Augmenting Diffs With Runtime Information

Khashayar Etemadi, Aman Sharma, Fernanda Madeiral, Martin Monperrus

Abstract—Source code diffs are used on a daily basis as part of code review, inspection, and auditing. To facilitate understanding, they are typically accompanied by explanations that describe the essence of what is changed in the program. As manually crafting high-quality explanations is a cumbersome task, researchers have proposed automatic techniques to generate code diff explanations. Existing explanation generation methods solely focus on static analysis, i.e., they do not take advantage of runtime information to explain code changes. In this paper, we propose COLLECTOR-SAHAB, a novel tool that augments code diffs with runtime difference information. COLLECTOR-SAHAB compares the program states of the original (old) and patched (new) versions of a program to find unique variable values. Then, COLLECTOR-SAHAB adds this novel runtime information to the source code diff as shown, for instance, in code reviewing systems. As an evaluation, we run COLLECTOR-SAHAB on 584 code diffs for Defects4J bugs and find it successfully augments the code diff for 95% (555/584) of them. We also perform a user study and ask eight participants to score the augmented code diffs generated by COLLECTOR-SAHAB. Per this user study, we conclude that developers find the idea of adding runtime data to code diffs promising and useful. Overall, our experiments show the effectiveness and usefulness of COLLECTOR-SAHAB in augmenting code diffs with runtime difference information.

Publicly-available repository: https://github.com/ASSERT-KTH/collector-sahab.

Index Terms—Code diff, dynamic program analysis, runtime differencing, code review.

---

1 INTRODUCTION

A software program evolves based on a series of changes to its source code. Developers are the gatekeepers of this evolution. Typically, they read, analyze, and ensure the quality of the code difference (hereafter, code diff) between two versions of a program, a best practice known as code review [1]. Listing 1 shows an example code diff that fixes a bug in Java, with the typical highlighting provided by IDEs and code review systems.

To facilitate the code reviewing process, various forms of explanations may be added to code diffs, such as commit messages [2], code comments, and pull request descriptions [3]. As manually crafting a high-quality explanation is time-consuming and may be neglected by developers [4], [5], [6], researchers have proposed techniques to automatically generate code diff explanations. Currently, explanation generation mostly focuses on static information, that is extracted from the code change without running the program [2], [4], [7]. However, runtime information is a great source of data for explaining code diffs [8]. For example, to properly understand the change in Listing 1, developers may need to know the impact of the condition change on the runtime value of variable `numericalVariance`.

In this paper, we present a study on using runtime information for explaining code diffs. This is a challenging task because there is an overwhelming number of events and values happening at runtime. This endeavor requires three components: first, a proper algorithm to monitor and select

```
1  public double getNumericalVariance() {
2  -    if (!numericalVarianceIsCalculated) {
3  +    if (!(sampleSize <= numberOfSuccesses)) {
4        numericalVariance =  calculateNumericalVariance();
5        numericalVarianceIsCalculated = true;
6      }
7      return numericalVariance;
8  }
```

Listing 1: A typical diff for a bug-fixing code diff, with state-of-the-art highlighting as found e.g. on Github.

data during the execution of both the original and patched versions of a program; second, an efficient algorithm to extract runtime differences between the collected traces; third, the extracted runtime differences should be integrated into code diffs in a useful manner for code reviewers.

In this paper, we propose COLLECTOR-SAHAB, a novel system to augment code diffs with runtime information. Given a code diff and a test case written in Java, COLLECTOR-SAHAB runs the test on both versions and uses Java bytecode instrumentation to collect *program states* during the execution, where program states consist of values assigned to specific, relevant variables. After collecting program states, COLLECTOR-SAHAB employs an original algorithm to find the variable values that are unique, in the sense that they only occur in one execution of the program under test. The first relevant unique variable value is added to the code diff to obtain an *augmented diff*, that is a diff with both static and dynamic differences. COLLECTOR-SAHAB is the first tool that augments code diffs with relevant runtime difference information.

To assess the effectiveness of COLLECTOR-SAHAB, we run it on 584 code diffs from the bug benchmark Defects4J [9]. Our results show that COLLECTOR-SAHAB detects a unique program state for 95% (555/584) of the code diffs in the benchmark. This outperforms the most

- *K. Etemadi, A. Sharma, and M. Monperrus are with the KTH Royal Institute of Technology, Stockholm, Sweden*
  Email: {khaes, amansha, monperrus}@kth.se

- *F. Madeiral is with the Vrije Universiteit Amsterdam, Amsterdam, Netherlands*
  Email: fer.madeiral@gmail.com

related work, DIDIFFFF [10]. DIDIFFFF is a web application to view the runtime value differences of a Java program. To assess COLLECTOR-SAHAB's applicability, we perform a field study on 50 code diffs from a diverse set of 10 real-world open-source projects. COLLECTOR-SAHAB successfully analyzes 96% (48/50) of these code diffs. We conduct a manual analysis on 30 of those real-world code diffs for which COLLECTOR-SAHAB augments the diff with runtime differences. This manual analysis confirms the usefulness of COLLECTOR-SAHAB in terms of correctness, understandability, and causal relation with the intended runtime behavioral change. Finally, we perform a user study and ask eight participants to rate them in terms of their usefulness, clarity, and novelty. The results confirm that developers find augmenting code diffs with runtime data a useful and promising practice.

To sum up, COLLECTOR-SAHAB is a novel tool that goes beyond existing explanation tools that only focus on static data [7], by integrating runtime differences into code diffs. COLLECTOR-SAHAB is the first approach that can provide developers with a concise set of runtime differences, usable for code review.

To summarize, we make the following contributions:

- We introduce COLLECTOR-SAHAB, a new system for augmenting code diffs with runtime information based on advanced execution monitoring, effective runtime data differencing, and well-integrated rendering for developers. COLLECTOR-SAHAB is publicly available at https://github.com/ASSERT-KTH/collector-sahab.
- We report original results of an evaluation of COLLECTOR-SAHAB's effectiveness in capturing runtime differences for 584 code diffs in Java. The results show that COLLECTOR-SAHAB reports a runtime difference for 95% (555/584) of the code diffs. COLLECTOR-SAHAB outperforms DIDIFFFF, the state-of-the-art tool in this domain, which only reports differences for 64.2% (377/584) of the cases.
- We perform a manual analysis on runtime differences detected by COLLECTOR-SAHAB for 30 code diffs from real-world, complex projects. This analysis confirms that COLLECTOR-SAHAB augmented diffs are correct, understandable, and successfully exclude random runtime differences.
- We report the results of an original user study on code diff augmentation with runtime data. The results of this study and our interview with the participants show that developers find COLLECTOR-SAHAB helpful.

The rest of the paper is organized as follows. In Section 2, we review the foundational concepts used in this paper. In Section 3, we explain how COLLECTOR-SAHAB works. Section 4 and Section 5 present the protocol that we use for our experiments and their results. Section 6 discusses the threats to the validity of our results. In Section 7, we review the related work. Finally, in Section 8 we conclude the paper.

2 CORE CONCEPTS

In this section, we review three foundational concepts of our paper: *program state*, *state depth*, and *breakpoint*.

The first core concept that is used in this work is program state. This concept is approached in different forms in previous work. Zeller defines program state as the set of program variables and their values [11], which is used by other researchers as well [12]. There are also more broad definitions for program states. For example, Xu et al. [13] also consider the call stack and program counter to be a part of the program state. Meinicke et al. [14] take heap objects as a part of the program state as well.

In the literature, there also exist program state definitions that are more narrow and tailored to the task under consideration. For example, the set of variables considered in the program state can be limited to variables used at a particular location [15], [16], can be fixed [17], can be in-scope variables at a particular execution point [18], can be variables assigned at a certain point in program [19], or finally variables that are read or written before a certain location [20]. Considering those various definitions from the literature, we attempt to generically define "program state" as follows.

Definition 1 (Program State – Generic). A program state s is the set of all variables V_l and their values accessed at a certain line l of a program p that is being executed.

There are two points to be noted regarding our definition of program state. First, since a line of a program may be executed multiple times during one single execution, it is possible to collect multiple program states for a given program line. Secondly, the read or written variables can refer to either primitive or non-primitive datatypes. This means the program state may include objects. The boundary of what constitutes an object at runtime is ambiguous. To fully capture the content of an object, we have to go inside the object and read the content of its fields, recursively. This leads us to the definition of state depth, as follows.

Definition 2 (State Depth). A program state s with depth d collected at line l consists of all values for primitive variables inside objects accessed at l that can be reached with at most d steps.[1]

For example, consider Listing 2, which shows a class "Student" and an object of that class that the variable "student1" refers to. The Student class contains a student's name, her supervisor's name, and her supervisor's education information. The object "student1" is serialized in JSON-like format at the bottom of the listing. If the state depth is set to 3, we collect all the information inside the object presented in Listing 2. If depth is set to 2, we do not collect the data inside "student1.supervisor.education". If depth is set to 1, we do not collect the data inside "student1.supervisor". Finally, if depth is set to 0, we do not collect any data inside "student1". In fact, when depth is set to 0, the program state only includes the value of primitive variables and excludes non-primitive objects such as "student1".

To collect the program states during execution, we have to specify the lines at which the states should be collected. This corresponds to the well-known debugging concept of "breakpoint". In debugging, breakpoints are used to suspend the program at a certain point of execution and explore the program state [21].

---

1. In our prototype implementation for Java, we consider the nine following datatypes as primitive: int, byte, short, long, float, double, boolean, char, and string.

```
1  Datatypes in Java-like syntax:
2  class Student {
3      String name;
4      Supervisor supervisor;
5  }
6  class Supervisor {
7      String name;
8      Education education;
9  }
10 class Education {
11     String university;
12     String city;
13 }
14
15 An object of the student datatype, serialized in JSON-like
       syntax:
16 {
17    "student1": {                              depth ≥ 1
18      "name": "Alice",                         depth ≥ 1
19      "supervisor": {                          depth ≥ 2
20        "name": "Bob",                         depth ≥ 2
21        "education": {                         depth ≥ 3
22          "university": "KTH Institute",       depth ≥ 3
23          "city": "Stockholm"                  depth ≥ 3
24        }
25      }
26    }
27 }
```

Listing 2: A program state is the runtime information associated to some datatype. In this example, the "Student" datatype is shown together with an object of this class "student1" with its content. The student object contains more or less information, depending on the observation depth, in color.

Definition 3 (Breakpoint). A breakpoint is a certain line l of a program p to be watched during an execution of p.

In this paper, we use breakpoints to collect program states, without human intervention. As opposed to manual debugging, no developer is involved. When a breakpoint is reached, the program execution is resumed right after the automated program state collection.

3 DESIGN AND IMPLEMENTATION

3.1 Overview

We design and implement COLLECTOR-SAHAB to generate a concise and useful report of runtime differences caused by a code change. This report puts the extracted execution data into a user interface (UI) to present code diffs. We call this report the *augmented diff* as it augments the code diff with dynamic information. Figure 1 represents an overview of how COLLECTOR-SAHAB works.

COLLECTOR-SAHAB gets an original and patched version of a program as well as a test case to execute them Section 3.2. In the first step, COLLECTOR-SAHAB executes the test on the original and patched versions of the program and compares the program states that occur during these two executions. To collect the program states, COLLECTOR-SAHAB first identifies breakpoints where the data should be collected (see Section 3.3) and collects the execution traces at those lines (see Section 3.4).

Then, COLLECTOR-SAHAB selects and reports a set of relevant values that occur in the program states of the original version but are absent in the patched version, and vice versa (see Section 3.5). This set of values represents the relevant runtime differences between two versions of a program. We focus on the relevant values, to make the data presented to code reviewers more precise [22].

Identifying values that occur only in program states of one of the versions can require excessive resources, as the number and size of collected program states can be very large. To address this issue, COLLECTOR-SAHAB adopts an effective algorithm to extract the fine-grained data in all program states and compare them against each other in massive numbers, see Section 3.5.

Finally, we integrate the extracted relevant execution differences into the code diffs typically used in code review (see Section 3.6). The result is an augmented diff, it enables developers to see runtime differences, together with the code change.

3.2 Input

Our tool takes in 3 inputs: the original version of the project, the patched version of the project, and a list of **covering tests**. A covering test is a test that runs at least part of the code that is changed. The challenge is to collect enough data and filter the interesting bits when running a covering test on the original and patched versions.

3.3 Breakpoint Identification

Given the original and patched versions, COLLECTOR-SAHAB first statically analyzes the code to identify the lines of the program that should be monitored as breakpoints. Next, a dynamic analysis is performed on these lines to collect the execution data that is used for runtime difference detection. For this, COLLECTOR-SAHAB finds a mapping between the lines of the original and patched versions. We call them the *matched lines* as defined below.

Definition 4 (Diff and Unchanged Lines). Given an original and a patched version of a program source file, *diff* is the set of lines that Myer's [23] algorithm detects as changed lines between the original and patched versions, as implemented in git diff. The diff contains a set of code lines deleted from the original and a set of code lines inserted into the patched. The remaining lines are the *unchanged lines*.

Definition 5 (Intra-function Matched Lines). In this paper, matched lines are all tuples $< l_o, l_p >$ where l_o and l_p meet two requirements. First, l_o is a line from the original version that is unchanged and mapped to l_p from the patched version. Second, l_o should be from a method that contains at least one line of the diff.

For example, Listing 1 shows an example of a patch with diff at lines 2 and 3. Lines 4, 5, 6, 7, and 8 are matched lines, for this example. Each of these lines is unchanged in the code diff and appears both in the original (l_o) and patched (l_p) versions. Also, these lines all appear in the getNumericalVariance() method which contains lines 2 and 3 as the diff. Therefore, lines 4, 5, 6, 7, and 8 meet both requirements and form the set of *matched lines*.

COLLECTOR-SAHAB registers breakpoints at all matched lines in the changed method. We focus on the matched lines so that we can compare the data between the two versions.
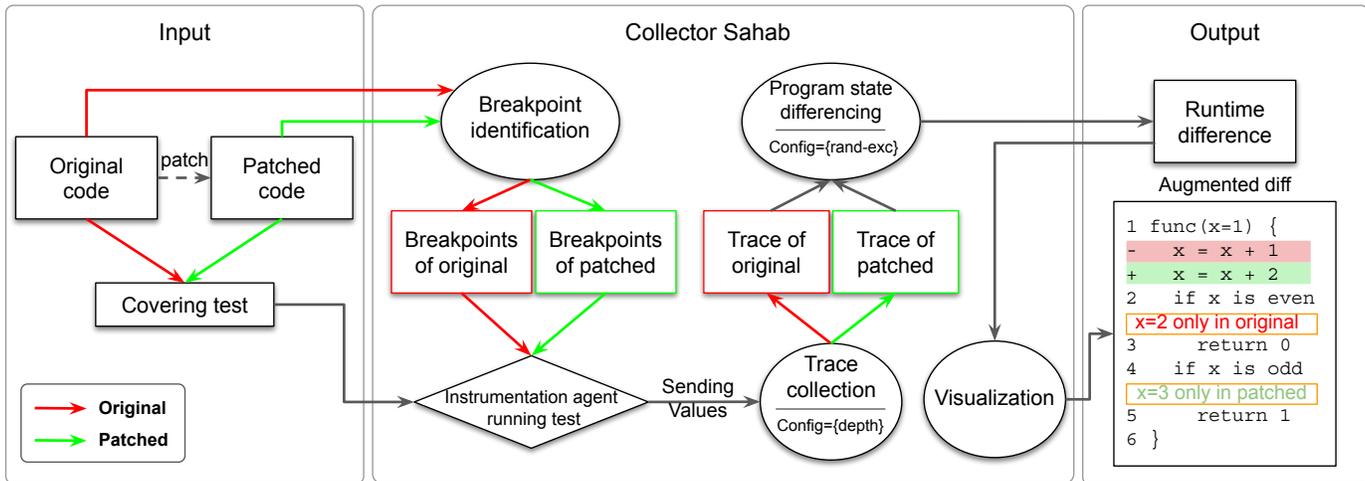
Figure 1: Overview of COLLECTOR-SAHAB.

```
1  {
2    "breakpoint": [
3      {
4        "file": "foo/BasicMath.java",
5        "lineNumber": 5,
6        "stackFrameContexts": [
7          {
8            "positionFromTopInStackTrace": 1,
9            "location": "foo.BasicMath:5",
10           "stackTrace": [
11             "add:5, foo.BasicMath",
12             "test_add:11, foo.BasicMathTest"
13           ],
14           "runtimeValueCollection": [
15             {
16               "kind": "LOCAL_VARIABLE",
17               "name": "x",
18               "type": "int",
19               "value": 23,
20               "fields": null,
21               "arrayElements": null
22             },
23             {
24               "kind": "LOCAL_VARIABLE",
25               "name": "y",
26               "type": "int",
27               "value": 2,
28               "fields": null,
29               "arrayElements": null
30             }
31           ]
32         }
33       ]
34     }
35   ]
36 }
```

Listing 3: Example runtime data obatined with COLLECTOR-SAHAB's trace collection component.

In theory, the runtime difference can be observed only after the change. However, COLLECTOR-SAHAB also monitors the lines before the change for two reasons. First, it enables us to better exclude random changes, see Section 3.5. Second, we note that the changed method can be executed multiple times; hence, a program state difference at a changed line can affect what happens above a code change in the method at later executions.

3.4 Trace Collection

After registering the breakpoints, COLLECTOR-SAHAB

runs the covering test on both original and patched versions and collects the execution trace as defined below.

Definition 6 (Program State – in COLLECTOR-SAHAB)). The program state s that COLLECTOR-SAHAB collects once a breakpoint b at line l is executed consists of the set of all variables V_l visible at l and their values, as well as the method calls leading to the execution of b. In practice, the program state object that COLLECTOR-SAHAB collects contains the following information:

- *file* is source file where b is located.
- *lineNumber* is the line number where b is located.
- *stack frame context* includes : the *stack trace* and the *runtime value collection* as follows.
- *stack trace* is the list of program lines that called a series of methods before reaching b during the execution.
- *runtime value collection* is the set of all variables and their values that are visible when the execution reaches b. This includes all local variables and class fields visible at b. If the variable is an array, all of its elements are also collected, and if it is a non-primitive variable, all the fields inside the object that it refers to are also stored. We note that COLLECTOR-SAHAB considers strings as primitive variables.

Note that COLLECTOR-SAHAB considers all visible variables at a line to compute the program state. This means even the variables that refer to objects of classes in third-party libraries are monitored. Therefore, COLLECTOR-SAHAB also detects runtime differences involving library objects.

Definition 7 (Execution Trace). In this paper, an execution trace is defined as the sequence of all program states that COLLECTOR-SAHAB collects during the execution of a test covering a change. Note that since every inserted breakpoint can be executed multiple times, the trace may include one or more program states per breakpoint.

The collection of execution traces is done as follows. Before the initiation of test execution, we instrument the Java classes to collect the local variables, field, and return values

at breakpoints. During the execution of the tests, the instrumented code is invoked and we save the runtime value as Java objects in memory. To avoid exceeding memory, the values of variables are collected only up to a state depth that is given as part of COLLECTOR-SAHAB's configuration (see Section 2). Eventually, when the test execution finishes, we output a serialized version of program states to disk. The program states are stored in the output file in the same order as they occur during the execution of the program.

Listing 3 shows the execution trace collected during the execution of test `BasicMathTest::test_add`. In this example, one program state is collected. This program state is for line 5 of the file `foo/BasicMath.java`. There are two variables whose values are stored in this trace: `x=23` and `y=2`.

3.5 Program State Differencing

We now explain how COLLECTOR-SAHAB computes the runtime difference between two versions. We devise a novel algorithm whose overarching goal is to produce a concise runtime difference that makes sense to developers. For this, we first define our key concept of *state value*, before its qualified version as relevant state value, and unique relevant state value.

Definition 8 (State value). For a given program state s corresponding to a breakpoint b located at line l, a state value is a tuple $< l, path, value >$. In this tuple, l is the line of b, $path$ is the path to a primitive variable inside *runtime value collection* of s, and $value$ is the value of that primitive variable.

The *path* in a state value specifies how a primitive variable can be accessed at the given line. Path appears in form of a list of variable names concatenated with dots such as $v_1.v_2...v_n$. v_1 is an object accessible at l, v_n is a primitive variable, and v_i is a field in the object referred by v_{i-1}. Note that we only consider paths primitive variables in state values. The value for a non-primitive variable is an object reference which is not understandable for developers, hence we do not take them as a part of the state value.

State value is the most fine-grained data that can be extracted from an execution trace. We compute all state values inside the collected execution trace to compute the runtime difference.

In Listing 2, if "student1" is the variable whose value is collected in program state s at a breakpoint at line l, the following are the state values that we extract for s:

```
<l,student1.name,"Alice">
<l,student1.supervisor.name,"Bob">
<l,student1.supervisor.education.university,"KTH Institute">
<l,student1.supervisor.education.city,"Stockholm">
```

Definition 9 (Relevant State value). A state value $< l, var_1..., val >$ is considered *relevant*, if and only if var_1 is accessed on line l.

Definition 10 (Unique Relevant State value). A state value is said to be unique if it only appears on the original or patched version. $< l, var.**, val >$ is a unique relevant state value if it is a relevant state value in the original version, and $< l', var.**, val >$ is not a relevant state value in

---

Algorithm 1 Algorithm for program state differencing.

Inputs:
   ot: The trace collected for the original version
   pt: The trace collected for the patched version
   osrc: The source of the original version
   psrc: The source of the patched version
Outputs:
   ousv: List of unique relevant state values in the original trace
   pusv: List of unique relevant state values in the patched trace

```
1:  oes ← program_states(ot)
2:  pes ← program_states(pt)
3:  osv ← GET_STATE_VALUES(osrc, oes)
4:  psv ← GET_STATE_VALUES(psrc, pes)
5:  ousv ← GET_UNIQUE_STATES(osv, psv)
6:  pusv ← GET_UNIQUE_STATES(psv, osv)
7:  return (ousv, pusv)
8:
9:  function GET_STATE_VALUES(src, ess)
10:     SV ← []
11:     for each execution_state es in ess do
12:         line ← es.line
13:         rv ← relevant_vars(src, line)
14:         for each variable_data vd in es do
15:             if rv.contains(vd.name) then
16:                 cur_SV ← EXTRACT_SVS(line, "", vd)
17:                 SV.insert(cur_SV)
18:     return SV
19: function EXTRACT_SVS(line, prefix, vd)
20:     SV ← []
21:     if is_primitve(vd) then
22:         sv_id = line + prefix + vd.name + vd.value
23:         SV.insert(sv_id)
24:         return SV
25:     for each field_data ivd in vd do
26:         SV.insert(EXTRACT_SVS(line, rv, prefix + vd.name, ivd))
27:     return SV
28: function GET_UNIQUE_STATES(left_sv, right_sv)
29:     SV ← []
30:     for each state_value sv in left_sv do
31:         if !right_sv.contains(sv) then
32:             SV.insert(sv)
33:     return SV
```

---

the patched version where l' is the line mapper to l in the matched lines. Conversely, $< l', var.**, val >$ is a unique relevant state value if it is a relevant state value in the execution trace of the patched version and $< l, var.**, val >$ is not among the relevant state values of the execution trace of the patched version.

Algorithm Algorithm 1 describes the algorithm we employ to compute the execution differences. The inputs to the algorithm are the collected traces for the original (*ot*) and patched (*pt*) versions as well as the original (*osrc*) and patched (*psrc*) versions of the source code. The algorithm outputs the unique relevant state values for the original (*outsv*) and patched (*pusv*) versions.

The algorithm first extracts the program states saved in collected traces for each version (lines Algorithm 1 and Algorithm 2). Note that each collected trace contains the list of all program states that the trace collector has encountered

while executing the matched lines. These program states are saved in the same order as they occur during the execution.

Next, all state values are computed for each version (lines Algorithm 3 and Algorithm 4). The GET_STATE_VALUES method at line Algorithm 9 traverses program states one by one (line Algorithm 11). For each state, it gets the program line where that state has happened (line Algorithm 12) and the relevant variables at that program line (line Algorithm 13). Note for each line, the relevant variables are the ones that are accessed to be read on that line. Then, for each variable data *vd* that corresponds to a relevant variable (see line Algorithm 15), we get all state values linked to it by calling the EXTRACT_SVS method (line Algorithm 16).

The EXTRACT_SVS method (line Algorithm 19) discovers all state values that can be extracted from a given variable, recursively. For a given primitive variable, the algorithm adds its variable name and value to the state values set (lines Algorithm 21-Algorithm 24) and returns it. For non-primitive variables, the algorithm looks into the fields of the object that the variable is referring to one by one (line Algorithm 25). Per each field in the referred object, EXTRACT_SVS is called recursively (line Algorithm 26). When the state values are extracted for all the fields of the referred object, the resulting collection of state values is returned (line Algorithm 27).

After running lines Algorithm 3 and Algorithm 4, the state values that occur in the original and patched versions are saved in *osv* and *psv*, respectively. To compute the unique relevant state values for each version, we iterate over the extracted state values and identify the ones that are not included in the state values of the opposite version. The GET_UNIQUE_STATES method that is called on lines Algorithm 5 and Algorithm 6 does exactly this. The result of calling this method is then saved in *ousv* and *pusv*, which form the returned result of the program state differencing algorithm.

Note that Algorithm 1 explains how unique state values are extracted. COLLECTOR-SAHAB also computes unique program states in three steps. First, it maps each program state of each execution trace to the hashed version of that program state. For this, the Java string hashing algorithm is used. Second, the hashes are compared to each other to detect unique program state hashes. Finally, COLLECTOR-SAHAB looks back to the mappings to find the original program states mapped to the detected unique hashes. The result is the list of unique program states in each of the original/patched execution traces.

As Figure 1 shows, COLLECTOR-SAHAB also has a `random-exclusion` configuration for the program state differencing state, noted as `Config=rand-exc` in the figure. When `random-exclusion` is set to `true`, COLLECTOR-SAHAB excludes state values that show random behavior, i.e. not having identical values accross runs. Using this setting enables the users to focus on runtime differences that are causally related to the code changes. For this purpose, COLLECTOR-SAHAB takes two steps. First, it does not consider values of variables that refer to file paths. This step is taken because our experiments on real-world commits show that, most of the time, file paths change on different system environments. This means a runtime

Figure 2: The augmented diff generated by COLLECTOR-SAHAB for a code diff for Math-80, showing important runtime information.

difference in file paths is usually because of a change in the environment, not a change in the logic of the program. In the second step, COLLECTOR-SAHAB runs each version of the program three times and only keeps state values that appear in all three runs. This step excludes values that reflect flaky behavior of a program. Excluding random state values generates a code diff augmentation that shows the crucial behavioral changes caused by the code changes.

3.6 Output & Visualization

The output of the tool is a **runtime difference** in two forms: a textual file and a visualization. The textual file contains all unique program states and unique state values. It is intended to be used by other tools, i.e. it is not user-facing. For example, such an output can be used by program repair tools to suggest non-overfitting patches based on runtime data. The other output is aimed at developers, it is a visualization of the runtime difference (see Figure 2). This visualization, which we call the *augmented diff*, is only generated when a unique state value is identified, and detecting a unique program state is not sufficient to produce an augmented diff.

Textual Output. The main textual output of COLLECTOR-SAHAB is a file containing all unique program states and unique relevant state values for both the original and patched versions. Also, all the traces collected at the trace collection step (Section 3.4) are stored in a separate file. These textual outputs can be consumed by other tools based on their own needs. We believe that COLLECTOR-SAHAB will be used as a foundation for future research on dynamic analysis.

Visual Output. COLLECTOR-SAHAB uses the unique state values extracted in the previous step to produce a user-friendly user interface (UI), which we call the **augmented diff**. This UI adds runtime information to the typical code diff view, e.g. that of GitHub. This new graphical component shows the first unique relevant state values for both the original and patched versions. We consider the first unique state values as the program state differences for two reasons. First, there are usually many unique state values, which is overwhelming for developers. Second, the literature has shown that the first point of execution divergence is the most relevant for code reviewers [24].

Figure 2 shows an example of an augmented diff for a code diff for the Math-20 bug. This generated UI clearly indicates that after changing how variable `j` is computed in line 1135, the variable is assigned a new value in the patched

version. As shown in Figure 2, a code reviewer directly sees that the execution of test `testMathpbx02` reveals that `j=24` only occurs in the patched version and `j=27` only occurs in the original version. By displaying those concrete runtime differences caused by the code change, COLLECTOR-SAHAB helps code reviewers to make an informed decision about the appropriateness of a code diff.

3.7 Implementation

COLLECTOR-SAHAB is an advanced tool made of 5,800 sophisticated lines of Java code, built using Java 11 and relying on ASM and Byte Buddy for instrumentation. All of the tool's source code is publicly available[2].

4 EXPERIMENTAL DESIGN

4.1 Research Questions

In this paper, we study the following research questions.

- RQ1 (effectiveness): What is the effectiveness of COLLECTOR-SAHAB to detect runtime differences in bug fixing code diffs? We run COLLECTOR-SAHAB on a large dataset of bug fixing APR code diffs from the DRR [25] dataset. We compare COLLECTOR-SAHAB effectiveness with DIDIFFFF [10] as baseline.
- RQ2 (applicability): To what extent does COLLECTOR-SAHAB perform on augmenting diffs from real-world projects? We assess COLLECTOR-SAHAB applicability on real-world commits. This improves the external validity of our evaluation beyond the projects of DRR.
- RQ3 (quality): What is the quality of the diffs augmented with runtime differences generated by COLLECTOR-SAHAB? To assess the quality of COLLECTOR-SAHAB reports for code reviewers, we conduct a manual analysis to evaluate their quality with respect to two criteria: correctness and understandability. We also manually check the accuracy of COLLECTOR-SAHAB in excluding random runtime differences. This analysis is performed on the real-world commits also used in the RQ2 experiment.
- RQ4 (user study): How do software developers assess COLLECTOR-SAHAB's augmented diffs? We conduct a user study and ask developers to rate COLLECTOR-SAHAB augmented diffs in terms of usefulness, understandability, and novelty. To compare COLLECTOR-SAHAB against the state of the art, we also ask the participants to give a score to DIDIFFFF outputs. In this study, we use the real-world commits collected for the RQ2 and RQ3 experiment.

4.2 Datasets

We use two different datasets for our experiments: MVN-DRR and RW3.

MVN-DRR is used for assessing the effectiveness of our tool (RQ1). This dataset contains a subset of code diffs from Ye et al.'s previous research [25]. Ye et al.'s DRR is a collection of 638 code diffs generated by 15 program repair tools for bugs in the Defects4J dataset [9]. Defects4J includes

bugs from five different projects: "jfreechart", "closure-compiler", "commons-lang", "commons-math", and "joda-time". In MVN-DRR, we select DRR code diffs that meet two conditions. First, the project should be a maven project with a "pom.xml". This means we can successfully build both the original and code diffs versions by running a "mvn compile". Second, the code diff changes should all appear in one single java method. We focus on single-method code diffs to make sure there is a single, well-scoped behavioral change. This means the code diffs that change multiple files, multiple methods, class declarations, method declarations, or imports are ignored. The result of these two filters is that MVN-DRR contains 584 code diffs from two projects: "commons-math" and "joda-time", appropriate for our experiments.

Our second dataset is RW3, which is used to study COLLECTOR-SAHAB applicability in the field, on a diverse set of complex real-world projects (RQ2 & RQ3 & RQ4). RW3 contains human-made code diffs[3] from GitHub repositories. To create this dataset, we first start by filtering repositories on GitHub. We select public Java repositories that are not archived and have a Codecov or Coveralls configuration file at the root of the repository. We consider projects using Codecov or Coveralls as they tend to be well-tested, which makes them appropriate for our experiments. To make sure the chosen repositories are well-established we select repositories that have 50 stars or more. These filters return a list of 158 repositories. Then, we scan these repositories for code diffs, excluding merge code diffs, from 2022-08-13 to 2022-09-12 which change exactly one file that meets two requirements. First, the changed file is a Java file, and second, the whole change lies inside a single method, to be consistent with RQ1. We perform manual analysis on each code diff to eliminate those where the code diff is not covered by a test, the project fails to compile, or we fail to execute the tests that cover the change. The resulting RW3 dataset contains 50 code diffs from 10 GitHub repositories as shown in Table 1.

In Table 1, we give the total number of commits, modules, and lines of code (KLOC) per considered project to give a sense of their complexity. As shown in the table, the median number of code diffs, modules, and lines of code are 6,069, 8, and 175K. This indicates that the projects in RW3 are established repositories with real-world complexity. These statistics confirm that the RW3 code diffs are selected from complex, real-world repositories. If COLLECTOR-SAHAB works on them, it is a good indicator of the quality of our approach and its prototype.

4.3 Baseline

There are very few runtime differencing tools that are available. In our work, we focus on Java, and there is one single open-source runtime differencing tool for Java, DIDIFFFF (pronounced di-di-fu-fu) [10].

Consequently, in our experimental evaluation, we consider DIDIFFFF as the baseline. DIDIFFFF runs two versions of a program and illustrates the difference between the values assigned to program variables in a web interface. DIDIFFFF uses SELOGGER [26] to record all the values each

---

2. https://github.com/ASSERT-KTH/collector-sahab

3. We consider each commit in GitHub as a code diffs.

Table 1: Descriptive statistics of open-source projects included in benchmark RW3.

Repository	#commits	#modules	KLOC
codingapi/springboot-framework	178	5	3
nosan/embedded-cassandra	2,149	1	12
infinitest/infinitest	489	7	25
jhipster/jhipster-lite	6,413	1	90
schemacrawler/SchemaCrawler	8,198	23	148
alibaba/nacos	4,224	18	203
apache/dubbo	5,725	27	229
cdk/cdk	17,263	9	502
apache/iotdb	7,364	43	591
JMRI/JMRI	74,997	1	4,741
Min	178	1	3
Median	6,069	8	175
Max	74,997	43	4,741

variable of a program holds during an execution. Using these values, DIDIFFFF creates two lists of values for each variable access that is not part of the code change. The first list represents the values assigned to the variable at that point of access in the original version. The second list represents values assigned to the variable in the patched version of the program. Finally, DIDIFFFF compares the two lists of values for each variable access. The result of this comparison is presented to code reviewers in form of a graphical user interface (GUI). This GUI shows the two lists of values assigned to a variable and highlights if there is a difference between these two lists. An example output of DIDIFFFF can be seen at https://bit.ly/3N7TKEb.

DIDIFFFF reports two types of runtime differences: 1) the difference in the content of the list of values assigned to primitive variables, and 2) the difference in the length of the list of values assigned to non-primitive variables. Regarding non-primitive variables, DIDIFFFF only counts how many values are assigned to them, it does not consider what is inside those values. That is because DIDIFFFF does not look deep into the objects referred by non-primitive variables, which makes it significantly different from COLLECTOR-SAHAB.

4.4 RQ1: Effectiveness on Benchmark

To answer **RQ1**, we run COLLECTOR-SAHAB on all 584 code diffs in MVN-DRR with four different state depths: depth=0,1,2,3. This maximum depth of three has been found empirically, because collecting data for higher depths requires resources beyond our experimental setup without improving effectiveness.

The key metric is the number of code diffs for which COLLECTOR-SAHAB augments the diff. We split the remaining code diffs into three groups and count the number of code diffs in each group. Code diffs for which no unique state value is detected, code diffs that cause a memory failure, and code diffs that make COLLECTOR-SAHAB exceed the time limit.

We also run DIDIFFFF on the MVN-DRR code diffs. We count the number of code diffs for which DIDIFFFF detects a runtime difference contrasting it with the number of code diffs for which COLLECTOR-SAHAB augments the diff.

We perform our experiment on a machine with sixteen Intel(R) Core(TM) i9-10980XE processors, each running at 3.00GHz and having 18 cores, and eight 16GB RAMs each of type DDR4 with 3600 MT/s clock speed.

4.5 RQ2: Applicability in the Field

To answer **RQ2**, we run COLLECTOR-SAHAB on all code diffs in RW3. Then, we compute the set of code diffs for which COLLECTOR-SAHAB successfully runs and lists all the unique state values (if any). These are the *successful code diffs*. Next, we get all the repositories corresponding to these code diffs and call the set of these repositories as *successful repositories*.

Finally, we compute the ratio $AC = \frac{|successfull\_code\_diffs|}{|all\_code\_diffs|}$, which represents the fraction of code diffs on which COLLECTOR-SAHAB works successfully and accurately. A high value of AC and a large set of successful repositories indicate how applicable our tool is on real-world projects and, from a scientific perspective, measure the external validity of our results.

4.6 RQ3: Manual Analysis on Augmentation Quality

In answer to **RQ3**, we conduct a detailed manual analysis of the outputs of running COLLECTOR-SAHAB on RW3 code diffs. In this analysis, we consider all the successful code diffs in RW3. This consists of the code diffs for which COLLECTOR-SAHAB detects a unique state value, and code diffs for which COLLECTOR-SAHAB does not detect any unique state values. Note that all code diffs in RW3 are covered by a test, but the covering test is not required to fail on the original version and pass on the patched version. This means the test may not reveal any runtime difference.

For code diffs with no detected unique state value, we manually verify that the code change does not entail any unique state value in neither the original nor the patched version. For each code diff with a detected unique state value, we study the quality of COLLECTOR-SAHAB augmented diff from three aspects: correctness, understandability, and randomness exclusion accuracy of the detected runtime difference. In case of negative results for any of the three mentioned criteria, we investigate the reasons behind them.

To assess correctness, we manually run the covering test line by line in a debugger twice, first on the original and second on the patched versions of the program under study. We ensure the first unique relevant state value reported by COLLECTOR-SAHAB matches what we see in the debugger value explorer.

Regarding understandability, we analyze whether a code reviewer could understand the reported runtime behavioral change in a reasonable amount of time (typically 5-30 minutes, per typical code review practices at the function level).

Finally, we evaluate the randomness exclusion accuracy of COLLECTOR-SAHAB to see if its `random-exclusion` configuration (as defined in Section 3.5) works accurately and effectively. We perform this evaluation in two steps as follows.

First, we divide the commits into two groups. Group 1 contains commits where the unique state value is a non-random value and Group 2 contains commits wherein

the unique state value is reported only because of differences in random value. We filter Group 2 commits by setting `random-exclusion=false` and then running Collector-Sahab again on these commits with `random-exclusion=true`. If the unique state value reported disappears, we consider such commits under Group 2. If after changing `random-exclusion` a unique state value is still reported, we consider the commit under Group 1.

In the second step, we manually perform two checks. In the first check, we investigate if the unique state values reported for commits in Group 2 are actually random and no other non-random unique state value exists. If our manual analysis confirms this, it means Collector-Sahab has performed accurately by not identifying any unique state value for these commits when `random-exclusion=true`. Second, we check if the unique state values reported for commits in Group 1 are not random, when `random-exclusion=true`. Again, if our manual analysis confirms this, it indicates the accuracy of randomness exclusion by Collector-Sahab. A high accuracy of randomness exclusion on the commits means that the `random-exclusion` configuration of Collector-Sahab is an accurate and highly applicable tool for excluding runtime differences that are caused by flaky behavior of programs.

This analysis is conducted by two of the authors, both experts in Java programming. In case of any discrepancies between the participants, they meet and discuss to resolve them and reach a conclusion.

4.7 RQ4: User Study

To answer **RQ4**, we perform a user study. In this study, we ask experienced developers to comment on the quality of Collector-Sahab and DiDiffff reports. For this experiment, we have two groups of participants: a group of four participants working in industry, and a group of four participants studying at our university. All the participants are experienced in programming with Java language and creating and merging GitHub pull requests.

For each participant, the experiment is performed in three steps as follows. First, in a live presentation, we use some examples to introduce the relevant concepts of runtime differencing and the two tools (Collector-Sahab and DiDiffff) to the participant. Second, we present a selected set of commits and their augmentations with both tools. For each commit, the participant gives a score between 1 and 5 to each augmented code diff in terms of **usefulness**, **clarity**, and **novelty**. Usefulness means whether the report generated by Collector-Sahab helps in understanding the changes made to the program. Clarity means how comprehensible the generated report is. Novelty measures how novel is the information provided by Collector-Sahab in the sense that it cannot be obtained by looking at the plain code diff. In the third and final step, we hold a short discussion with the participant when they have finished rating the augmented code diffs. In the final discussion, we ask for an overall feedback regarding code diff augmentation with runtime data, how the tools compare against each other, and potential improvements in each tool.

To select the commits that are shown to participants, we consider the RW3 commits that cause a runtime difference per our manual analysis in RQ3. We exclude three types of commits from the RQ3 experiment. First, the commits that suffer from random runtime differences. Second, if there are multiple commits changing the same line of a program, we only show one of them to a participant. Third, we also exclude commits that change a test method. After applying these filters, we select a total of 8 commits. We split these commits into two groups (industry and students) so that each group gets 4 commits to analyze. This is is an appropriate number as it fits well within the one-hour time limit that we set for each participant.

In answer to each question, the participants give a score between 1 and 5 to the augmented code diff. To have an overall assessment of the given scores, we first compute the average score a participant has given to the augmented diffs of each tool in answer to each question. Then, we calculate the median score given by participants in each group (industry/students) to obtain the overall assessment in relation to each criterion.

Note that for the commits for which only one of the tools detects a runtime difference, the scores cannot be compared between the two tools. Therefore, we consider these commits separately.

In our final live discussion with participants, we ask them three questions. First, if they think augmenting code diffs with runtime data can be a useful practice. Second, we ask them which tool they prefer overall. Finally, we ask whether they have any specific suggestion for improving each tool. The answers given by participants help us to see if this path of code diff augmentation is promising and what are the best ways to improve the state of the art.

5 Experimental Results

5.1 Collector-Sahab Effectiveness (RQ1)

Table 2 shows the results of running Collector-Sahab on 584 code diffs from our benchmark MVN-DRR. In each row, "Diff Tool" shows the used tool and its configuration. More specifically, we run Collector-Sahab with four different depths from zero to three and DiDiffff on the dataset (see Section 2 and Section 3.4). The second column "Augmented Diff" indicates the number of code diffs for which a runtime difference is detected and the diff is augmented with a unique state value. Each of the remaining three columns explains the failure modes. "No Diff Detected" indicates the number of code diffs for which no unique state value in either of the versions is found (see Section 3.6). Finally, "Memory Failure" and "Time Limit" show the number of code diffs on which Collector-Sahab faces memory and time limit during execution, respectively.

Consider "Collector-Sahab (depth=1)" in Table 2 as an example. The table shows that in this case, Collector-Sahab augments the diff with a detected unique state value for 95% (555/584) code diffs. For most code diffs, Collector-Sahab is reliable to identify an execution difference between two versions of a program. With depth set to 1, Collector-Sahab faces the memory limit for 1.5% (9/584) of the code diffs and fails due to exceeding the

Table 2: Results of running COLLECTOR-SAHAB on 584 code diffs in MVN-DRR.

Diff Tool	Augmented Diff	No Diff Detected	Memory Failure	Time Limit
DIDIFFFF	377 (64.5%)	207 (35.5%)	–	–
COLLECTOR-SAHAB (depth=0)	543 (93.0%)	32 (4.9%)	9 (1.5%)	0 (0.0%)
COLLECTOR-SAHAB (depth=1)	555 (95.0%)	18 (3.2%)	9 (1.5%)	2 (0.3%)
COLLECTOR-SAHAB (depth=2)	554 (94.8%)	16 (2.8%)	13 (2.2%)	1 (0.2%)
COLLECTOR-SAHAB (depth=3)	544 (93.1%)	22 (3.9%)	18 (3.0%)	0 (0.0%)

time limit only for 0.3% (2/584) of the code diffs, which represents a tiny minority. For the remaining 3.2% (18/584) of the diffs, COLLECTOR-SAHAB runs successfully but does not detect any unique state value in either of the versions. This shows for most of the code diffs the differentiating test triggers a unique state value that is identified and reported by COLLECTOR-SAHAB.

As shown in Table 2, COLLECTOR-SAHAB outperforms the state-of-the-art tool DIDIFFFF. The number of code diffs for which we have a runtime difference is higher for all depths, from 93% (543/584) to 95% (555/584) for different depths, while it is 64.5% (377/584) for DIDIFFFF. COLLECTOR-SAHAB detects and reports a runtime difference in notably more cases than DIDIFFFF.

Next, we manually analyze the code diffs for which COLLECTOR-SAHAB or DIDIFFFF fail to generate a diff report. Our analysis indicates that there are three reasons why COLLECTOR-SAHAB outperforms DIDIFFFF in terms of identifying runtime differences. First, DIDIFFFF only considers differences between values of primitive types, while COLLECTOR-SAHAB also considers differences between the value of non-primitive variables that refer to objects. Second, DIDIFFFF only considers differences between local variable values, while COLLECTOR-SAHAB also detects the difference between the values returned by methods (inside a return statement without an explicit variable). The third and last reason that COLLECTOR-SAHAB outperforms DIDIFFFF is related to DIDIFFFF not logging the variable values inside inner classes. The latter case is present in our dataset. These three reasons explain why COLLECTOR-SAHAB is able to detect more runtime differences compared to DIDIFFFF.

The column "No Diff Detected" shows the number of code diffs for which COLLECTOR-SAHAB does not find a unique state vlue. As shown in Table 2, generally, the number of these cases decreases when depth increases: the number of code diffs with no diff detected is 32 for depth=0; this number decreases to 16 for depth=2. This indicates that COLLECTOR-SAHAB's power to look deep inside non-primitive objects adds to its effectiveness in detecting runtime differences. There is an exception for depth=3, where the number of code diffs without a detected diff increases to 22. Per our detailed manual analysis of the logs, the reason behind this is a problem with one of the libraries that COLLECTOR-SAHAB uses. This library [4] fails serialize very large Java object into JSON format. As the size of trace object collected for depth=3 is as large as 10GB in some cases, this leads to COLLECTOR-SAHAB failure to detect a runtime difference. This shows increasing the depth has its own cost and we should carefully select an optimal depth for configuring COLLECTOR-SAHAB in a given setup.

4. https://github.com/FasterXML/jackson

In the current version, COLLECTOR-SAHAB stores all the collected data in RAM during execution and prints the whole trace of the program under test at the end. As the entire trace may be very large, it causes memory failure in some cases. The "Memory Failure" column of Section 5.1 shows that facing a memory limit during COLLECTOR-SAHAB execution is rare. When the depth is set to zero or one, we face memory limit issues only for 1.5% (9/584) code diffs. When depth is set to two or three, the memory limit causes an issue for 2.2% (7/584) and 3.0% (18/584) of code diffs, respectively. There is room for further improvement in this regard, as memory failures can be partially avoided by printing the execution trace step-by-step. This means there is a significant potential to improve COLLECTOR-SAHAB memory usage.

COLLECTOR-SAHAB puts breakpoints at all matched lines Section 3.3, which means the instrumenter collects the program state at many places. This adds an overhead during the runtime and makes the execution take longer than when the code is not instrumented. With respect to time, the "Time Limit" column shows that we rarely face cases where COLLECTOR-SAHAB fails due to time issues. We exceed the time limit only with depth=1 and depth=2 and only for 0.3% (2/584) and 0.2% (1/584) of the code diffs. This means COLLECTOR-SAHAB is fast enough to finish its analysis on most code diffs. We note that COLLECTOR-SAHAB exceeds the time limit for more cases when depth=1, compared to when depth=2 and depth=3. This happens while increasing the depth should make the tool take more time to collect the trace and compute the diff. The reason is that when the depth is increased, we may face memory failure or the library issue mentioned above even before exceeding the time limit. Therefore, the "Time Limit" case turns into "Memory Failure" or "No Diff Detected" in such cases.

Overall, we notice that COLLECTOR-SAHAB performs the best when the depth is set to one by augmenting the diff for 95% (555/584) of code diffs. Depth=1 outperforms other depths as it provides an accurate balance between looking at details (better than depth=0) and collecting too much data (better than depth=2,3).

Answer to RQ1: What is the effectiveness of COLLECTOR-SAHAB to detect runtime differences in bug fixing code diffs?
Our experiment on benchmark MVN-DRR shows that COLLECTOR-SAHAB is effective at detecting runtime differences. COLLECTOR-SAHAB augments diffs with runtime information for 95% (555/584) of the code diffs in our benchmark, outperforming DIDIFFFF, which detects a runtime difference in only 64.5% (377/584) of the code diffs. It is clear that COLLECTOR-SAHAB enhances the state of the art of runtime differencing.

Table 3: Results of running COLLECTOR-SAHAB on RW3.

Repository	|ALL|	Success	AC	Failure
JMRI/JMRI	13	12	0.92	1
alibaba/nacos	2	2	1.00	0
apache/dubbo	4	4	1.00	0
apache/iotdb	8	7	0.88	1
cdk/cdk	5	5	1.00	0
codingapi/springboot-framework	1	1	1.00	0
infinitest/infinitest	1	1	1.00	0
jhipster/jhipster-lite	7	7	1.00	0
nosan/embedded-cassandra	1	1	1.00	0
schemacrawler/SchemaCrawler	8	8	1.00	0
Total	50	48	0.96	2

5.2 COLLECTOR-SAHAB Applicability (RQ2)

In our study of COLLECTOR-SAHAB applicability, we run COLLECTOR-SAHAB on code diffs from benchmark RW3. The results of this experiment are shown in Table 3. In this table, the left column indicates the repository id. The number of all code diffs from each real-world repo is shown in the "|ALL|" column, per our selection criteria. As defined in Section 4.5, "Success" represents the number of all code diffs successfully processed by COLLECTOR-SAHAB. "AC" is the ratio of code diffs successfully processed by COLLECTOR-SAHAB and finally, "Failure" indicates the number of code diffs that COLLECTOR-SAHAB fails to process.

Consider the repository jhipster/jhipster-lite as an example (eighth row). RW3 consists of 7 code diffs from this repository. COLLECTOR-SAHAB successfully performs on all 7 code diffs, giving AC equal to 1 (7/7).

In total, our selection criteria yield 50 code diffs from real-world complex projects. They are sampled from a set of ten diverse open-source projects. COLLECTOR-SAHAB successfully analyzes 48 code diffs out of a total of 50. For all repositories, we have at least one successfully processed code diff. This indicates that COLLECTOR-SAHAB is successful in handling diverse Java projects from RW3. With 50 total code diffs and 48 successfully analyzed ones, the corresponding value of AC is 0.96 (48/50). Notably, the median AC is 1, because COLLECTOR-SAHAB handles all code diffs for 6 projects. Such a high value of AC attests the real applicability of COLLECTOR-SAHAB.

Finally, as shown in the "Failure" column, there are 2 commits where COLLECTOR-SAHAB fails to process the code diff to collect a runtime diff. Both of them are due to out-of-memory exception. This can be attributed to the huge number of states collected. In these cases, either the tool runs out of heap space while generating the execution trace (step 'trace collection') or while iterating through states to find a diff (step 'program state differencing'). It is known that fine-grain monitoring is very memory intensive [27].

Answer to RQ2: To what extent does COLLECTOR-SAHAB perform on augmenting diffs from real-world projects?

To analyze applicability in the wild, we build a well-formed benchmark of 50 code diffs from 10 complex open-source Java repositories. COLLECTOR-SAHAB successfully runs on 48 code diffs, which yields a high success ratio of 0.96. COLLECTOR-SAHAB is able to handle real-world Java projects, incl. large ones with multi-module build systems.

5.3 COLLECTOR-SAHAB Augmentation Quality (RQ3)

Per the RQ3 protocol, we run COLLECTOR-SAHAB with two configurations (`random-exclusion=true,false`) and use a Java debugger to analyze the RW3 augmented code diffs.

First, we analyze the 18 code diffs for which COLLECTOR-SAHAB does not detect a unique state value even with `random-exclusion=false`. Our manual analysis through the debugger verifies what COLLECTOR-SAHAB states about these code diffs: they do not entail any unique state value. This indicates the correctness of COLLECTOR-SAHAB results on this set of code diffs.

More importantly, there are 30 code diffs for which COLLECTOR-SAHAB detects a unique state value and generates an augmented diff. The list of these code diffs is presented in Table 4. As shown in the table, for 17 of the code diffs, COLLECTOR-SAHAB outputs a unique state value with `random-exclusion=true`. For the remaining 13 code diffs, COLLECTOR-SAHAB detects a unique state value only with `random-exclusion=false`. Therefore, COLLECTOR-SAHAB labels 40% (13/30) of the considered code diffs as causing random runtime differences. This indicates the notable importance of the random value exclusion feature on real-world commits.

Table 4 also presents the results of our manual analysis on the correctness, understandability, and randomness exclusion accuracy of the runtime difference detected for the code diffs. We define these metrics as follows. Column "Correctness" shows if COLLECTOR-SAHAB's augmented diff matches the difference we witness by debugging the code. The "Understandability" column shows if it is possible to understand the augmented diff in a reasonable amount of time.

Let's take the code diff #7 as an example, from project apache/iotdb. COLLECTOR-SAHAB generates the augmented diff for this code diff as shown in Figure 3. The message of this commit reads as "Call start() method before FI really executing" and the plain text code diff also shows `instanceContext.start()` is called on line 73 of the patched version. However, this does not show how calling `instanceContext.start()` affects the runtime behavior of the program. Here COLLECTOR-SAHAB comes in and claims that when this method is not called in the original version, `startNanos.value` of `instanceContext` is set to 0 and calling this method in the patch changes this value. This gives the code reviewer a very concrete and detailed runtime difference information that cannot be easily obtained just by looking at the textual diff. Our

Table 4: Manual analysis of the quality of COLLECTOR-SAHAB's augmented diffs on real-world code changes.

	ID	Code diff	Correctness	Understandability
rand-exc=true	1	SchemaCrawler@07c7368	✓	✗
	2	SchemaCrawler@55a55de	✓	✓
	3	SchemaCrawler@f950519	✓	✓
	4	JMRI@0b53ea8	✓	✓
	5	JMRI@4a30042	✓	✓
	6	iotdb@5faa9da	✓	✗
	7	iotdb@a3559e5	✓	✓
	8	iotdb@e5e4f17	✓	✓
	9	springboot-framework@70b0d12	✓	✓
	10	jhipster-lite@510a975	✓	✓
	11	jhipster-lite@5d2a6b8	✓	✓
	12	jhipster-lite@69a7d5a	✓	✓
	13	jhipster-lite@9162791	✓	✓
	14	jhipster-lite@e2fc61f	✓	✓
	15	jhipster-lite@f5b592a	✓	✓
	16	cdk@d500be0	✓	✓
	17	dubbo@29e4e42	✓	✓
rand-exc=false	18	embedded-cassandra@dcc5818	✓	✓
	19	SchemaCrawler@5955cd0	✓	✓
	20	SchemaCrawler@c06131a	✓	✓
	21	SchemaCrawler@451a5a5	✓	✓
	22	JMRI@3e24272	✓	✓
	23	JMRI@d21837e	✓	✓
	24	dubbo@bb7f942	✓	✓
	25	iotdb@35541f4	✓	✓
	26	iotdb@94657cf	✓	✓
	27	iotdb@e5e4f17	✓	✓
	28	jhipster-lite@2b38799	✓	✓
	29	dubbo@6b095f1	✓	✓
	30	nacos@0cf9c24	✓	✓
		Total	100% (30/30)	93% (28/30)

manual debugging of the program verifies the correctness of COLLECTOR-SAHAB's report for this code diff. This reported runtime difference in this augmented diff is also easy to understand.

Overall, the execution difference reported by COLLECTOR-SAHAB is correct for all 30 code diffs in our benchmark. Generating a correct augmented diff for 100% of the code diffs indicates that COLLECTOR-SAHAB is a reliable tool for execution difference detection. Our manual analysis of the randomness exclusion accuracy also shows that COLLECTOR-SAHAB correctly determines whether unique state values are random or not for 100% of the code diffs. In other words, COLLECTOR-SAHAB is able to correctly exclude unique state values that are flaky or related to file paths as explained in Section 3.5. This indicates that COLLECTOR-SAHAB's augmented diffs help developers focus on non-random and runtime differences causally related to the change.

We also see that the COLLECTOR-SAHAB output for code diff #6 is not easy to understand. By carefully analyzing this code diff, we see that the reported runtime difference is hard to understand because the code change is a multi-line change and all changed lines affect the reported unique state value. To understand the augmented diff, we need to know the value of many variables involved in the code diff. Recall that we designed COLLECTOR-SAHAB to only show the first unique relevant state value. This demonstrates the trade-off between conciseness and comprehensiveness of the augmented diff. To summarize, COLLECTOR-SAHAB's report for the code diff #6 is not easy to understand because more variable values are necessary for this particular case.

The code diff #1 is another case where the augmentation by COLLECTOR-SAHAB is not easily understandable. This augmented code diff is shown in Figure 4. The detected unique state value is the `false` value of the `printStackPropertiesSet` field in the `method` variable. The issue with this runtime data is that it does not state if `printStackPropertiesSet` is actually a static field. Consequently, the user may imply that `method` is referring to a not-null object whose `printStackPropertiesSet` is set to `true`, while in fact the value of `method` is null in this example. Based on this observation, we envision a future improvement in COLLECTOR-SAHAB that adds the type (static/non-static) of the fields to the reported runtime differences.

> **Answer to RQ3: What is the quality of the diffs augmented with runtime differences generated by COLLECTOR-SAHAB?**
> The manual analysis of 30 code diffs in the benchmark RW3 shows that 100% of augmented diffs are correct. The randomness of the detected unique state values is correctly determined in 100% as well. For 93% (28/30) of the code diffs, the reported runtime difference is easily understandable. Overall, based on our experiments, we conclude that COLLECTOR-SAHAB is a reliable and useful tool to understand code changes, beyond purely syntactic code diffs.

Figure 3: Augmented diff for apache/iotdb@a3559e5, accurately pointing to the first unique relevant state value in the runtime data.

Figure 4: The augmented diff for SchemaCrawler@07c7368 shows a runtime diff that is hard to understand, due to reporting a static field value for a null object.

5.4 COLLECTOR-SAHAB User Study (RQ4)

After applying the filters mentioned in Section 4.7, we end up with eight commits shown in Table 5. It shows, per commit, the tools that detects a runtime difference. Note that COLLECTOR-SAHAB detects a runtime difference for all commits, while DIDIFFFF fails to do this for C4 and C8. The last column of the table shows the group that the commit is assigned to. As explained in Section 4.7, the commits are split into two groups, one group for participants from industry and one group for students. This gives us four commits for each group.

Table 6 summarizes the results of the user study. As shown in the table, P1-P4 are the participants from industry and P5-P8 are the student participants. For each participant, the table shows the average score that the participant gives to the reports of each tool per each considered criterion. As explained in Section 4.7, three criteria are considered: "Usefulness", "Clarity", and "Novelty". For each criteria, the "CS" column represents the scores given to COLLECTOR-SAHAB and "DD" represents the scores given to DIDIFFFF. The last column on the right shows the participant preference overall per the discussion at the end of the study. Note that the average score per participant is computed for the scores for which both COLLECTOR-SAHAB and DIDIFFFF give an output: C1-C3 for participants from industry and C5-C7 for students. This means the tools are tested on the same set of commits per participant group and a comparison

Table 5: Commits considered for the RQ4 experiment, their augmenting tool(s), and assigned group.

ID	Commit	COLLECTOR-SAHAB	DIDIFFFF	Group
C1	cdk@d500be0	✓	✓	Industry
C2	iotdb@5faa9da	✓	✓	Industry
C3	SchemaCrawler@55a55de	✓	✓	Industry
C4	jhipster-lite@9162791	✓	✗	Industry
C5	SchemaCrawler@07c7368	✓	✓	Students
C6	iotdb@e5e4f17	✓	✓	Students
C7	SchemaCrawler@f950519	✓	✓	Students
C8	jhipster-lite@f5b592a	✓	✗	Students

Table 6: Scores given by each participant to augmented diffs generated by COLLECTOR-SAHAB and DIDIFFFF. "CS" and "DD" represent COLLECTOR-SAHAB and DIDIFFFF, respectively.

	Particpant ID	Usefulness CS	DD	Clarity CS	DD	Novelty CS	DD	Preference
Industry	P1	3.0	2.3	3.3	3.0	4.3	4.0	COLLECTOR-SAHAB
	P2	4.3	2.0	5.0	1.6	4.3	5.0	COLLECTOR-SAHAB
	P3	3.6	2.3	3.6	2.6	2.0	2.3	COLLECTOR-SAHAB
	P4	2.6	1.3	3.3	1.0	4.3	1.0	COLLECTOR-SAHAB
	median	**3.3**	2.3	**3.3**	2.1	**4.3**	3.1	COLLECTOR-SAHAB
Students	P5	2.0	2.6	3.3	3.0	4.0	3.6	COLLECTOR-SAHAB
	P6	3.0	3.3	2.6	3.6	3.6	2.6	DIDIFFFF
	P7	2.6	3.3	3.6	3.3	3.6	4.0	COLLECTOR-SAHAB
	P8	2.0	3.6	3.0	4.3	2.0	3.0	DIDIFFFF
	median	2.3	**3.3**	3.1	**3.4**	3.6	**3.8**	–

between their scores is fair.

For example, consider participant P1. This participant has given COLLECTOR-SAHAB average scores of 3.0, 3.3, and 4.3 for usefulness, clarity, and novelty, respectively. All these scores are greater than or equal to 3, which means this participant's view about COLLECTOR-SAHAB is more positive than negative with respect to all criteria. The average scores given to DIDIFFFF by this participant are 2.3, 3, and 4.0 for usefulness, clarity, and novelty, respectively. This means this participant prefers COLLECTOR-SAHAB in terms of all criteria, as they have given a higher average score to COLLECTOR-SAHAB for each of them. Finally, P1 prefers COLLECTOR-SAHAB overall as shown in the last column.

All the other median scores for COLLECTOR-SAHAB are above 3, except for median usefulness score given by students. This indicates the positive attitude of participants about the tool. Regarding the median usefulness score by students, we ask the participants what can be improved to make the augmentation more useful. Per our discussion, we notice that for two of the commits (C5 and C7) they also want to see the type of the variable that is taking a unique value in one version. This is a concrete suggestion for future improvement of our tool.

We also note that median score that industry participants give to COLLECTOR-SAHAB is at least one point higher that the median score they give to DIDIFFFF in terms of each criteria. This indicates that people from industry find COLLECTOR-SAHAB more helpful. Per our discussion with these participants (P1-P4), we realize that they think using COLLECTOR-SAHAB is more practical, as the data it represents is more concise and its UI is integrated into GitHub. In contrast, DIDIFFFF presents all values for all variables, which takes too long to understand. For the same reasons, all participants from industry prefer COLLECTOR-SAHAB as

mentioned in the last column of Table 6.

Students give a higher score to DIDIFFFF in term of usefulness. As mentioned above, the reason is that for two of the analyzed commits they want to see the type of the variable with unique value as well. However, in terms of clarity and novelty, they give the tools very close scores. The close score for the two tools is also reflected in the last column for students, as two people prefer COLLECTOR-SAHAB (P5 and P7), and two prefer DIDIFFFF (P6 and P8).

In the live discussion after the scoring task, we ask the participants what they think about the idea of augmenting code diffs with runtime data. All participants have a very positive view on this idea. The positive view comes from the fact that these tools can "ease the painful process of code review" as P4 says, and "help [us] catch unexpected behavior" as P2 states. This all indicates there is a promising future for code diff augmentation with runtime data.

We also ask the participants to score the augmented diff for commits that only COLLECTOR-SAHAB can detect its runtime difference: C4 and C8. As a result, participants from industry give a median score of 3, 5, and 2 to C4 in terms of usefulness, clarity, and novelty, respectively. Students also give a median score of 4, 4.5, and 2 to C8 in terms of usefulness, clarity, and novelty, respectively. The novelty score is low for these two commits because these commits are adding/modifying literal strings, which means the runtime value for the invovled objects can be seen just by looking at the static code diff.

```
1 void print100() {
2   for (int i = 1; i <= 100; i++) {
3   for (int i = 100; i >= 1; i--) {
4     System.out.println(i);
5   }
6 }
```

Listing 4: A code diff that does not produce any new program state.

Answer to RQ4: How do software developers assess COLLECTOR-SAHAB's augmented diffs?

We have performed a user study involving four participants from industry and four students, who all scored COLLECTOR-SAHAB and DIDIFFFF outputs. Six of the eight participants prefer COLLECTOR-SAHAB and seven participants have a positive view about code diff augmentation with runtime data.

---

6 DISCUSSION

6.1 External Validity

Effectiveness on different targets: In this paper, we focus on code changes that only modify one method of a Java program. Accordingly, as explained in Section 4.4, we evaluate the effectiveness on 584 single-method code diffs from MVN-DRR. This is a threat to the external validity of this study, as it means the effectiveness results may not be generalizable to bigger changes in Java programs. Beyond size, we emphasize that we carefully build a diverse dataset, namely RW3, for the RQ2 and RQ3 experiments, from 10 various sizable Java projects. The satisfactory performance of COLLECTOR-SAHAB on this dataset indicates the arguably wide applicability of COLLECTOR-SAHAB on complex Java programs.

6.2 Internal Validity

Certainty of manual analysis results: As described in Section 4.5, we conduct a manual analysis to assess the usefulness of COLLECTOR-SAHAB augmented diffs. This manual analysis may lead to results that are subjective or uncertain. To address this, two participants with expertise in Java programming both perform the analysis and compare their assessments. In case of discrepancies, they have a thorough discussion about the results until they make a final conclusion. This protocol is meant to mitigate the threats created by the subjective nature of manual analysis.

6.3 Code Changes with No Augmentation

There are two types of code diffs that are never augmented by COLLECTOR-SAHAB. First, code diffs that do not cause any behavioral changes on matched lines. For example, refactoring the program by renaming a variable does not cause runtime differences. Second, there are code diffs that may cause some behavioral changes but do not produce any new state values. This happens when the code diff only changes the order in which state values appear, while COLLECTOR-SAHAB works with set-based abstractions.

For example, Listing 4 shows the `print100()` method that runs a for loop and prints the natural numbers less than 100. The code diff in this example reverses the for loop iteration. In this original version (line 2), `i` goes from 1 to 100 and in the patched version (line 3), it goes from 100 to 1. Consequently, in the patched version, line 4 prints the numbers in reverse order. Even though the patch changes the behavior of the program as explained, the set of state values at the matched (line 4) is the same, they just appear in a different order. Therefore, COLLECTOR-SAHAB does not catch the runtime difference in this code diff.

The experiments in this study show that many of the important code diffs, such as bug-fixes, produce runtime differences going beyond value ordering, and are consequently augmented by COLLECTOR-SAHAB.

6.4 Diff in Runtime Generated Code

In Java, there are different options to generate and add code to the program at runtime, such as dependency injection, annotation processing and metaprogramming [28]. In COLLECTOR-SAHAB, the code added at runtime is not considered, this is because COLLECTOR-SAHAB adds the breakpoints before the execution, by statically matching source code lines and comparing the two versions, as explained in Section 3.3. Therefore, COLLECTOR-SAHAB considers a runtime difference caused by the code added at runtime only if the difference is reflected in the state values in matched lines, later in the execution. As future engineering work, one can improve COLLECTOR-SAHAB by considering the code added at runtime, with sophisticated runtime analysis. This however poses UI challenges with respect to presenting the information to the developers in the diff.

6.5 Scalability of COLLECTOR-SAHAB

In the experiments of this paper, we use COLLECTOR-SAHAB to augment code diffs that change a single method in different programs. In answer to RQ2, we see this already applies to a significant number of real-world commits. However, to make COLLECTOR-SAHAB as applicable as possible and run it on all types of commits, incl. commits that change many different methods, future research will study scalability in terms of time and memory resources it needs. As shown in answer to RQ1 and RQ2, facing time limits is very rare for COLLECTOR-SAHAB. This is because it just adds an extra computation statement at each line to extract state values, which does not change the order or magnitude of the execution time complexity. However, there are a notable number of code changes studied in RQ1 and RQ2 on which COLLECTOR-SAHAB fails due to exceeding the memory limit. In this regard, we first note that the state depth configuration provided by COLLECTOR-SAHAB helps users configure the tool to get the most detailed trace, while avoiding memory failure. Nevertheless, our careful inspection of the experiments reveal that we can avoid memory failures even with a very high state depth by making one technical improvement: COLLECTOR-SAHAB could write and read the execution trace in small batches. This engineering improvement would enable COLLECTOR-SAHAB to compute runtime differences without saving the whole trace in RAM.

6.6 Test Suite Inputs

As mentioned in Section 3.2, COLLECTOR-SAHAB needs a differentiating test as input to be able to detect runtime differences. This means COLLECTOR-SAHAB does not augment commits with low quality test suites. For example, in our search for real-wolrd commits in the RW3 dataset, we see that we collect 212 commits and 26% (56/212) of them satisfy all the requirements except having a test that covers the code diff. Moreover, in Section 5.2, 18 commits have tests that covers the code diff but none of them shows a runtime behavioral difference. This entails that a significant number of real-world commits do not have a strong enough test suite that includes a differentiating test. A number of tools have been proposed to amplify test suites and generate such differentiating tests [29]. In the future, COLLECTOR-SAHAB can take advantage of such tools to find a runtime difference for commits without a strong test suite.

7 RELATED WORK

As explained in Section 2, previous studies consider different segments of the execution data as part of the program state, from all variable values [11], [12] or some of the variable values [15], [16], [17], [18], [19], [20] to the call stack and program counter [13].

In the following, we first review studies on collecting execution data and using this data for detecting the behavioral difference between two versions of a program. Next, we look at works focused on improving code diffs by augmenting them with execution data.

7.1 Collecting Execution Data

Lewis [30] defines the concept of omniscient debugger as a tool that "works by collecting events at every state change and every method call in a program". Different tools have been created that fully or partially implement this concept of omniscient debuggers [31]. The state-of-the-art omniscient debugger is introduced by Shimari et al. [26], [32]. They introduce SELOGGER, which records events, such as method entry, method exit, return values, and reading or writing a value into a variable. In contrast with COLLECTOR-SAHAB, SELOGGER does not record values that are stored deep inside visible program objects. This means, compared to SELOGGER, the logging mechanism of COLLECTOR-SAHAB provides more information for identifying execution differences.

Previous studies have also proposed execution trace collection tools for other programming languages, such as C/C++. Millnert et al. [33] propose DMCE which adds probe statements after changed lines of C/C++ programs. After executing the code diff, they check which probe statements are executed. Based on the probe statement execution data, they determine which parts of the code diff are covered.

Magalhaes et al. [34] introduce WHIRO, which collects program state information in C/C++ programs. WHIRO is built based on LLVM compiler and instruments the program at the compiler intermediate representation level. This tool can be customized to work at different levels of granularity. It can collect data allocated statically, in the stack or in the heap. Also, it can be set up to save the graph of relations between pointers up to a specified depth. Finally, it can be set up to collect data at different points: only before the main return, only before function returns, or after each data store statement. What makes Whiro novel and useful is how it collects data in an uncooperative environment, like in C/C++, where the program data does not have a datatype attached to it. It has to find the relation between low-level data at the heap and high-level program variables to do so. In contrast with DMCE and WHIRO, COLLECTOR-SAHAB collects execution data for Java programs.

Orton and Mycroft propose REHYPE which collects the executed methods of a Java program [35]. Based on the collected data they output potential inefficient code in the program. They also introduce SCOPDA which uses REHYPE output and suggests concrete improvements to the program [36]. These two tools work with execution traces of Java programs, however their goal is different from COLLECTOR-SAHAB: COLLECTOR-SAHAB uses execution differences to explain code diffs, while SCOPDA analyzes execution trace to improve the efficiency of programs.

7.2 Execution Differencing

Execution logs are one of the sources that can be used to model the execution behavior [37]. Goldstein et al. [38] present an algorithm to visualize behavioral differences using execution logs in four steps. First, they remove noisy lines of the logs that do not play an important role related to the behavior. Second, they create a finite state automaton (FSA) for each original/patched log using the KTAILS algorithm [39]. Third, they find the differences between FSAs and finally, visualize the difference. This algorithm has been used in real-world setups in later studies [40], [41]. These studies show the effectiveness of Goldstein et al.'s algorithm in helping developers identify program changes. The main difference between COLLECTOR-SAHAB and these works is that in contrast with COLLECTOR-SAHAB, they do not consider state values to detect execution differences. This makes COLLECTOR-SAHAB able to identify more fine-grained differences compared to these log-based differencing approaches.

Test suites may not execute programs exactly the same as they are executed in the field [42], [43]. Nevertheless, it is common to utilize the traces produced during the execution of test suites to understand program changes [42], [44]. Some researchers have introduced novel methods to generate tests that can differentiate between two versions of a program. Danglot et al. [45] propose DCI. This approach amplifies existing tests to generate new tests that can detect behavioral changes in the continuous integration pipeline. Johnson et al. [46] also use EVOSUITE [47] to generate tests for a given program. Then, they find two similar tests that one of them passes and the other one fails on the program. These tests are utilized to detect the root causes of defects in a program. These works are complementary to COLLECTOR-SAHAB. They can create tests that can be used as the input covering test in COLLECTOR-SAHAB.

Many studies compare execution traces to perform different tasks, such as fault localization [48], [49] and failure root cause analysis [24], [50]. The most advanced techniques

for comparing execution traces are TRACESIM [51] and S3M [52]. These two works compare stack traces that are collected as ordered lists of lines that are executed before a program crash. TRACESIM performs the comparison by computing measures such as edit distance and tf-idf between traces and their elements. On the other hand, S3M uses deep learning to determine the similarity between two traces. For this, they present traces as vectors and then utilize the siamese architecture [53] to compute the similarity between them. All these works take the list of executed statements as the execution trace, while COLLECTOR-SAHAB also considers state values. This can make COLLECTOR-SAHAB more precise in detecting differences in executions.

There are also several studies that compare variable values to detect what causes differences between two executions [11], [54], [55]. Cleve and Zeller [55] use *delta debugging* [54], [56] to narrow down differences between program states in two executions of a program. The result is a small set of variable value differences that make one execution of a program pass and another execution fail. Other researchers improve the precision of delta debugging by mixing it with other techniques [50], such as execution coverage [57], observation-based slicing [58], dual slicing [59], and hierarchical information [60]. COLLECTOR-SAHAB is different from these tools mainly because it is focused on improving the code diff by adding a concise execution difference between two versions of the program. On the contrary, these tools are designed to identify the root cause of a failure.

Abramson et al. have introduced *relative debugging* [61], [62], [63], [64], [65], [66]. Relative debugging is implemented in their tool, called GUARD. GUARD runs two programs and reports the differences between the values assigned to their variables at runtime. This is close to what COLLECTOR-SAHAB does, however GUARD requires the users to specify the data structures, the variables, the breakpoints, and the mappings between the variables of two versions. In COLLECTOR-SAHAB, these steps are all done automatically. Consequently, COLLECTOR-SAHAB is able to generate one augmentation for a given code diff, while GUARD acts similar to a traditional debugger, in which the user interacts with the tool to collect runtime differences and explore them.

The tool that is most similar to COLLECTOR-SAHAB is DIDIFFFF [10]. As explained in Section 4.3, DIDIFFFF lists values assigned to each variable during execution and compares the lists for two versions of a program. DIDIFFFF is built on top of SELOGGER. There are two main differences between COLLECTOR-SAHAB and DIDIFFFF. COLLECTOR-SAHAB explores the data inside complex Java objects, while DIDIFFFF only considers values of primitive variables. Also, COLLECTOR-SAHAB only reports a concise execution difference in its UI output, while DIDIFFFF reports all the extracted differences.

7.3 Improving Code Diffs

Previous studies have shown how integrating runtime data into the developers' working environment helps them better perform tasks, such as bug fixing and code comprehension [67], [68]. In this section, we review studies focused on improving code diffs and making them more understandable. This can be done in different ways, such as splitting

code changes into smaller sets of related changes [69] or by improving how code changes are presented. Our focus is on the diff presentation.

Some tools improve the code diff presentation by making it more fine-grained, such as MERGELY [70] and GUMTREE [71]. For this, MERGELY considers changes at the level of code elements inside lines, instead of looking at whole line differences. GUMTREE enhances existing AST differencing algorithms, such as CHANGEDISTILLER [72], and proposes an algorithm that computes minimum changes that should be made on the original AST to reach the patched one. GUMTREE marks these fine-grained AST level changes on the code diff to help developers easily see what is the exact change to the program.

Decker et al. [73] introduce a new syntactic differencing tool, called SRCDIFF. This tool is built on top of SRCML which represents source code in an XML format and annotates the code with the syntax information. Based on manual and statistical analysis, they propose several heuristics to determine if a change is actually a modification or a complete replacement. Hence, in contrast with GUMTREE, the goal of SRCDIFF is not producing an optimal diff. SRCDIFF intends to produce a diff that is more similar to the changes performed by developers. The difference between COLLECTOR-SAHAB and tools like GUMTREE and SRCDIFF is that COLLECTOR-SAHAB adds execution data to the diff, while they solely use static data to improve the code change presentation.

Another way of improving code diff presentation is based on the idea of literate programming. literate-diff-viewer [74] is a tool where the description of the source code changes is provided in natural language and then the diff generated shows how the application behavior varies, together with the description. Applied to video game diffs, one see how the games evolve as we scroll through the diff description.

In another line of code, researchers try to integrate new information into the integrated development environments (IDE) [75]. This can include putting code change information with data related to developers' actions in the IDE [76] or production data [77]. For example, Cito et al. [78] add the time spent to run a method at production to the IDE code editor. When the code is changed, they predict the new production time for the changed code and show it to the developer in IDE. As a result, developers better detect performance problems in their code changes. These works are different from COLLECTOR-SAHAB as they do not add program state data to the code diff.

Bohnet et al. [79] combine execution traces with code changes to help developers identify the root cause of a failure in C/C++ programs. They define the execution trace as the sequence of executed functions. When a program faces a failure after a code change, the execution trace is collected and shown to the developer. Next, the developer interacts with the tool to detect the failure causing behavioral change. This work is different from COLLECTOR-SAHAB as it does not consider state values as a part of the execution trace. In contrast with COLLECTOR-SAHAB, it is also a semi-automatic tool requiring interaction with the developer.

Table 7 summarizes the closest related work to COLLECTOR-SAHAB. COLLECTOR-SAHAB is the first open-

Table 7: Collector-Sahab compared to the related work.

Tool	Language	Collected Runtime Data	Computed Runtime Diff	Presentation	Open-source
SELogger [32]	Java	Variable & Return Values	–	Log File	✓
DMCE [33]	C/C++	Covered Lines	–	Log File	✓
Whiro [34]	C/C++	Heap Data	–	Graphical Visualization	✓
VISSOFT [40]	–	–	Detects Outlier Logs	Graphical Visualization	✓
DCI [45]	Java	Covered Lines	Coverage Difference	Amplified Test	✓
TraceSim [51]	–	Stack Trace	Stack Trace Similarity	Similarity Score	✓
DeltaDebugging [56]	C/C++	Variable Values	Failure-inducing Changes	Text	✓
Guard [61]	C/Fortran	Variable Values	Variable Value Difference	Text & Visualization	✗
didifff [10]	Java	Variable Values	Variable Value Difference	Ad hoc GUI	✓
Collector-Sahab (this paper)	Java	State & Return Values	State & Return Value Diff	Augmented Code Review Diff	✓

source execution differencing tool for Java that explores deep inside the value of all non-primitive types and computes state values. Collector-Sahab detects and reports a unique state value in the original or patched version of a program to help developers understand the changed behavior.

8 Conclusion

In this paper, we introduce Collector-Sahab, a novel tool that detects runtime differences between two versions of a program and augments the code diff with it. The augmentation is done with unique variable and return values that occur during the execution of a test case. This new information is important for developers to better understand behavioral changes caused by a code change. We demonstrate the effectiveness of Collector-Sahab by running it on 584 code diffs for Defects4J bugs. Collector-Sahab produces an execution difference for 95% (555/584) of the code diffs, better than the closest related work. The high quality of Collector-Sahab's augmented diffs is also confirmed by a manual analysis of 30 commits from real-world open-source projects. Our user study shows the participants find the idea of code diff augmentation with runtime useful for software development process.

In the future, research is needed to detect and mitigate spurious program state differences. Such research relates to the currently active research field on flaky tests. Moreover, while this work focuses on bug fixes, the augmentation done by Collector-Sahab could also help explain other types of changes, e.g., refactorings. Additional studies are required to investigate this potential application. Finally, Collector-Sahab could be integrated into a fully-automated pipeline, e.g., continuous integration, so that developers receive its augmented code diff directly after proposing changes for code review.

9 Acknowledgments

References

[1] C. Sadowski, E. Söderberg, L. Church, M. Sipko, and A. Bacchelli, "Modern code review: a case study at google," in *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Practice*, 2018, pp. 181–190.

[2] S. Jiang, A. Armaly, and C. McMillan, "Automatically generating commit messages from diffs using neural machine translation," in *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*. IEEE Press, 2017, pp. 135–146.

[3] Z. Liu, X. Xia, C. Treude, D. Lo, and S. Li, "Automatic generation of pull request descriptions," in *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2019, pp. 176–188.

[4] Z. Liu, X. Xia, A. E. Hassan, D. Lo, Z. Xing, and X. Wang, "Neural-machine-translation-based commit message generation: how far are we?" in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 373–384.

[5] R. Dyer, H. A. Nguyen, H. Rajan, and T. N. Nguyen, "Boa: A language and infrastructure for analyzing ultra-large-scale software repositories," in *2013 35th International Conference on Software Engineering (ICSE)*. IEEE, 2013, pp. 422–431.

[6] P. Liu, L. Li, Y. Yan, M. Fazzini, and J. Grundy, "Identifying and characterizing silently-evolved methods in the android api," in *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 2021, pp. 308–317.

[7] S. Liu, C. Gao, S. Chen, N. L. Yiu, and Y. Liu, "Atom: Commit message generation based on abstract syntax tree and hybrid ranking," *IEEE Transactions on Software Engineering*, 2020.

[8] J. Galenson, P. Reames, R. Bodik, B. Hartmann, and K. Sen, "Codehint: Dynamic and interactive synthesis of code snippets," in *Proceedings of the 36th International Conference on Software Engineering*, 2014, pp. 653–663.

[9] R. Just, D. Jalali, and M. D. Ernst, "Defects4j: A database of existing faults to enable controlled testing studies for java programs," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, 2014, pp. 437–440.

[10] T. Kanda, K. Shimari, and K. Inoue, "didiffff: A viewer for comparing changes in both code and execution traces," in *2022 IEEE/ACM 30th International Conference on Program Comprehension (ICPC)*. IEEE, 2022, pp. 528–532.

[11] A. Zeller, "Isolating cause-effect chains from computer programs," *ACM SIGSOFT Software Engineering Notes*, vol. 27, no. 6, pp. 1–10, 2002.

[12] L. H. Pham, J. Sun, L. L. T. Thi, J. Wang, and X. Peng, "Learning likely invariants to explain why a program fails," in *2017 22nd International Conference on Engineering of Complex Computer Systems (ICECCS)*. IEEE, 2017, pp. 70–79.

[13] G. Xu, A. Rountev, Y. Tang, and F. Qin, "Efficient checkpointing of java software using context-sensitive capture and replay," in *Proceedings of the the 6th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering*, 2007, pp. 85–94.

[14] J. Meinicke, C. Wong, C. Kästner, and G. Saake, "Understanding differences among executions with variational traces," *CoRR*, vol. abs/1807.03837, 2018. [Online]. Available: http://arxiv.org/abs/1807.03837

[15] T.-Y. Huang, P.-C. Chou, C.-H. Tsai, and H.-A. Chen, "Automated fault localization with statistically suspicious program states," *Acm Sigplan Notices*, vol. 42, no. 7, pp. 11–20, 2007.

[16] R. Gupta, "Deep learning for bug localization and program repair," Ph.D. dissertation, PhD thesis, 2020.

[17] M. Schäf, D. Schwartz-Narbonne, and T. Wies, "Explaining inconsistent code," in *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering*, 2013, pp. 521–531.

[18] T. Xie and D. Notkin, "Checking inside the black box: Regression testing based on value spectra differences," in *20th IEEE International Conference on Software Maintenance, 2004. Proceedings.* IEEE, 2004, pp. 28–37.

[19] R. Abou Assi, W. Masri, and C. Trad, "Substate profiling for effective test suite reduction," in *2018 IEEE 29th International Symposium on Software Reliability Engineering (ISSRE).* IEEE, 2018, pp. 123–134.

[20] B. Mehne, H. Yoshida, M. R. Prasad, K. Sen, D. Gopinath, and S. Khurshid, "Accelerating search-based program repair," in *2018 IEEE 11th International Conference on Software Testing, Verification and Validation (ICST).* IEEE, 2018, pp. 227–238.

[21] C. Zhang, J. Yang, D. Yan, S. Yang, and Y. Chen, "Automated breakpoint generation for debugging." *J. Softw.*, vol. 8, no. 3, pp. 603–616, 2013.

[22] R. Santelices, Y. Zhang, S. Jiang, H. Cai, and Y.-j. Zhang, "Quantitative program slicing: Separating statements by relevance," in *2013 35th International Conference on Software Engineering (ICSE).* IEEE, 2013, pp. 1269–1272.

[23] E. W. Myers, "AnO(ND) difference algorithm and its variations," *Algorithmica*, vol. 1, no. 1-4, pp. 251–266, Nov. 1986.

[24] Y. Zhang, K. Rodrigues, Y. Luo, M. Stumm, and D. Yuan, "The inflection point hypothesis: a principled debugging approach for locating the root cause of a failure," in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, 2019, pp. 131–146.

[25] H. Ye, M. Martinez, and M. Monperrus, "Automated patch assessment for program repair at scale," *Empirical Software Engineering*, vol. 26, 2021.

[26] K. Shimari, T. Ishio, T. Kanda, and K. Inoue, "Near-omniscient debugging for java using size-limited execution trace," in *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME).* IEEE, 2019, pp. 398–401.

[27] E. T. Barr and M. Marron, "Tardis: Affordable time-travel debugging in managed runtimes," *ACM SIGPLAN Notices*, vol. 49, no. 10, pp. 67–82, 2014.

[28] R. Pawlak, M. Monperrus, N. Petitprez, C. Noguera, and L. Seinturier, "Spoon: A Library for Implementing Analyses and Transformations of Java Source Code," *Software: Practice and Experience*, vol. 46, pp. 1155–1179, 2015. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01078532/document

[29] B. Danglot, O. L. Vera-Pérez, B. Baudry, and M. Monperrus, "Automatic test improvement with dspot: a study with ten mature open-source projects," *Empirical Software Engineering*, vol. 24, pp. 2603–2635, 2019.

[30] B. Lewis, "Debugging backwards in time," *arXiv preprint cs/0310016*, 2003.

[31] T. Matsumura, T. Ishio, Y. Kashima, and K. Inoue, "Repeatedly-executed-method viewer for efficient visualization of execution paths and states in java," in *Proceedings of the 22nd International Conference on Program Comprehension*, 2014, pp. 253–257.

[32] K. Shimari, T. Ishio, T. Kanda, N. Ishida, and K. Inoue, "Nod4j: Near-omniscient debugging tool for java using size-limited execution trace," *Science of Computer Programming*, vol. 206, p. 102630, 2021.

[33] V. Millnert, M. Templing, and P. Åberg, "Dmfe: did my function execute?" in *Proceedings of the 14th IEEE/ACM International Conference on Utility and Cloud Computing*, 2021, pp. 1–10.

[34] J. W. d. S. Magalhães, C. Liao, and F. M. Q. Pereira, "Automatic inspection of program state in an uncooperative environment," *Software: Practice and Experience*, vol. 52, no. 12, pp. 2727–2758, 2022.

[35] I. Orton and A. Mycroft, "Refactoring traces to identify concurrency improvements," in *Proceedings of the 23rd ACM International Workshop on Formal Techniques for Java-like Programs*, 2021, pp. 16–23.

[36] ——, "Source code patches from dynamic analysis," in *Proceedings of the 23rd ACM International Workshop on Formal Techniques for Java-like Programs*, 2021, pp. 1–8.

[37] M. Gupta, A. Mandal, G. Dasgupta, and A. Serebrenik, "Runtime monitoring in continuous deployment by differencing execution behavior model," in *International Conference on Service-Oriented Computing.* Springer, 2018, pp. 812–827.

[38] M. Goldstein, D. Raz, and I. Segall, "Experience report: Log-based behavioral differencing," in *2017 IEEE 28th International Symposium on Software Reliability Engineering (ISSRE).* IEEE, 2017, pp. 282–293.

[39] A. W. Biermann and J. A. Feldman, "On the synthesis of finite-state machines from samples of their behavior," *IEEE transactions on Computers*, vol. 100, no. 6, pp. 592–597, 1972.

[40] C. Deknop, K. Mens, A. Bergel, J. Fabry, and V. Zaytsev, "A scalable log differencing visualisation applied to cobol refactoring," in *2021 Working Conference on Software Visualization (VISSOFT).* IEEE, 2021, pp. 1–11.

[41] C. Deknop, "Log-based behavioural differencing applied to an industrial case," in *THE 1ST EARLY CAREER RESEARCHERS WORKSHOP COLLOCATED WITH ECCS 2021*, 2021, p. 72.

[42] Q. Wang, Y. Brun, and A. Orso, "Behavioral execution comparison: Are tests representative of field behavior?" in *2017 IEEE International Conference on Software Testing, Verification and Validation (ICST).* IEEE, 2017, pp. 321–332.

[43] J. Cabrera-Arteaga, M. Monperrus, and B. Baudry, "Scalable comparison of javascript v8 bytecode traces," in *Proceedings of the SPLASH workshop on Virtual Machines and Language Implementations*, 2019. [Online]. Available: http://arxiv.org/pdf/1910.03478

[44] K. D. Sherwood and G. C. Murphy, "Reducing code navigation effort with differential code coverage," *Department of Computer Science, University of British Columbia, Tech. Rep*, 2008.

[45] B. Danglot, M. Monperrus, W. Rudametkin, and B. Baudry, "An approach and benchmark to detect behavioral changes of commits in continuous integration," *Empirical Software Engineering*, vol. 25, no. 4, pp. 2379–2415, 2020.

[46] B. Johnson, Y. Brun, and A. Meliou, "Causal testing: understanding defects' root causes," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 87–99.

[47] G. Fraser and A. Arcuri, "Whole test suite generation," *IEEE Transactions on Software Engineering*, vol. 39, no. 2, pp. 276–291, 2012.

[48] L. Guo, A. Roychoudhury, and T. Wang, "Accurately choosing execution runs for software fault localization," in *International Conference on Compiler Construction.* Springer, 2006, pp. 80–95.

[49] H. Wang, Y. Lin, Z. Yang, J. Sun, Y. Liu, J. Dong, Q. Zheng, and T. Liu, "Explaining regressions via alignment slicing and mending," *IEEE Transactions on Software Engineering*, vol. 47, no. 11, pp. 2421–2437, 2019.

[50] Q. Yi, Z. Yang, J. Liu, C. Zhao, and C. Wang, "A synergistic analysis method for explaining failed regression tests," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 1. IEEE, 2015, pp. 257–267.

[51] I. M. Rodrigues, A. Khvorov, D. Aloise, R. Vasiliev, D. Koznov, E. R. Fernandes, G. Chernishev, D. Luciv, and N. Povarov, "Tracesim: An alignment method for computing stack trace similarity," *Empirical Software Engineering*, vol. 27, no. 2, pp. 1–41, 2022.

[52] A. Khvorov, R. Vasiliev, G. Chernishev, I. M. Rodrigues, D. Koznov, and N. Povarov, "S3m: Siamese stack (trace) similarity measure," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR).* IEEE, 2021, pp. 266–270.

[53] S. Chopra, R. Hadsell, and Y. LeCun, "Learning a similarity metric discriminatively, with application to face verification," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 539–546.

[54] A. Zeller and R. Hildebrandt, "Simplifying and isolating failure-inducing input," *IEEE Transactions on Software Engineering*, vol. 28, no. 2, pp. 183–200, 2002.

[55] H. Cleve and A. Zeller, "Locating causes of program failures," in *Proceedings. 27th International Conference on Software Engineering, 2005. ICSE 2005.* IEEE, 2005, pp. 342–351.

[56] A. Zeller, "Yesterday, my program worked. today, it does not. why?" *ACM SIGSOFT Software engineering notes*, vol. 24, no. 6, pp. 253–267, 1999.

[57] K. Yu, M. Lin, J. Chen, and X. Zhang, "Practical isolation of failure-inducing changes for debugging regression faults," in *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*, 2012, pp. 20–29.

[58] D. Binkley, N. Gold, M. Harman, J. Krinke, and S. Yoo, "Observation-based slicing," *RN*, vol. 13, no. 13, p. 13, 2013.

[59] W. N. Sumner and X. Zhang, "Comparative causality: Explaining the differences between executions," in *2013 35th International Conference on Software Engineering (ICSE).* IEEE, 2013, pp. 272–281.

[60] G. Misherghi and Z. Su, "Hdd: hierarchical delta debugging," in *Proceedings of the 28th international conference on Software engineering*, 2006, pp. 142–151.

[61] R. Sosič and D. Abramson, "Guard: A relative debugger," *Software: Practice and Experience*, vol. 27, no. 2, pp. 185–206, 1997.

[62] D. Abramson and R. Sosic, "A debugging and testing tool for supporting software evolution," *Automated Software Engineering*, vol. 3, pp. 369–390, 1996.

[63] D. Abramson, M. N. Dinh, D. Kurniawan, B. Moench, and L. DeRose, "Data centric highly parallel debugging," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, 2010, pp. 119–129.

[64] M. N. Dinh, D. Abramson, and C. Jin, "Scalable relative debugging," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 3, pp. 740–749, 2013.

[65] D. Abramson, C. Chu, D. Kurniawan, and A. Searle, "Relative debugging in an integrated development environment," *Software: Practice and Experience*, vol. 39, no. 14, pp. 1157–1183, 2009.

[66] A. Searle, J. Gough, and D. Abramson, "Duct: an interactive define-use chain navigation tool for relative debugging," *arXiv preprint cs/0311037*, 2003.

[67] D. Rothlisberger, M. Harry, W. Binder, P. Moret, D. Ansaloni, A. Villazón, and O. Nierstrasz, "Exploiting dynamic information in ides improves speed and correctness of software maintenance tasks," *IEEE Transactions on Software Engineering*, vol. 38, no. 3, pp. 579–591, 2012.

[68] X. Liu and R. Holmes, "Exploring developer preferences for visualizing external information within source code editors," in *2020 Working Conference on Software Visualization (VISSOFT)*. IEEE, 2020, pp. 27–37.

[69] M. Dias, "Supporting software integration activities with first-class code changes," Ph.D. dissertation, Laboratoire d'Informatique Fondamentale de Lille, 2015.

[70] Mergely. (2022) Diff text documents online with mergely, an editor and html5 javascript library. [Online]. Available: https://www.mergely.com/

[71] J.-R. Falleri, F. Morandat, X. Blanc, M. Martinez, and M. Monperrus, "Fine-grained and accurate source code differencing," in *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*, 2014, pp. 313–324.

[72] B. Fluri, M. Wursch, M. PInzger, and H. Gall, "Change distilling: Tree differencing for fine-grained source code change extraction," *IEEE Transactions on software engineering*, vol. 33, no. 11, pp. 725–743, 2007.

[73] M. J. Decker, M. L. Collard, L. G. Volkert, and J. I. Maletic, "srcdiff: A syntactic differencing approach to improve the understandability of deltas," *Journal of Software: Evolution and Process*, vol. 32, no. 4, p. e2226, 2020.

[74] Abagames, "Abagames/literate-diff-viewer: Show source code description with literate diff way." [Online]. Available: https://github.com/abagames/literate-diff-viewer

[75] M. Sulír, M. Bačíková, S. Chodarev, and J. Porubän, "Visual augmentation of source code editors: A systematic review," *Computer Languages, Systems & Structures*, 2018.

[76] S. Proksch, S. Nadi, S. Amann, and M. Mezini, "Enriching in-ide process information with fine-grained source code history," in *2017 IEEE 24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2017, pp. 250–260.

[77] J. Winter, M. Aniche, J. Cito, and A. v. Deursen, "Monitoring-aware ides," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2019, pp. 420–431.

[78] J. Cito, P. Leitner, M. Rinard, and H. C. Gall, "Interactive production performance feedback in the ide," in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 2019, pp. 971–981.

[79] J. Bohnet, S. Voigt, and J. Döllner, "Projecting code changes onto execution traces to support localization of recently introduced bugs," in *Proceedings of the 2009 ACM symposium on Applied Computing*, 2009, pp. 438–442.