



# MuRS: Mutant Ranking and Suppression using Identifier Templates

Zimin Chen\*

zimin@kth.se

KTH Royal Institute of Technology  
Sweden

Goran Petrović

goranpetrovic@google.com

Google  
Switzerland

Małgorzata Salawa

magorzata@google.com

Google  
Switzerland

Marko Ivanković

markoi@google.com

Google  
Switzerland

Manushree Vijayvergiya

manushree@google.com

Google  
Switzerland

René Just\*

rjust@cs.washington.edu

University of Washington  
USA

## ABSTRACT

Diff-based mutation testing is a mutation testing approach that only mutates lines affected by a code change under review. This approach scales independently of the code-base size and introduces test goals (mutants) that are directly relevant to an engineer’s goal such as fixing a bug, adding a new feature, or refactoring existing functionality. Google’s mutation testing service integrates diff-based mutation testing into the code review process and continuously gathers developer feedback on mutants surfaced during code review. To enhance the developer experience, the mutation testing service uses a number of manually-written rules that suppress not-useful mutants—mutants that have consistently received negative developer feedback. However, while effective, manually implementing suppression rules requires significant engineering time.

This paper proposes and evaluates MuRS, an automated approach that groups mutants by patterns in the source code under test and uses these patterns to rank and suppress future mutants based on historical developer feedback on mutants in the same group. To evaluate MuRS, we conducted an A/B testing study, comparing MuRS to the existing mutation testing service. Despite the strong baseline, which uses manually-written suppression rules, the results show a statistically significantly lower negative feedback ratio of 11.45% for MuRS versus 12.41% for the baseline. The results also show that MuRS is able to recover existing suppression rules implemented in the baseline. Finally, the results show that statement-deletion mutant groups received both the most positive and negative developer feedback, suggesting a need for additional context that can distinguish between useful and not-useful mutants in these groups. Overall, MuRS is able to recover existing suppression rules and automatically learn additional, finer-grained suppression rules from developer feedback.

\*Work done at Google.



This work is licensed under a Creative Commons Attribution 4.0 International License.

ESEC/FSE '23, December 3–9, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0327-0/23/12.

<https://doi.org/10.1145/3611643.3613901>

## CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging.**

## KEYWORDS

Mutation Testing, Developer Feedback, Code Review

### ACM Reference Format:

Zimin Chen, Małgorzata Salawa, Manushree Vijayvergiya, Goran Petrović, Marko Ivanković, and René Just. 2023. MuRS: Mutant Ranking and Suppression using Identifier Templates. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '23)*, December 3–9, 2023, San Francisco, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3611643.3613901>

## 1 INTRODUCTION

Software testing is an essential part of software development that validates and verifies the software under test. A well tested program is an indication of a reliable program [13]. But how do we assess whether a program is well tested? One way is to use mutation testing [1], which systematically introduces small changes into the program and checks whether the program’s tests detect the changes. Mutation testing has been widely studied [8] and sees increasing adoption in industry [2, 19].

Mutation testing systematically applies mutation operators that create program variants called mutants. Each mutant differs from the original program by a small change such as a changed literal, a changed operator, or a deleted statement. Even a small number of mutation operators, when systematically applied to an entire program, can generate vast numbers of mutants, and evaluating all of them is computationally expensive [23]—prohibitively so for an industry-level code base. To solve this problem, *diff-based mutation testing* [19] incrementally mutates, during code review, only source code lines that are affected by the code change under review.

Google has successfully deployed diff-based mutation testing for a code base that sees more than 40 000 code commits per day to more than a billion lines of code [19]. The corresponding code-review tool offers an option for authors and reviewers to provide feedback for mutants surfaced during code review (see Figure 1). If a reviewer clicks the “Please fix” button, the code-review tool auto-generates a review comment, asking the author to resolve the mutant. Additionally, both authors and reviewers can click “Thumbs up” or “Thumbs down” to give positive or negative feedback for

a surfaced mutant. The negative feedback ratio is regularly measured and serves as an indicator for effective false positives [22]. To lower that ratio, the existing mutation testing service implements a number of static suppression rules that prevent mutants that are likely to receive negative feedback from being surfaced or even generated, and it prioritizes mutation operators based on the historical effectiveness of the mutants they generate. Over the years, the manually developed suppression and prioritization rules have reduced the negative feedback ratio from over 80% to well below 15% [19]. However, manually implementing rules is labor-intensive, the rules require regular maintenance, and prioritization based on mutation operators may be too coarse-grained.

This paper presents MuRS (**M**utant **R**anking & **S**uppression) that aims to tackle these problems and further reduce the negative feedback ratio. MuRS uses identifier templates—an abstraction over the unified diff between the mutant and the original program. Specifically, MuRS takes, for a given mutant, the unified diff and replaces (1) literals with their type names (e.g., `int` or `String`) and (2) identifiers with `IDENTIFIER`. MuRS optionally preserves the most common literals and identifiers, using a configurable vocabulary size. Similarly, the number of context lines in the unified diff considered when building identifier templates is another hyper-parameter. MuRS then groups all historical mutants and aggregates their developer feedback by identifier template. The aggregated feedback for each identifier template is then used to rank and suppress newly generated mutants. The intuition behind this approach is that mutants whose identifier templates have largely received positive feedback should be ranked highly, whereas mutants whose identifier templates have received negative feedback should be ranked lower or suppressed altogether. MuRS scales very well to a large code base: our analyses computed millions of identifier templates on a machine with commodity hardware and only 16 GB of RAM.

We conducted an A/B testing study to evaluate MuRS' performance, randomly assigning 50% of code reviews at Google during the evaluation period to the experiment group (MuRS) and all other code reviews to a control group (existing mutation testing service). During the evaluation period, a total of 666 143 mutants surfaced across the five languages supported by MuRS (Python, Java, C++, Go, and TypeScript). The key results are as follows:

- The overall negative feedback ratio for the experiment group is 11.45%, which is statistically significantly lower than the negative feedback ratio of 12.41% for the control group.
- The correlation coefficient (Kendall's Tau-b) between MuRS's ranking score and the ranked mutants' perceived usefulness (developer feedback) is 0.2615, which indicates a moderate correlation. Further, the top 50% of ranked mutants only contain 16.39% of mutants with negative feedback. This suggests that MuRS's ranking of mutants effectively decreases the likelihood of surfacing not-useful mutants.
- When applied to all surfaced mutants in the control group, MuRS would have suppressed 49% of mutants that received negative feedback, and it would have suppressed 12% of mutants that received positive feedback. This means that, when deploying MuRS, choosing a suppression threshold involves a trade-off between correctly suppressing not-useful mutants, which waste engineers' time, and incorrectly suppressing useful mutants, which likely lead to additional tests.



**Figure 1:** A mutant surfaced in the code review tool. A reviewer can click “Please fix”, indicating that the author should resolve the mutant. Both author and reviewers can click “Thumbs up” and “Thumbs down” to give feedback.

- For all supported languages, the identifier templates that receive the most controversial feedback (i.e., the most positive and the most negative feedback) are all statement removal mutants. Prior research has found that statement removal mutants are coupled with real faults more often than other types of mutants [10]. This finding suggests that although statement removal mutants are useful in general, further investigation is necessary to identify the specific subset of statement removal mutants that are not useful.

We additionally performed a retrospective analysis, investigating what identifier templates MuRS would have learned and scored at different points in the past six years. A manual analysis of the most common identifier templates that consistently received negative feedback showed that MuRS is able to recover suppression rules implemented in the existing mutation testing service.

## 2 TERMINOLOGY

This section defines the terminology used throughout the paper:

- **Changelist:** A set of changes made to the code base, which are submitted for review (similar to a merge request).
- **Snapshot:** A specific point in time, capturing the state of a changelist as it evolves during the code review process. Each changelist has one or more snapshots.
- **Generated mutants:** All mutants generated by the mutation testing service.
- **Surviving mutants:** Mutants that are not detected by any of the existing tests—a subset of all generated mutants.
- **Surfaced mutants:** Mutants displayed in the code review tool—a subset of surviving mutants. The number of surfaced mutants is limited to avoid overwhelming developers.
- **Positive feedback:** A surfaced mutant receives positive feedback when the “Please Fix” or “Thumbs Up” button is clicked in Figure 1; we deem such mutants useful.
- **Negative feedback:** A surfaced mutant receives negative feedback when the “Thumbs Down” button is clicked in Figure 1; we deem such mutants not useful.

### 3 MUTATION TESTING AT GOOGLE

Integrating mutation testing into a large code base, like Google’s, poses significant challenges related to scalability and workflow integration. Generating and testing all possible mutants is computationally infeasible, and reporting all surviving mutants to developers would be overwhelming. To overcome these challenges, Google’s mutation testing service implements *diff-based mutation testing*: it only generates mutants for the lines of code that have been changed in a given changelist. It also implements *arid node detection* and *probabilistic mutation-operator selection* to surface mutants that developers mostly consider useful in practice [19].

In Figure 2, “Mutation Testing Service” illustrates the entire process, from submitting a changelist for code review to surfacing mutants as findings in the code review tool. A detailed report [19] provides a comprehensive description of the mutation testing service. This section provides a high-level overview:

- (1) A developer submits a changelist for code review.
- (2) The mutation testing service identifies all changed lines that are covered by at least one test.
- (3) The mutation testing service builds the abstract syntax tree of each affected file and marks arid nodes that are not eligible for mutation. Mutating arid nodes produces non-actionable mutants, such as String mutations in logging statements, for which developers justifiably would not write tests.
- (4) The mutation testing service uses *probabilistic mutation-operator selection* that prioritizes mutation operators based on their historical performance until at most one mutant is generated for each affected line. The historical performance is computed for each mutation operator and programming language, which is relatively coarse grained.
- (5) The mutation testing service identifies surviving mutants by executing the existing tests against each generated mutant.
- (6) The mutation testing service surfaces a random sample of surviving mutants in the code review tool, based on per file and per changelist thresholds.

MuRS aims to address two limitations of the current mutation testing service, related to the implementation of arid nodes (step 3) and the probabilistic mutation-operator selection (step 4). While arid-node detection is effective, it currently requires human effort to identify and implement a static rule for suppressing a certain group of mutants. This implementation is specific to each programming language and must be manually maintained as the code base evolves. As for probabilistic mutation-operator selection, the problem is that the selection is coarse-grained and context-agnostic. Currently, the mutation testing service implements five mutation operators: AOR (Arithmetic operator replacement), LCR (Logical connector replacement), ROR (Relational operator replacement), UOI (Unary operator insertion), and SBR (Statement block removal). However each mutation operator corresponds to a number of possible mutations, such as replacing addition with subtraction, division, or multiplication. In contrast, MuRS considers the actual mutation and the code before and after as additional context.

### 4 MURS: MUTANT RANKING & SUPPRESSION

MuRS aims to reduce the negative feedback ratio of surfaced mutants. The **negative feedback ratio** is defined as the number of

**Table 1: The attributes collected in the *Data Collection* stage.**

Attribute	Explanation
<i>CL</i>	The changelist ID
<i>Filename</i>	The name of the mutated file
<i>Diff</i>	The diff between the original file and the mutant in unified format with $N$ context lines
<i>PosFeedList</i>	A list of type boolean (one per snapshot) indicating whether a mutant received a positive feedback at that snapshot
<i>NegFeedList</i>	A list of type boolean (one per snapshot) indicating whether a mutant received a negative feedback at that snapshot
<i>KilledList</i>	A list of type boolean (one per snapshot) indicating whether a mutant was killed at that snapshot

surfaced mutants with negative feedback divided by the total number of surfaced mutants with positive or negative feedback.

MuRS implements a template-based approach for five languages (C++, Java, Go, Python and TypeScript) that ranks and suppresses mutants based on historical developer feedback. It consists of two phases, as shown in the upper part of Figure 2:

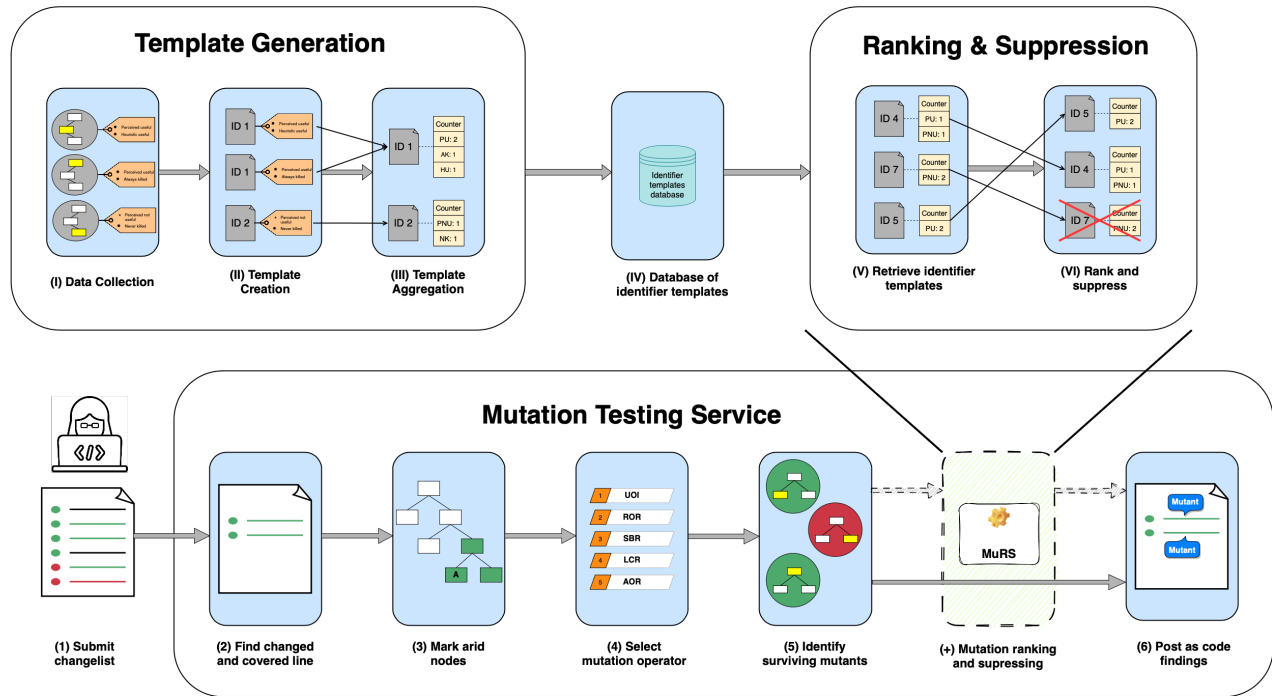
- (1) **Template Generation:** During this phase, MuRS generates identifier templates from all past mutants, associating each identifier template with a list of metrics such as number and type of received feedback, and number of times it was generated and killed. This phase is intended to run periodically to collect identifier templates with all statistics until that point in time and store them in a database. The time interval for template re-computation can be chosen based on needs and resource constraints, balancing data freshness with an acceptable resource consumption. In our case, this phase runs every two months.
- (2) **Ranking & Suppression:** During this phase, MuRS uses the latest identifier templates stored in a database to rank and suppress mutants in production. Specifically, it uses the list of metrics to represent the past performance associated with each identifier template to rank and suppress surviving mutants that match an identifier template.

There are different ways to generate the identifier templates and to rank and suppress mutants. Sections 4.1 and 4.2 discuss these configuration options in detail, and section 4.3 describes the configuration that we chose based on a preliminary evaluation. All of MuRS’ identifier templates can be stored in a lookup table in memory, using less than 16GB of RAM.

#### 4.1 Template Generation

Template Generation consists of three stages, as shown in Figure 2:

- (I) **Data Collection:** Gather all previously generated mutants along with their associated attributes (see Table 1).
- (II) **Template Creation:** Create identifier templates for all previously generated mutants.
- (III) **Template Aggregation:** Aggregate mutants with the same identifier templates.



**Figure 2: Mutation Testing Service (bottom) shows the end-to-end workflow from submitting a changelist for code review to surfacing mutants as code findings in the code review tool:** (1) A changelist is submitted for review by a developer. Red and green dots indicate that a line is (not) covered by existing tests. A green, red, or black line indicates that it is added, removed, or unchanged, respectively. (2) Lines that are both changed and covered are identified. (3) Arid nodes, which are not eligible for mutation, are marked with ‘A’. (4) Mutation operators are applied in order to generate at most one mutant per changed line. (5) Test cases are run and surviving mutants are identified. (6) A random subset of surviving mutants are surfaced as code findings. MuRS, is added between (5) and (6): it ranks and suppresses surviving mutants as opposed to randomly sampling from all surviving mutants. MuRS consists of two phases, Template Generation and Ranking & Suppression. Template Generation consists of three stages: (I) Historical mutants and their feedback are collected. (II) Identifier templates are generated for each mutant. (III) Mutants are grouped by identifier template and feedback scores are aggregated. (IV) The identifier templates are stored in a database. Ranking & Suppression consists of two stages: (V) The aggregated feedback score is retrieved from the database for each surviving mutant’s identifier template. (VI) Surviving mutants are ranked and suppressed based on their feedback score.

**4.1.1 Data Collection.** In this stage, MuRS gathers mutants and their attributes from a database of all previously generated mutants. MuRS then collects various attributes (Table 1) for each mutant and assigns two kinds of labels to each mutant, based on said attributes: *Perceived feedback* and *Killed status*.

*Perceived feedback* represents a mutant’s usefulness as perceived by developers. It is derived from *PosFeedList* and *NegFeedList*:

- **Perceived useful:** *PosFeedList* contains at least one *True* and *NegFeedList* is all *False*.
- **Perceived not-useful:** *PosFeedList* is all *False* and *NegFeedList* contains at least one *True*.
- **Mixed feedback:** Both *PosFeedList* and *NegFeedList* contain at least one *True*.
- **No feedback:** Both *PosFeedList* and *NegFeedList* only contain *False*.

*Killed status* is a label representing whether a mutant was killed (or not) throughout the code review. It is derived from *KilledList*:

- **Always killed:** All entries in *KilledList* are *True*.
- **Never killed:** All entries in *KilledList* are *False*.
- **Eventually killed:** The *KilledList* is a sequence of *False*, followed by a sequence of *True*, which means that a mutant survived up to a certain point in the code review process and was consistently killed afterwards.
- **Mixed killed:** None of the above conditions are met. This can happen when the entries in *KilledList* alternate between *False* and *True*.

At the end of this stage, each mutant is assigned an appropriate *Perceived feedback* and *Killed status* label.

**4.1.2 Template Creation.** In this stage, MuRS generates identifier templates for all labeled mutants based on the *Diff* attribute.

We define an abstraction, which we refer to as an *identifier template*, over the diff (in unified format) between the original and the mutated code. This abstraction is line-based and removes comments from the code before processing the diff, including context

```
def is_even(x):
- return x % 2 == 0
+ return x % 2 != 0

def is_odd(x):
```

### Listing 1: An example of a mutant in Python.

lines (where applicable). Note that this abstraction is only one point in the design space and a limitation is that it only matches syntactic similarities but disregards semantic similarities. We chose this abstraction because it is relatively fast to compute and allows for quick inference, even in the presence of millions of mutants. The template’s specificity can be controlled using different abstraction levels and various parameters, which we subsequently explain using the mutant in Listing 1 as a running example.

**Abstraction levels** that we explored are:

- **Original code template:** Use each changed line as is.

```
- return x % 2 == 0
+ return x % 2 != 0
```

- **Typed template:** Replace all literals with their corresponding type name, and identifiers with *IDENTIFIER*. For instance, integer literals are replaced with *INT*. However, all keywords such as *while*, *for* and *return* are kept.

```
- return IDENTIFIER % INT == INT
+ return IDENTIFIER % INT != INT
```

- **Indexed typed template:** Similar to Typed template, but index each type name. This means that each name is followed by a number, in order of appearance, to differentiate between identifiers of the same type. For example, the two integers, 2 and 0, are replaced with *INT\_0* and *INT\_1*, respectively.

```
- return IDENTIFIER % INT_0 == INT_1
+ return IDENTIFIER % INT_0 != INT_1
```

**Identifier template parameters** that we explored are:

- **Context size:** The number of lines preceding and succeeding the mutated line(s). Larger context sizes lead to more specificity, because more context lines are included and the corresponding identifier template matches fewer, but more specific types of mutants. The following identifier template corresponds to an Indexed typed template with a context size of 1.

```
def IDENTIFIER_0(IDENTIFIER_1):
- return IDENTIFIER_1 % INT_0 == INT_1
+ return IDENTIFIER_1 % INT_0 != INT_1
```

- **Vocabulary size:** The number of the most common variable names and literals to keep without abstraction. Larger vocabulary sizes lead to more specificity. For example, if *numpy* is part of the vocabulary, mutants in the context of the identifier *numpy* can be ranked and suppressed differently than similar mutants in other contexts. The following identifier

template corresponds to an Indexed typed template with a vocabulary containing *x*, but not containing 2 and 0.

```
- return x % INT_0 == INT_1
+ return x % INT_0 != INT_1
```

**4.1.3 Template Aggregation.** Similar mutants result in the same identifier template, depending on the abstraction level and identifier template parameters. In this stage, MuRS aggregates all labeled identifier template instances—counting the number of *Perceived feedback* and *Killed status* labels. *Perceived feedback*: PU, PNU, MF, and NF count perceived useful, perceived not-useful, mixed feedback and no feedback; *Killed status*: AK, NK, EK, and MK count always killed, never killed, eventually killed and mixed killed.

All identifier templates and their corresponding counters are then stored in a database, shown as (IV) in Figure 2. These identifier templates are then used to populate a lookup table and efficiently rank and suppress mutants in the Ranking & Suppression phase.

## 4.2 Ranking & Suppression

Ranking & Suppression consumes the pre-computed identifier templates and their associated counters, as derived from the Template Generation phase, to rank and suppress mutants. For example, a mutant whose identifier template has a high PU count may be considered more useful than one whose identifier template has a low PU count, and hence should be ranked higher. Similarly, a mutant whose identifier template has a high PNU count may be considered not useful, and hence should be suppressed. Mutant ranking and mutant suppression are two independent procedures, as the former only gives a relative ordering between mutants and the highest ranked mutant may still not be worth surfacing to a developer.

**4.2.1 Ranking.** MuRS ranks identifier templates and matching mutants based on a usefulness score, which comprises two distinct scores: (1) the developer-feedback score, derived from the *Perceived feedback* label, and (2) the killed score, derived from the *Killed status* label. MuRS ranks the identifier templates based on the developer-feedback score, and breaks ties based on the killed score.

We explored two definitions for the developer-feedback score:

- **Usefulness score:** The ratio of PU over PU+PNU:  $\frac{PU}{PU+PNU}$ .
- **Bayes usefulness score:** It uses the Usefulness score and the average Usefulness score across all identifier templates:

$$w * Usefulness\ score + (1 - w) * Average\ Usefulness\ score$$

where the weight *w* is defined as  $\frac{PU+PNU}{PU+PNU+m}$  and *m* is the average of *PU* + *PNU* across all identifier templates. Compared to the Usefulness score, this score accounts for noise when *PU* + *PNU* is low—that is, when there is insufficient feedback from developers for a given identifier template.

We also explored two definitions for killed score:

- **Kill-ratio score:** The ratio of mutants killed over mutants generated for a given identifier template.
- **Kill-counter score:** It is a tuple with (AK, EK, MK, NK) counters. The ordering between the counters is based on their correlation with the developer feedback.

Since there are 2 developer-feedback scores and 2 killed scores, there are 4 different combinations of usefulness scores.

**Table 2: Hyper-parameter tuning for MuRS. Mutants on the left are surfaced without using MuRS. Mutants on the right are surfaced using MuRS with a specific hyper-parameter setting that suppresses mutant C. In this case, the original negative feedback ratio is  $\frac{2}{4} = 0.5$ , and it is  $\frac{1}{3} \approx 0.33$  for MuRS. The best hyper-parameter setting is the one that achieves the lowest negative feedback ratio.**

Original		MuRS	
Mutant	Label	Mutant	Label
A	Positive feedback	A	Positive feedback
B	No feedback	B	No feedback
C	Negative feedback	D	Negative feedback
D	Negative feedback		

4.2.2 *Suppression.* MuRS aims to suppress mutants that are likely to receive negative feedback. When designing the suppress function, we chose to implement the following two desired properties:

- (1) If a mutant’s identifier template is not in the database (i.e., the mutant has not been encountered before), that mutant will not be suppressed. This allows for gathering feedback for new identifier templates.
- (2) Mutant suppression should be conservative and driven by developer feedback. Since suppressing an identifier template means that any associated mutant will be suppressed, we conservatively opted to suppress mutants only based on explicit developer feedback—that is, the PU and PNU counters.

We explored the following three suppression functions:

- **No suppression:** Do not suppress any mutants. This is added as an experimental control.
- **Average threshold:** Suppress mutants whose Usefulness score is lower than the average Usefulness score across all identifier templates.
- **Probabilistic:** Use the p-value of the z-test to probabilistically suppress mutants, as the Usefulness score follows a normal distribution. For example, if the z-score of a mutant’s Usefulness score is  $-1$ , then the corresponding p-value is 0.1587: the mutant will be suppressed with probability  $p = 1 - 0.1587 = 0.8413$ . This suppression function is only applied when the Usefulness score is lower than the average Usefulness score across all identifier templates.

At the end of Ranking & Suppression, MuRS has ranked all generated mutants and suppressed mutants that are likely to receive negative feedback. At this point, the mutation testing service continues to surfacing non-suppressed mutants as findings in the code review tool.

### 4.3 Hyperparameter Tuning

MuRS’s design space involves five dimensions: (1) template type (original code template, typed template or indexed typed template), (2) vocabulary size (0, 1000, 5000, or 10000), (3) context size (0 or 1), (4) ranking function (4 combinations of developer-feedback score and killed score), and (5) suppression function (no suppression, average threshold, or probabilistic).

We conducted a preliminary evaluation of all possible hyperparameter settings using historical mutants. Specifically, we used the mutants generated before July 2022 to generate identifier templates and computed the hypothetical negative feedback ratio for each hyperparameter setting on the mutants generated in July 2022. We then selected the setting with the lowest ratio. The example in Table 2 illustrates how the hypothetical negative feedback ratio is calculated. On the left we have mutants with their labels. Note that we have the labels because they have already been generated and shown to developers. On the right is an example of how one setting might rank and suppress the mutants, where mutant C is suppressed. In this example, the original negative feedback ratio is  $\frac{2}{4} = 0.5$ , the hypothetical negative feedback ratio for MuRS is  $\frac{1}{3} \approx 0.33$ . The best hyperparameter setting is the one that achieves the lowest hypothetical negative feedback ratio.

The hyperparameter tuning identified the following setting as the best: (1) **Indexed type template**, (2) **vocabulary size = 0**, (3) **context size = 0** (4) **Bayes usefulness score with kill-counter score**, and (5) **Probabilistic suppression function**. We consistently used this hyperparameter setting to answer our research questions in section 5.

## 5 EVALUATION METHODOLOGY

To evaluate MuRS, we compared the current mutation testing service, which randomly selects  $n$  mutants and surfaces them as code findings in the code-review tool, to an enhanced variant, which surfaces the top- $n$  mutants after ranking and suppressing mutants with MuRS (see Figure 2). Note that both variants use existing optimizations, such as arid-node suppression and probabilistic mutation-operator selection [19]. Moreover, the mutation testing service limits the number of surfaced mutants during code review to avoid overwhelming developers: it surfaces at most three mutants per file and at most ten mutants per changelist. We kept these thresholds constant for both variants.

We followed a standard A/B testing methodology for changelists subject to mutation by the mutation testing service: each changelist was randomly assigned to the experiment group with probability  $p = 0.5$ ; all other changelists formed the control group. Regardless of group, each mutant that surfaced in the code review tool was subject to developer feedback (i.e., positive feedback, negative feedback, or no feedback at all). The A/B testing allowed us to answer the following two research questions:

**RQ1** Does MuRS decrease the overall negative feedback ratio?

**RQ2** Is MuRS’s ranking associated with developer feedback?

The existing mutation testing service has been deployed for over seven years, with domain experts manually developing arid-node heuristics to suppress mutants that will likely receive negative feedback [19]. The majority of these heuristics were implemented early during development and deployment, others were added over time in response to user feedback. Considering historical mutant data, we additionally answered the following research question:

**RQ3** Do templates identified by MuRS as not-useful correspond to suppression rules developed by domain experts?

At a high level: RQ1 uses an extrinsic evaluation and determines whether MuRS can improve mutant suppression over a manually

tuned baseline; RQ2 uses an intrinsic evaluation and determines the degree to which MuRS’s ranking of mutants is correlated with their perceived usefulness. RQ3 uses case studies and determines whether MuRS’s templates and their usefulness statistics can be linked to manually developed suppression rules.

### 5.1 RQ1: Does MuRS Decrease the Overall Negative Feedback Ratio?

During the A/B testing evaluation period<sup>1</sup>, 666 143 mutants surfaced across 84 827 changelists. For the control group, 365 542 mutants surfaced for 45 210 changelists. For the experiment group, 300 601 mutants surfaced for 39 618 changelists. For each group we computed the overall negative feedback ratio and compared the results. Additionally, we conducted a two-sample hypothesis test and computed the effect size, which quantifies the magnitude of the difference. Specifically, we used the chi-square test of independence [17], testing the null hypothesis that developer feedback is independent of group association (i.e., experiment vs. control group). We set the significance threshold to 0.05. We also computed the odds ratio [6] (effect size) and its 95% confidence interval.

### 5.2 RQ2: Is MuRS’s Ranking Associated with Developer Feedback?

The mutation testing service may generate more surviving mutants than can be surfaced in the code-review tool. Thus, it is important to rank mutants that are likely to receive positive feedback higher. To determine to what extent MuRS succeeds in properly ranking mutants, we computed the rank of all surfaced mutants in the control group and correlated that rank with the mutants’ received feedback. Specifically, we computed Kendall’s Tau-b (accounting for expected ties in the dichotomous feedback variable) between MuRS’s ranking score and developer feedback. We again set the significance threshold to 0.05. We repeated the same analysis for all mutants in the experiment group. Additionally, we computed the ratios of mutants that already received positive or negative feedback that would have been suppressed in the control group, investigating the trade-off between suppressing not-useful mutants while retaining useful ones.

### 5.3 RQ3: Do Templates Identified by MuRS as Not-useful Correspond to Suppression Rules Developed by Domain Experts?

We wished to better understand whether identifier templates with consistent not-useful feedback are linked to existing suppression rules in the mutation testing service. Since not-useful mutants can no longer be observed once a suppression rule is implemented, we resorted to historical mutant data to determine whether MuRS, if applied prior to suppression, would have identified and suppressed said mutants. Specifically, we compiled a data set with all surfaced mutants generated since 2017. For each mutant, we computed the corresponding identifier template and aggregated template data per month. We then performed a case study on a sample of 10

<sup>1</sup>Due to industry confidentiality reasons, we cannot disclose the exact time range, but it was a span of multiple weeks starting in July 2022.

**Table 3: Mutants with feedback during A/B testing.**

	Positive	Negative	Total
<b>Experiment group</b>	9688	1253	10 941
<b>Control group</b>	11 233	1592	12 825
<b>Total</b>	20 921	2845	23 766

templates that produced a sufficiently large number of mutants and consistently received negative feedback.

## 6 EVALUATION RESULTS

This section presents the results for our three research questions.

### 6.1 RQ1: Does MuRS Decrease the Overall Negative Feedback Ratio?

During the A/B testing evaluation period, a total of 23 766 mutants received developer feedback. Table 3 shows the distribution among the two groups and positive vs. negative feedback. The negative feedback ratio for mutants in the experiment group is 11.45% (1253 out of 10 941 mutants). In contrast, the negative feedback ratio for mutants in the control group is 12.41% (1592 out of 12 825 mutants). The Chi-squared test of independence for the two variables in Table 3 yields a p-value of 0.0242. Since this value is lower than our predefined significance threshold of 0.05, we can reject the null hypothesis. This means that the observed differences in feedback are not independent of the group—that is, the lower negative feedback ratio of 11.45% in the experiment group is statistically significant. The odds ratio for the contingency table in Table 3 is 0.91. This means that a mutant in the experiment group has a lower chance of receiving negative feedback than a mutant from the control group. The 95% confidence interval for the odds ratio is (0.84, 0.99) is strictly smaller than 1, which is consistent with the observed statistical significance.


### 6.2 RQ2: Is MuRS’s Ranking Associated with Developer Feedback?

As shown in Table 3, the control group had 12 825 mutants with feedback, and the experiment group had 10 941.


For the control group, the rank correlation coefficient (Kendall Tau-b) between MuRS’s usefulness score and the developer feedback is 0.2658 ( $p < 0.0001$ ). For the experiment group the correlation coefficient is 0.2615 ( $p < 0.0001$ ). These coefficients suggest a weak to moderate correlation and there is no significant difference between the control and experiment groups.

Figure 3 displays the cumulative number of mutants with negative feedback among the top x% ranked mutants in the control group. The top 50% of the ranked mutants only contain 16.39% (261 out of 1592) of all mutants that received negative feedback. The negative feedback ratio for the top x% mutants is shown in Figure 4. It reveals that the negative feedback ratio remains around 4% for the top half of the ranked mutants, but it steadily increases to 12% when including all mutants. Both figures suggest that higher ranked mutants are less likely to receive negative feedback.

Figure 4 shows a trade-off: while aggressive mutant suppression (e.g., the bottom 40% of ranked mutants) might lower the negative



**Figure 3: Accumulated number of not-useful mutants among the top x % ranked mutants in the control group.**




**Figure 4: Negative feedback ratio for top x % ranked mutants in the control group.**

feedback ratio to about 6%, it would likely suppress useful mutants as well. We determined how many mutants with positive or negative feedback would have been suppressed in the control group. Since suppression is probabilistic, we computed the expected values:

- 772.4159 out of 1592 mutants with negative feedback (49%) would have been suppressed. Since these mutants indeed received negative feedback, they are correctly suppressed.
- 1338.1659 out of 11 233 mutants with positive feedback (12%) would have been suppressed. Since these mutants received positive feedback, they are incorrectly suppressed.

### 6.3 RQ3: Do Templates Identified by MuRS as Not-useful Correspond to Suppression Rules Developed by Domain Experts?

Figure 5 shows the results of applying MuRS to all mutants surfaced since mid 2017 and filtering the resulting set of identifier templates



**Figure 5: Sampled not-useful templates for which at least 50 live mutants surfaced per month on average.**

such that each retained template (1) corresponds to at least 50 surfaced mutants per month on average and (2) consistently received negative feedback (i.e., its PU counter is 0 and its PNU counter is strictly greater than 0). This filtering step retains templates whose mutants MuRS would suppress with very high probability (sufficient feedback and consistently negative feedback). Each row in this plot corresponds to a retained template and each dot corresponds to a month during which a template was observable. A line connects multiple months of the same template, indicating the overall time period a template was observable.

From the identifier templates in Figure 5, we sampled 10 templates for manual inspection. We arbitrarily sampled two templates per language that were no longer observable after 2019. Since most suppression rules were added during initial development of the mutation testing service for mutants that received strong negative feedback, MuRS should recover this information in the form of templates that correspond to said not-useful mutants. Our manual analysis resulted in two key insights.

First, for most identifier templates, we were able to identify a corresponding suppression rule in the existing mutation testing service that suppresses the corresponding mutants. As one concrete example, the following error-handling pattern corresponds to not-useful mutants of a common error-handling pattern:



---

```
-if GO_IDENTIFIER_0 != GO_IDENTIFIER_1 { return
  GO_IDENTIFIER_0 }
```

---

This template corresponds to statement removal mutants that delete `if err != nil { return err }`. Since these statements are ubiquitous but the mutants are not generally worth testing, the negative developer feedback provided a very strong signal for suppression.

Second, the identifier templates are finer-grained than the existing suppression rules, meaning that the identifier templates only correspond to a subset of the mutants that are suppressed by the existing rules. As an example, the following Java template refers to specific mutations in annotations:

---

```
-@JAVA_IDENTIFIER_0(JAVA_IDENTIFIER_1 = JAVA_IDENTIFIER_2,
  JAVA_IDENTIFIER_3, JAVA_IDENTIFIER_4 = true)
+@JAVA_IDENTIFIER_0(JAVA_IDENTIFIER_1 = JAVA_IDENTIFIER_2,
  JAVA_IDENTIFIER_3, JAVA_IDENTIFIER_4 = !(true))
```

---

This template corresponds to mutants that change a boolean attribute of an annotation from `true` to `!(true)`. Annotations in Java provide metadata or compiler directives, or are used for meta-programming. Common examples include the `@Override` annotation, which indicates that a subclass method is overriding a parent class method, the `@Deprecated` annotation, which indicates that a method is deprecated, or the `@Test` annotation for JUnit tests. However, annotations are rarely, if ever, subject to (unit) testing. The clear negative sentiment of developers towards these mutants prompted the addition of a suppression rule, which in contrast to the identifier template, suppresses mutations in annotations more generally.

In addition to the 10 sampled identifier templates, we also inspected templates with positive and mixed feedback. Below are two concrete examples.

Among the C++ templates, the mutants corresponding to the following template are considered useful by many developers:

---

```
-} else if (CPP_IDENTIFIER_0 == CPP_IDENTIFIER_1) {
+} else if (true) {
```

---

This template captures mutants that replace a relational operator in the condition of an `else if` statement with the constant value `true`.

The following Python template has both the highest total PU and PNU count:

---

```
-PYTHON_ID_0(PYTHON_ID_1)
```

---

This template captures mutants that remove a function call with a single argument. Just et al. found that statement deletion mutants are often associated with real faults, more so than other types of mutants [10]. Therefore, it is not surprising that many of these mutants are considered useful. At the same time, there are likely many function calls that match this pattern that are not worth testing. One common class of function calls that matches this pattern are telemetry statements, e.g. `log(id)`. Developers typically do not consider these to be directly impacting user experience and are less interested in testing them specifically. Templates like this are the most controversial ( $\text{controversy} = PU * (1 - PU) * (PU + PNU)$ ).

By computing the controversy score for all identifier templates, we found that the most controversial templates for Python, Java, C++, Go, and TypeScript are all statement removal mutants. One of the reasons is that statement removal mutants represent 68% of all

mutants [19], which makes them also more likely to receive more feedback. This observation calls for more research to understand the differences between useful statement removal mutants that are associated with real bugs and not-useful statement removal mutants. For example, future research can inform additional context that should be considered to distinguish these.

## 6.4 Discussion

It is important to put the results in perspective when assessing their practical significance. First, we are comparing MuRS against a strong baseline, which employs many suppression heuristics that have lowered the negative feedback ratio from over 80%, during initial development, to well below 15% now. Second, reducing the negative feedback ratio even by 1% substantially reduces negative interactions across all developers. While each individual negative interaction has a small cost, the total cost reduction across tens of thousands of developers that participate in the code review process multiple times per day is significant. Third, the negative feedback ratio is a metric derived from developer feedback, which is subjective. A mutant can receive negative feedback for many reasons—some mutants represent infeasible test goals, others are simply not worth satisfying [16, 20]. While many not-useful mutants consistently receive negative feedback, others receive no or even inconclusive feedback. For example, the very same mutant can receive different feedback from the author and reviewers during a code review. As a result, the lower bound for the negative feedback ratio is very likely greater than zero. Our baseline is already approaching a ratio in the single digits, and the true lower bound, and thus the best possible result, could be anywhere between 0% and 10%.

MuRS suppresses about 50% mutants with negative feedback, at the cost of suppressing about 12% mutants with positive feedback. One of the reasons why MuRS does not suppress more not-useful mutants is that the corresponding identifier templates have not yet received enough negative feedback, and as a result the suppression function conservatively retains them. Indeed, we have observed that about 80% of all templates correspond to at most ten surfaced mutants—that is, the templates are too specific. We conjecture that this specificity is also the reason why the best setting for the vocabulary and context sizes are both 0 (see subsection 4.3). Increasing either size would make the identifier templates even more specific, grouping even fewer syntactically similar mutants. At the same time, the current templates do not consider enough context to retain all useful mutants and to disambiguate mutants with mixed feedback—that is, the templates are too generic. We leave a deeper investigation and possible refinements for future work.

An alternative mutant suppression approach can use static rules to prevent mutants from being generated in the first place (e.g., mutations of logging statements). The existing mutation testing service already implements this using arid nodes, which are AST nodes that are identified through pattern matching and that are never mutated. While effective, MuRS offers several advantages over the arid node approach. It automatically creates an identifier template that is used to suppress a particular type of mutant, eliminating the need for manual implementation of static rules. In fact, we have discovered multiple identifier templates that correspond to our existing static rules, demonstrating that MuRS is able to recover

these rules. Overall, bootstrapping mutant suppression with arid-node heuristics and using MuRS to automatically learn additional, finer-grained suppression rules is effective.

## 7 RELATED WORK

The closest related work on industry-scale mutation testing is Beller et al.'s study of mutation testing at Facebook [2]. They also focused on diff-based mutation testing, but a key difference is that their approach uses semi-automatic learning on common Java bug patterns to create few mutation operators that yield mostly useful mutants by design. In contrast, our approach uses a predefined set of mutation operators and suppresses not-useful mutants. The two approaches could potentially be combined for greater effect.

### 7.1 Mutant Selection and Prioritization

A number of studies have shown that a small percentage of mutants are sufficient, that there is substantial redundancy among generated mutants, and that many generated mutants are not useful [5, 14, 18]. We refer the reader to the literature review by Pizzoleto et al. for more comprehensive related work [21], especially PG-1 (*Reducing the number of mutants*) and PG-5 (*Avoiding the creation of certain mutants*) mentioned in the paper for the closest related work regarding mutation selection or prioritization.

Gopinath et al. compared multiple strategies for reducing the number of mutants against random sampling [7]. They found that none of the strategies yield an effective advantage larger than 5%, when compared with random sampling. Given the result, they warn against adopting mutant reduction techniques without adequate reason. In an industrial context, however, a 1% reduction in mutants with negative feedback could be worthwhile (it saves valuable developer time) as long as mutants with positive feedback are retained.

Brown et al. introduced an approach that extracts mutation operators from the revision history of software projects [3]. The approach is based on the observation that if a commit corrects a bug, then its reversal is potentially introducing a bug, making it a viable mutation operator. Using this approach, they mined mutation operators from the revision history of the top-50 most forked C projects on GitHub. Their findings showed that mutants generated from the mined mutation operators were just as challenging to kill as mutants generated using traditional mutation operators. However, the mined mutation operators exhibited greater diversity in the types of changes than traditional mutation operators.

### 7.2 Mutant Usefulness

Just et al. studied the degree to which program context, defined over the abstract syntax tree, is correlated with a mutant's expected usefulness [11]. In this work, usefulness was termed mutant utility and measured along three dimensions: equivalence, triviality, and dominance. The results showed that mutant utility is context-dependent: the same mutation may lead to a high-utility mutant in one context but not necessarily in another.

Kaufman et al. proposed test completeness advancement probability (TCAP) as a measure to define how useful a mutant is [12]. TCAP is a probability that if a mutant is presented as a test goal, it will elicit a test to improve the test completeness. Evaluated on

9 projects from the Defects4J benchmark [9], they concluded that TCAP can be predicted from program context and that TCAP-based mutant prioritization improves test completeness faster than the previous state-of-the-art, which is random prioritization.

MuRS and its evaluation differ from prior work in two key aspects. First, MuRS uses an abstraction over the diff, which encodes both the mutation and the surrounding context. Second, mutant usefulness is derived from developer feedback as opposed to from test or mutant characteristics.

### 7.3 Diff-Based Mutation Testing

Cachia et al. focused on industrial adoption of mutation testing, and suggested incremental mutation testing that limits the scope of mutation testing to the changed code [4]. In their evaluation, they found that incremental mutation testing reduced the number of generated mutants and the execution time of mutation testing.

Ma et al. studied the relationship between commit-aware and traditional mutation testing [15]. They found that the commit-aware mutation score and the traditional mutation score are only weakly correlated, and that mutants from traditional mutation testing have a 30% lower chance of revealing faults introduced by the commit.

## 8 CONCLUSIONS

This paper presents MuRS, an approach designed to enhance the user experience of mutation testing by reducing the negative feedback ratio of mutants surfaced during code review. MuRS ranks and suppresses newly generated mutants based on identifier templates that group similar mutants and aggregate historical developer feedback for said mutants.

Based on an A/B testing study, we found that the negative feedback ratio for MuRS (11.45%) is statically significantly lower, compared to a manually tuned baseline. The results also show that mutants higher ranked by MuRS have a lower probability of receiving negative feedback. Additionally, we observe that MuRS would have suppressed about 50% of not-useful mutants but also 12% of useful mutants in the control group, highlighting a trade-off. Finally, we identified statement deletion mutants as the type of mutants whose identifier templates frequently receive both positive and negative feedback. Going forward, we aim to investigate what distinguishes useful statement deletion mutants from not-useful ones despite their syntactical similarities w.r.t. identifier templates. Furthermore, we would like to understand whether MuRS should apply different abstractions for the mutated lines vs. context lines.

One important aspect of MuRS in an industrial setting is its lower cost of development and maintenance, compared to manually implementing suppression rules. Our results provide evidence that MuRS produces identifier templates that correspond to static suppression rules implemented in the past, but future work should further evaluate to what extent MuRS can fully replace suppression rules or to what extent the two approaches are complementary.

## REFERENCES

- [1] Allen T Acree, Timothy A Budd, Richard A DeMillo, Richard J Lipton, and Frederick G Sayward. 1979. *Mutation Analysis*. Technical Report. Georgia Inst of Tech Atlanta School of Information And Computer Science.
- [2] Moritz Beller, Chu-Pan Wong, Johannes Bader, Andrew Scott, Mateusz Machalica, Satish Chandra, and Erik Meijer. 2021. What it would take to use mutation testing in industry—a study at facebook. In *International Conference on Software*

- Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 268–277. <https://doi.org/10.1109/ICSE-SEIP52600.2021.00036>
- [3] David Bingham Brown, Michael Vaughn, Ben Liblit, and Thomas Reps. 2017. The Care and Feeding of Wild-Caught Mutants. In *Proceedings of the Joint Meeting of the European Software Engineering Conference and the Symposium on the Foundations of Software Engineering (ESEC/FSE) (ESEC/FSE 2017)*. Association for Computing Machinery, New York, NY, USA, 511–522. <https://doi.org/10.1145/3106237.3106280>
- [4] Mark Anthony Cachia, Mark Micallef, and Christian Colombo. 2013. Towards incremental mutation testing. *Electronic Notes in Theoretical Computer Science (ENTCS)* 294 (2013), 2–11. <https://doi.org/10.1016/j.entcs.2013.02.012>
- [5] Yiqun T. Chen, Rahul Gopinath, Anita Tadakamalla, Michael D. Ernst, Reid Holmes, Gordon Fraser, Paul Ammann, and René Just. 2020. Revisiting the Relationship Between Fault Detection, Test Adequacy Criteria, and Test Set Size. In *Proceedings of the International Conference on Automated Software Engineering (ASE)*. 237–249. <https://doi.org/10.1145/3324884.3416667>
- [6] Ronald A Fisher. 1935. The logic of inductive inference. *Journal of the royal statistical society* 98, 1 (1935), 39–82. <https://doi.org/10.2307/2342435>
- [7] Rahul Gopinath, Iftekhar Ahmed, Mohammad Amin Alipour, Carlos Jensen, and Alex Groce. 2017. Mutation reduction strategies considered harmful. *IEEE Transactions on Reliability* 66, 3 (2017), 854–874. <https://doi.org/10.1109/TR.2017.2705662>
- [8] Yue Jia and Mark Harman. 2010. An analysis and survey of the development of mutation testing. *IEEE Transactions on Software Engineering (TSE)* 37, 5 (2010), 649–678. <https://doi.org/10.1109/TSE.2010.62>
- [9] 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 International Symposium on Software Testing and Analysis (ISSTA)*. 437–440. <https://doi.org/10.1145/2610384.2628055>
- [10] René Just, Darioush Jalali, Laura Inozemtseva, Michael D Ernst, Reid Holmes, and Gordon Fraser. 2014. Are mutants a valid substitute for real faults in software testing?. In *Proceedings of the Symposium on the Foundations of Software Engineering (FSE)*. 654–665. <https://doi.org/10.1145/2635868.2635929>
- [11] René Just, Bob Kurtz, and Paul Ammann. 2017. Inferring Mutant Utility from Program Context. In *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA)*. 284–294. <https://doi.org/10.1145/3092703.3092732>
- [12] Samuel J Kaufman, Ryan Featherman, Justin Alvin, Bob Kurtz, Paul Ammann, and René Just. 2022. Prioritizing mutants to guide mutation testing. In *Proceedings of the International Conference on Software Engineering (ICSE)*. 1743–1754. <https://doi.org/10.1145/3510003.3510187>
- [13] Pavneet Singh Kochhar, Ferdian Thung, and David Lo. 2015. Code coverage and test suite effectiveness: Empirical study with real bugs in large systems. In *IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 560–564. <https://doi.org/10.1109/SANER.2015.7081877>
- [14] Bob Kurtz, Paul Ammann, Jeff Offutt, Márcio E Delamaro, Mariet Kurtz, and Nida Gökçe. 2016. Analyzing the validity of selective mutation with dominator mutants. In *Proceedings of the Symposium on the Foundations of Software Engineering (FSE)*. 571–582. <https://doi.org/10.1145/2950290.2950322>
- [15] Wei Ma, Thomas Laurent, Miloš Ojđanić, Thierry Titchou Chekam, Anthony Ventresque, and Mike Papadakis. 2020. Commit-aware mutation testing. In *International Conference on Software Maintenance and Evolution*. IEEE, 394–405. <https://doi.org/10.1109/ICSM46990.2020.00045>
- [16] A Jefferson Offutt and Jie Pan. 1997. Automatically detecting equivalent mutants and infeasible paths. *Software Testing, Verification and Reliability (JSTVR)* 7, 3 (1997), 165–192. [https://doi.org/10.1002/\(SICI\)1099-1689\(199709\)7:3<165::AID-STVR143>3.0.CO;2-U](https://doi.org/10.1002/(SICI)1099-1689(199709)7:3<165::AID-STVR143>3.0.CO;2-U)
- [17] Karl Pearson. 1900. X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 50, 302 (1900), 157–175. <https://doi.org/10.1080/14786440009463897>
- [18] Goran Petrović, Marko Ivanković, Gordon Fraser, and René Just. 2021. Does mutation testing improve testing practices?. In *Proceedings of the International Conference on Software Engineering (ICSE)*. 910–921. <https://doi.org/10.1109/ICSE43902.2021.00087>
- [19] Goran Petrovic, Marko Ivankovic, Gordon Fraser, and René Just. 2021. Practical mutation testing at scale: A view from Google. *IEEE Transactions on Software Engineering (TSE)* (2021). <https://doi.org/10.1109/TSE.2021.3107634>
- [20] Goran Petrovic, Marko Ivankovic, Bob Kurtz, Paul Ammann, and René Just. 2018. An industrial application of mutation testing: Lessons, challenges, and research directions. In *Proceedings of the International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. 47–53. <https://doi.org/10.1109/ICSTW.2018.00027>
- [21] Alessandro Viola Pizzoleto, Fabiano Cutigi Ferrari, Jeff Offutt, Leo Fernandes, and Márcio Ribeiro. 2019. A systematic literature review of techniques and metrics to reduce the cost of mutation testing. *Journal of Systems and Software (JSS)* 157 (2019), 110388. <https://doi.org/10.1016/j.jss.2019.07.100>
- [22] Caitlin Sadowski, Edward Aftandilian, Alex Eagle, Liam Miller-Cushon, and Ciera Jaspán. 2018. Lessons from building static analysis tools at Google. *Communications of the ACM (CACM)* 61, 4 (2018), 58–66. <https://doi.org/10.1145/3188720>
- [23] W Eric Wong and Aditya P Mathur. 1995. Reducing the cost of mutation testing: An empirical study. *Journal of Systems and Software (JSS)* 31, 3 (1995), 185–196. [https://doi.org/10.1016/0164-1212\(94\)00098-0](https://doi.org/10.1016/0164-1212(94)00098-0)

Received 2023-05-18; accepted 2023-07-31