

# Using Defect Prediction to Improve the Bug Detection Capability of Search-Based Software Testing

Anjana Perera  
Anjana.Perera@monash.edu  
Faculty of Information Technology  
Monash University  
Melbourne, Australia

## ABSTRACT

Automated test generators, such as search based software testing (SBST) techniques, replace the tedious and expensive task of manually writing test cases. SBST techniques are effective at generating tests with high code coverage. However, is high code coverage sufficient to maximise the number of bugs found? We argue that SBST needs to be focused to search for test cases in defective areas rather in non-defective areas of the code in order to maximise the likelihood of discovering the bugs. Defect prediction algorithms give useful information about the bug-prone areas in software. Therefore, we formulate the objective of this thesis: *Improve the bug detection capability of SBST by incorporating defect prediction information*. To achieve this, we devise two research objectives, i.e., 1) Develop a novel approach (SBST<sub>CL</sub>) that allocates time budget to classes based on the likelihood of classes being defective, and 2) Develop a novel strategy (SBST<sub>ML</sub>) to guide the underlying search algorithm (i.e., genetic algorithm) towards the defective areas in a class. Through empirical evaluation on 434 real reported bugs in the Defects4J dataset, we demonstrate that our novel approach, SBST<sub>CL</sub>, is significantly more efficient than the state of the art SBST when they are given a tight time budget in a resource constrained environment.

## ACM Reference Format:

Anjana Perera. 2020. Using Defect Prediction to Improve the Bug Detection Capability of Search-Based Software Testing. In *35th IEEE/ACM International Conference on Automated Software Engineering (ASE '20), September 21–25, 2020, Virtual Event, Australia*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3324884.3415286>

## 1 INTRODUCTION

Search based software testing (SBST) has been gaining maturity not only in research, but also in the industry [19, 21]. Harman and Jones [20] coined the emerging research area, search based software engineering (SBSE), which uses search algorithms like genetic algorithm (GA) to solve software engineering problems. SBST is a sub-area of SBSE, which specialises in solving software testing problems like test data generation. Research on SBST dates

back to 1976 [31], and since then it has been seen a growing interest by the research community [19]. At the same time, SBST has made its way to the industry as well, i.e., Sapienz at Facebook [2, 30].

Previous work shows that SBST techniques are effective at achieving high code coverage [33, 34]. However, the more important question is ‘How does SBST perform in finding real bugs?’. Both Shamshiri et al. [39] and Almasi et al. [1] showed that SBST outperformed other approaches in finding real bugs on open source and industrial projects respectively. However, overall the results of these studies suggest that SBST methods are not as effective in finding real bugs. We find that high code coverage is not sufficient to maximise the number of bugs found.

Defect prediction algorithms predict the areas in software that are likely to be buggy. They give predictions at file [8, 9, 28], class [4] or method [6, 15, 22] level. Previous work shows the effectiveness of defect predictors in locating the buggy areas in software [6, 32, 35]. An effective defect predictor can be used to guide developers to efficiently test a software system to find bugs [8].

The objective of this thesis is to improve the bug detection performance of SBST by incorporating defect prediction information. We argue that SBST needs to generate more test cases covering the defective areas in the code in order to increase the likelihood of finding bugs. Therefore, we propose to leverage defect prediction information to inform the search process of the defective areas. To accomplish our main research objective, we formulate two research objectives; 1) Develop a novel approach (SBST<sub>CL</sub>) that allocates time budget to classes based on the likelihood of classes being defective, and 2) Develop a novel strategy (SBST<sub>ML</sub>) to guide the underlying search algorithm (i.e., genetic algorithm) towards the defective areas in a class. Finally, the expected contributions of this research are 1) a time budget allocation approach that improves bug detection capability of SBST, 2) an analysis of the bug detection performance improvement of SBST by the time budget allocation approach when varying the performance of the defect predictor, 3) a defectiveness-aware search algorithm guidance strategy to improve the bug detection capability of SBST, and 4) an analysis of the bug detection performance improvement of SBST by the defectiveness-aware search algorithm guidance strategy when varying the performance of the defect predictor.

Through empirical evaluation on 434 real bugs from 6 open source java projects in the Defects4J dataset [25], we demonstrate the improved efficiency of SBST by our time budget allocation approach in a resource constrained environment. Further analysis to the results suggests the improvement of SBST<sub>CL</sub> comes from its ability to find more unique bugs which cannot be found otherwise.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ASE '20, September 21–25, 2020, Virtual Event, Australia

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6768-4/20/09...\$15.00

<https://doi.org/10.1145/3324884.3415286>

## 2 MOTIVATING EXAMPLE

The following example illustrates the limitation of SBST focusing only on high code coverage when using it for finding real bugs. Figure 1a shows a classic example of a bug due to an integer overflow. The `if` condition at line 412 is expected to evaluate to `true` only if either `u` or `v` is zero. However, if the method is called with the following inputs `u = 1073741824` and `v = 1032`, then the `if` condition at line 412 is evaluated to `true` since the multiplication of `u` and `v` causes an integer overflow to zero. Fixed code shown in Figure 1b rectifies this issue by individually checking if `u` or `v` are zero. Thus, in order to detect this bug, a test should pass non-zero arguments `u` and `v` such that their multiplication causes an integer overflow to zero. A search algorithm [34] focusing on high code coverage is more likely to find a test case with both/either `u` and/or `v = 0` and consider this code as covered than find a test case with a pair of `u` and `v` causing an integer overflow to zero. In contrast, we propose to 1) use defect prediction information to identify bug-prone areas in the code and 2) guide the search algorithm to find more test cases covering the bug-prone code to increase the likelihood of finding the bug.

|  |  |
|--|--|
| <pre> 411 public static int gcd(int u, int v) { 412   if (u * v == 0) { 413     return (Math.abs(u) + Math.abs(v)); 414   } 415   ... 416 } </pre> | <pre> 411 public static int gcd(int u, int v) { 412   if ((u == 0)    (v == 0)) { 413     return (Math.abs(u) + Math.abs(v)); 414   } 415   ... 416 } </pre> |
| (a) Buggy code   | (b) Fixed code   |

Figure 1: MathUtils class from Math-94 Bug in Defects4J

## 3 RELATED WORK

Shamshiri et al. [39] and Almasi et al. [1] evaluated EvoSuite [12], a state of the art SBST tool, against other test generation techniques on open source and industrial software projects respectively. While the results of these studies showed that EvoSuite outperformed the other techniques, overall EvoSuite is not as effective in finding real bugs. In particular, EvoSuite only found on average 23% bugs from the Defects4J dataset [39]. Both studies used only 100% branch coverage criterion for the test case search. We argue that targeting 100% code coverage is not sufficient to find real bugs. Instead, we propose to guide SBST to search for test cases covering defective areas in code to maximise the bug detection. Contrary to these works, we leverage defect prediction information to guide the search algorithm to extensively explore the search space for test cases covering defective areas.

Defect predictors employ statistical approaches [6, 32], learning techniques [15, 18] or risk estimation algorithms [28, 29, 35] etc. to estimate the likelihood of a file, class or method is defective. Based on this likelihood and a threshold (e.g., 0.5), they classify if the component (i.e., class, file or method) is defective or not. As a result of their efficacy, defect predictors are used in the industry to assist developers in code reviewing [29] and guide the developers to efficiently test defective components [8]. Defect predictors can be useful not only for humans, but also for other automated tools. Paterson et al. [35] used a defect predictor to inform a test case prioritisation strategy, G-clef, of the classes that are likely to be

defective and found it is promising. While the predictions are beneficial for developers, the inaccuracies in the estimations (e.g., false positives) can become costly and inefficient. Our approaches take this into account when using defect predictions to inform SBST.

## 4 RESEARCH OBJECTIVES

RO1: Develop a novel approach (SBST<sub>CL</sub>) that allocates time budget to classes based on the likelihood of classes being defective

SBST techniques such as the ones implemented in EvoSuite [12] employ a search algorithm such as genetic algorithm (GA) [11] to generate unit tests for each class independently by spending a given time budget. This time budget has to be tuned carefully as it is a stopping criterion for the GA and depends on the available computational resources in the organisation. Allocating a higher time budget is beneficial for the search because it allows the search algorithm to extensively explore the search space of test cases, thereby increasing the likelihood of finding bugs.

Previous work that studied SBST in terms of finding real bugs allocated a fixed time budget to all the buggy classes [1, 14, 39]. Ideally, buggy classes should receive more time budget than the non-buggy ones to maximise the likelihood of finding bugs. Allocating different time budgets to classes has been explored in previous research. For example, Campos et al. [7] allocated time budget to classes based on the complexity of them in order to maximise branch coverage. In contrast, we propose to allocate time budgets based on the likelihood of a class being buggy, i.e., classes that are more likely to be buggy receive a higher time budget, in order to increase the number of bugs found.

In the first research objective, we simulate defect predictor outputs at different performance levels (e.g., Matthews Correlation Coefficient (MCC)) including an ideal defect predictor to investigate how the bug detection performance of SBST<sub>CL</sub> changes with varying performance levels of defect predictors. An actual defect predictor comes with some degree of inaccuracy. Thus, we investigate the bug detection performance of SBST<sub>CL</sub> when using a real defect predictor at class level.

RO2: Develop a novel strategy (SBST<sub>ML</sub>) to guide the underlying search algorithm (i.e., genetic algorithm) towards the defective areas in a class

SBST techniques receive a limited time budget to generate unit tests for a class. Hence, spending search resources on covering non-defective parts in the class will deteriorate the bug detection performance of SBST. GA uses search heuristics (e.g., fitness function and selection operator) to guide the search to the given objective. Current state of the art search algorithm, *DynaMOSA* [34], employs a preference sorting algorithm to sort the best test cases to be selected to the next generation in GA. We find that preference sorting treats each target (i.e., coverage goal) is equally important and that leads to sub-optimal test suites in terms of finding bugs. Thus, we plan to use defect prediction at method level to inform the underlying GA of the buggy methods, hence it can prioritise the coverage of buggy code over non-buggy code.

As in the RO1, we simulate defect predictor outputs to investigate how the bug detection performance of  $SBST_{ML}$  changes when varying the performance of the defect predictor. Then, we investigate the performance of  $SBST_{ML}$  when using an actual defect predictor at method level.

## 5 CURRENT STATE OF WORK




Figure 2:  $SBST_{CL}$  Overview

### 5.1 RO1: Develop a novel approach ( $SBST_{CL}$ ) that allocates time budget to classes based on the likelihood of classes being defective

**5.1.1 Motivation.** Real-world projects are often very large. For instance, a modern car has millions of lines of code and thousands of classes [5]. Thus, running test generation on each individual class in the project would take significant amount of resources. At the same time, the available computational resources are limited in practice [7]. In order for the test generation tool to be adapted in the continuous integration (CI) system [10] without causing an impact to the existing processes (e.g., making other jobs idle), the test generation tool should use as little resources possible. Therefore this prompts the question, ‘How can we optimally utilise the available computational resources (e.g., time budget) to generate tests for the whole project with maximal bug detection?’.

**5.1.2 Approach.** We propose our novel approach,  $SBST_{CL}$ , consisting of three main modules (see Figure 2); i) Defect Predictor, ii) Budget Allocation based on Defect Scores (BADS), and iii) SBST. The defect predictor estimates the likelihood of defectiveness for each class in the project. We call this likelihood a defect score. BADS takes these defect scores as input and decides how to allocate the available time budget to classes based on these scores. At high level, BADS allocates more time budget for more likely to be defective classes and less time budget for less likely to be defective classes. BADS employs an exponential function to highly favour the budget allocation for a few number of highly likely to be defective classes. In order to account for inaccuracies present in defect predictors, we introduce a lower bound for time budget allocation. This lower bound ensures that every class gets a time budget allocated regardless of its defect score. Finally, the SBST module takes these allocated time budgets and runs test generation for each of the classes. Given its maturity, we use EvoSuite as the SBST tool and *DynaMOSA* as the search algorithm.

**5.1.3 Evaluation.** We answer the following research questions in order to evaluate  $SBST_{CL}$  in terms of efficiency in finding real bugs and effectiveness of discovering unique bugs, i.e., bugs that are not found by the baseline approach.

*RQ1.1: Is  $SBST_{CL}$  with simulated defect predictors more efficient in finding bugs compared to the state of the art?*

To investigate the bug detection performance of  $SBST_{CL}$  when employing defect predictors with different performance levels, we simulate the defect predictor outputs. Studies on defect prediction use various metrics to evaluate and compare the models’ performance, i.e., recall, precision, accuracy, MCC and AUC etc. [18, 23]. We choose MCC as the performance measure to simulate the defect predictions as it is insensitive to class imbalance unlike the other metrics [40]. Class imbalance prevails in defect prediction as there are only a few number of buggy classes (or methods) compared to non-buggy ones. We simulate defect predictor outcomes ranging from  $MCC = 1.0$  (i.e., ideal defect predictor) to  $MCC = 0.0$  (i.e., random classification).

We choose Defects4J as a suitable benchmark to evaluate our approach as it is widely used in research [27, 35, 36, 39], and hence allows us to compare the results to existing work. It contains 835 real bugs from 17 open source java projects [24]. The baseline for comparison is the state of the art SBST, i.e., *DynaMOSA*, with fixed time budget allocation. We conduct experiments for 2 cases of total time budgets ( $T$ );  $15 * N$  and  $30 * N$  seconds, where  $N$  is the number of classes in the project. We choose these time budgets for the experiments as *DynaMOSA* is capable of converging to the final branch coverage quickly, sometimes with a time budget like 20 seconds [34]. Also, faster test generation makes SBST more suitable to fit into the CI/CD pipeline without affecting the existing processes.

For each bug, we run test generation with the baseline and  $SBST_{CL}$  using each simulated defect predictor. We repeat the test generation of each approach for 20 runs to account for randomness of SBST. Once the test cases are generated, we check if they find the bugs in the programs, and compare the results of  $SBST_{CL}$  with different simulated predictors against the baseline. We use two-tailed non-parametric Mann-Whitney U-Test with the significance level ( $\alpha$ ) 0.05 [3] and Vargha and Delaney’s  $\hat{A}_{12}$  statistic [42] to check the statistical significance of the differences and effect size of the results (i.e., number of bugs found).

*RQ1.2: Does the efficiency of  $SBST_{CL}$  change when using defect predictors with different performance levels?*

To answer this research question, we compare  $SBST_{CL}$  with simulated defect predictors ( $MCC < 1.0$ ) against  $SBST_{CL}$  with the ideal defect predictor (i.e.,  $MCC = 1.0$ ).  $SBST_{CL}$  considers the inaccuracies present in the actual defect predictors when allocating time budgets to classes. However, the defect predictors having performance closer to a random classifier (i.e.,  $MCC = 0.0$ ) may deteriorate the performance of  $SBST_{CL}$ .

*RQ1.3: Does  $SBST_{CL}$  with simulated defect predictors find more unique bugs?*

In this research question, we investigate if  $SBST_{CL}$  is capable of finding more unique bugs which cannot be revealed by the baseline approach. This is an important measure since it indicates the testing technique is capable to find bugs which cannot be discovered otherwise within the given time budget [17]. The increased efficiency of  $SBST_{CL}$  could also be due to its robustness, which is measured by the success rate. Thus, we also analyse how often a bug is found over 20 runs.

*RQ1.4: Is SBST<sub>CL</sub> with a real defect predictor more efficient and effective in finding bugs when compared to the state of the art?*

The first three research questions (i.e., RQ1.1-3) investigate the bug detection performance of SBST<sub>CL</sub> using simulated defect predictors. In this research question, we investigate the performance of SBST<sub>CL</sub> when using an actual defect predictor.

We use Schwa [9, 13], which uses three metrics that have been shown to be effective in the literature [16, 26, 32, 38], as the defect predictor. Schwa produces defect scores for each class in the project by using a time weighted risk formula, whereas the other approaches require to train a classifier using training data. This makes it ideal to use in an industrial scenario where such training data is not available. In addition, Paterson et al. [35] successfully leveraged Schwa in G-clef to inform a test case prioritisation strategy of the buggy classes.

We run experiments for SBST<sub>CL</sub> with Schwa (SBST<sub>CL+Schwa</sub>) as outlined in RQ1.1. Then, we compare the efficiency and effectiveness (i.e., finding unique bugs) of SBST<sub>CL+Schwa</sub> against the state of the art SBST (i.e., baseline approach).

**5.1.4 Results.** We are yet to complete the experiments for RQ1.1-3. For RQ1.4, the empirical evaluation demonstrates that SBST<sub>CL+Schwa</sub> is significantly more efficient than state of the art SBST when they are given a tight time budget ( $T = 15 * N$  seconds) in a usual resource constrained scenario. In particular, SBST<sub>CL+Schwa</sub> finds an average of 13.1% more bugs than the baseline method. However, when there is sufficient time budget ( $T = 30 * N$  seconds), SBST<sub>CL+Schwa</sub> is more effective than the state of the art SBST 67% of the time (effect size = 0.67), which is significant given how hard it is to find failing test cases [17]. Further analysis to the results tells us that the superior performance of SBST<sub>CL+Schwa</sub> is supported by both its ability to find unique bugs that the baseline could not find and robustness in finding bugs. More details can be found in our publication [37].

## 5.2 RO2: Develop a novel strategy (SBST<sub>ML</sub>) to guide the underlying search algorithm (i.e., genetic algorithm) towards the defective areas in a class

**5.2.1 Motivation.** We establish in RO1 that the limited resources (i.e., time budget) available in practice can be allocated optimally for test generation of classes in order to maximise the number of bugs found. However, does the search algorithm spend the allocated resources (i.e., time budget) to optimally search for test cases that can reveal the bugs? We find that the current state of the art search algorithm, *DynaMOSA*, treats each target (i.e., coverage goal) is equally important to cover. However, the buggy code could exist in one or few methods, hence we argue that spending the search resources on covering the non-buggy methods may lead to sub-optimal test suites in terms of finding bugs.

**5.2.2 Approach.** We use a defect predictor that works at the method level. It estimates the likelihood of defectiveness (i.e., defect score) for each method in a class. We choose state of the art *DynaMOSA* as the search algorithm in SBST<sub>ML</sub>. First, SBST<sub>ML</sub> uses the defect scores to identify the buggy methods that need to be prioritised in the search. Then, all the targets (e.g., branches) in each buggy method are considered equally important. We introduce a

defectiveness-aware preference sorting algorithm that prioritises the selection of test cases that are closer to cover the buggy targets to the next generation in GA. This way SBST<sub>ML</sub> can generate more test cases covering the buggy targets in the class, hence increasing the likelihood of finding the bug.

**5.2.3 Evaluation.** We mainly evaluate SBST<sub>ML</sub> in terms of efficiency in finding real bugs and effectiveness of revealing unique bugs against the state of the art, *DynaMOSA*. Similar to RO1, we plan to carry out the investigation in two parts. First, we study if the performance of SBST<sub>ML</sub> changes when using simulated defect predictors with varying performance levels. Second, we investigate the performance of SBST<sub>ML</sub> when using an actual defect predictor that works at method level.

We use the Defects4J dataset as the experimental subjects. Similar to the experimental protocol outlined in RO1, we run test generation for each bug with the the baseline and SBST<sub>ML</sub> using simulated and actual defect predictors. We repeat the test generations for 20 runs and conduct statistical tests to check for significance of the differences of results (i.e., number of bugs found).

**5.2.4 Results.** We conducted a pilot study on 60 bugs from Apache-commons Lang project [24]. Our preliminary results show that SBST<sub>ML</sub> with an ideal defect predictor is significantly more efficient than *DynaMOSA*. In particular, SBST<sub>ML</sub> finds an average of 11.13%, 20.5% and 21% more bugs than *DynaMOSA* at 15, 30 and 60 seconds respectively.

## 6 FUTURE WORK

We have three future works pertaining to RO1 (RQ1.1-3) and RO2. In the first research objective, we need to perform the experiments for SBST<sub>CL</sub> with simulated defect predictors and evaluate its performance as outlined in Section 5.1.3.

Second and third future works aim at achieving the second research objective. The outcome of the pilot study suggests the merits of using defect prediction at method level to guide the search algorithm to find more bugs. Our next step is to incorporate the lessons learned from the pilot study. This is followed by an empirical evaluation as outlined in Section 5.2.3. Upon completing the second future work, we apply an actual defect predictor that works at the method level. The existing work on method level defect prediction suggests several metrics that are effective in producing predictions [15, 22, 41]. We plan to use these metrics to build a model and use it in the SBST<sub>ML</sub>.

## 7 CONCLUSION

In this thesis, we aim at achieving the main research objective; *Improve the bug detection capability of SBST by incorporating defect prediction information*. Through empirical evaluation on 434 real bugs in Defects4J dataset, we demonstrate the increased efficiency of SBST by our time budget allocation approach in a resource constrained environment. Preliminary results of the second research objective shed a positive light on guiding the search algorithm using the likelihoods of methods being defective in order to find more bugs. The outcome of RQ1.4 has been accepted at the International Conference on Automated Software Engineering (ASE) [37].

## REFERENCES

- [1] M Moein Almasi, Hadi Hemmati, Gordon Fraser, Andrea Arcuri, and Jānis Bene-felds. 2017. An industrial evaluation of unit test generation: Finding real faults in a financial application. In *Proceedings of the 39th International Conference on Software Engineering: Software Engineering in Practice Track*. IEEE Press, 263–272.
- [2] Nadia Alshahwan, Xinbo Gao, Mark Harman, Yue Jia, Ke Mao, Alexander Mols, Taijin Tei, and Ilya Zorin. 2018. Deploying search based software engineering with Sapienz at Facebook. In *International Symposium on Search Based Software Engineering*. Springer, 3–45.
- [3] Andrea Arcuri and Lionel Briand. 2014. A Hitchhiker’s guide to statistical tests for assessing randomized algorithms in software engineering. *Software Testing, Verification and Reliability* 24, 3 (2014), 219–250.
- [4] Victor R Basili, Lionel C. Briand, and Walcécio L Melo. 1996. A validation of object-oriented design metrics as quality indicators. *IEEE Transactions on software engineering* 22, 10 (1996), 751–761.
- [5] Manfred Broy, Ingolf H Kruger, Alexander Pretschner, and Christian Salzmann. 2007. Engineering automotive software. *Proc. IEEE* 95, 2 (2007), 356–373.
- [6] Bora Caglayan, Burak Turhan, Ayse Bener, Mayy Habayeb, Andriy Miransky, and Enzo Cialini. 2015. Merits of organizational metrics in defect prediction: an industrial replication. In *Proceedings of the 37th International Conference on Software Engineering—Volume 2*. IEEE Press, 89–98.
- [7] José Campos, Andrea Arcuri, Gordon Fraser, and Rui Abreu. 2014. Continuous test generation: enhancing continuous integration with automated test generation. In *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*. ACM, 55–66.
- [8] Hoa Khanh Dam, Trang Pham, Shien Wee Ng, Truyen Tran, John Grundy, Aditya Ghose, Taeksu Kim, and Chul-Joo Kim. 2019. Lessons learned from using a deep tree-based model for software defect prediction in practice. In *Proceedings of the 16th International Conference on Mining Software Repositories*. IEEE Press, 46–57.
- [9] Paulo André Faria de Freitas. 2015. Software Repository Mining Analytics to Estimate Software Component Reliability. (2015).
- [10] Martin Fowler and Matthew Foemmel. 2006. Continuous integration.
- [11] Gordon Fraser and Andrea Arcuri. 2011. Evolutionary generation of whole test suites. In *2011 11th International Conference on Quality Software*. IEEE, 31–40.
- [12] Gordon Fraser and Andrea Arcuri. 2012. Whole test suite generation. *IEEE Transactions on Software Engineering* 39, 2 (2012), 276–291.
- [13] André Freitas. 2015. schwa. <https://github.com/andrefreitas/schwa> Last accessed on 16/09/2019.
- [14] Gregory Gay. 2017. Generating effective test suites by combining coverage criteria. In *International Symposium on Search Based Software Engineering*. Springer, 65–82.
- [15] Emanuel Giger, Marco D’Ambros, Martin Pinzger, and Harald C Gall. 2012. Method-level bug prediction. In *Proceedings of the 2012 ACM-IEEE International Symposium on Empirical Software Engineering and Measurement*. IEEE, 171–180.
- [16] Todd L Graves, Alan F Karr, James S Marron, and Harvey Siy. 2000. Predicting fault incidence using software change history. *IEEE Transactions on software engineering* 26, 7 (2000), 653–661.
- [17] Andrew Habib and Michael Pradel. 2018. How many of all bugs do we find? a study of static bug detectors. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 317–328.
- [18] Tracy Hall, Sarah Beecham, David Bowes, David Gray, and Steve Counsell. 2011. A systematic literature review on fault prediction performance in software engineering. *IEEE Transactions on Software Engineering* 38, 6 (2011), 1276–1304.
- [19] Mark Harman, Yue Jia, and Yuanyuan Zhang. 2015. Achievements, open problems and challenges for search based software testing. In *2015 IEEE 8th International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 1–12.
- [20] Mark Harman and Bryan F Jones. 2001. Search-based software engineering. *Information and software Technology* 43, 14 (2001), 833–839.
- [21] Mark Harman and Peter O’Hearn. 2018. From start-ups to scale-ups: Opportunities and open problems for static and dynamic program analysis. In *2018 IEEE 18th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 1–23.
- [22] Hideaki Hata, Osamu Mizuno, and Tohru Kikuno. 2012. Bug prediction based on fine-grained module histories. In *2012 34th international conference on software engineering (ICSE)*. IEEE, 200–210.
- [23] Seyedrebrvar Hosseini, Burak Turhan, and Dimuthu Gunarathna. 2017. A systematic literature review and meta-analysis on cross project defect prediction. *IEEE Transactions on Software Engineering* 45, 2 (2017), 111–147.
- [24] René Just. 2019. Defects4J - A Database of Real Faults and an Experimental Infrastructure to Enable Controlled Experiments in Software Engineering Research. <https://github.com/rjust/defects4j> Last accessed on: 02/10/2019.
- [25] René Just, Darioush Jalali, and Michael D Ernst. 2014. Defects4J: A database of existing faults to enable controlled testing studies for Java programs. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis*. ACM, 437–440.
- [26] Sunghun Kim, Thomas Zimmermann, E James Whitehead Jr, and Andreas Zeller. 2007. Predicting faults from cached history. In *Proceedings of the 29th international conference on Software Engineering*. IEEE Computer Society, 489–498.
- [27] Xuan Bach D Le, David Lo, and Claire Le Goues. 2016. History driven program repair. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, Vol. 1. IEEE, 213–224.
- [28] Chris Lewis, Zhongpeng Lin, Caitlin Sadowski, Xiaoyan Zhu, Rong Ou, and E James Whitehead Jr. 2013. Does bug prediction support human developers? findings from a google case study. In *Proceedings of the 2013 International Conference on Software Engineering*. IEEE Press, 372–381.
- [29] Chris Lewis and Rong Ou. 2011. Bug Prediction at Google. <http://google-engtools.blogspot.com/2011/12/> Last accessed on: 16/09/2019.
- [30] Ke Mao, Mark Harman, and Yue Jia. 2016. Sapienz: Multi-objective automated testing for Android applications. In *Proceedings of the 25th International Symposium on Software Testing and Analysis*. 94–105.
- [31] Webb Miller and David L. Spooner. 1976. Automatic generation of floating-point test data. *IEEE Transactions on Software Engineering* 3 (1976), 223–226.
- [32] Nachiappan Nagappan, Brendan Murphy, and Victor Basili. 2008. The influence of organizational structure on software quality. In *2008 ACM/IEEE 30th International Conference on Software Engineering*. IEEE, 521–530.
- [33] Annibale Panichella, Fitsum Meshesha Kifetew, and Paolo Tonella. 2015. Reformulating branch coverage as a many-objective optimization problem. In *2015 IEEE 8th international conference on software testing, verification and validation (ICST)*. IEEE, 1–10.
- [34] Annibale Panichella, Fitsum Meshesha Kifetew, and Paolo Tonella. 2017. Automated test case generation as a many-objective optimisation problem with dynamic selection of the targets. *IEEE Transactions on Software Engineering* 44, 2 (2017), 122–158.
- [35] David Paterson, Jose Campos, Rui Abreu, Gregory M Kapfhammer, Gordon Fraser, and Phil McMinn. 2019. An Empirical Study on the Use of Defect Prediction for Test Case Prioritization. In *2019 12th IEEE Conference on Software Testing, Validation and Verification (ICST)*. IEEE, 346–357.
- [36] Spencer Pearson, José Campos, René Just, Gordon Fraser, Rui Abreu, Michael D Ernst, Deric Pang, and Benjamin Keller. 2017. Evaluating and improving fault localization. In *Proceedings of the 39th International Conference on Software Engineering*. IEEE Press, 609–620.
- [37] Anjana Perera, Aldeida Aleti, Marcel Böhme, and Burak Turhan. 2020. Defect Prediction Guided Search-Based Software Testing. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*. ACM. <https://doi.org/10.1145/3324884.3416612>
- [38] Foyzur Rahman, Daryl Posnett, Abram Hindle, Earl Barr, and Premkumar Devanbu. 2011. BugCache for inspections: hit or miss?. In *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*. ACM, 322–331.
- [39] Sina Shamshiri, Rene Just, Jose Miguel Rojas, Gordon Fraser, Phil McMinn, and Andrea Arcuri. 2015. Do automatically generated unit tests find real faults? an empirical study of effectiveness and challenges (t). In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 201–211.
- [40] Martin Shepperd, David Bowes, and Tracy Hall. 2014. Researcher bias: The use of machine learning in software defect prediction. *IEEE Transactions on Software Engineering* 40, 6 (2014), 603–616.
- [41] Thomas Shippey, David Bowes, and Tracy Hall. 2019. Automatically identifying code features for software defect prediction: Using AST N-grams. *Information and Software Technology* 106 (2019), 142–160.
- [42] András Vargha and Harold D Delaney. 2000. A critique and improvement of the CL common language effect size statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (2000), 101–132.