.g., `NullPointerException`). For the builds with a failing test case only, we collect information about the type of the failure. However, differently from crashing test cases, failures do not have an explicit type information assigned. We thus classified failures based on a manual analysis of the test and of the failed assertion in particular.

Table 2 shows the categories that we identified. The first column (*Failing Assertion Category*) specifies the category of the assertion failure, the second column (*Definition*) defines the category, and finally the third column (*Example of potentially failing test code for this reason*) contains a sample failure assertion extracted from our benchmark to exemplify the category.

Wrong Values generally represents all those cases in which there is a difference between the expected value and the one generated by the program under test. The code snippet shows an example of a test case whose aim is to verify if the status code of the HTTP response is 200.

Mocking Verification Failure represents a failure reported by a Mock framework, for instance because a specific method is not called the expected number of times. The code snippet shows the case of a failure reported by Mockito because the method `process` with the argument `k1` is called a number of times different from 2.

Exception Difference represents failures caused by a difference between the observed exception and the expected exception. The code snippet shows the case of a test case that fails because the expected exception of type `ParseException` is not generated.

Timeout represents tests that fail due to the timeout of an operation. The code snippet shows an example of a test case that fails because the `Subscriber` does not receive a notification within a second from the time `Observable` finished its job.

Finally, *Environment Misconfiguration* represents failures caused by problems in the testing environment. The software environment includes every entity external to the program, such as configuration files, system variables, external programs. The code snippet shows a test case that fails because the username and password associated with the Maven environment have not been properly set up.

3.4 Analysis of Human Patches and Automated Code-removal Patches

The third part of the methodology consists of a quantitative and qualitative analysis of both the generated code-removal patches and the ground truth human patches. In the rest of this section, we present our analysis and the identified categories.

Classification of Code-removal Patches We analyze code-removal patches to first determine their correctness. To this end, similarly to previous studies (Martinez et al. 2016; Qi et al. 2015), we consider a patch to be correct if it is either identical or semantically-equivalent to the corresponding human patch.

Table 2 Classification of reasons for failing test cases

Failing Assertion Category	Definition	Example of potentially failing test code for this reason
Wrong Values	The value generated by the program is not acceptable according to the expected value in the test assertion.	<pre> CloseableHttpResponse response = httpClient.execute(httpGet); assertEquals(200, response.getStatusLine().getStatusCode()); Processor<String> mockProc = mock(Processor.class); verify(mockProc, times(2)).process(eq("k1")); @Test (expected = ParseException.class) public void whenGivenBadScnlThrowHelpfulExcept () { parser.parse("not a SCNL"); } TestObserver<String> test = rxResponse.test(); test.awaitTerminalEvent(1, TimeUnit.SECONDS); assertNotNull("ensure \${env.CI_OPT_MVN_CENTRAL_USER} and \${env.CI_OPT_MVN_CENTRAL_PASS} is set.", plainText); </pre>
Mocking Verification Failure	The test execution does not pass a check performed on a mocked component.	
Exception Difference	The failure is caused by the generation of an exception of the wrong type or with the wrong message, or by a missing exception.	
Timeout	The failure is caused by the program not generating an output in the maximum allowed time.	
Environment Misconfiguration	The failure is caused by an incorrect execution environment (e.g., an environment variable is not set).	

In Listing 2, we exemplify the case of a code-removal patch generated by jKali for the failing Travis CI build with id 322406277¹² associated with the project pac4j. This patch is semantically-equivalent to the human patch, and is thus considered correct. The code-removal patch changes the `if` statement using `false` as condition, which forces the execution of the `else` branch, and consequently forces the method to always return `null`. The human patch removes the overridden method shown in Listing 2¹³. The program without the overridden method has the same behavior as the program with the code-removal patch. Thus, even though the code-removal patch does not entirely delete the overridden method **internalConverter(Object)**, it generates a program with the same behavior of the one that includes the human patch.

```

1      @java.lang.Override
2      protected String internalConvert(final Object attribute) {
3      -      if (null != attribute) {
4      +      if (false) {
5          return attribute.toString();
6      } else {
7          return null;

```

Listing 2: Example of code-removal patch generated by jKali for failing Travis CI build 322406277.

If the code-removal patch is not correct, we classify it according to its nature. In particular, we identified four different potential issues affecting the test cases that caused the acceptance of wrong code-removal patches: Weak Test Suite, Buggy Test Case, Rottening Test, and Flaky Test. Table 3 summarizes these cases.

Well known in the literature, a *weak test suite* might be responsible of the acceptance of a wrong code-removal patch. We confirmed this by showing that we can manually add assertions to existing tests or add new test cases that discard the code-removal patch.

For example, a `NullPointerException` is thrown by the program during the execution of the test case shown in Listing 3. This `NullPointerException` is generated because Hibernate saves a new record with an ID that differs from the one expected in the test case at line 8. The code-removal patch reported in Listing 4 works because it removes the instruction that assigns the ID and, since the test case does not check if the ID of the new record is not null, jKali is able to create the patch. Adding for example the assertion `assertNotNull(avantage.getId());` avoids the generation of the code-removal patch.

Sometime there are *buggy test cases*, that is, a wrong code-removal patch is not discarded because the test expects the wrong behavior from the program. This is revealed and conformed by manual fixes made to the test cases after the failure.

¹² <https://travis-ci.org/github/pac4j/pac4j/builds/322406277>

¹³ <https://github.com/pac4j/pac4j/pull/1076>

Table 3 Classification of code-removal patches. While the literature has focused on WT, reasons CP, BT, RT and FT have never been studied before

Category	Definition
CP: Correct Patch	The code-removal patch is identical or semantically-equivalent to the human patch.
Wrong Patch	
WT: Weak Test Suite	The available test suite does not cover the program well enough, better assertions and better tests might be needed.
BT: Buggy Test Case	The code-removal patch works due to a fault in a test case (e.g. the expected value in the test case is not correct).
RT: Rottening Test	The code-removal patch disables the execution of the failing assertion, that is located in a control flow statement. This means that the code-removal patch affects the return values used in the expression of the control flow statement containing the failing assertion, thus avoiding its execution.
FT: Flaky Test	A code-removal patch is accepted due to a flaky test that now passes.

```

1  @Autowired
2  private Avantage avantage;
3  ...
4  @Test
5  public void test_sauvegarder_lister_mettre_a_jour() {
6      ...
7      Avantage av = avantageRepository.findOne(avantage.getId());
8      assertThat(avantage.equals(av));
9      ...
10 }

```

Listing 3: Example of weak test case for failing Travis CI build 384760371.

For example, in Listing 5, it is reported the change made to the test case after the failure of build 368867994. The code-removal patch generated for this build is shown in Listing 6. The code-removal patch works because the test case expects that two `IniSection` objects are equal (line 10 of Listing 5), but the patch removes the `return false;` instruction executed when two objects are different (line 6 of Listing 6), resulting in pairs of objects that are always equal, even though they are different. For this reason, there is not the possibility to recognise the error in the test case and the code-removal patch works.

```

1      public void setId(Integer id) {
2  -          this.id = id;
3      }

```

Listing 4: Code-removal patch generated for failing Travis CI build 384760371.

```

1  public void VehicleFaultyTest(){
2      ...
3      IniSection result = vehicle.generateReport(1);
4      IniSection correct = new IniSection("vehicle_report");
5      ...
6      correct.setValue("time", "2");
7  -   correct.setValue("speed", "10");
8  +   correct.setValue("speed", "0");
9      ...
10     assertEquals(correct, result);

```

Listing 5: Wrong test case for failing Travis CI build 368867994.

A *rottening test* is a test that contains assertions that are executed only based on some conditions (Delplanque et al. 2019). If the code-removal patch changes the program in a way that the execution of the assertion is skipped, the test is now green (because the assertion is not executed) and the patch is accepted even if the program is still faulty. This is the first paper showing that such test issues affect the acceptability of program repair patches.

An example of this scenario is presented in Listing 7. The test case throws an exception when it tries to execute the cast at line 7. This instruction is executed only if the JSON object created at line 3 has the property “priority”. The code-removal patch generated for this build and shown in Listing 8 removes the instruction to add the property “priority” to the JSON object, and in this way the program passes the test case.

```

1  @Override
2  public boolean equals(Object obj) {
3      ...
4      for (java.lang.String key : this.getKeys()) {
5          if (!this.getValue(key).equals(other.getValue(key))) {
6  -             return false;
7              }
8          }
9      return true;

```

Listing 6: Code-removal patch for failing Travis CI build 368867994.

```

1  @org.junit.Test
2  public void serializeTest() {
3      JSONObject ds = todo.serialize();
4      assertEquals( todo.getTitle(), (String) ds.get("title") );
5      ...
6      if (ds.has("priority")) assertEquals( todo.getPriority().toString(),
7                                          (String) ds.get("priority"));
8  }

```

Listing 7: Example of rottening test for failing Travis CI build 88971125.

Finally, the presence of flaky tests (Luo et al. 2014) may let jKali accept a patch. When a flaky test stops failing, it has actually no causal relation with the generated patch, and the patch is then wrongly assumed as being correct. For example, the test case shown in Listing 9 can fail only due to a timeout that not always occurs. The generated code-removal patch shown in Listing 10 changes a piece of code that is not executed by the failing test case, and thus it confirms that the behavior of the program is not influenced by the code-removal patch, and

Table 4 Classification of human patches when a code-removal patch exists

Patch Type	Category	Definition
Fix in Test	Fix Test Code	The patch fixes the logic of one or more test cases to properly reflect the expected behavior.
	Fix Test Data	The patch fixes the data used in one or more test cases.
Statement-Level Change	Change Condition	The patch modifies a condition used in the program, for example in a <code>if-statement</code> or in a <code>while</code> .
	Add if-else Statement	The patch adds a new <code>if-else</code> statement to conditionally execute part of the code.
Method-Level Change	Change Method Implementation	The patch modifies multiple statements inside the same method.
	Override Method	The patch introduces a method that overrides another method present in the program.
Code Removal	Remove Variable Assignment	The patch removes an assignment statement from the program.
	Remove Variable Annotation	The patch removes an annotation from the program.
	Revert	The patch reverts the code to a previous commit.
Not Available	No Change	The patch includes no changes, or if they are present, they are not related to the failure, typically because the build failure is due to the presence of flaky tests.
	Not Found	It is not possible to determine the changes that make the patch, for instance because they are mixed with many other changes not related to the removed fault.

that test case is flaky.

```

1 public JSONObject serialize() {
2     JSONObject ds = new JSONObject();
3     ds.put("title", this.getTitle());
4     ...
5 - ds.put("priority", this.getPriority());
6     return ds;
7 }

```

Listing 8: Code-removal patch generated for failing Travis CI build 88971125.

```

1 @Mock
2 private Timer.Context context;
3 ...
4 @Test
5 public void timerIsStoppedCorrectly() throws Exception {
6     HttpHost host = startServerWithGlobalRequestHandler(STATUS_OK);
7     HttpGet get = new HttpGet("/?q=anything");
8
9     ...
10    verify(context, timeout(100).times(1)).stop();
11 }

```

Listing 9: Failing test case for failing Travis CI build 374587117.

```

1 @java.lang.Override
2 public void failed(java.lang.Exception ex) {
3     timerContext.stop();
4 -     if (callback != null) {
5 +         if (true) {

```

Listing 10: Code-removal patch for failing Travis CI build 374587117.

Classification of Human Patches In our study, for 32 out of 48 cases (66.67%) the human patch corresponding to a code-removal patch is found. For each of them, we classify the patch based on the type of changes implemented by the developers. The classification takes into

account the size of the change (e.g., if changes are localized in a statement or in a method) and the target of the changes (e.g., if the changes target the program or the tests). We also look for code-removal changes made by developers which match the type of changes produced by jKali. Table 4 shows a summary of all the categories we found by analyzing the patches.

The first two categories, *Fix Test Code* and *Fix Test Data*, correspond to patches implemented by changing the code of the test, while distinguishing between changes to the logic of the tests and changes to the data used in the tests. A number of categories capture the case of actual modifications implemented in the code of the faulty program (excluding code-removal only patches). These changes might be at the level of the *individual statements* or at the level of the *methods*. Changes to individual statements involved either conditions or *if-else* statements. Although other types of changes to individual statements are possible, we do not list patterns for which no code-removal patch exists in our dataset. Changes to methods involved either a method or the addition of an overridden version of a method.

Interestingly, we have a number of human patches that consist of *Code Removal* operations. We report cases in which the developers removed assignments, annotations or revert a code change in Section 4.4.1.

Finally, sometimes the human patch is not available, either because the failed build has been intentionally not fixed (e.g., because the failure was caused by a flaky test) or because our procedure was not able to uniquely identify the human patch, as already described in Section 3.2.

3.5 Summary

In this section, we have presented a novel categorization of failures and code-removal patches. In particular, the taxonomy of failures (Table 2), the classification of code-removal patches based on the causes that make them work (Table 3), and the categorization of ground-truth human patches (Table 4) are contributions that can provide a solid foundation for future studies in program repair.

4 Experimental Results

This section presents the results of our analysis and provides the answers to our research questions.

4.1 What is the relation between assertion failures and the generation of test-suite-adequate code-removal patches? (RQ1)

In this research question, we analyze the relation between the category of assertion failures and the proportion of generated code-removal patches, per the methodology of Section 3. Table 5 reports the information about the number of builds per category of assertion failures and the corresponding code-removal patches generated by jKali.

The first column (*Failing Assertion Category*) lists the different categories of assertion failures, the second column (*Occurrences*) shows the number of builds that have a failing test case for each category of assertion failure, the third column (*# Builds with Patch*) indicates the number of builds that have a code-removal patch generated by jKali, and the fourth column (%)

Table 5 Relation between the test failure categories and code-removal patches

Failing Assertion Category	Occurrences	# Builds with Patch	% Builds
Wrong Values	842	22	2.61%
Exception Difference	39	3	7.69%
Mocking Verification Failure	20	2	10.00%
Timeout	46	1	2.17%
Environment Misconfiguration	2	0	0%
Total	949	28	2.95%

Builds) indicates the percentage of builds with a code-removal for every failing assertion category and over all 949 builds. The data are presented in descending order by the number of patched builds.

4.1.1 Analysis of the Results

The most frequent category of assertion failure is *Wrong Values* with 842 occurrences, while the least frequent category is *Environment Misconfiguration* with 2 occurrences.

jKali generates a patch for all failure categories, with the exception of *Environment Misconfiguration*. This is due to both the few occurrences in this category, but also to the missing capability of changing the environment configuration files in jKali. In fact, jKali is designed to create patches that change the source code of the target program, and cannot make any change to the environment.

Wrong Values is the category of assertion failure with the highest number of occurrences (842) and the highest number of code-removal patches (22). The high number of patches is only due to the high number of failing builds in that category. The assertion failures with the highest percentage of code-removal patches are *Mocking Verification Failure* (10.00%) and *Exception Difference* (7.69%). While this percentage is high, it is still a rare event and the absolute number of patches is still low (2 for *Mocking Verification Failure* and 3 for *Exception Difference*). This event rarity prevents us to make strong claims that those failures types are more amenable to code-removal patches.

To study if the generation of code-removal patches is dependent on the category of the assertion failure, we apply Fisher's exact test on the number of patched builds per assertion failure category. The null hypothesis of our test is that *the number of patched builds is independent of the assertion failure category*. We observe that the p-value is 0.0928, that is greater than the significance level α set to 0.05. This means that the null hypothesis cannot be rejected, and the capability to generate code-removal patches is likely independent of the category of assertion failure.

4.1.2 Comparison of the Results with Previous Studies

Another interesting aspect is that the proportion of generated code-removal patches is significantly lower than in previous studies. Indeed, in our case, the likelihood is 2.95%, while in the previous studies is 36.23% Long and Rinard (2016), 25.71% Qi et al. (2015), and 9.28% Martinez et al. (2016). Our best explanation is that it is due to the small size of the datasets used in former experiments, which consist of 69 cases for Long and Rinard (2016), 105 cases for Qi et al. (2015), and 224 cases for Martinez et al. (2016). A second explanation is that those

previous studies did not sample over builds but over commits. A third explanation is that the benchmarks of those studies used some kind of selection, which results in a biased sampling.

Yet, our result is aligned with Durieux et al. (2019), in which the proportion of patches generated by jKali on the bugs of Bears is about 3.0%. Unlike the other benchmarks, Bears is a benchmark which uses CI builds to identify buggy and patched program version candidate, and it contains bugs associated with more different programs (72 projects) compared to the other ones (8 for ManyBugs benchmark (Le Goues et al. 2015), and 5 for Defects4J benchmark (Just et al. 2014)). Both Bears and this study sample builds, which explains the strong consistency.

What is the relation between assertion failures and the generation of test-suite-adequate code-removal patches? (RQ1)

jKali has been able to create a patch for 28 out of 949 (2.95%) builds having only one failing test case. We obtained one patch for every category of assertion failure with the exception of Environment Misconfiguration, which is out of scope for current generators of code-removal patches. Our results show that the generation of a code-removal patch is independent of the category of assertion failure. Our analysis suggests that former studies tended to over-estimate the prevalence of code-removal patches, because of the selection criteria considered. Our results are useful for program repair researchers, they give a better understanding of the somewhat limited repair capability of code-removal patches.

4.2 What is the relation between crashing tests and the generation of test-suite-adequate code-removal patches? (RQ2)

In this research question, we analyze the relation between the type of exceptions and the proportion of generated code-removal patches for crashing tests, per the methodology of Section 3. Table 6 shows the details about the number of available builds, divided per exception type, and the corresponding number and percentage of patches produced by jKali.

The first column (*Exception Type*) indicates the name of the exception, the second column (*Occurrences*) reports the number of builds having a crashing test case with the specific type of exception, the third column (*# Builds with Patch*) indicates the number

Table 6 Relation between the types of crashing exceptions and code-removal patches

Exception Type	Occurrences	# Builds with Patch	% Builds
NullPointerException	124	9	7.26%
Exception	66	2	3.03%
OutOfMemoryError	5	2	40.00%
ClassCastException	7	1	14.29%
FileNotFoundException	27	1	3.70%
IllegalArgumentException	54	1	1.85%
IllegalStateException	241	1	0.41%
javax..PersistenceException	2	1	50.00%
rocketmq..MQClientException	1	1	100.00%
RuntimeException	62	1	1.61%
org.apache.dubbo.rpc.RpcException	20	3	15.00%
ConnectorStartFailedException	1	1	100.00%
Other exceptions	359	–	–
Total	969	24	2.48%

of builds that have a code-removal patch generated by jKali, while the fourth column (*% Builds*) reports the percentage of code-removal patches for every type of exceptions and over all 969 builds. We report the data in descending order by the number of patched builds including only the exceptions for which there is at least one code-removal patch. In total, jKali is able to create a patch for 24 out of 969 builds with a crashing test case, with a success rate of 2.48%.

4.2.1 Comparison of the Results between Builds with Failing and Crashing Test Cases

The success rate reported for crashing failures is in line with the success rate reported for builds with failing test cases (2.95%). To study if the nature of the failure, produced by either a crashing or a failing test case, has an impact on the capability to produce a code-removal patch, we applied the Fisher's exact test considering the following null hypothesis: *the number of patched builds is independent of the type of test case that reveals the failure (either a failing or a crashing test case)*. We observe a p-value equals to 1.196954, that is higher than the significance level α set to 0.05. This means that the null hypothesis cannot be rejected, and conclude that the generation of code-removal patches is likely independent on the type of test case that reveals the failure.

4.2.2 Analysis of the Results

The most frequent type of exception is `IllegalStateException`¹⁴ with 241 occurrences, but only one build has a code-removal patch. This type of exception occurs when a method has been invoked at an inappropriate time, thus a code-removal patch has a low likelihood to make pass a crashing test. Indeed, a code-removal patch can remove the wrong method call, but to avoid that the exception is thrown, it is usually also necessary to replace the removed call with the right piece of code (e.g., the invocation of another method), in order to create a legal program state. However, this is outside the scope of code-removal patches, that can only remove, but not add code.

A larger number of code-removal patches for `NullPointerException` (9) and the generic `Exception`¹⁵ (2) is reported. We explain this result by the fact that these exception types are prevailing in the benchmark. Furthermore, we manually analyzed them. We observe that jKali successfully patches programs producing a `NullPointerException` by removing the usage of the object that is null. Regarding the generic exception `java.lang.Exception`, the 2 code-removal patches are related to timeout errors and they remove the piece of code that causes the timeout. For example, the code-removal patch of the Travis CI build for Apache Twill (id 356030973) removes the call to method `java.net.InetAddress.getLoopbackAddress()`, whose execution could generate a deadlock.

Some exceptions only sporadically present in the benchmark have been successfully patched, that is the case for `MQClientException`¹⁶,

¹⁴ <https://docs.oracle.com/en/java/javase/15/docs/api/java.base/java/lang/IllegalStateException.html>

¹⁵ <https://docs.oracle.com/en/java/javase/15/docs/api/java.base/java/lang/Exception.html>

¹⁶ <https://rocketmq.apache.org/docs/quick-start/>

Table 7 Relation between the failing builds and code-removal patches

Patch reason	# Builds - fail. test	# Builds - crash. test	Total
Correct	1	1	2
Weak Test Suite	9	15	24
Buggy Test Case	7	2	9
Rotting Test	4	2	6
Flaky Test	7	4	11
Total	28	24	52

`ConnectorStartFailedException`¹⁷, and `PersistenceException`¹⁸. The number of builds with these types of exceptions is far too low to be able to generalize a finding. Finally, we report that `OutOfMemoryError`, `ClassCastException`, and `RpcException`¹⁹ have been patched with good frequency. The first one is an error that is thrown when there is insufficient memory for the program to work properly. The success rate of jKali is 40%. The second one is an exception that occurs when there is an instruction in the source code that tries to convert an object from one type to another, but they are incompatible (e.g., converting an Integer to a String). In this case, the success rate of jKali is 14.29%. The third one is a custom exception that occurs when remote procedure call fails. In this case, the success rate of jKali is 15%.

Finally, to study if the generation of code-removal patches is dependent on the exception type, we apply Fisher's exact test on the number of patched builds per exception type. The null hypothesis is that *the number of patched builds is independent of the crash category*. We report a p-value equals to 0.0004998, that is less than the significance level α set to 0.05, and thus this means that the type of exception influences the success rate of generating a code-removal patch.

What is the relation between crashing tests and the generation of test-suite-adequate code-removal patches? (RQ2) jKali has been able to create a patch for 24 out of 969 builds having one crashing test case only (2.48%). The most repairable type of common exception is `OutOfMemoryError` (5 occurrences), and `NullPointerException` is a prevalent exception with code-removal patches (9). Our experiment shows that the generation of code-removal patches can be influenced by the exception type in a statistically significant manner. This experiment shows that `IllegalStateException` is a very common kind of exception, for which we need specific program repair tools.

4.3 To what extent can code-removal patches, even if incorrect, give valuable information to developers? (RQ3)

In this research question we manually analyze why code-removal patches have been generated, and study if these patches can provide valuable information about the cause of the build failure.

¹⁷ <https://docs.spring.io/spring-boot/docs/1.5.7.RELEASE/api/org/springframework/boot/context/embedded/tomcat/ConnectorStartFailedException.html>

¹⁸ <https://docs.oracle.com/javase/7/api/javax/persistence/PersistenceException.html>

¹⁹ <https://grpc.github.io/grpc/csharp/api/Grpc.Core.RpcException.html>

Table 7 shows the different reasons that lead to the generation of a code-removal patch. The first column (*Patch reason*) lists the possible reasons, the second column (# Builds - fail. test) and the third column (# Builds - crash. test) report the corresponding number of builds with a failing and crashing test case, respectively. Finally the fourth column (*Total*) shows the total number of the patched builds for every patch reason.

4.3.1 Correct Patches

For only 2 out of 52 builds (3.85%), jKali managed to create a correct code-removal patch. This confirms previous research showing that code-removal patches are mostly incorrect (Qi et al. 2015; Martinez et al. 2016). For the other 50 builds, the generated patches are all different manifestations of the inadequacy in the test suite.

4.3.2 Weak Test Suite

For 24 out of 52 builds (46.15%), the problem is a *Weak Test Suite* that does not sufficiently assert the behavior of the program under test. For example, build 400611810 generates an assertion error when comparing the expected status code with the actual one. As shown in 11, the code-removal patch removes the instruction that adds the HTTP header, and in this way the resulting HTTP answer has the expected status code. The patch works because a test²⁰ checks only if the status code is correct, without checking if the HTTP request contains the right header.

```

1  --- /src/main/java/com/http/Request.java
2  +++ /src/main/java/com/http/Request.java
3  @@ -235,7 +235,6 @@
4      header.put("Accept-Encoding", "gzip, deflate, br");
5      header.put("Accept", "text/html,application/xhtml+xml,application/xml;
6                  q=0.9,image/webp,image/apng,*/*;q=0.8");
7      header.put("Connection", "Keep-Alive");
8  -   this.setHeader(header);
9  }
10
11 public java.lang.String getContent() {

```

Listing 11: Example of code-removal patch generated by jKali for build 400611810.

²⁰ <https://github.com/repairator/repairator-experiments/blob/bd4b41ef77dc3db1d7ac4c1ac991c5f214a8a84f/src/test/java/com/http/TestRequest.java#L11>

4.3.3 Buggy Test Case

For 9 out of 52 builds (17.31%), the code-removal patch reveals a *Buggy Test Case*, that is, a faulty test case that allows the acceptance of an incorrect patch. For example, build 35121194 of our dataset fails after the addition of a change that trims the output associated with the result of a command execution, but the test case is not updated to support this change. The code-removal patch works because it removes exactly the new instruction that trims the output (i.e., the patch undoes the change), and so the test case passes. To our knowledge, this is the first ever report of this phenomenon in the literature.

4.3.4 Rottening Test

We observed that the acceptance of incorrect test-suite-adequate patches has been caused by *Rottening Tests* in 6 out of 52 builds (11.54%). In such cases, the failing or crashing test case's assertion is no longer executed after the application of a code-removal patch. This result confirms that change in the application code can have an effect on the test execution (Martinez et al. 2020). For example, build 38897112 generates a `ClassCastException` exception when a test case checks the value associated with a property of a JSON object under test. Since the check is executed only if the object has the property, and since the code-removal patch removes the instruction to set the property of the JSON object, the patch passes all tests because the assertion is not executed. To our knowledge, this is the first ever report of this phenomenon in the literature.

4.3.5 Flaky Test

Finally, another interesting category is *Flaky Test* with 11 out of 52 builds for which incorrect test-suite adequate patches (21.15%) have been created. In these cases, the patch is accepted because of a *Flaky Test*, i.e. the patch is not related to the test pass. These code-removal patches make irrelevant changes that do not modify the logic of the program, (e.g., printing of log information, as observed for build 40344741 on Travis CI), but end up being accepted as side effect of the intermittent failures generated by flaky tests (e.g., if the flaky test does not fail after an irrelevant change is introduced in the program, the change is reported as an acceptable patch). An easy way to mitigate this problem is to run the failing or crashing tests multiple times before accepting a code-removal patch.

4.3.6 Debugging Hints from Code-removal Patches

Overall, this evidence suggests a code-removal patch always tells something interesting to the developers. Developers can exploit code-removal patches even when they are incorrect. In fact, by looking at the changes in the code-removal patches, developers can focus on the instructions that are either removed or not executed anymore due to the code-removal patch and understand if there is any bug in that code or if there is an error in the failing test case. As reported in Table 9, Column *Correlation Type*, when the human patch is not applied to a test case, we observe that the locations changed by the code-removal patches and human patches are different most of the time. Only in 4 out of 14 cases (28.57%) there is partial relation between the locations changed by the code-removal and human patches. Thus, the code-

Table 8 Strategies actually used by developers to fix builds patched by jKali

Patch Type	Category	# Builds fail. test	# Builds crash. test	Tot per Category	Tot per Patch Type
Statement-Level Change	Change Condition	1	0	1	2
	Add if-else Statement	0	1	1	
Method-Level Change	Change Method Implementation	2	3	5	6
	Override Method	0	1	1	
Code Removal	Remove Variable Assignment	0	2	2	6
	Remove Variable Annotation	0	1	1	
	Revert	2	1	3	
Fix in Test	Fix Test Code	9	4	13	19
	Fix Test Data	2	4	6	
Not Available	No Change	5	2	7	19
	Not Found	7	5	12	

removal patches should not be used to identify the precise locations that have to be changed to fix the bug, but as a way to obtain some hints for the debugging phase, to better understand the reason of the bug and how to fix it.

To what extent can code-removal patches, even if incorrect, give valuable information to developers? (RQ3) Our investigation reports that for only 2 out of 52 (3.85%) builds jKali managed to create a correct patch, showing that code-removal patches cannot be trusted. However, in all the cases where the patches are incorrect, the patches reveal different kinds of problems affecting the test suites that are relevant for the developers. Only in 4 out of 14 cases (28.57%) where the human patch is not applied to a test case, there is a partial relation between the locations changed by the code-removal and the human patches. Thus, the code-removal patches should not be used to identify the precise code location that must be changed to fix the program, but as a useful source of information to understand the reason of the failure. This result is relevant to researchers, since it opens new ways of exploiting code-removal patches, for instance as means to automatically improve test suites. This result is also interesting for practitioners, since our experiments suggest that code-removal patches carry useful information to identify fixes and weak test suites.

4.4 How do developers fix the failed builds associated with a test-suite-adequate code-removal patch? (RQ4)

In this research question, we investigate the relation between the fixes produced by developers and the automatically generated code-removal patches. To this end, we first retrieve and

analyze the fixes produced by the developers and then we relate these fixes to the code-removal patches.

4.4.1 Fixes Produced by Developers

Table 8 shows the details about the relation between the failing builds for which there is a code-removal patch and the types of human fixes. In particular, the first column (*Patch Type*) lists the different types of human fixes associated with the builds for which jKali is able to create a patch, the second column (*Category*) shows the specific categories of every patch type, the third column (*# Builds - fail. test*), and the fourth column (*# Builds - crash. test*) indicate the number of builds with a failing or crashing test case respectively fixed according to a specific category of human fix, the fifth column (*Tot per Category*) reports the total number of builds whose fixes belong to a specific category, and finally the sixth column (*Tot per Patch Type*) reports the total number of builds fixed by a specific type of human fix.

Fixes patching a single program statement is quite rare in our benchmark: it happens in just 2 cases (3.85%).

A more significant number of fixes span entire methods (6 cases, 11.54%). Method level changes do not follow specific patterns because they introduce or change pieces of logic in the program, they are far more complex than code-removal patches.

Interestingly, a non-trivial number of programmer fixes are actually code-removal patches (6 cases, 11.54%) that remove specific program elements (e.g., assignments and annotations) or revert changes (e.g., by undoing commits or closing pull requests without merging changes). The action of reverting a change is considered like a removal of code, because the code associated with the changes is deleted with that action. This result shows that revert-based repair is relevant, while this has been little researched in academia (Tan and Roychoudhury 2015).

Unexpectedly, the most frequent type of human fixes target the test cases and not the application code (19 cases, 36.54%). Human developer either fix the test code or the test data used by a test (e.g., a JSON file used by a test). This result reinforces the finding of RQ3 that incorrect code-removal patches can be exploited to improve test suites, including fixing wrong test cases. Moreover, it calls for more repair techniques able to generate fixes for test cases and not only programs (Daniel et al. 2009).

Finally, per our methodology, the human fixes are sometimes not available. In a significant number of cases, 11 (21.15%), this is because the build failure is due to flaky tests. Indeed, we notice that in 4 out of 11 cases (36.37%), the status of the build on Travis CI became `passed` after the original failure detected by Repairator. This is a piece of evidence that techniques are needed to make sure that the build failures to be repaired are indeed not due to flakiness.

4.4.2 Relation between Code-removal Patches and Developers Fixes

Table 9 relates code-removal patches to developer fixes. The first column (*Build ID*) contains the Travis CI IDs of the builds. The second column (*Build Type*) indicates if a build has a failing (FT) or crashing test case (CT). The third column (*# Code-Removal Patches*) indicates the number of code-removal patches generated by jKali for a specific build. The fourth column (*Code-Removal Patch Reason*) shows why a code-removal patch has been generated for a specific build. The fifth column (*Human Fix Category*) shows which is the category of fix that developers implemented to fix a bug in a particular build (WT is Weak Test Suite, BT is

Table 9 Comprehensive Data of the 52 Builds with Code-removal Patches

Build ID	Build Type	# Code-removal Patches	Code-removal Patches Reason	Human Fix Category	Correlation Type
322406277	FT	1	CP	Revert	Partial
365170225	FT	1	WT	Not found	None
397786068	FT	1	WT	Fix Test Data	None
353457987	FT	1	WT	Fix Test Code	None
368867994	FT	9	WT	Fix Test Code	None
400611810	FT	6	WT	Fix Test Code	None
249918159	FT	1	WT	Change Method Impl.	Disjoint
380634197	FT	1	WT	Change Method Impl.	Partial
372495757	FT	1	WT	Change Condition	Partial
413754623	FT	2	WT	Not found	None
354875355	FT	1	RT	Not found	None
403087258	FT	1	RT	Not found	None
351075282	FT	2	RT	Fix Test Data	None
378592651	FT	1	RT	Not found	None
351211949	FT	1	BT	Fix Test Code	None
408694507	FT	3	BT	Fix Test Code	None
390335750	FT	4	BT	Fix Test Code	None
349620528	FT	5	BT	Fix Test Code	None
363986485	FT	1	BT	Fix Test Code	None
387671228	FT	1	BT	Fix Test Code	None
396857150	FT	1	BT	Not found	None
214962527	FT	1	FT	Not found	None
403447416	FT	1	FT	No Change	None
415750114	FT	1	FT	Revert	Disjoint
374587117	FT	9	FT	No Change	None
402096641	FT	1	FT	No Change	None
387846982	FT	1	FT	No Change	None
415477949	FT	6	FT	No Change	None
384713759	CT	2	CP & WT	Remove Assignment	Same-location & Disjoint
356030973	CT	1	WT	Change Method Impl.	Partial
348327780	CT	1	WT	Fix Test Data	None
348335601	CT	1	WT	Fix Test Data	None
348337755	CT	1	WT	Fix Test Data	None
372415239	CT	1	WT	Fix Test Data	None
389668297	CT	1	WT	Revert	Partial
386721415	CT	2	WT	Fix Test Code	None
384760371	CT	1	WT	Remove Annotation	Disjoint
354919174	CT	4	WT	Add if-else Statement	Disjoint
346537408	CT	1	WT	Not found	None
373018834	CT	1	WT	Change Method Impl.	Disjoint
373043004	CT	1	WT	Change Method Impl.	Disjoint
367766867	CT	4	WT	Not found	None
418325841	CT	5	WT	Not found	None
358186949	CT	1	WT	Fix Test Code	None
388971125	CT	1	RT	Override Method	Disjoint
388975720	CT	1	RT		Disjoint

Table 9 (continued)

Build ID	Build Type	# Code-removal Patches	Code-removal Patches Reason	Human Fix Category	Correlation Type
				Remove Assignment	
385681821	CT	1	BT	Fix Test Code	None
363526725	CT	1	BT	Fix Test Code	None
421420531	CT	1	FT	No Change	None
415654258	CT	1	FT	No Change	None
407166687	CT	11	FT	Not found	None
415796275	CT	19	FT	Not found	None

Buggy Test Case, RT is Rottening Test Case, FT is Flaky Test Case, and CP is Correct Patch). The sixth column (*Correlation Type*) indicates which type of correlation exists between the changes performed by developers fixes and code-removal patches. We define three different types of correlations: 1) *Same-location*, when the code-removal patch and the human fix change exactly the same statements of source code, 2) *Partial*, when the code-removal patch and the human fix have in commons at least one line of code that is changed, and 3) *Disjoint*, when the code-removal patch and the human fix change different points of the source code and they do not have anything in common.

Correct Patches Notably, there are two cases in which the code-removal patch is correct, build 322406277 that fails due to a failing test case, and build 384713759, that fails due to a crashing test case. To fix the bug in build 322406277, the developer closed the pull request refusing the change that overrides a method. The corresponding code-removal patch changes this method, forcing the execution of a specific branch, whose behaviour is the same of the original method without the overriding. Thus, in this case there is a partial relation between the human fix and the code-removal patch. For build 384713759, the developer removed an assignment to a variable, and the code-removal patch (ID 1) does exactly the same change. This is the only case in which the developer and the code-removal patch change exactly the same location of the source code.

Weak or Incorrect Test Cases Considering the 24 builds for which jKali is able to generate a code-removal patch because of a weak test suite, there are 4 cases in which the code-removal patches and the human patches are partially related. For build 380634197, there is a partial relation between the changes applied by the developer and the corresponding code-removal patch, because the code-removal patch deletes an `else` branch of the same method fixed by developers. For the build 372495757, the human fix and the code-removal patch are partially related, because they change the same method, but in different parts. For the build 356030973, the human fix and the code-removal patch are partially related, in this case, the code-removal patch avoids the execution of the `if` branch that is changed by developers. In the case of build 389668297, the changes that introduce the bug have been reverted, while the corresponding code-removal patch avoids the execution of one of the paths of the new faulty method introduced by developers, indicating a partial relation between human and automated fix.

There are also five cases (249918159, 384760371, 354919174, 373018834, and 373043004) where the code-removal patches and human fixes are disjoint because they

change different points of the source code. In particular, for the build 384760371, the human fix and the code-removal patch are disjoint because they change different points of the source code, but they are semantically related because both changes influence the same value used by the program to save records in a database.

Interestingly, considering both the 24 code-removals patches that work due to weak test suites and the 9 ones that work due to buggy test cases (33 builds in total), in 18 out of 33 cases (54.55%) the human fix precisely consists in fixing the test code (13 cases) or the test data (5 cases). This finding confirms the relation between the quality of the test suites and the ability of jKali to generate code-removal patches.

Flaky Test When a code-removal patch works because of a flaky test, in 7 out of 8 cases (87.5%) there are no changes applied by developers to fix the failure, which is consistent.

The generation of patches that trivially alter, or do not alter at all, the semantics of the program are good indicators of failures caused by flaky tests. In fact, for the builds 403447416 and 421420531, the code-removal patches simply force logging, without introducing any other change to the logic of the programs. For the builds 374587117 and 415654258, the flakiness is related to timeouts. For the build 374587117, the code-removal patch removes a piece of code not executed by the failing test case, while for the build 415654258, the code-removal patch removes the instruction that closes a `Dispatcher` object. For the builds 402096641 and 387846982, the code-removal patches remove an assignment instruction, while for the build 41547794, the code-removal patch removes a method call, but apparently these actions do not influence the behavior of the program. For the remaining build 415750114, the flakiness is related to a rare race condition that causes the failure of the test case. The corresponding code-removal patch forces the execution of a specific `if` statement, without skipping the execution of other parts of code because it is not associated with an alternative (`else`) branch.

Rotting Test The generation of code-removal patches accepted because of rotting tests can provide useful information to improve the tests, for instance by tracking the tests and the assertions that are not anymore executed after the patch. Indeed, in 2 cases (builds 354875355 and 403087258) the code-removal patches remove the code that enables the execution of the failing test. Indeed, the test case fails only when a certain value is higher than a specified threshold. Since the code-removal patch avoids the increase of that value, the test case does not fail anymore. In other case, build 388971125, the developers fixed the source code overriding the method tested in the crashing test case. In these cases, the changes performed by developers and code-removal patch are disjoint.

For the build 378592651, the code-removal patch forces the execution of a specific branch, changing a value that is used in a condition of a test case to execute certain assertions. Since the condition checks if the value is different from a given threshold before a certain time, and the code-removal patch changes that value also when it should not happen, the value satisfies the condition, and the failing assertion is not executed anymore.

In the remaining case associated with the build 351075282, while the developer fixed the test data, the code-removal patch drops an entire block of code that influences the execution of the failing assertion in the test.

Problem of Fault Localization Overall, there are 14 builds for which the human patch changes the source code. We observe that in 8 out of 14 cases (57.14%), the code-removal patch is at a totally different location compared to the human patch. In other terms, the fault localization

technique used²¹ has a poor effectiveness in those 8 cases. This is another piece of evidence that the state of the art of fault localization is under-optimal for program repair (Liu et al. 2019).

4.4.3 Analysis of Builds with More Than One Code-removal Patch

Finally, we now consider the cases for which jKali is able to create more than one code-removal patch. Considering the builds with one failing test case, it was possible to generate more than one code-removal patch for only 9 out of 949 (0.95%). Considering the builds with one crashing test case, jKali created more than one code-removal patch for only 7 out of the 969 (0.72%) builds. We now analyze these cases in details.

Incorrect Patches Some patches are simply incorrect. For instance, the second patch (ID 2) generated for build 384713759 is incorrect since it prevents the program from throwing the `javax.persistence.PersistenceException`, by removing the instruction that stores the data in the database. Thus, the crashing test case is turned into a passing test case by removing a functionality.

Weak or Incorrect Test Cases Some additional code-removal patches are incorrect because of problems in the test suite. For example, considering build 400611810, jKali manages to create 6 code-removal patches, all incorrect. This suggests there is a problem in the tests. Indeed, all the patches, in different ways, avoid to set the header value. In this way, the request has the correct status code expected by the test case, but the header is empty, making the program incorrect, although it passes the test cases since no test checks the content of the header.

Another example is build 408694507 for which jKali creates 3 code-removal patches. These patches, in different ways, avoid the execution of a `return` statement introduced by the developer in a recent change. The developer did not update the test case to support the new behavior of the program, and for this reason a test failed. Thus, the 3 code-removal patches, although incorrect, indicate that there is a problem with the last change made by the developer.

Finally, for build 367766867, jKali creates 4 code-removal patches. All of them focus on the same piece of code that is related to an `if`-statement that contains a `for`-loop. This case is interesting because only the fourth patch (ID 4) deletes a single instruction of that fragment of code, while the others avoid the execution of the entire `if`-statement. A developer could focus only on the single instruction deleted by Patch 4, forgetting about the other patches, to understand how to fix the bug.

Rotting Test Some cases are determined by rotting tests. For instance, the two patches generated for build 351075282 are incorrect. The correct human patch modifies the tests by adding some missing test data. Both code-removal patches are related to the same part of the program with the difference that patch with ID 1 deletes only a subset of the piece of code removed by the patch with ID 2. This means that the developer could focus only on the logic of the code removed by Patch 1, since the extra piece of code deleted only by Patch 2 does not change the final result.

Flaky Test It happens that code-removal patches change the program in an opposite way. Indeed, in 3 out of 4 cases (builds 374587117, 407166687, and 415796275) jKali creates this type of patches. An example is reported in Listing 12 and Listing 13, that show two code-

²¹ Ochiai in Astor/jKali, as presented in Section 2.1

removal patches generated by jKali for build 407166687. In the first case, Patch 8, the `if` condition is changed with the keyword `true` forcing the execution of the `then`-branch, while in the second case, Patch 9, the `if` condition is changed with the keyword `false`, forcing the execution of the `else`-branch. When such a situation occurs, this should be interpreted by the developer as a clear sign that the result of a test case is independent of the change applied to the `if` condition. Indeed, the program passes the test cases both without the `then`-branch and with the `then`-branch. The failing test is likely flaky and should be revised.

Additional Details About Code-removal Patches Further details about the code-removal patches generated by jKali are contained in a public repository²².

How do developers fix the failed builds associated with a testsuite-adequate code-removal patch? (RQ4) In 33 out of 52 cases (63.46%), the builds for which code-removal patches exist are caused by problems in the test suite rather than problems in the program. In particular, there are 13 out of 52 builds (25%) for which the developers fixed the test code, and 5 out of 52 builds (9.62%) for which the developers fixed the data used by the test cases. This is a novel observation in the literature and important for the research field: it shows that the presence of code-removal patches is a good signal about problems in tests, further confirming our results from RQ3. Also, our experiment clearly shows that fault localization often does not point to the right location to change (8 out of 14 cases, 57.14%). These results are significant for the program repair research community: this is a need to research on using the presence of code-removal patches as test adequacy criterion, and there is also a need for more research on fault localization.

```

1  --- /src/main/java/org/apache/dubbo/rpc/protocol/dubbo/DubboProtocol.java
2  +++ /src/main/java/org/apache/dubbo/rpc/protocol/dubbo/DubboProtocol.java
3  @@ -183,183 @@
4  -         if (isStubSupportEvent && (!isCallbackService)) {
5  +         if (true) {

```

Listing 12: Code-removal patch (ID 8) for build 407166687.

```

1  --- /src/main/java/org/apache/dubbo/rpc/protocol/dubbo/DubboProtocol.java
2  +++ /src/main/java/org/apache/dubbo/rpc/protocol/dubbo/DubboProtocol.java
3  @@ -183,183 @@
4  -         if (isStubSupportEvent && (!isCallbackService)) {
5  +         if (false) {

```

Listing 13: Code-removal patch (ID 9) for build 407166687.

²² <https://github.com/repairator/open-science-repairator>

4.5 Discussion about the Results

Given the results associated with the RQs, here we discuss the main findings and describe possible future research in this field.

4.5.1 Success Rate of Code-removal Patches

Benchmark Overfitting The results reported in Sections 4.1 and 4.2 show that the success rate of code-removal patches is low and it is less than the one reported in previous studies (Long and Rinard 2016; Qi et al. 2015; Martinez et al. 2016). This applies both to builds with one failing test case (2.95% of success rate) and builds with one crashing test case (2.48% of success rate).

This result is important since it questions the existing knowledge. Our new experiments suggest that the benchmarks used for evaluation may influence the effectiveness of code-removal patches. This is along the line of Durieux et al., who showed that program repair techniques can overfit a benchmark (Durieux et al. 2019).

Type of Bugs Our data show that the generation of code-removal patches is independent of the category of assertion failure, but it is dependent on the type of exception that characterizes the crashing test case. For example, a `NullPointerException` is more likely to be fixed with a code-removal patch than an `IllegalStateException`. Analyzing the builds with one crashing test case, we also noticed that certain types of exceptions are more frequent than others, e.g., `NullPointerException` and `IllegalStateException`. This suggests that the success rate of code-removal patches depends on the considered failure.

Quality of the Test Suite Finally, the quality of test cases can influence the likelihood of creating code-removal patches. This is in line with the results reported by Qi et al., stating that the reason why a high number of code-removal patches has been generated is related to the weak test suites associated with the programs included in the benchmarks used in their experiment (Qi et al. 2015).

Future Research Future research should keep into consideration the characteristics of the benchmarks used for the evaluation in order to better explain the results of a certain repair technique. Additionally, program repair techniques should have strategies to guide the repair process based on the type of bugs. In this way, repair techniques can reduce the search space in terms of locations to consider and the ingredients to use to fix a bug. For example, for an `IllegalStateException`, the repair process should focus only on finding the correct location or conditions to call the method that makes the program throw the exception.

4.5.2 Code-removal Patches and Test Cases Problems

Based on the results reported in Section 4.3, we know that when a code-removal patch works, it is often due to problems in test cases. There are different types of problems: well-known ones, such as weak assertions (Qi et al. 2015), but also new problems never discussed in the context of code-removal patches. These cases include bugs in the test cases, tests that are executed only if some conditions are satisfied (rottening test cases), and flakiness. For only 1 out of 28 builds with one failing test case, we found a correct code-removal patch, and similarly for only 1 out of 24 builds with one crashing test case, the code-removal patch was correct.

This is in line with the results reported in previous works in which the higher number of code-removal patches reported relate to problems affecting the test suite (Qi et al. 2015; Long and Rinard 2016; Durieux et al. 2019). The studies by Long and Rinard (2016) and Qi et al. (2015) consider code-removal patches in C programs. Both studies use the ManyBugs benchmark (Le Goues et al. 2015), which is the one used to evaluate the original implementation of GenProg, but Long and Rinard restrict their investigation to 69 bugs, since they classified the rest of the bugs as not being actual detects. Their results show that most of the code-removal patches exist because of a weak test suite.

In the experiment conducted by Martinez et al. (2016), there are 22 code-removal patches, but only one is correct. They also report that the test cases influence the results about the generation and correctness of a patch. Indeed, if the test cases insufficiently cover the program functionality, program repair tools can generate patches that make the program pass all the test cases, but they can also introduce some bugs not revealed by the test suite. For example, considering the Math-32 bug of the Defects4J benchmark (Just et al. 2014), the code-removal patch shown in Listing 14 forces the execution of the `else`-branch, but since the test case does not correctly test the `then`-branch, the patch makes the program pass the test case, even though the patch is not correct.

```

1  - if ((Boolean) tree.getAttribute()) {
2  + if (false) {
3      setSize(Double.POSITIVE_INFINITY);
4      setBarycenter(Vector2D.NaN);
5  } else {
6      setSize(0);
7      setBarycenter(Vector2D(0 , 0));

```

Listing 14: Example of code-removal patch generated by jKali for the bug Math-32.

Future Research Future research can exploit code-removal patches to improve the quality of patches generated by program repair techniques. Indeed, code-removal patches could be used to detect problems in the test suite, and avoid the generation of patches that are test-suite-adequate, but actually incorrect. Moreover, it would be interesting to investigate the possibility to create fixes directly in the test case and not only in the source code of the program.

4.5.3 Code-removal Patches used as Debugging Hints

Putting in relation the type of change performed by the code-removal patches and the human patches, the results reported in Sections 4.3 and 4.4 show that for the 14 builds for which the human patch is available, in 4 cases the code-removal and human patches are partially related, in 8 cases they are disjoint.

Based on these data, code-removal patches are not good indicators to precisely localize the correct point where to apply the fix. However, they give hints to developers: they can be used to identify the methods in the program that need to be analyzed by the developers in order to understand the failure and localize the bug.

Future Research Future research can leverage code-removal patches as a first pass of coarse grain fault localization, or to refine the results of an existing fault localization technique. More generally, studies could focus on trying to improve fault localization of bugs by exploiting not only test case traces, but program patches, akin to mutation-based fault-localization (Papadakis and Le Traon 2015).

5 Threats to Validity

A threat to the validity of the results is about their generalization. Indeed, given the width of possible code-removal patches, the study considers faults revealed by either one failing or one crashing test case only, since it is a situation commonly encountered in practice. For instance, the Bears benchmark (Madeiral et al. 2019), which is a benchmark of 251 reproducible bugs from 72 different projects, has 71.32% of builds with a single failing (38.65% of the total) or crashing (32.67% of the total) test case. Thus, it is necessary to conduct further studies to understand if the results obtained are generalizable also to builds that have more than one failing or crashing test case. Another threat to validity is related to the execution time chosen for the repair process, that was set to 100 minutes. To address this threat, the setup of the experiment was based on previous research (Durieux et al. 2019). Since we never encountered a timeout case in our experiment, we consider this threat not likely to affect our results.

Another risk is about the correctness of the implementations used in the experiments. jKali, the Java implementation of Kali (Qi et al. 2015), was used to generate the code-removal patches. To mitigate this threat, both the tool and the results were made publicly available. Moreover, jKali was already used in other previous studies (Martinez and Monperrus 2016; Durieux et al. 2019; Ye et al. 2021).

Finally, a threat to validity is related to the manual analysis and classification of patches. To mitigate this threat, we relied on the information associated with the build using the Travis CI API and the changes classified as human patches were also applied in a local environment to make sure that the failing build was not able to pass the test cases. Moreover, we chose to be maximally conservative with the analysis of the build. When there were too many changes between the failed and the passed builds, we decided to label the human patch as “Not found”, mitigating the risk of making unsound claims.

6 Related Work

The studies most relevant to our work concern with the analysis of patches, with the quality of tests, with the effectiveness of repair strategies, and with the interplay between automated repair techniques and human activities. In this section, we discuss how they relate to our work.

6.1 Analysis of Patches

A primary concern about program repair techniques is the quality of the generated patches. Indeed, several studies show that plausible patches are often unsatisfactory for developers. For instance, the study by Qi et al. shows that the majority of the patches generated by tools like GenProg simply delete functionality, resulting in incorrect code changes (Qi et al. 2015). This

paper extensively studies code-removal patches, investigating their causes, and the information that can be extracted from the patches, even if incorrect.

The study by Martinez et al. (2016) based on the empirical comparison of three repair techniques (jGenProg Martinez and Monperrus (2016), jKali Martinez and Monperrus (2016), and Nopol Xuan et al. (2017)) points out the problem of overfitting patches, which are test-suite-adequate patches that fail to repair the target bugs. Similarly, the study by Ye et al. (2021) investigates the effectiveness of program repair tools on the QuixBugs (Lin et al. 2017) benchmark, confirming the challenge of overfitting patches as one of the main challenges for program repair techniques. In particular, the study shows that jKali is able to generate three patches for two different programs, and only one of them is correct. This result is aligned with our findings, where only 2 out of the 52 (3.85%) generated patches are correct.

One of the biggest experiment that considers 2,141 bugs belonging to 5 different benchmarks and 11 repair tools (Durieux et al. 2019) analyzes the effectiveness of program repair techniques identifying factors that prevent the generation of patches. In particular, the study shows that jKali is able to fix 52 out of 2,141 bugs (2.43%). This success rate is aligned with our experiment (2.71% of success rate) that considers 1,918 bugs, and it is different from the results reported in previous studies with small-scale benchmarks, where the success rate is 36.23% Long and Rinard (2016) over 69 bugs, 25.71% Qi et al. (2015) over 105 bugs, and 9.28% Martinez et al. (2016) over 224 bugs. We can conclude that diversity in terms of numbers and types of bugs involved in a study is indeed an important factor to be taken into consideration when evaluating program repair techniques.

The study by Wang et al. (2019) reports the analysis of the patches generated by automatic program repair techniques on the Defects4J benchmark in comparison to human fixes. Results show that 25.4% of the correct patches are syntactically different from the patches implemented by developers, especially when the developers patches are large. The study suggests that it is not mandatory for program repair techniques to fix bugs as developers do. In this sense, code-removal patches, although not perfectly matching human fixes, can still be useful.

6.2 Quality of Test Cases

Since test cases represent an important part in the process of generate-and-validate APR systems, researchers studied the relation between the quality of test suites and the quality of the generated patches. For instance, the study conducted by Yi et al. (2018) shows that improving test suite related metrics may increase the reliability of repair tools. In line with this finding, we also reported that the presence of a code-removal patch is a possible signal of problems in the tests. In fact, it is often possible to remove incorrect code-removal patches implementing additional test cases that cover untested behavior.

Other studies consider the relation between tests and specific types of patches. For instance, Dziuranski et al. (2020) reported that the quantity of tests can be reduced without any negative impact on the generation of 1-edit degree patches, since test suites usually include several positive tests unrelated to the points that have to be changed to fix the bug.

The study by Jiang et al. (2019) analyzes 50 real-world bugs from Defects4J benchmark in order to understand if the low performance of program repair tools relates to the presence of weak test suites. They discover that several defects can be fixed even when test suites do not cover every case, and they propose different strategies that can be encoded to improve the effectiveness of program repair tools. To reduce the problem of overfitting patches, Yang et al. (2017) propose to enhance the test suites with fuzz testing.

Finally, the study by Martinez et al. (2016) shows that test suites are often too weak to effectively support program repair techniques. Our analysis complements these findings with additional evidence of the issues that may affect test suites, and consequently program repair techniques. These issues include rotting tests, buggy test cases, and flaky tests.

6.3 Effectiveness of Repair Strategies

Program repair techniques can use different strategies to generate patches. In particular, there are studies that define repair actions based on the analysis of human patches. Martinez and Monperrus (2015) show that it is possible to mine repair actions from patches written by developers, and use probabilistic models to reason on the search space of the possible patches. A similar study shows that expressions in a statement, assignments, and variable declarations are more likely to be changed than other types of program elements to fix a bug Soto and Goues (2018). This aspect explains why a program repair technique that just drops functionality, like jKali, has less likelihood to create a correct patch. To increase the effectiveness of repair models, Zhong and Mei (2018) propose to put them in relation with the different categories of bugs based on the type of exception, and the results show that these models are better than the generic ones.

Liu et al (2020) focus on the efficiency of patch generation based on the number of patch candidates that are created before finding a valid one. They show that fault localization used by the repair techniques influences the result, and when it is wrong, it increases the chances of producing overfitting patches, as in our case. Indeed, in 8 out of 14 cases (57.14%) where the human patch is applied to the source code, the corresponding code-removal patch is in a completely different location.

The study by Motwani et al. (2018) considers the applicability of repair techniques and the characteristics of the defects that the techniques can repair. The study shows that program repair techniques have less likelihood to produce a patch for bugs for which developers need to apply many changes or that have many different failing test cases. For this reason, we focused our study on code-removal patches on cases with a single failing or crashing test case.

6.4 Code-removal in Other Patch Generation Techniques

Some program repair techniques although not exclusively focusing on code-removal patches, as Kali does, can be able to generate them. We now discuss the presence of code-removal patches in other program repair techniques.

Ye et al. (2021) collected patches for the bugs in Defects4J benchmark generated by 14 different repair techniques. In this dataset there is only one bug, Math-50, for which different repair tools (e.g., jGenProg Martinez and Monperrus (2016) and ELIXIR Saha et al. (2017)) generate a correct code-removal patch, and there are 12 bugs for which a plausible code-removal patch has been generated.

In a recent study on another benchmark, QuixBugs Ye et al. (2021), the authors analyzed the effectiveness of the repair strategies implemented in different tools. According to the reported data, jGenProg generates 8 code-removal patches out of 164 patches (4.88%), while jKali creates 3 code-removal patches. Note that 6 out of 8 code-removal patches generated by jGenProg are for the same bug.

There are also other techniques that leverage templates extracted from human-written fixes in order to automatically create new patches. Liu et al. analyzed the recurrently-used fix

patterns in automatic program repair tools, and they reported that 4 out of the 35 identified patterns implement delete actions (Liu et al. 2019). In particular, among them, only 2 out of 4 can be exploited to create code-removal patches as jKali does, because the other two patterns remove arguments from method invocations if the method has overloaded methods and remove some of the expressions in conditional statements. They also implemented the repair patterns in a tool called TBar, and they tested them on the Defects4J benchmark. Based on the data reported in the repository associated with the experiment, 20 bugs have been fixed with a code-removal patch.

SOFix is another template-based technique in which the repair templates have been extracted from the answers written by users in the Stack Overflow forum²³. In particular, they mined 13 repair templates, and among them 2 are for generating code-removal patches. In their evaluation on the Defects4J benchmark, SOFix fixes 2 bugs with a code-removal patch.

As TBar and SOFix, also jKali has been evaluated on Defects4J benchmark, and it fixed 27 bugs (Durieux et al. 2019).

Thus, based on these data, although different techniques may generate code-removal patches, they actually implement additions and code replacements most of the time. Only TBar, which has specific patterns for deleting code, is able to create a number of code-removal patches similar to jKali. On the other hand, jGenProg and its variants/successors generate code-removal patches only sometimes.

6.5 Software Debloating

Deleting code is a pattern that can be used to create code-removal patches, but it is a practice that can be exploited also in other fields, like software debloating. Software debloating aims to improve the security and performance of software by deleting library code and features that are not necessary for the end user (Brown and Pande 2019; Ramanathan et al. 2020). Thus, code-removal patches and software debloating are related.

For example, Piranha is a code refactoring tool designed to automatically create differential revisions to identify stale feature flags (Ramanathan et al. 2020). Although jKali and Piranha work in different field, they both try to improve the software by deleting code. A code-removal patch can prevent the program to execute a specific path to make the program pass the test case, and a Piranha deletion identifies stale flags that make the program more complex to maintain. The experimental results associated with Piranha show that it was possible to reduce 17.3% of the flags of the examined projects and reduce the codebase size by 1%, indicating the usefulness of such a tool.

The problem of software bloat can affect also the external dependencies of a program, and not only the source code (Soto-Valero et al. 2021b, a). For example, Soto-Valero et al. conducted a study to investigate bloated dependencies, that are libraries packaged with the program's compiled source code, but they are actually not needed to build and run the programs (Soto-Valero et al. 2021b). After developing a tool called DepClean, they analyzed 9,639 Java artifacts hosted on Maven Central, and they notice that 2.7% of the dependencies directly declared are bloated, 15.4% of the inherited dependencies are bloated, and finally 57% of the transitive dependencies of the studied artifacts are bloated. Thus, deleting code is not useful only to fix and improve the codebase of a program as jKali and Piranha do respectively,

²³ <https://stackoverflow.com>

but it can be used also to better manage the external dependencies, removing the ones that are not necessary.

Humans tend to add features instead of removing them Meyvis and Yoon (2021), and we believe developers are the same when it comes to software engineering. Yet, deleting code is clearly an option to manage code ranging from program repair to software debloating.

6.6 Empirical Studies Investigating how Humans Use Automatic Patches

Finally, there are studies to understand how humans work with the patches generated automatically by program repair techniques. Indeed, program repair techniques may produce patches that are not accepted by developers (Tan et al. 2016).

The early study by Fry et al. (2012) measures the maintainability of patches, and it shows that machine-generated patches are slightly less maintainable than the human-written ones. Ryan et al. (2019) show that developers in general tend to trust more in patches produced by humans than automatic program repair techniques. Similar results are the ones obtained by Alarcon et al. (2020), indeed, they report that programmers find human repairs more trustworthy than the ones generated by GenProg. The study by Cambronero et al. (2019) investigates how humans are used to use the patches automatically generated by program repair techniques. The study shows that providing only patches generated by program repair techniques is not enough to ensure that developers integrate them in the code base. On the other hand, Monperrus et al. (2019) developed a bot called Repairator, and they demonstrated that it is possible to produce patches that are accepted by developers.

Liu et al. proposed an approach based on convolutional neural networks to automatically identify patterns related to security vulnerabilities or bad programming practices (Liu et al. 2021). The approach is based on mining common code patterns for each type of violation, and extract common fix patterns for each type of fixed violation. The study shows that humans trust the automatically generated fixes because they accepted 69 out 116 fixes, showing that fixes created following human strategies have a good success rate to be integrated in the codebase by humans.

Kim et al. proposed PAR, a program repair tool that generates patches following patterns learned from existing patches implemented by developers, and they conducted a user study to investigate if humans accept or not the patches generated by PAR (Kim et al. 2013). The user study involved 72 students and 96 developers who were proposed patches generated by PAR and GenProg (Le Goues et al. 2012). The results show that humans accepted more patches generated by PAR than GenProg, confirming also in this case that humans tend to trust patches that follow strategies already used by developers in the past more than other types of patches, that follow strategies different from the human reasoning, like genetic programming.

Finally, a study conducted by Tan et al. (2016) shows that enforcing a program repair tool such as GenProg (Le Goues et al. 2012) to use anti-patterns allows to have less deletion of program functionalities. Moreover, the generated patches have better quality than the ones of a customized version of GenProg that prohibits deletions of code. Since the patches produced in this way have better quality, developers might more likely accept them.

As reported in Section 4.3, only two code-removal patches are classified as correct in our experiment, thus it is very unlikely that developers accept them. However, code-removal patches, although incorrect, can be proposed to developers, that can exploit them to discover some problems in the test suite or to simplify the debug phase.

7 Conclusion

In this paper, we focused on code-removal patches, and we used jKali to generate them. In particular, we analyzed 1,918 failed builds with only one failing test case (949) or only one crashing test case (969) to determine the information that can be extracted from the code-removal patches and their relation with human fixes. Moreover, we proposed a comprehensive taxonomy of code-removal patches that can be exploited to better understand the current limitations of program repair techniques.

Our results show that code-removal patches are often insufficient to fix bugs, contrarily to previous studies (Long and Rinard 2016; Qi et al. 2015; Martinez et al. 2016) where the effectiveness of code-removal patches is higher. Moreover, while other approaches generically explain the presence of code-removal (or plausible) patches with the presence of a weak test suite, our study provides detailed evidence about issues that may affect test suites, such as rotting tests, buggy test cases, and flaky tests. The relation between code-removal patches and human fixes provides additional insights about the meaning of code-removal patches. Finally, our study provides evidence that code-removal patches could be exploited to automatically improve test suites, opening new opportunities for the studies in the field of program repair.

Funding Open access funding provided by Università degli Studi di Milano - Bicocca within the CRUI-CARE Agreement. This work was supported in part by the ERC Consolidator Grant 2014 Program through the EU H2020 Learn Project under ERC Grant 646867. This work was partially supported by the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation, and by the Swedish Foundation for Strategic Research (SSF).

Declarations The computations were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC).

Conflicts of interest The authors have no conflicts of interests to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alarcon G, Walter C, Gibson A, Gamble R, Capiola A, Jessup S, Ryan T (2020) Would you fix this code for me? effects of repair source and commenting on trust in code repair. In: Systems vol 8, p 8. <https://doi.org/10.3390/systems8010008>
- Britton T, Jeng L, Carver G, Cheak P, Katzenellenbogen T (2013) Reversible debugging software: quantify the time and cost saved using reversible debuggers. Tech. rep., <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.444.9094&rep=rep1&type=pdf>
- Brown MD, Pande S (2019) Carve: Practical security-focused software debloating using simple feature set mappings. In: Proceedings of the 3rd ACM workshop on forming an ecosystem around software

- transformation, FEAST'19. London, United Kingdom: Association for Computing Machinery, pp 1–7. <https://doi.org/10.1145/3338502.3359764>
- Cambronero JP, Shen J, Cito J, Glassman E, Rinard M (2019) Characterizing developer use of automatically generated patches. In: Smith J, Bogart C, Good J, Fleming SD (eds) 2019 IEEE symposium on visual languages and human-centric computing, VL/HCC 2019, Memphis, Tennessee, USA, October 14–18 2019, IEEE Computer Society, pp 181–185. <https://doi.org/10.1109/VLHCC.2019.8818884>,
- Campos J, Ribeiro A, Perez A, Abreu R (2012) Gzoltar: an eclipse plug-in for testing and debugging. In: 2012 Proceedings of the 27th IEEE/ACM international conference on automated software engineering, pp 378–381. <https://doi.org/10.1145/2351676.2351752>
- Colanzi TE, McMinn P (2018) Search-based software engineering - 10th international symposium, SSBSE 2018, Montpellier, France, September 8–9, 2018, Proceedings, Lecture Notes in Computer Science, vol 11036. Springer. <https://doi.org/10.1007/978-3-319-99241-9>
- Coralogix (2015) This is what your developers are doing 75% of the time, and this is the cost you pay, Tech. rep. <https://coralogix.com/log-analytics-blog/this-is-what-your-developers-are-doing-75-of-the-time-and-this-is-the-cost-you-pay/>
- Daniel B, Jagannath V, Dig D, Marinov D (2009) ReAssert: Suggesting repairs for broken unit tests. In: Proceedings of the 24th IEEE/ACM international conference on automated software engineering
- Delplanque J, Ducasse S, Polito G, Black AP, Etien A (2019) Rotten green tests. In: 2019 IEEE/ACM 41st International conference on software engineering (ICSE), pp 500–511
- Durieux T, Madeiral F, Martinez M, Abreu R (2019) Empirical review of java program repair tools: A large-scale experiment on 2,141 bugs and 23,551 repair attempts. In: Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering, ESEC/FSE 2019. Tallinn, Estonia: Association for Computing Machinery, pp 302–313. <https://doi.org/10.1145/3338906.3338911>
- Durieux T, Goues CL, Hilton M, Abreu R (2020) Empirical study of restarted and flaky builds on travis ci. [arxiv: 2003.11772](https://arxiv.org/abs/2003.11772) [cs.SE]
- Dziuranski P, Gerasimou S, Kolovos D, Matragkas N (2020) Empirical analysis of 1-edit degree patches in syntax-based automatic program repair. In: 2020 IEEE Congress on evolutionary computation (CEC), pp 1–10
- Fry ZP, Landau B, Weimer W (2012) A human study of patch maintainability. In: In international symposium on software testing and analysis, pp 177–187
- Gazzola L, Micucci D, Mariani L (2019) Automatic software repair: A survey. *IEEE Transactions on Software Engineering* 45(1):34–67
- Jiang J, Xiong Y, Xia X (2019) A manual inspection of defects4j bugs and its implications for automatic program repair. In: Science china information sciences, vol 62. <https://doi.org/10.1007/s11432-018-1465-6>
- Just R, Jalali D, Ernst MD (2014) Defects4j: A database of existing faults to enable controlled testing studies for java programs. In: Proceedings of the 2014 international symposium on software testing and analysis, ISSTA 2014. San Jose, CA, USA: Association for Computing Machinery, pp 437–440. <https://doi.org/10.1145/2610384.2628055>,
- Kim D, Nam J, Song J, Kim S (2013) Automatic patch generation learned from human-written patches. In: Proceedings of the 2013 international conference on software engineering, ICSE '13. San Francisco, CA, USA: IEEE Press, pp 802–81
- Krasner H (2018) The cost of poor quality software in the US: A 2018 report. Tech. rep., <https://www.it-cisq.org/the-cost-of-poor-quality-software-in-the-us-a-2018-report/The-Cost-of-Poor-Quality-Software-in-the-US-2018-Report.pdf>
- Langr J, Hunt A, Thomas D (2015) Pragmatic unit testing in java 8 with JUnit, 1st edn. Pragmatic Bookshelf
- Le Goues C, Nguyen T, Forrest S, Weimer W (2012) Genprog: A generic method for automatic software repair. *IEEE Transactions on Software Engineering* 38(1):54–72
- Le Goues C, Holtschulte N, Smith EK, Brun Y, Devanbu P, Forrest S, Weimer W (2015) The manybugs and introclass benchmarks for automated repair of c programs. *IEEE Transactions on Software Engineering* 41(12):1236–1256
- Lin D, Koppel J, Chen A, Solar-Lezama A (2017) Quixbugs: A multi-lingual program repair benchmark set based on the quixey challenge. In: Proceedings companion of the 2017 ACM SIGPLAN international conference on systems, programming, languages, and applications: software for humanity. SPLASH Companion 2017, Vancouver, BC, Canada: Association for Computing Machinery, pp 55–56. <https://doi.org/10.1145/3135932.3135941>,
- Liu K, Koyuncu A, Bissyandé TF, Kim D, Klein J, Le Traon Y (2019) You cannot fix what you cannot find! an investigation of fault localization bias in benchmarking automated program repair systems. In: 2019 12th IEEE conference on software testing, validation and verification (ICST), pp 102–113. <https://doi.org/10.1109/ICST.2019.00020>

- Liu K, Koyuncu A, Kim D, Bissyandé TF (2019) Tbar: Revisiting template-based automated program repair. In: Proceedings of the 28th ACM SIGSOFT international symposium on software testing and analysis. ISSTA 2019. Beijing, China: Association for Computing Machinery, pp 31–42. <https://doi.org/10.1145/3293882.3330577>
- Liu K, Wang S, Koyuncu A, Kim K, Wu P, Klein J, Mao X, Le Traon Y, Bissyandé T, Kim D (2020) On the efficiency of test suite based program repair: A systematic assessment of 16 automated repair systems for java programs. <https://doi.org/10.1145/3377811.3380338>
- Liu K, Kim D, Bissyandé TF, Yoo S, Le Traon Y (2021) Mining fix patterns for findbugs violations. IEEE Transactions on Software Engineering 47(1):165–188. <https://doi.org/10.1109/TSE.2018.2884955>
- Long F, Rinard M (2016) Automatic patch generation by learning correct code. SIGPLAN Not 51(1):298–312. <https://doi.org/10.1145/2914770.2837617>
- Luo Q, Hariri F, Eloussi L, Marinov D (2014) An empirical analysis of flaky tests. In: Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering. FSE 2014. Hong Kong, China: Association for Computing Machinery, pp 643–653. <https://doi.org/10.1145/2635868.2635920>
- Madeiral F, Urli S, Maia M, Monperrus M (2019) Bears: An extensible java bug benchmark for automatic program repair studies. In: Proceedings of the 26th IEEE international conference on software analysis, evolution and reengineering (SANER '19). <https://arxiv.org/abs/1901.06024>
- Martinez M, Monperrus M (2015) Mining software repair models for reasoning on the search space of automated program fixing. Empirical Softw Engg 20(1):176–205. <https://doi.org/10.1007/s10664-013-9282-8>
- Martinez M, Monperrus M (2016) Astor: A program repair library for java. In: Proceedings of ISSTA, <https://doi.org/10.1145/2931037.2948705>
- Martinez M, Durieux T, Sommerard R, Xuan J, Monperrus M (2016) Automatic Repair of Real Bugs in Java: A Large-Scale Experiment on the Defects4J Dataset. In: Springer empirical software engineering <https://doi.org/10.1007/s10664-016-9470-4>, <https://hal.archives-ouvertes.fr/hal-01387556/document>
- Martinez M, Etien A, Ducasse S, Fuhrman C (2020) Rtlj: A java framework for detecting and refactoring rotten green test cases. In: Proceedings of the ACM/IEEE 42nd international conference on software engineering: companion proceedings. ICSE '20. Seoul, South Korea: Association for Computing Machinery, pp 69–72. <https://doi.org/10.1145/3377812.3382151>
- Meyvis T, Yoon H (2021) Adding is favoured over subtracting in problem solving
- Monperrus M (2018) Automatic software repair: a bibliography. In: ACM Comput. Surv. 51.1. <https://doi.org/10.1145/3105906>
- Monperrus M, Urli S, Durieux T, Martinez M, Baudry B, Seinturier L (2019) Repairator patches programs automatically. In: Ubiquity (July 2019). <https://doi.org/10.1145/3349589>
- Motwani M, Sankaranarayanan S, Just R, Brun Y (2018) [Journal First] do automated program repair techniques repair hard and important bugs? In: 2018 IEEE/ACM 40th International conference on software engineering (ICSE), pp 25–25
- Papadakis M, Le Traon Y (2015) Metallaxis-fl: mutation-based fault localization. Software Testing, Verification and Reliability 25(5–7):605–628
- Qi Z, Long F, Achour S, Rinard M (2015) An analysis of patch plausibility and correctness for generate-and-validate patch generation systems. In: Proceedings of the 2015 international symposium on software testing and analysis. ISSTA 2015. Baltimore, MD, USA: Association for Computing Machinery, pp 24–36. <https://doi.org/10.1145/2771783.2771791>
- Ramanathan MK, Clapp L, Barik R, Sridharan M (2020) Piranha: Reducing feature flag debt at uber. In: Proceedings of the ACM/IEEE 42nd international conference on software engineering: software engineering in practice. ICSE-SEIP '20. Seoul, South Korea: Association for Computing Machinery, pp 221–230. <https://doi.org/10.1145/3377813.3381350>
- Ryan TJ, Alarcon GM, Walter C, Gamble RF, Jessup SA, Capiola AA, Pfahler MD (2019) Trust in automated software repair - the effects of repair source, transparency, and programmer experience on perceived trustworthiness and trust. In: Moallem A (ed) HCI for Cybersecurity, Privacy and Trust - First International Conference, HCI-CPT 2019, Held as Part of the 21st HCI International Conference, HCII 2019, Orlando, FL, USA, July 26-31, 2019, Proceedings, Springer, Lecture Notes in Computer Science, vol 11594, pp 452–470. https://doi.org/10.1007/978-3-030-22351-9_31
- Saha RK, Lyu Y, Yoshida H, Prasad MR (2017) Elixir: Effective object-oriented program repair. In: 2017 32nd IEEE/ACM international conference on automated software engineering (ASE), pp 648–659. <https://doi.org/10.1109/ASE.2017.8115675>
- Undo Software (2014) Increasing software development productivity with reversible debugging. white paper, https://undo.io/media/uploads/files/Undo_ReversibleDebugging_Whitepaper.pdf
- Soto M, Goues CL (2018) Common statement kind changes to inform automatic program repair. In: Proceedings of the 15th international conference on mining software repositories, MSR '18. Gothenburg, Sweden: Association for Computing Machinery, pp 102–105. <https://doi.org/10.1145/3196398.3196472>

- Soto-Valero C, Durieux T, Baudry B (2021a) A longitudinal analysis of bloated java dependencies. [arxiv:2105.14226](https://arxiv.org/abs/2105.14226)
- Soto-Valero C, Harrand N, Monperrus M, Baudry B (2021) A comprehensive study of bloated dependencies in the maven ecosystem. *Empir Softw Eng* 26(3):45. <https://doi.org/10.1007/s10664-020-09914-8>
- Tan SH, Roychoudhury A (2015) Relifix: Automated repair of software regressions. In: 37th International conference on software engineering, IEEE
- Tan SH, Yoshida H, Prasad MR, Roychoudhury A (2016) Anti-patterns in search-based program repair. In: Proceedings of the 2016 24th ACM SIGSOFT international symposium on foundations of software engineering. FSE 2016. Seattle, WA, USA: Association for Computing Machinery, pp 727–738. <https://doi.org/10.1145/2950290.2950295>,
- Urli S, Yu Z, Seinturier L, Monperrus M (2018) How to design a program repair bot? insights from the repairator project. In: Proceedings of the 40th international conference on software engineering
- Wang S, Wen M, Chen L, Yi X, Mao X (2019) How different is it between machine-generated and developer-provided patches?: An empirical study on the correct patches generated by automated program repair techniques. In: 2019 ACM/IEEE International symposium on empirical software engineering and measurement (ESEM). <https://doi.org/10.1109/esem.2019.8870172>
- Weimer W, Fry ZP, Forrest S (2013) Leveraging program equivalence for adaptive program repair: Models and first results. In: Proceedings of the 28th IEEE/ACM international conference on automated software engineering (ASE), IEEE Press
- Xuan J, Martinez M, DeMarco F, Clement M, Marcote SL, Durieux T, Le Berre D, Monperrus M (2017) Nopol: Automatic repair of conditional statement bugs in java programs. *IEEE Trans Softw Eng* 43(1):34–55. <https://doi.org/10.1109/TSE.2016.2560811>
- Yang J, Zhikhartsev A, Liu Y, Tan L (2017) Better test cases for better automated program repair. In: Proceedings of the 2017 11th joint meeting on foundations of software engineering. ESEC/ FSE 2017. Paderborn, Germany: Association for Computing Machinery, pp 831–841. <https://doi.org/10.1145/3106237.3106274>
- Ye H, Martinez M, Durieux T, Monperrus M (2021) A comprehensive study of automatic program repair on the quixbugs benchmark. *Journal of Systems and Software* 171:110825. <https://doi.org/10.1016/j.jss.2020.110825>
- Ye H, Martinez M, Monperrus M (2021) Automated patch assessment for program repair at scale. *Empir Softw Eng* 26(2):20. <https://doi.org/10.1007/s10664-020-09920-w>
- Yi J, Tan SH, Mechtaev S, Böhme M, Roychoudhury A (2018) A correlation study between automated program repair and test-suite metrics. In: Proceedings of the 40th international conference on software engineering. ICSE '18. Gothenburg, Sweden: Association for Computing Machinery, p 24. <https://doi.org/10.1145/3180155.3182517>
- Yu Z, Martinez M, Danglot B, Durieux T, Monperrus M (2017) Test case generation for program repair: A study of feasibility and effectiveness. [arXiv:1703.00198](https://arxiv.org/abs/1703.00198)
- Yuan Y, Banzhaf W (2020) Arja: Automated repair of java programs via multi-objective genetic programming. *IEEE Transactions on Software Engineering* 46(10):1040–1067. <https://doi.org/10.1109/TSE.2018.2874648>
- Zhong H, Mei H (2018) Mining repair model for exception-related bug. *Journal of Systems and Software* 141: 16–31