TERMINATOR: Better Automated UI Test Case Prioritization

Zhe Yu

Fahmid Fahid

Tim Menzies

Gregg Rothermel

zyu9, ffahid, gerothier@ncsu.edu, timm@ieee.org

North Carolina State University

Raleigh, NC, USA

Kyle Patrick

Snehit Cherian

kyle.patrick@lexisnexis.com

snehit.cherian@lexisnexis.com

LexisNexis Legal & Professional

Raleigh, NC, USA

ABSTRACT

Automated UI testing is an important component of the continuous integration process of software development. A modern web-based UI is an amalgam of reports from dozens of microservices written by multiple teams. Queries on a page that opens up another will fail if any of that page's microservices fails. As a result, the overall cost for automated UI testing is high since the UI elements cannot be tested in isolation. For example, the entire automated UI testing suite at LexisNexis takes around 30 hours (3-5 hours on the cloud) to execute, which slows down the continuous integration process.

To mitigate this problem and give developers faster feedback on their code, test case prioritization techniques are used to reorder the automated UI test cases so that more failures can be detected earlier. Given that much of the automated UI testing is "black box" in nature, very little information (only the test case descriptions and testing results) can be utilized to prioritize these automated UI test cases. Hence, this paper evaluates 17 "black box" test case prioritization approaches that do not rely on source code information. Among these, we propose a novel TCP approach, that dynamically re-prioritizes the test cases when new failures are detected, by applying and adapting a state of the art framework from the total recall problem. Experimental results on LexisNexis automated UI testing data show that our new approach (which we call TERMINATOR), outperformed prior state of the art approaches in terms of failure detection rates with negligible CPU overhead.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging**; • **Information systems** → **Learning to rank**.

KEYWORDS

automated UI testing, test case prioritization, total recall

ACM Reference Format:

Zhe Yu, Fahmid Fahid, Tim Menzies, Gregg Rothermel, Kyle Patrick, and Snehit Cherian. 2019. TERMINATOR: Better Automated UI Test Case Prioritization. In *Proceedings of the 27th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '19)*, August 26–30, 2019, Tallinn, Estonia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3338906.3340448>

1 INTRODUCTION

Complex web applications must be validated using a variety of methods. Early in the lifecycle, unit and integration tests can be used to validate many aspects of an application. However, these isolated tests cannot realistically simulate the behavior of a modern web page that combines together many services. For this reason, user interface (UI) testing are applied and automated. Compared to unit tests, automated UI tests are more expensive to write and maintain. Also, they usually have a higher execution time. For example, the automated UI test suite at LexisNexis takes around 30 hours (3-5 hours on the cloud) to execute. Hence:

- Testing has become the key component in continuous integration. If a system can only be rebuilt and tested (say) twice in an eight-hour work day, that also determines how often developers can ship new (or fixed) features to clients.
- Developers are less agile when testing forces them to wait for hours to get feedback on their latest changes.

Worse still, since automated UI tests are expressed in terms of actions taken by a browser user agent, failures do not have a straightforward relationship to the underlying application code or architecture. As shown in Figure 1, a modern web-based UI is an amalgam of reports from dozens of microservices written by multiple teams. If a query on one webpage opens up another, then that query interacts with code from multiple teams. The code that generated Figure 1 comes from seven different teams which, in turn, use "under-the-hood" components from dozens of other teams. Consequently, much of the testing is "black box" in nature since most of the code associated with subpages will be opaque to the query author. Hence, it is hard to determine what part of the code is being tested by which test case.

One possible solution to this problem is test case prioritization (TCP). TCP is a topic widely studied in software testing, where test cases are prioritized so that test cases revealing faults are executed earlier, allowing software developers to identify and fix problems earlier [17, 25, 34, 35, 39]. A problem with many of the existing TCP techniques is that they rely on coverage information extracted from the source code, such as what branches/statements/functions of the source code are covered by which test cases [4, 26, 30, 40, 41].

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 ACM 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.

ESEC/FSE '19, August 26–30, 2019, Tallinn, Estonia

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5572-8/19/08...\$15.00

<https://doi.org/10.1145/3338906.3340448>

Figure 1: A modern web page is an amalgam of many microservices. This figure shows a real example from LexisNexis where a single click invokes multiple components of the software. By our estimates, the code on this page comes from seven different teams which, in turn, use “under-the-hood” components from dozens of other teams. The failure of any component will lead to a test failure. That failure may be opaque to the author of the test since it may result from the failure of components written by many other teams. As a result, when a team is prioritizing their own tests, all they can reason about is the test description information (e.g see Figure 2(b)) and prior results from their own testing results (time to run the test; whether or not that prior test failed).

Figure 2: To test a page shown at the left (a), programmers write a test description (b) which is converted to test code (c).

Unfortunately, such techniques are not applicable to automated UI testing since the coverage information is not available—as discussed above, the relationship between an automated UI test case and the source code being tested is blurred.

This paper explores TCP techniques for higher failure detection rates utilizing only “black box” information; i.e. test descriptions and results from their own prior (and current) test runs (time to run the test, whether or not that test failed). Our central insight is that this kind of test case prioritization problem belongs to a class of information retrieval problems called the *total recall* problem [20] (defined in §3.1). To test this idea, in this work we apply total recall methods to “read” the description and history results of the test cases to build a classifier that predicts which tests might fail sooner. Specifically, we incrementally update a support vector machine (SVM) model with the results of executed tests, then reflect on the SVM model to decide which test to run next. Formally speaking, this is an *active learning* [46] based framework.

We compared our proposed approach, which we call TERMINATOR, with 12 state of the art TCP algorithms and 2 baseline algorithms on 54 consecutive runs of automated UI testing (data from three months of testing at LexisNexis) to answer the following two research questions:

- **RQ1: can the proposed approach TERMINATOR achieve significantly higher failure detection rates than other TCP algorithms when prioritizing automated UI tests from LexisNexis?** Our results show that TERMINATOR significantly outperformed other TCP algorithms in terms of failure detection rates. This suggests that techniques from the total recall problem can be successfully applied and adapted to address the automated UI test case prioritization problem.
- **RQ2: what is the computational overhead of TERMINATOR?** Our results show that TERMINATOR had 50% more computational overhead than the simple history-based TCP algorithms because it recursively updates the SVM model and dynamically adjusts the order of the unexecuted tests. However, TERMINATOR’s overhead was still negligible compared to the runtime of the test cases (0.33% of the runtime of the test suite). Therefore, we conclude that it is practical to apply TERMINATOR for prioritizing automated UI tests in LexisNexis.

The main contributions of this paper are:

- (1) Identification of an under-explored problem: prioritizing automated UI tests. Features of this problem are: (a) the prioritization target is to reach higher failure detection rates and (b) source code information is unavailable. By our reading of the literature, most existing TCP approaches utilize source code information and prioritize for fault detection rates. Only 27 of 239 approaches do not rely on source code information, among which we derived 12 approaches that can prioritize automated UI tests.
- (2) A dataset for automated UI testing is provided¹ for reproducing and improving this work. This is an important contribution since, in the research literature, there are very few examples of real-world industrial test results from large systems.
- (3) The lesson learned that test case prioritization for failure detection rates can be generalized as the total recall problem.

- (4) By applying and adapting the total recall problem, we propose a new TCP algorithm that significantly outperformed other TCP algorithms in terms of failure detection rates.
- (5) The proposed algorithm can be easily applied or tested on other TCP problems since it requires minimal information.

The rest of this paper is structured as follows. Background and related work on automated UI testing and test case prioritization are discussed in §2. The total recall problem is introduced and the proposed new TCP algorithms are presented in §3. The proposed TCP approach TERMINATOR is compared against other TCP algorithms in §4. Impacts of this work to LexisNexis are reported in §5 while conclusion and future work are provided in §6.

2 BACKGROUND AND RELATED WORK

LexisNexis is a corporation providing computer-assisted legal research (CALR) as well as business research and risk management services [1, 49]. During the 1970s, LexisNexis pioneered the electronic accessibility of legal and journalistic documents [38]. As of 2006, the company had the world’s largest electronic database for legal and public-records related information [38].

LexisNexis provides regulatory, legal, business information and analytics to the legal community. Legal and research professionals use the Lexis Advance platform to find relevant information more easily and efficiently [2]. It helps them prepare legal cases and drive better legal outcomes. The Lexis Advance platform is maintained by a set of automated UI tests. Those automated UI tests simulate user behaviours on the interface of the platform and detect potential failures of the underlying microservices whenever the system is modified and rebuilt.

2.1 Automated UI Testing

Automated UI testing is an important component of the continuous integration process of software development. To automatically test the user interface of a complex system such as the Lexis Advance platform, the test team of LexisNexis designs a large set of automated UI test cases, each of which simulates one action in a specific scenario. As an example, Figure 2 shows how one automated UI test case is designed for a simple search of “1+1”. In this example, the test designer wants to test the search function by (a) verifying that when a user inputs “1+1” and hits the search button, a result of “2” will show up. To automate this UI test, the test designer would first (c) define the test code for a set of scenarios, then (b) write down the automated UI test case with the pre-defined scenarios and expected input and output. In this way, the test designer does not need to know what code will be executed when an automated UI test is executed, and the pre-defined scenarios can be reused for designing other automated UI test cases.

2.2 Test Case Prioritization

The goal of test case prioritization is to schedule test cases for execution in an order that attempts to increase their effectiveness at meeting some performance goal [44]. Here we briefly introduce different types of TCP algorithms in terms of what they prioritize for and what information they use.

2.2.1 What to Prioritize For. Three different performance goals have been targeted by TCP algorithms:

¹<https://github.com/ai-se/Data-for-automated-ui-testing-from-LexisNexis>

- **Coverage:** Some approaches aim to order test cases so that higher coverage can be achieved earlier. Here the coverage can refer to requirement coverage [24], statement coverage [30], decision coverage [30], block coverage [30], branch coverage [44], etc.
- **Fault detection rate:** Most approaches prioritize the test cases to achieve higher fault detection rates. An improved rate of fault detection during testing can provide faster feedback on the system under test and let software engineers begin correcting faults earlier than might otherwise be possible [44]. The most popular performance metric for evaluating early fault detection is average percentage of faults detected (APFD) proposed by Rothermel et al. [44] in 2001. APFD computes the area under the curve (AUC) of the the gain in the percentage of detected faults as follows:

$$APFD = 1 - \frac{TF1 + TF2 + \dots + TFm}{nm} + \frac{1}{2n} \quad (1)$$

where TFi is the first test case that reveals fault i , m is the total number of faults revealed by the test suite, and n is the total number of test cases in the test suite. Ranging from 0% to 100%, higher APFD values mean better ordering of test cases in terms of early fault detection.

In APFD, each test case is considered to be of the same cost and the same severity. However in practice, some test cases can take much longer to run than others while some faults can cause more damage than others if not detected. To this end, Elbaum et al. [15] proposed the metric of average percentage of faults detected with cost (APFDc), which takes into consideration the cost related to the resources required to execute and validate each test case and the severity of each fault. Let T be a test suite containing n test cases with costs $t1, t2, \dots, tn$. Let F be a set of m faults revealed by the test suite and let $f1, f2, \dots, fm$ be the severities of those faults. Let TFi be the first test case that reveals fault i . APFDc is calculated as follows:

$$APFDc = \frac{\sum_{i=1}^m (fi \times (\sum_{j=TFi}^n tj - 0.5tTFi))}{\sum_{i=1}^n ti \times \sum_{i=1}^m fi} \quad (2)$$

In order to measure APFD or APFDc, researchers use manually seeded faults [58], mutation faults [12, 47] or real fault detection records [33, 42]. Source code information is required to manually seed faults or to generate mutation faults while failure to fault mapping information (i.e. which failures reveal which faults) is required for real fault detection records.

- **Failure detection rate:** Some studies prioritize test cases to achieve higher failure detection rates. This is a compromise prioritization target when failure to fault mapping information is not available and is often seen when real world data is applied to evaluate TCP performance. Similar to fault detection rate, Liang et al. [31] apply APFDc to measure failure detection rates. Here it becomes average percentage of failures detected with the same equation as (2), except that TFi represents i th failed test case.

2.2.2 What Information is Used. Different TCP algorithms utilize different information to decide the execution order of test cases:

- **Coverage information:** Most test case prioritization techniques rely on coverage information extracted from source code [44],

changes in source code [37], requirements [32], previously detected faults [14], etc. Usually, coverage information requires access to the source code and takes time to compute.

- **History information:** Some algorithms learn the fault/failure exposing potential of each test case from its fault/failure detection history to prioritize test cases for current execution. Then, the test cases can be prioritized in descending order of their fault/failure exposing potential.
- **Cost information:** Some TCP algorithms take into consideration the cost of each test case to prioritize the test cases, e.g. execute all test cases in ascending order of their runtime estimated from the execution history [23].
- **Test description information:** Standard natural language processing techniques (stopword removal, stemming, bag-of-word, etc.) can be applied to extract features from test description information as shown in Figure 2(b). Some algorithms train a predictive model with independent variables from test description features and dependent variables extracted from failure/fault history. The test cases are then prioritized based on the predicted probability of failure/fault detection [22, 29].
- **Feedback information:** Some algorithms iteratively utilize execution results (pass/fail) in current test runs to reorder the rest of the unexecuted test cases in the test suite [5, 59]. The priorities of unexecuted test cases are adjusted dynamically based on two heuristics: (1) if one test case fails, increase the priority of unexecuted test cases that are similar/related to the failed one; (2) if one test case passes, decrease the priority of unexecuted test cases that are similar/related to the passed one.

2.3 Prioritizing Automated UI Test Cases

There are two characteristics of automated UI test case prioritization: (1) its prioritization target is to reach higher failure detection rate, and (b) source code information is unavailable. To analyze what existing TCP algorithms can be applied to prioritize automated UI test cases, we conducted a systematic literature review [55] and ended up identifying 239 TCP papers. Surprisingly, we found that there are few papers [3, 22, 48, 59] (4 out of 239) exploring similar scenarios as automated UI test case prioritization—prioritizing for failure detection rates without source code information. This suggests that although lots of work has been done in the TCP arena, the scenario where source code information and failure to faults mapping are not available is under-explored and many of the existing TCP algorithms cannot be applied to prioritizing automated UI tests.

3 METHODOLOGY

The central idea of this paper is that prioritizing automated UI tests can be generalized as the total recall problem [20] and techniques addressing the total recall problem can be applied and adapted to better prioritize the automated UI tests.

3.1 Total Recall

The aim of *total recall* is to optimize the cost for achieving very high recall (ideally, very close to 100%) with a human assessor in the loop [20]. More specifically, the total recall problem can be described as follows:

The Total Recall Problem:

Given candidates E with a small positive fraction $R \subset E$, each $x \in E$ can be inspected to label it positive ($x \in R$) or negative ($x \notin R$) at a cost. Starting with the labels $L = \emptyset$, the task is to inspect and label as few candidates as possible ($\min |L|$) while achieving very high recall ($\max |L \cap R|/|R|$).

State of the art solutions [56, 57] for the total recall problem use an active learning based framework. The key idea behind active learning

is that a machine learning algorithm can train faster (i.e. using less data) if it is allowed to choose the data from which it learns [46]. The experience in explored total recall problems is that such *active learners* outperform supervised and semi-supervised learners and can significantly reduce the effort required to achieve high recall [7–11, 19, 50–54, 56, 57]. To understand active learning, consider the decision plane between the failed and passed test cases of Figure 3. Suppose we want to find more failing test cases and we had access to the Figure 3 model. One tactic would be to execute test cases that fall into the region of red circles in this figure, as far as possible from the green squares (this tactic is called *certainty sampling*). Another tactic would be to verify the position of the boundary; i.e. execute test cases that are closest to the boundary (this tactic is called *uncertainty sampling*).

Figure 3: Separating failing (red circles) from passing (green squares) test cases.

3.2 TERMINATOR

The automated UI test case prioritization problem can be generalized to the total recall problem as follows:

- E : the test suite to be prioritized.
- R : the set of test cases that will fail if executed.
- L : the set of test cases already executed in the current run.
- $L_R = L \cap R$: the set of failed test cases in the current run.

Given the information available in automated UI testing, we extract three types of features:

- **Text feature:** the same text mining feature extraction used in the total recall approaches [56, 57] is applied to extract text features from the test case descriptions. Specifically, we:
  - (1) Tokenized the test case descriptions without stop/control word removal or stemming.
  - (2) Built a term frequency matrix.
  - (3) Normalized each row (feature vector for each test case) with their L2-norm².
- **History feature:** same as history-based algorithms, the testing result from previous runs (failed, passed, skipped) are vectorized as feature.
- **Hybrid feature:** concatenation of text and history features.

² L2-norm for a vector x is $\sqrt{x^T x}$, where T denotes “transpose”

Using the foregoing types of features, the proposed framework is described in Algorithm 1 with engineering choices of N_1, N_2 . N_1 is the batch size of the process. In this paper, we chose $N_1 = 10$ to simulate the scenario where test cases are executed in parallel on 10 computation nodes on the cloud. N_2 represents the threshold above which certainty sampling is applied instead of uncertainty sampling. In this paper, we chose $N_2 = 30$ as suggested by previous works on total recall [57].

Algorithm 1: Pseudo Code for TERMINATOR

```

Input :  $E$ , the test suite
         $R$ , test cases that will fail
         $N_1$ , batch size
         $N_2$ , threshold of query strategy
Output :  $L$ , list of executed test cases
         $L_R$ , list of failed test cases

1  $L \leftarrow \emptyset$ ;
2  $L_R \leftarrow \emptyset$ ;
// Keep reviewing until all the test cases are executed
3 while  $|L| < |E|$  do
    // Start training or not
4     if  $|L_R| \geq 1$  then
        // Presumptive non-relevant examples
5          $L_{pre} \leftarrow \text{Presume}(L, E \setminus L)$ ;
6          $CL \leftarrow \text{Train}(L_{pre})$ ;
        // Query next
7          $X \leftarrow \text{Query}(CL, E \setminus L, L_R)$ ;
8     else
        // Random Sampling
9          $X \leftarrow \text{Random}(E \setminus L)$ ;
        // Execute the selected test cases
10    foreach  $x \in X$  do
11         $L_R, L \leftarrow \text{Execute}(x, R, L_R, L)$ ;
12 return  $L, L_R$ ;

13 Function  $\text{Presume}(L, E \setminus L)$ 
    // Randomly sample  $|L|$  points from  $E \setminus L$ , presume those to be passed
    test cases
14    return  $L \cup \text{Random}(E \setminus L, |L|)$ ;

15 Function  $\text{Train}(L_{pre})$ 
    // Train linear SVM with Weighting
16     $CL \leftarrow \text{SVM}(L_{pre}, \text{kernel=linear, class\_weight=balanced})$ ;
17    if  $|L_R| \geq N_2$  then
        // Aggressive undersampling
18         $L_I \leftarrow L_{pre} \setminus L_R$ ;
19         $\text{tmp} \leftarrow L_I[\text{argsort}(CL.\text{decision\_function}(L_I))[: |L_R|]]$ ;
20         $CL \leftarrow \text{SVM}(L_R \cup \text{tmp}, \text{kernel=linear})$ ;
21    return  $CL$ ;

22 Function  $\text{Query}(CL, E \setminus L, L_R)$ 
23    if  $|L_R| \geq N_2$  then
        // Certainty Sampling (highest predicted probability of failing)
24         $X \leftarrow \text{argsort}(CL.\text{decision\_function}(E \setminus L))[: -1] : N_1$ ;
25    else
        // Uncertainty Sampling
26         $X \leftarrow \text{argsort}(\text{abs}(CL.\text{decision\_function}(E \setminus L)))[: N_1]$ ;
27    return  $X$ ;

28 Function  $\text{Execute}(x, R, L_R, L)$ 
    // Append selected test case to the list of executed ones
29     $L \leftarrow [L, x]$ ;
    // If the selected test case fails, append it to the list of failed ones
30    if  $x \in R$  then
31         $L_R \leftarrow [L_R, x]$ ;
32    return  $L_R, L$ ;

```

4 EMPIRICAL STUDY

We conducted an empirical study to answer the following research questions:

Table 1: A simple example of automated UI test case prioritization, where P, F, and S indicate passed, failed, and skipped testing results. For each test session, the runtime is $t1 < t2 < t3 < t4$.

Test	Description	Test Session			
		1	2	3	4
t1	Test Check Box in page A	P	F	S	F
t2	Test Radio Button in page A	F	F	P	F
t3	Test Check Box in page B	P	P	F	P
t4	Test Radio Button in page B	F	P	F	F

- **RQ1: can TERMINATOR achieve significantly higher failure detection rates than other TCP algorithms when prioritizing automated UI tests from LexisNexis?** Given that the goal of automated UI test case prioritization is to detect failures faster, TERMINATOR will only be considered useful if it can achieve significantly higher failure detection rates than other TCP algorithms.
- **RQ2: what is the computational overhead of TERMINATOR?** It is also important for a TCP algorithm to have low computational overhead, e.g., a TCP algorithm will not be considered useful if it can sort test cases in an effective order but requires even longer time than the total runtime of the test suite to compute that order.

4.1 Dataset

To conduct this study, we collected data for 54 consecutive runs of 2661 automated UI test cases in LexisNexis from September 27th to November 15th in 2018. The collected data contains the test description written using Gherkin syntax³ (as shown in Figure 2 (b)), test duration, outcome (passed/failed), and error message associated with each test case. To remove the effect of environmental errors, failures with "time out" messages are labeled with "Timeout" and are not considered failures. This dataset is available online⁴. For confidentiality, only the featurized vectors are provided for test descriptions while error messages are omitted.

The collected dataset is used to simulate the performance of different TCP algorithms. We simulate Run 6 to Run 54 by using different prioritization approaches and compare their performance⁵. When prioritizing for Run n , the execution results from Run 1 to Run $n-1$ are available as information for the prioritization algorithm.

Table 1 shows a small example of the dataset with 4 test cases and 4 consecutive runs. To prioritize for test cases in Run 4, all testing results from Run 1 to Run 3 and the description of each test case can be utilized.

4.2 Independent Variables

Three variants of the proposed active learning based TCP framework TERMINATOR are compared against 12 existing "black box" TCP techniques and 2 baselines. As shown in Table 2, in total 17

algorithms are tested in which Group A are baseline algorithms, Groups B, C, D, and E are the existing "black box" TCP techniques, and Group F are three variants of TERMINATOR that differs from what types of features they use. The 12 TCP techniques are chosen because they do not rely on source code information. During the simulation, three types of information can potentially be applied by these techniques to prioritize the test cases in Run n :

- **Execution history:** execution results of each test case from Run 1 to Run $n-1$. For example the execution history information for t1 in Run 4 would be {P, F, S} in Table 1.
- **Test case description:** the natural language description of the automated UI test case. For example the test case description information for t1 would be "Test Check Box in page A" in Table 1.
- **Feedback:** execution results of prior test cases in Run n . For example in Table 1, in Run 4, when t1 has been executed and other tests have not, the feedback information would be "t1 failed".

Note that some of the state of the art algorithms are not reproduced exactly as presented in the original paper, rather, they are partially implemented because of lack of information or differences in the workflow of automated UI testing.

4.2.1 Group A: Baselines. Group A shows baseline algorithms for comparison, in which

- **A1** is the simplest TCP algorithm utilizing no information. It executes all the test cases in a random order. A TCP algorithm has no value if it cannot perform better than A1. In the example of Table 1, A1 could place tests in an order such as {t1, t3, t2, t4} for Run 4.
- **A2** executes the test cases in the optimal order. It represents the upper bound of the TCP performance any algorithm can ever achieve. In the example of Table 1, A2 will execute the tests in the order {t1, t2, t4, t3} for Run 4 to maximize the failure detection rate. A2 is not applicable in practice because it requires knowledge of which test cases will fail to determine the optimal order, but it can be used in simulations to show how much room there is to improve from the TCP algorithms.

4.2.2 Group B: History-based Algorithms. Group B algorithms use metrics extracted from the execution history information, as described in §2.2.2 to order the test cases before each run:

- **B1:** time since last failure. Many algorithms assign higher priority to test cases with more recent failures [16, 21, 28]. To apply this metric to our dataset, it becomes the number of consecutive non-failures before Run n . For example, B1 metrics for Run 4 in Table 1 are {1, 1, 0, 0} for t1 to t4. Therefore, B1 will sort the test cases in the order {t3, t4, t1, t2} before Run 4.
- **B2:** failure rate (number of times failed/number of times executed). Fazlalizadeh et al. [17], Aman et al. [3], and Tsai et al. [48] apply this metric to determine the order of test cases (higher priority for higher value of B2 metrics). For example, B2 metrics for Run 4 in Table 1 are {1/2, 2/3, 1/3, 2/3} for t1 to t4. Therefore, B2 will sort the test cases in the order {t2, t4, t1, t3} before Run 4.
- **B3:** Exponential decay metrics [27]

$$P_0 = h_1$$

$$P_i = \alpha h_i + (1 - \alpha)P_{i-1}, 0 \leq \alpha \leq 1, i \geq 1 \quad (3)$$

where $h_i = 0$ if the test case passed or was skipped in Run i and $h_i = 1$ if the test case failed in Run i . α is the learning rate.

³<https://cucumber.io/docs/gherkin/>

⁴<https://github.com/ai-se/Data-for-automated-UI-testing-from-LexisNexis>

⁵Run 1 to Run 5 are not simulated since there are not enough execution history information available for those runs.

Table 2: Test Case Prioritization Algorithms and Information They Utilize

ID	Algorithm Description	Utilized Information		
		Execution history	Test case description	Feedback
A1	Execute test cases in random order.			
A2	Execute test cases in optimal order.			
B1	Execute test cases in ascending order of time since last failure.	✓		
B2	Execute test cases in descending order of number of times failed/number of times executed.	✓		
B3	Execute test cases in descending order of exponential decay metrics as in (3).	✓		
B4	Execute test cases in descending order of ROCKET metrics as in (4).	✓		
B5	Execute test cases in descending order of the Mahalanobis distance of each test case to the origin (0,0) when considering two metrics—time since last execution and failure rate.	✓		
C1	Execute test cases in ascending order of the estimated test case runtime.	✓		
D1	Supervised learning with Simple History (SH).	✓	✓	
D2	Supervised learning with All History (AH).	✓	✓	
D3	Supervised learning with Weighted History (WH).	✓	✓	
E1	Dynamic test case prioritization with co-failure information.	✓		✓
E2	Dynamic test case prioritization with flipping history.	✓		✓
E3	Dynamic test case prioritization with rules mined from failure history.	✓		✓
F1	TERMINATOR with text feature.		✓	✓
F2	TERMINATOR with history feature.	✓		✓
F3	TERMINATOR with hybrid feature.	✓	✓	✓

In our experiments, $\alpha = 0.9$ achieves better performance than other values ranging from 0.1 to 1. Higher priority is assigned for higher B3 metrics. For example, B3 metrics for Run 4 in Table 1 are {0.09, 0.1, 0.9, 0.99} for t1 to t4. Therefore, B3 will sort the test cases in the order {t4, t3, t2, t1} before Run 4.

- **B4: ROCKET metrics [36]**

$$P_i = \sum_{j=1}^{i-1} \omega_{i-j} h_j$$

$$\omega_k = \begin{cases} 0.7, & \text{if } k = 1 \\ 0.2, & \text{if } k = 2 \\ 0.1, & \text{if } k \geq 3 \end{cases} \quad (4)$$

where $h_j = 0$ if the test case passed or was skipped in Run j and $h_j = 1$ if the test case failed in Run j . Higher priority is assigned for higher B4 metrics. For example, B4 metrics for Run 4 in Table 1 are {0.2, 0.3, 0.7, 0.8} for t1 to t4. Therefore, B4 will sort the test cases in the order {t4, t3, t2, t1} before Run 4.

- **B5: time since last execution.** Some algorithms assign higher priority to test cases that have not been executed for a long time [3, 16]. To apply this metric to our dataset, we utilized the number of consecutive skips before Run n . For example, using this metric, the value for Run 4 in Table 1 are {1, 0, 0, 0} for t1 to t4. Given that there are only a few skipped test cases in our dataset, this metric is never applied alone. Therefore, we followed the work of Aman et al. [3] using time since last execution and failure rate (B5) to prioritize the test cases by the Mahalanobis distance of the two metrics to the origin:

$$d_M(\mathbf{x}, (0, 0)) = \mathbf{x}^T S^{-1} \mathbf{x} \quad (5)$$

where S is the variance-covariance matrix of all the data $\mathbf{x}$ and S^{-1} is its inverse. For example, B5 metrics for t1 to t4 in Run 4 are {(1, 1/2), (0, 2/3), (0, 1/3), (0, 2/3)}. Their variance-covariance

matrix and its inverse are:

$$S = \begin{bmatrix} 0.250 & -0.014 \\ -0.014 & 0.025 \end{bmatrix}, \text{ and } S^{-1} = \begin{bmatrix} 4.124 & 2.25 \\ 2.25 & 40.5 \end{bmatrix}$$

Their Mahalanobis distances to the origin are {17.625, 18, 4.5, 18}. Therefore, B5 will sort the test cases in the order {t2, t4, t1, t3} before Run 4.

4.2.3 Group C: Cost-based Algorithms. C1 executes test cases in ascending order of their costs [23]. Usually, the cost is the runtime of each test case (which can be estimated from the execution history). For example, in Table 1, C1 will sort the test cases in the order {t1, t2, t3, t4} before Run 4 based on their increasing estimated cost from previous runs.

4.2.4 Group D: Test Case Description-based Algorithms. Group D algorithms utilize execution history information as dependent variables and test case description information as independent variables to build regression models and predict the probability that each test case will fail. The hypothesis behind this is that, if two test cases have similar descriptions, they tend to also have the same outcomes (fail/pass together). Text features are extracted from the test case description as described in §3.2. There are three ways to construct the dependent variables [22]:

- **D1: Simple History (SH).** Assign value 0 to the passed and skipped test cases and value 1 to the failed test cases in the immediately previous run. For example, the dependent variables of t1 to t4 of Run 4 in Table 1 will be {0,0,1,1}. A regression model will be trained to learn that test cases are more likely to fail if “page B” is in their description. Therefore, D1 will sort the test cases in the order {t3, t4, t1, t2} before Run 4.
- **D2: All History (AH).** Assign the failure rate value (B2) as the dependent variable. This measure considers the entire history rather than just the immediately previous run (SH). For example, the dependent variables of t1 to t4 of Run 4 in Table 1 will be {1/2, 2/3, 1/3, 2/3}. A regression model will be trained to learn that test

cases with “Radio Button” are more likely to fail than test cases with “Check Box” in their descriptions and test cases with “page A” are more likely to fail than test cases with “page B” in their descriptions. Therefore, D2 will sort the test cases in the order {t2, t4, t1, t3} before Run 4.

- **D3: Weighted History (WH).** Assign weighted averages (as opposed to simple averages of AH measures) of SH values in all previous runs placing more importance on recent runs. For example, ω_k in (4) from B4 can be applied as the weights and the dependent variables of t1 to t4 of Run 4 in Table 1 will be {0.2, 0.3, 0.7, 0.8}. A regression model will be trained to learn that test cases with “page B” are more likely to fail than test cases with “page A” in their descriptions and test cases with “Radio Button” are more likely to fail than test cases with “Check Box” in their descriptions. Therefore, D2 will sort the test cases in the order {t4, t3, t2, t1} before Run 4.

4.2.5 Group E: Feedback-based Algorithms. Group E algorithms utilize execution history information to find correlations between test cases, then dynamically order the test cases based on feedback information. For example, if t1 failed, then increase the priority of t3 because t1 and t3 are related.

- **E1: co-failure.** The idea behind this is that, if test A and test B failed together in the past 75% of the time, then if we observe a failure of test A in the current run, we may prioritize the execution of test B since there would be a high probability that test B will fail as well [59]. Given a just finished test case result t_{finished} , for every unexecuted test case $t \in T_{\text{unexecuted}}$:

$$\text{Priority}(t) = \text{Priority}(t) + (P(t = \text{fail}, t_{\text{finished}}) - 0.5)$$

where $P(t = \text{fail}, t_{\text{finished}})$ is the probability that t fails given that t_{finished} has passed/failed. For example, E1 initializes the priority of each test case as 0 for Run 4. Then it executes t1 first, and observes that t1 fails. The priority of t2, t3, and t4 will be updated to {0.5, -0.5, -0.5}. Therefore, t2 will be executed next. Given that t2 also fails, the priority of t3 and t4 will be updated to [-0.5, -1]. Then t3 will be executed and pass. The priority of t4 will be updated to -1 and finally t4 will be executed.

- **E2: Flipping history.** AFSAC [5] analyzes the execution history of the test suite and builds a correlation matrix of test cases. Two test cases are correlated when their results are changed to the opposite status by one commit in two consecutive test sessions (flipped together). The correlation matrix is composed of values reflecting the accumulated number of flipped results. AFSAC [5] uses an improved ROCKET metric to find the first failure before reordering the test cases based on flipping history. For example, in Table 1, E2 initializes the order of Run 4 as {t4, t3, t2, t1}, then executes t4 and t4 fails. The priority of t1, t2, and t3 will be updated to {1, 1, 1} given that t1, t2, and t3 all flipped with t4 together once in the previous runs (runs 1 to 3). Therefore, there is no adjustment to the order and t3 is executed next. Since t3 passes, there is no update on the priorities and t2 will be executed. After t2 fails, the priority of t1 is updated to $\max\{1, 1\} = 1$ since t1 also flipped with t2 together once. Finally t1 is executed.
- **E3: Rule-based.** REMAP [43] mines fail and pass rules over 90% confidence with the Repeated Incremental Pruning to Produce Error Reduction (RIPPER) algorithm, and dynamically prioritizes

the unexecuted test cases. If a test case fails, its corresponding fail rule test cases will be executed next; if a test case passes, its corresponding pass rule test cases will be pushed back to the end of the execution queue. The initial order of test cases is calculated by their failure rate (B2) and the number of rules associated. For example, in Table 1, with minimum support set to 2, one pass rule (if t2 fails then t3 will pass) and one fail rule (if t3 passes then t2 will fail) can be mined from Run 1 to Run 3. An initial order will be {t2, t3, t4, t1} given that t2 and t3 each has one rule associated and their failure rates are {1/2, 2/3, 1/3, 2/3}. After t2 fails, according to the pass rule, t3 is pushed back to the end of the execution order {t4, t1, t3}. Since there are no rules associated with t4 or t1, the order will not change after their execution.

4.2.6 Group F: TERMINATOR. Group F includes three variants of the proposed active learning based framework TERMINATOR. As shown in Table 2, different variants of TERMINATOR utilize different information.

Consider the example of Table 1, with $N_1 = 1$ and $N_2 = 2$. Here, TERMINATOR first randomly selects t2 to execute. After t2 fails, t1 is selected (randomly) as a presumptive non-relevant example, and a model is trained with t2 being positive (failure) and t1 being negative (non-failure). Suppose the prediction probabilities of being positive for t1, t3, and t4 are {0.1, 0.8, 0.6}. Since $|R| < N_2$, uncertainty sampling is applied to select t4 for execution since the model is most uncertain about the prediction of t4 (closest to 0.5). After t4 fails, t3 and t1 are selected as presumptive non-relevant examples. After aggressive undersampling, a model is trained with t2 and t4 being positive, and with t1 and t3 being negative. Suppose the prediction probabilities of being positive for t1 and t3 are {0.3, 0.2}. Since $|R| \geq N_2$, certainty sampling is applied to select t1 for execution since t1 has a higher predicted probability for being positive than t3.

Note that in the above example, F1, F2, and F3 will have different features; thus the predictions will be different, providing a different dynamic ordering of test cases.

4.3 Dependent Variables

To answer the two research questions, we collected two performance metrics to evaluate the performance of each algorithm:

- **APFDc:** average percentage of failure detected with cost, as calculated in (2) with severities of each failure set to be the same ($f_i = 1$). APFDc is applied to evaluate the failure detection rates and answer **RQ1**.
- **Overhead:** computation time of the algorithm / total runtime of all test cases. Overhead is collected to evaluate the extra computational cost of each TCP algorithm and answer **RQ2**.

To test the significance of any improvements, Scott-Knott analysis is applied to cluster and rank the two performance metrics of each algorithm from Run 6 to Run 54. It clusters algorithms with little difference in performance together and ranks each cluster with the median performances [45]. Nonparametric hypothesis tests are applied to handle the non-normal distribution. Specifically, Scott-Knott decided that, the two algorithms are not of little difference if both bootstrapping [13], and an effect size test (Cliff’s delta) [6], agreed that the difference is statistically significant (99% confidence) and not a negligible effect (Cliff’s Delta ≥ 0.147).

Table 3: Results for the two metrics described in §4.3. Overhead is presented as percentages. Medians and IQRs show the 50th and (75-25)th percentile results for 49 simulations runs. Results are divided into “ranks” (shown in the left most columns). Results have the same rank if statistical tests showed no significant difference. Note that the result of E3 is not shown because it took too much time to mine the association rules from 2661 test cases. E3 ran for 30 hours and was still not finished. As a result, E3 has an overhead of more than 100% and is considered not useful no matter how high an APFDc it can achieve.

(a) APFDc					(b) Overhead				
Rank	Treatment	Median	IQR		Rank	Treatment	Median	IQR	
1	E2	0.49	0.08	●	1	A2	0.02	0.01	●
1	A1	0.50	0.01	●	2	A1	0.2	0.1	●
1	C1	0.50	0.03	●	2	C1	0.21	0.07	●
1	E1	0.52	0.05	●	2	B1	0.21	0.11	●
2	D1	0.60	0.15	●	2	D1	0.22	0.08	●
2	D3	0.61	0.11	●	2	B4	0.22	0.11	●
2	F1	0.61	0.07	●	2	B3	0.22	0.10	●
2	D2	0.62	0.10	●	2	B2	0.22	0.08	●
2	B1	0.63	0.19	●	2	D2	0.22	0.10	●
3	F2	0.66	0.18	●	2	B5	0.22	0.12	●
3	B5	0.67	0.11	●	2	D3	0.22	0.11	●
3	B2	0.67	0.11	●	3	F2	0.23	0.07	●
3	B4	0.67	0.17	●	3	E2	0.24	0.07	●
3	B3	0.67	0.18	●	4	F3	0.33	0.11	●
4	F3	0.73	0.13	●	4	F1	0.34	0.10	●
5	A2	0.95	0.09	●	5	E1	5.79	2.52	●

4.4 Threats to Validity

This section discusses validity threats [18] to the above design.

External validity concerns how well the conclusion can be applied outside. All the conclusions in this study are drawn from 54 automated UI testing runs from LexisNexis. When applied to other case studies, the following concerns might arise: 1) TERMINATOR might not be applicable if not all three types of information are not available; 2) other TCP algorithms might perform better than TERMINATOR when using more information, e.g. coverage.

Internal validity focuses on how sure we can be that the treatment caused the outcome. To enhance internal validity, we heavily constrained our simulations to the same dataset with the same workflow for the TCP algorithms we evaluate. However, some of the algorithms in Table 2 are not exact implementations or reproductions of the original algorithms. Rather, they are only partially implemented to fit the workflow and information availability of our data. This may be a threat to internal validity. Another threat to internal validity comes from flaky tests. Apart from removing all the “time out” failures, we did not identify other types of flaky tests. These flaky tests could have affected the model learned as well as the evaluations. We plan to explore ways to tackle flaky tests problem in our future work.

Construct validity focuses on the relation between the theory behind the experiment and the observation. We use default parameters for the 17 algorithms in our simulations. For some of the algorithms, e.g. B3 and D1, different parameter settings can make large differences in failure detection rates. Therefore, it is possible that our observation does not hold when different parameters are applied for each algorithm. We would consider hyper-parameter tuning as future work to alleviate this concern.

Conclusion validity focuses on the significance of the treatment. To enhance conclusion validity, we applied Scott-Knott tests to the collective results of 49 runs to see whether one algorithm was significantly better than another on certain performance metrics.

4.5 Results

Table 3 shows the results of the 17 algorithms presented in Table 2. For each algorithm, APFDc and overhead from 49 runs are recorded and their medians and iqr are presented. Two algorithms perform significantly different in terms of one metric if they are in different ranks. APFDc evaluates failure detection rates, with higher numbers being better. Overhead represents the extra cost of running the algorithms, with lower numbers being better.

First, we analyze the APFDc results from Table 3 for **RQ1**:

- F3, which is TERMINATOR with the hybrid feature, performed the best in terms of APFDc. Its failure detection rates are significantly higher than those of the other algorithms (9% higher than a Rank 3 algorithm like B4).
- Simple history-based algorithms (B2, B3, B4, B5) are the second best algorithms in terms of APFDc. The more complex algorithms from Groups D or E that utilized more information performed worse than the simple history-based algorithms from Group B. Even though 9% higher APFDc does not seem to be a large improvement, it actually increases the failure detection rates substantially. As an example, Figure 4 shows the failure detection curves for the algorithms of each rank. Here we can see that TERMINATOR (F3) finds many more failures earlier than the other TCP algorithms, e.g., when 20% of the time is spent, F3 finds 60% of the failures while B4 finds 30%. By comparing the A2 and F3 results from Table 3, we can offer a precise measure for the success of the proposed method:

RQ1: can TERMINATOR achieve significantly higher failure detection rates than other TCP algorithms when prioritizing automated UI tests from LexisNexis?

As seen in the **Rank** column of Table 3, F3 (TERMINATOR with the hybrid feature) performed significantly better than other TCP methods. Further, TERMINATOR performed within $73/95 = 75\%$ of the optimal (A2) in median.

Figure 4: Failure detection curve of an example run at 02:02:34, October 24th, 2018. Here recall is the number of failures detected divided by the total number of failures in the test suite. An algorithm is considered better than another if it achieves higher recall with less time spent.

Second, we analyze the overhead results from Table 3 for RQ2:

- Even though F3 has higher overhead (50% higher than simple history-based algorithms from Group B), its overhead (35 seconds in median, which is only 0.33% of the total runtime of the test suite) is negligible compared to the cost it can potentially save.
- Algorithms utilizing feedback information (Groups E and F) tend to have higher overhead since the orders of test cases are dynamically adjusted based on the execution results.

RQ2: what is the computational overhead of TERMINATOR?

The computational overhead of TERMINATOR is 0.33% of the total runtime of the test suite, which is negligible compared to the cost it can potentially save by achieving higher failure detection rates.

The reasons for the superior performance of TERMINATOR are:

- (1) TERMINATOR has an advanced framework adapted from the solutions of the total recall problem which utilizes different data balancing techniques and query strategies.
- (2) TERMINATOR utilizes more information (test history, test case description, and current feedback) than other algorithms.

That said, there are some limitations to the current results. Specifically, compared to the optimal result, the APFDc of F3 is still lower than that of A2. That is, there is still room to improve this result.

5 ORGANIZATIONAL IMPACTS

With higher failure detection rates provided by TERMINATOR, LexisNexis teams can build more versions of tested software within one eight hour shift. This will have three major effects:

- *Higher responsiveness to customer feedback.* The faster LexisNexis can build and test systems, the faster their customers can see new features (or fixes to old features).
- *Greater programmer adaptability.* The faster programmers can get feedback on their systems, the better they can become in software adaption and extension.

- *Increased recruitment and retention of programmers.* Here at NC State, many of our graduates work in the large and local information technology industry. Based on feedback from that population, we can assert that organizations that ship products faster can also recruit and retain more software developers.

6 CONCLUSION AND FUTURE WORK

Automated UI testing is an important component of the continuous integration process of software product development. Faster automated UI testing leads to shorter development cycles and better quality. This paper studies the specific problem of how to prioritize automated UI test cases. By analyzing what information is available when prioritizing automated UI test cases, 12 state of the art TCP algorithms are found to be applicable to automated UI test case prioritization. In addition, prioritizing test cases for failure detection rates can be generalized as a total recall problem. A new TCP algorithm is proposed by adapting the active learning based framework from the total recall problem to the TCP problem. Utilizing the information of execution history, test case description, and feedback, TERMINATOR (F3) outperformed all other TCP algorithms in terms of failure detection rates (APFDc) on a dataset of 54 consecutive runs of real world automated UI testing in LexisNexis. Lessons learned from this work are summarized as follows:

- Automated UI testing is widely applied in industry and consumes a large portion of time and resources. Optimizing the execution of automated UI testing deserves much more attention in research.
- Black box test case prioritization (given no information on source code) for failure detection rates is under-explored in the research arena. This is strange since, at least in our experience, this kind of testing is more often than not the only practical approach for large organizations like LexisNexis.
- The test case prioritization for failure detection rates problem can be generalized as a total recall problem. Techniques from the total recall problem can be applied to improve performance on the TCP problem. This also suggests that future improvements in TCP can in return be applied to help solve other total recall problems like citation screening, vulnerability inspection, and static warning identification.

Given the validity threats and limitations of current work, our future work includes:

- (1) Hyper-parameter tuning for different TCP algorithms including the proposed active learning based framework to see if better configurations can be found.
- (2) Identify flaky test cases and avoid their impact on the learned active learning model.
- (3) It is tedious and time-consuming to classify each test failure to a specific fault caused by some piece of the source code. The study of fault localization might help address this problem.
- (4) Apply TERMINATOR to other TCP problems.
- (5) Follow up with LexisNexis on how to integrate the proposed framework with their current automated UI testing system to obtain higher failure detection rates.

ACKNOWLEDGEMENT

This work was partially funded by a gift from LexisNexis managed by Phillippe Poignant and an NSF CCF grant #1703487.

REFERENCES

- [1] 1994. Company News; A Name Change is Planned for Mead Data Central. *The New York Times*. (1994). <https://www.nytimes.com/1994/12/02/business/company-news-a-name-change-is-planned-for-mead-data-central.html?src=pm>
- [2] 2019. LEGAL. *www.relx.com* (2019). <https://www.relx.com/our-business/market-segments/legal>
- [3] Hirohisa Aman, Yuta Tanaka, Takashi Nakano, Hideto Ogasawara, and Minoru Kawahara. 2016. Application of Mahalanobis-Taguchi method and 0-1 programming method to cost-effective regression testing. In *2016 42th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 240–244.
- [4] Árpád Beszéd, Tamás Gergely, Lajos Schrettnér, Judit Jász, László Langó, and Tibor Gyimóthy. 2012. Code coverage-based regression test selection and prioritization in WebKit. In *2012 28th IEEE international conference on software maintenance (ICSM)*. IEEE, 46–55.
- [5] Younghwan Cho, Jeongho Kim, and Eunseok Lee. 2016. History-based Test Case Prioritization for Failure Information. In *2016 23rd Asia-Pacific Software Engineering Conference (APSEC)*. IEEE, 385–388.
- [6] Norman Cliff. 1993. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin* 114, 3 (1993), 494.
- [7] Gordon V Cormack and Maura R Grossman. 2014. Evaluation of machine-learning protocols for technology-assisted review in electronic discovery. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*. ACM, 153–162.
- [8] Gordon V Cormack and Maura R Grossman. 2015. Autonomy and Reliability of Continuous Active Learning for Technology-Assisted Review. *arXiv preprint arXiv:1504.06868* (2015).
- [9] Gordon V Cormack and Maura R Grossman. 2016. Engineering Quality and Reliability in Technology-Assisted Review. (2016), 75–84.
- [10] Gordon V Cormack and Maura R Grossman. 2016. Scalability of Continuous Active Learning for Reliable High-Recall Text Classification. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. ACM, 1039–1048.
- [11] Gordon V Cormack and Maura R Grossman. 2017. Navigating Imprecision in Relevance Assessments on the Road to Total Recall: Roger and Me. In *The International ACM SIGIR Conference*. 5–14.
- [12] Hyunsook Do and Gregg Rothermel. 2006. On the use of mutation faults in empirical assessments of test case prioritization techniques. *IEEE Transactions on Software Engineering* 32, 9 (2006), 733–752.
- [13] Bradley Efron. 1982. *The jackknife, the bootstrap and other resampling plans*. SIAM.
- [14] Nandhini Elanthiraiyan and Chamundeswari Arumugam. 2014. Parallelized ACO algorithm for regression testing prioritization in hadoop framework. In *Advanced Communication Control and Computing Technologies (ICACCCT), 2014 International Conference on*. IEEE, 1568–1571.
- [15] Sebastian Elbaum, Alexey Malishevsky, and Gregg Rothermel. 2001. Incorporating varying test costs and fault severities into test case prioritization. In *Proceedings of the 23rd International Conference on Software Engineering*. IEEE Computer Society, 329–338.
- [16] Sebastian Elbaum, Gregg Rothermel, and John Penix. 2014. Techniques for Improving Regression Testing in Continuous Integration Development Environments. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE 2014)*. ACM, New York, NY, USA, 235–245. <https://doi.org/10.1145/2635868.2635910>
- [17] Y Fazlalizadeh, A Khalilian, M Abdollahi Azgomi, and S Parsa. 2009. Prioritizing test cases for resource constraint environments using historical test case performance data. In *Computer Science and Information Technology, 2009. ICCSIT 2009. 2nd IEEE International Conference on*. IEEE, 190–195.
- [18] Robert Feldt and Ana Magazinius. 2010. Validity Threats in Empirical Software Engineering Research—An Initial Survey.. In *SEKE*. 374–379.
- [19] M. R. Grossman and G. V. Cormack. 2013. The Grossman-Cormack glossary of technology-assisted review with foreword by John M. Facciola, U.S. Magistrate Judge. *Federal Courts Law Review* 7, 1 (2013), 1–34.
- [20] Maura R Grossman, Gordon V Cormack, and Adam Roegiest. 2016. TREC 2016 Total Recall Track Overview. In *TREC*.
- [21] Hadi Hemmati, Zhihan Fang, and Mika V Mantyla. 2015. Prioritizing manual test cases in traditional and rapid release environments. In *Software Testing, Verification and Validation (ICST), 2015 IEEE 8th International Conference on*. IEEE, 1–10.
- [22] Hadi Hemmati and Fatemeh Sharifi. 2018. Investigating NLP-Based Approaches for Predicting Manual Test Case Failure. In *Software Testing, Verification and Validation (ICST), 2018 IEEE 11th International Conference on*. IEEE, 309–319.
- [23] Kim Herzig, Michaela Greiler, Jacek Czerwinka, and Brendan Murphy. 2015. The Art of Testing Less without Sacrificing Quality. In *IEEE/ACM IEEE International Conference on Software Engineering*.
- [24] Yen-Ching Hsu, Kuan-Li Peng, and Chin-Yu Huang. 2014. A study of applying severity-weighted greedy algorithm to software test case prioritization during testing. In *Industrial Engineering and Engineering Management (IEEM), 2014 IEEE International Conference on*. IEEE, 1086–1090.
- [25] Passant Kandil, Sherin Moussa, and Nagwa Badr. 2015. A methodology for regression testing reduction and prioritization of agile releases. In *Information & Communication Technology and Accessibility (ICTA), 2015 5th International Conference on*. IEEE, 1–6.
- [26] Jonathan Miller Kauffman and Gregory M Kapfhammer. 2012. A framework to support research in and encourage industrial adoption of regression testing techniques. In *Software Testing, Verification and Validation (ICST), 2012 IEEE Fifth International Conference on*. IEEE, 907–908.
- [27] Jung-Min Kim and Adam Porter. 2002. A history-based test prioritization technique for regression testing in resource constrained environments. In *Proceedings of the 24th international conference on software engineering*. ACM, 119–129.
- [28] Jung-Hyun Kwon and In-Young Ko. 2017. Cost-Effective Regression Testing Using Bloom Filters in Continuous Integration Development Environments. In *Asia-Pacific Software Engineering Conference (APSEC), 2017 24th*. IEEE, 160–168.
- [29] Remo Lachmann, Sandro Schulze, Manuel Niek, Christoph Seidl, and Ina Schaefer. 2016. System-Level Test Case Prioritization Using Machine Learning. In *Machine Learning and Applications (ICMLA), 2016 15th IEEE International Conference on*. IEEE, 361–368.
- [30] Zheng Li, Mark Harman, and Robert M Hierons. 2007. Search algorithms for regression test case prioritization. *IEEE Transactions on software engineering* 33, 4 (2007).
- [31] Jingjing Liang, Sebastian Elbaum, and Gregg Rothermel. 2018. Redefining Prioritization: Continuous Prioritization for Continuous Integration. In *Proceedings of the 40th International Conference on Software Engineering (ICSE '18)*. ACM, New York, NY, USA, 688–698. <https://doi.org/10.1145/3180155.3180213>
- [32] Wenhong Liu, Xin Wu, WeiXiang Zhang, and Yang Xu. 2014. The research of the test case prioritization algorithm for black box testing. In *Software Engineering and Service Science (ICSESS), 2014 5th IEEE International Conference on*. IEEE, 37–40.
- [33] Yiling Lou, Dan Hao, and Lu Zhang. 2015. Mutation-based test-case prioritization in software evolution. In *Software Reliability Engineering (ISSRE), 2015 IEEE 26th International Symposium on*. IEEE, 46–57.
- [34] Yafeng Lu, Yiling Lou, Shiyang Cheng, Lingming Zhang, Dan Hao, Yangfan Zhou, and Lu Zhang. 2016. How does regression test prioritization perform in real-world software evolution?. In *Software Engineering (ICSE), 2016 IEEE/ACM 38th International Conference on*. IEEE, 535–546.
- [35] Surendra Mahajan, Shashank D Joshi, and V Khanaa. 2015. Component-based software system test case prioritization with genetic algorithm decoding technique using java platform. In *Computing Communication Control and Automation (ICCCUBEA), 2015 International Conference on*. IEEE, 847–851.
- [36] Duscica Marijan, Arnaud Gotlieb, and Sagar Sen. 2013. Test case prioritization for continuous regression testing: An industrial case study. In *Software Maintenance (ICSM), 2013 29th IEEE International Conference on*. IEEE, 540–543.
- [37] Lijun Mei, Yan Cai, Changjiang Jia, Bo Jiang, Wing Kwong Chan, Zhenyu Zhang, and TH Tse. 2015. A subsumption hierarchy of test case prioritization for composite services. *IEEE Transactions on Services Computing* 8, 5 (2015), 658–673.
- [38] Stephen Miller. 2012. For Future Reference, a Pioneer in Online Reading. *The Wall Street Journal* (2012). <https://www.wsj.com/articles/SB10001424052970203721704577157211501855648?KEYWORDS=lexisnexis>
- [39] Tim Miller et al. 2013. Using dependency structures for prioritization of functional test suites. *IEEE transactions on software engineering* 39, 2 (2013), 258–275.
- [40] Siavash Mirarab and Ladan Tahvildari. 2008. An empirical study on bayesian network-based approach for test case prioritization. In *2008 International Conference on Software Testing, Verification, and Validation*. IEEE, 278–287.
- [41] Reetika Nagar, Arvind Kumar, Gaurav Pratap Singh, and Sachin Kumar. 2015. Test case selection and prioritization using cuckoo search algorithm. In *Futuristic Trends on Computational Analysis and Knowledge Management (ABLAZE), 2015 International Conference on*. IEEE, 283–288.
- [42] David Paterson, Gregory M Kapfhammer, Gordon Fraser, and Phil McMinn. 2018. Using Controlled Numbers of Real Faults and Mutants to Empirically Evaluate Coverage-Based Test Case Prioritization. In *Proceedings of the International Workshop on Automation of Software Test (AST 2018)*. IEEE.
- [43] Dipesh Pradhan, Shuai Wang, Shaikat Ali, Tao Yue, and Marius Liaaen. 2018. REMAP: Using Rule Mining and Multi-objective Search for Dynamic Test Case Prioritization. In *Software Testing, Verification and Validation (ICST), 2018 IEEE 11th International Conference on*. IEEE, 46–57.
- [44] Gregg Rothermel, Roland H. Untch, Chengyun Chu, and Mary Jean Harrold. 2001. Prioritizing test cases for regression testing. *IEEE Transactions on software engineering* 27, 10 (2001), 929–948.
- [45] ANDREW JHON Scott and M Knott. 1974. A cluster analysis method for grouping means in the analysis of variance. *Biometrics* (1974), 507–512.
- [46] Burr Settles. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 6, 1 (2012), 1–114.
- [47] Chang-ai Sun, Cuiyang Fan, Zhen Wang, and Huai Liu. 2017. dμReg: a path-aware mutation analysis guided approach to regression testing. In *Proceedings of the 12th International Workshop on Automation of Software Testing*. IEEE Press, 59–64.

- [48] Wei-Tek Tsai, Xiaoying Bai, Yinong Chen, and Xinyu Zhou. 2005. Web service group testing with windowing mechanisms. In *IEEE International Workshop on Service-Oriented System Engineering (SOSE'05)*. IEEE, 213–218.
- [49] Ashlee Vance. 2010. Legal Sites Plan Revamps as Rivals Undercut Price. *The New York Times*. (2010). https://www.nytimes.com/2010/01/25/technology/25westlaw.html?_r=1&ref=reedelsevier
- [50] Byron C Wallace and Issa J Dahabreh. 2012. Class probability estimates are unreliable for imbalanced data (and how to fix them). In *Data Mining (ICDM), 2012 IEEE 12th International Conference on*. IEEE, 695–704.
- [51] Byron C Wallace, Issa J Dahabreh, Kelly H Moran, Carla E Brodley, and Thomas A Trikalinos. 2013. Active literature discovery for scoping evidence reviews: How many needles are there. In *KDD workshop on data mining for healthcare (KDD-DMH)*.
- [52] Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. 2010. Active learning for biomedical citation screening. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 173–182.
- [53] Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. 2011. Who Should Label What? Instance Allocation in Multiple Expert Active Learning.. In *SDM*. SIAM, 176–187.
- [54] Byron C Wallace, Thomas A Trikalinos, Joseph Lau, Carla Brodley, and Christopher H Schmid. 2010. Semi-automated screening of biomedical citations for systematic reviews. *BMC bioinformatics* 11, 1 (2010), 1.
- [55] Zhe Yu. [n.d.]. PhD thesis: Total Recall and Software Engineering. ([n. d.]).
- [56] Zhe Yu, Nicholas A Kraft, and Tim Menzies. 2018. Finding better active learners for faster literature reviews. *Empirical Software Engineering* 23, 6 (2018), 3161–3186.
- [57] Zhe Yu and Tim Menzies. 2019. FAST2: An intelligent assistant for finding relevant papers. *Expert Systems with Applications* 120 (2019), 57 – 71.
- [58] Lingming Zhang, Ji Zhou, Dan Hao, Lu Zhang, and Hong Mei. 2009. Prioritizing JUnit test cases in absence of coverage information. In *Software Maintenance, 2009. ICSM 2009. IEEE International Conference on*. IEEE, 19–28.
- [59] Yuecai Zhu, Emad Shihab, and Peter C Rigby. 2018. Test Re-prioritization in Continuous Testing Environments. In *2018 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 69–79.