

# Automated Evolution of Feature Logging Statement Levels Using Git Histories and Degree of Interest

Yiming Tang<sup>a,\*</sup>, Allan Spektor<sup>b</sup>, Raffi Khatchadourian<sup>b</sup>, Mehdi Bagherzadeh<sup>c</sup>

<sup>a</sup>*Department of Computer Science, City University of New York (CUNY) Graduate Center,  
365 5th Ave, New York, NY 10016 USA*

<sup>b</sup>*Department of Computer Science, City University of New York (CUNY) Hunter College,  
695 Park Avenue, Room HN 1008, New York, NY 10065 USA*

<sup>c</sup>*Department of Computer Science & Engineering, Oakland University, Rochester, MI  
48309 USA*

---

## Abstract

Logging—used for system events and security breaches to describe more informational yet essential aspects of software features—is pervasive. Given the high transactionality of today’s software, logging effectiveness can be reduced by information overload. Log levels help alleviate this problem by correlating a priority to logs that can be later filtered. As software evolves, however, levels of logs documenting surrounding feature implementations may also require modification as features once deemed important may have decreased in urgency and vice-versa. We present an automated approach that assists developers in evolving levels of such (feature) logs. The approach, based on mining Git histories and manipulating a degree of interest (DOI) model<sup>1</sup>, transforms source code to revitalize feature log levels based on the “interestingness” of the surrounding code. Built upon JGit and Mylyn, the approach is implemented as an Eclipse IDE plug-in and evaluated on 18 Java projects with ~3 million lines of

---

\*Corresponding author.

*Email addresses:* [ytang3@gradcenter.cuny.edu](mailto:ytang3@gradcenter.cuny.edu) (Yiming Tang),  
[allan.spektor03@myhunter.cuny.edu](mailto:allan.spektor03@myhunter.cuny.edu) (Allan Spektor),  
[raffi.khatchadourian@hunter.cuny.edu](mailto:raffi.khatchadourian@hunter.cuny.edu) (Raffi Khatchadourian),  
[mbagherzadeh@oakland.edu](mailto:mbagherzadeh@oakland.edu) (Mehdi Bagherzadeh)

*URL:* <http://linkedin.com/in/gracetang1993> (Yiming Tang),  
<https://www.linkedin.com/in/allan-spektor> (Allan Spektor),  
<http://cs.hunter.cuny.edu/~Raffi.Khatchadourian99> (Raffi Khatchadourian),  
<http://mbagherz.bitbucket.io> (Mehdi Bagherzadeh)

<sup>1</sup>Degree of interest model (DOI) was proposed by Kersten and Murphy [1] to gauge the degree of developers’ interests in program elements.

code and  $\sim 4K$  log statements. Our tool successfully analyzes 99.22% of logging statements, increases log level distributions by  $\sim 20\%$ , and increases the focus of logs in bug fix contexts  $\sim 83\%$  of the time. Moreover, pull (patch) requests were integrated into large and popular open-source projects. The results indicate that the approach is promising in assisting developers in evolving feature log levels.

*Keywords:* logging, software evolution, software repository mining, software transformation, source code analysis, degree of interest

---

## 1. Introduction

Modern software typically includes logging, which documents useful information about a system’s behavior at run-time and facilitates system understanding. Logs help diagnose run-time issues and can be used to monitor processes [2], transfer knowledge [3], and detect errors [4–6]. Other logs, feature logs, may be more informational yet essential as they describe aspects of features the surrounding code implements.

However, the high transactionality of today’s software can cause logging to be less effective due to information overload. The sheer number of logs emitted can make it challenging to debug during development; logs pertaining to auxiliary features may be tangled with those features under current development. Also, parsing necessary information from logs to understand system behavior, how features interact, and diagnosing problems can be challenging.

To help alleviate these problems, logging frameworks and libraries empower developers to write logging statements consisting of several parts dictating how the log should be emitted, if at all. A logging statement is comprised of a particular log *object*, each of which is associated with a *run-time* level and other attributes. A logging *method* is invoked on the log object; one parameter is a log priority *level*. Log levels are ordered, and—during execution—the log *message* is emitted iff the log statement level is greater than or equal to the log object run-time level. Messages are typically dynamically constructed with static text

and dynamic contexts, such as the contents of one or more variables [7]. For example, the following statement outputs system-health information iff the run-time level of `logger` is  $\leq$  `FINER` [8]: `logger.log(Level.FINER, "Health is: " + DiagnosisMessages.systemHealthStatus())`. Controlling the log run-time level affords developers the ability to limit the types of log information emitted either for particular environments (e.g., development, deployment) or other factors.

As software evolves, however, levels of logging statements correlated with surrounding feature implementations may also need to be modified. Such *feature* logging statements could, for example, serve as algorithm checkpoints, where critical variables are outputted for validation and progress is ensured. Ideally, levels of *feature* logs would *evolve with* systems as they are developed, with higher log levels (e.g., `INFO`) being assigned to logs corresponding to features with more current stakeholder interest than those with less (e.g., `FINEST`). However, as developers tend not to manually change log levels [9], feature log levels may become stale, causing irrelevant logs to accumulate, increased information overload, and tangling of relevant feature logs with those not currently being developed, thereby complicating debugging. Furthermore, *manually* maintaining log levels can be tedious and error- and omission-prone as logging statements are highly scattered [5]. Moreover, developers may not use the full spectrum of available levels.

Existing approaches [3, 7, 9–11] focus on either new logging statements or messages. Logger hierarchies [8, 12] may be useful but still require manual maintenance. To the best of our knowledge, there is currently no automated solution for logging statement level evolution. Therefore, we present an automated approach that assists developers in evolving feature logging statement levels. The approach first mines Git repositories to discover developers’ “interestingness” of code surrounding feature logging statements by adapting the degree of interest (DOI) model of Mylyn [1]. Mylyn [13] is a standard Eclipse Integrated Development Environment (IDE) [14] plug-in that facilitates software evolution by focusing graphical IDE components so that only artifacts related to the currently active task are revealed [15]. Mylyn manipulates DOI so that artifacts

(e.g., files) with more interaction are more prominently displayed in the IDE than those less recently used. Each program element is associated with a float value named DOI value that is used to gauge the developers' degree of interest in it. The approach later correlates "interestingness" of code surrounding feature logging statements with feature log levels. If our approach detects a mismatch between feature log levels and feature interests, it could suggest appropriate log levels to the feature logging statements.

We programmatically manipulate DOI using modifications made in source code repositories. Our approach transforms code to reinvigorate feature logging statement levels, pulling those related to features whose implementations are worked on more and more recently to the forefront, while pushing those worked on less and less recently to the background. Our goal is information overload reduction and improved debugging by automatically bringing more relevant features to developers' attention and vice-versa throughout system evolution.

Logging levels are often used to differentiate various logging *categories*, i.e., levels having special semantics that are not on a "sliding scale." Altering such levels may violate the preservation of the log's intended semantics. In this work, we focus on the levels *feature* logs, i.e., those highly related to feature implementations, as feature interests vary over time and whose related logging statements may benefit from aligning levels correspondingly. Thus, to distinguish feature logs from those that are more categorical, e.g., those conveying more critical information (e.g., errors, security), a series of novel heuristics, mainly derived from first-hand developer interactions, are introduced. On the other hand, the heuristics also account for less-critical debugging logs, e.g., tracing, using a keyword-based technique. This effort focuses our approach on only manipulating logging statements tied to features to better coordinate them with developers' current interests.

Our approach is implemented as an open-source plug-in to the Eclipse IDE, though it may be used with other IDEs via popular build systems. It supports two popular logging frameworks and integrates with JGit [16] and Mylyn. The evaluation involved 18 Java projects of varying sizes and domains with a total of

~3 million lines of code and ~4K logging statements. Our study indicates that (i) given its ability to process a significant number and size of Git changesets, the fully-automated analysis cost is viable, with an average running time of 10.66 secs per logging statement and 0.89 secs per thousand lines of code changed, (ii) developers do not actively think about how their logging statement levels evolve with their software, motivating an automated approach, and (iii) our approach is promising in evolving feature log levels.

This work’s contributions are summarized as follows:

**Approach design.** We present an automated approach that programmatically manipulates a Mylyn DOI model using Git histories to evolve feature logging statement levels to better align with the current features of interest. Widespread manual log level modification is alleviated, information overload is reduced, and more relevant events are underscored, potentially exposing bugs.

**Heuristic formulation.** Heuristics—based on first-hand developer feedback—to distinguish between feature logs and those with more critical information are proposed.

**Implementation & experimental evaluation.** To ensure real-world applicability, we implemented our approach as an open-source Eclipse IDE plug-in built upon Mylyn and JGit and used it to study 18 Java projects. Our technique successfully analyzes 99.22% of logging statements, increases log level distributions by ~20%, and increases the focus of logs in bug fix contexts at a rate of ~83%. Furthermore, several pull requests were integrated into large and popular open-source projects.

## 2. Motivating Example

Lst. 1 portrays a hypothetical code snippet [17] that uses `java.util.logging` (JUL) [18] having log levels that include—in ascending order—FINEST, FINER, FINE, INFO, WARNING, and SEVERE. A `Wombat` class starts at line 1 and has a `logger`

---

**Listing 1** Hypothetical logging usage example [17].

---

```
1 public class Wombat {
2     private static final Logger logger = Logger.getLogger("global");
3     logger.setLevel(Level.FINE); // Only logs  $\geq$  FINE.
4     private double temp; private double oldTemp;
5
6     public void setTemp(double val) {
7         this.oldTemp = temp; this.temp = val;
8         logger.log(Level.FINER, "Temp set to: " + this.temp);
9         logger.finer("Old temperature was: " + this.oldTemp);}
10
11     public static void main(String[] args) {
12         Wombat w = new Wombat();
13         Scanner scanner = new Scanner(System.in);
14
15         System.out.println("Enter a temperature: ");
16         double input = scanner.nextDouble(); w.setTemp(input);
17
18         try { // send to file.
19             logger.fine("Writing to file.");
20             Files.writeString("output.txt", w.toString(), WRITE);
21         } catch (IOException e) { // Fatal error.
22             logger.severe("Couldn't open file for writing.");
23             throw e;}}
```

---

(line 2) and current and previous `temp` values (line 4). The `logger` is configured so that only logs with levels  $\geq$  `FINE` are emitted to reduce information overload.

A mutator for `temp` begins on line 6. On line 7, old `temp` values are cached. Then, new and old temperatures are logged on lines 8 and 9, respectively. Both statements log at the `FINER` level. Since `logger` has been previously configured not to emit logs with levels  $\leq$  `FINER`, the statements have no effect.

When creating `Wombats` (line 12), the user is asked for a temperature (line 16). A string representation of the `Wombat` not shown in this paper is then saved to a file (lines 18–23). Line 19 logs that the writing has commenced, and since the level is `FINE`, the statement emits a log. The actual file writing takes place on line 20. Because `Files.writeString()` possibly throws an `IOException`, the call is surrounded by a try/catch block. Line 22 executes when the specified

---

**Listing 2** Rejecting invalid temperatures.

---

```
1 @@ -23,11 +23,15 @@ public void setTemp(double val) {
2 + if (val > 0) {
3     this.oldTemp = temp; this.temp = val;
4     logger.log(Level.FINER, "Temp set to: " + this.temp);
5     logger.finer("Old temperature was: " + this.oldTemp);}
6 + else throw new IllegalArgumentException("Invalid: " + val);
7
8 @@ -38,7 +42,17 @@ public static void main(String[] args)
9 + while (true) {
10 + try {
11     w.setTemp(input);
12 + break; // succeeded.
13 + } catch (IllegalArgumentException e) {
14 + // Not a fatal error. Log the exception and retry.
15 + logger.log(Level.INFO, "Invalid input: " + input, e);
16 + System.out.println("Invalid temp. Please retry.");}}
```

---

---

**Listing 3** Warning about drastic temperature changes.

---

```
1 @@ -30,6 +30,9 @@ public void setTemp(double val) {
2     else throw new IllegalArgumentException("Invalid:" + val);
3 +
4 + if ((this.temp - this.oldTemp) / this.oldTemp > 0.05)
5 +     logger.warning("Temperature has risen above 5%.");
```

---

exception has been caught. This log message is emitted since `SEVERE`  $\geq$  `FINE`.

Lst. 2 depicts a `changeset`<sup>2</sup> where invalid temperatures (i.e., negative temperatures) are rejected by guarding lines 3–5 and throwing an exception on line 6. As a result, client code (lines 9–16) is modified to handle the exception, looping until valid input is entered. On line 15, a log is issued when the exception is caught, documenting the retry. Because the error is non-fatal, `INFO` is used. An ensuing `changeset` (lst. 3) logs a warning (line 5) when temperatures increase by more than 5%.

Lst. 4 shows an abbreviated result, containing only relevant parts of lst. 1 with lst. 2 and 3 applied and transformations made to feature logging statement

---

<sup>2</sup>Although only additions are shown, similar issues may arise with deletions.

---

**Listing 4** Resulting “reinvigorated” logging levels.

---

```
5 public void setTemp(double val) {
6     if (val > 0) {
7         this.oldTemp = temp; this.temp = val;
8         logger.log(Level.FINER, "Temp set to: " + this.temp);
9         logger.finer("Old temperature was: " + this.oldTemp);}
10    else throw new IllegalArgumentException("Invalid: " + val);
11
12    if ((this.temp - this.oldTemp) / this.oldTemp > 0.05)
13        logger.warning("Temperature has risen above 5%.");}
14
15    public static void main(String[] args) { // ...
16        } catch (IllegalArgumentException e) {
17            // Not a fatal error. Log the exception and retry.
18            logger.log(Level.INFO, "Invalid input: "+ input, e);
19            System.out.println("Invalid temp. Please retry.");//}
20
21    try { // send to file.
22        logger.finest("Writing to file.");
23        Files.writeString("output.txt", w.toString(), WRITE);
24    } catch (IOException e) { // Fatal error.
25        logger.severe("Couldn't open file for writing."); // ...
```

---

levels. Recent changes to nearby code of feature logging statements may indicate that features, e.g., temperature management, implemented at this part of the code are of a higher “interest.” As such, the levels at lines 8–9 have *increased* from FINER to FINE, potentially helping developers debug the new feature code. In this example, the transformed logs will now emit.

Contrarily, as file writing is *not* being actively developed—the feature logging statement level at line 22 *decreased* from FINE to FINEST. Recent changesets did *not* include edits to this region, thus resulting in log suppression. While code recently edited may have regressions, the (non-feature) logging statement at line 25 did *not* have its level lowered and remains useful in finding possible regressions. Likewise, the non-feature logging statement at line 18, although non-fatal, was *not* lowered *despite* recent non-local edits.

As logging is pervasive [5], manually managing feature logging statement levels can be overwhelming. Even this simple example demonstrates that logging






Figure 1: Logging level revitalization approach overview.

statements can be scattered and tangled with code implementing core functionality. Automatically evincing information related to features that are developed more and more often to the forefront, while gradually suppressing those less and less frequently developed may enhance focus and help expose potential bugs.

### 3. Approach

#### 3.1. Overview

Our automated approach (fig. 1) programmatically manipulates a Mylyn DOI model [1] by mining Git repositories to evolve feature logging statement levels to better align with the current features of interest. Mylyn [15] has traditionally been used to associate IDE edit and graphical interaction events to a particular *task* to support *task context switching*. The two branches between yellow diamonds could be executed concurrently. We adapt this model to track the “interest” of code elements surrounding feature logging statements (step 4) to find *mismatches* (step 7) between log levels and feature interests for feature logging statements. Furthermore, to aggregate interest levels across a development team, edits are extracted from Git histories (step 3) instead of IDE interactions. The DOI model is then programmatically manipulated using the extracted edits (step 5) and finally partitioned (step 6) and compared to the current feature logging statement (step 1) levels (step 2). If a mismatch is

found, the level is transformed in the code (step 9). To distinguish feature logging statements from other logging that could be more critical and to guide the transformation, a set of heuristics that we define are used (step 8).

### 3.2. Feature Logging Statement Level Extraction

Logging statements—later used for (i) correlating the degree of interest of surrounding code when finding mismatches (step 7), (ii) applying heuristics (step 8), and (iii) as potential transformation sites (step 9)—are extracted in step 1. Current levels are extracted from the statements (step 2) and later used in mismatch identification (step 7). Depending on the API used, level extraction is performed in two ways. For instance, JUL has both convenience and standard logging APIs. Convenience APIs have method names that correspond to the desired log level (e.g., line 9, lst. 1). On the other hand, standard APIs have logging levels passed as parameters (e.g., line 8, lst. 1). In both cases, Abstract Syntax Trees (ASTs) are extracted from the underlying source code. Whereas the convenience case is straight-forward, in the standard case, our current implementation only extracts levels represented as literals. Using data-flow analysis is a subject of future work; however, we analyzed 99.22% of logging statements during our study successfully despite this limitation.

### 3.3. Mylyn DOI Model Manipulation

#### 3.3.1. Background

Mylyn [13] maintains focused contexts of entities relevant to a particular task using a DOI model. A context comprises the relevant elements (e.g., classes, methods, fields), along with how *interesting* the elements are to the related task. The more a developer *interacts* with an element (e.g., navigates to a file, edits a file) when working on a task, the more interesting the element is deemed to be, and vice-versa. Elements also *decay*, i.e., as other elements increase in DOI, elements that were once “interesting” decrease. Mylyn then alters the IDE’s behavior so that interesting elements are displayed more prominently throughout its views.

### 3.3.2. Repository Mining

The Mylyn context is adapted to ascertain the interest levels of code surrounding logging statements. Traditionally, Mylyn is used for *context switching*; i.e., relevant elements are stored in task contexts. That way, developers can easily switch between tasks without losing their focus on related code elements. It is confined to a single developer’s workspace; however, code modifications made by our approach are *global*—affecting all project developers. As such, the context is “expanded” to include *all* developers’ interests by mining Git repositories in step 3.

Mylyn can record a wide range of interaction events between developers and program elements, such as element selection, element editing, etc. However, our approach only takes into account the interaction event of element editing because (1) edits can be easily saved, for example, through Git histories (2) edits can more directly express developers’ interest than other interaction events, such as element selection.

### 3.3.3. Converting Code Changes to Interaction Events

A central notion of Mylyn are “interaction events,” which dictate how the DOI model is manipulated. The more interaction a particular element has, the larger its DOI and vice-versa. Although Mylyn has a broad spectrum of interaction event types, we focus on “edit” events as we mine Git code changesets.

In converting Git code changes to Mylyn interaction events (step 4), we mainly focus on changes to method and constructor bodies. While edits to other kinds of code elements, e.g., fields, could be near logging statements, this is currently not supported and is not a representative use case. Furthermore, Git edits are parsed instead of AST differencing, which can be computationally expensive. Moreover, AST differencing does not include whitespace changes, which are desirable as they may indicate interest.<sup>3</sup> Interaction events are then processed by Mylyn as if they had emanated from the IDE in step 5.

---

<sup>3</sup>We also consider changes to non-source lines, e.g., comments.

*Rename Refactorings & Copying.* Program elements (e.g., methods) altered in Git repositories may no longer exist in the current version where level transformations would take place. Such historical elements that were removed are ignored as they do not exist in the current version. However, elements that have undergone rename refactorings need to be considered as they will have a counterpart in the current version. To this end, during repository mining, we maintain a data structure that associates rename relationships between program elements, e.g., method signatures. Before converting changesets to interaction events, a lookup is first performed to retrieve the changed element’s signature in the current version. Unfortunately, handling refactorings necessitates two history traversals, one to create the renaming data structure and the other to process changesets. However, in our implementation, we have a performance improvement where code changes are cached during the renamings detection, and only code changes are traversed subsequently rather than the entire Git history.

Because only rename refactorings are needed, instead of more advanced approaches [19], our current implementation uses lightweight approximations, such as basic method signature and body similarity comparison. Nevertheless, during our evaluation, we were able to successfully analyze the change history of code surrounding 99.22% of ~4K logging statements across 18 projects. In the future, we will explore integrating more advanced techniques.

For copying, we use the copy detection features of Git at the file level. If Git detects a copied file, any DOI values associated with the original file serve as the starting values for elements in the new file. We will explore integrating more advanced copy detection, e.g., for methods [20], in the future.

#### *3.4. DOI-Feature Logging Level Association*

Step 5 results in a rich DOI model where the most and most recently edited code is correlated with the highest DOI and vice-versa. Final DOI values that are non-negative reals are then partitioned so that DOI ranges can be associated with log levels (step 6). The association is then used to discover mismatches

between interest and feature logging statement levels (step 7).

Partitions are created by subtracting the largest DOI value by the smallest and dividing the result by the number of available levels, producing a DOI partition range size. Then, each DOI range is associated with levels by order. For example, the least logging level (e.g., `FINEST`) is associated with the first DOI partition (e.g.,  $[0, 2.54)$ ). However, this scheme can be influenced by the log category heuristic, i.e., treating `WARNING` and `SEVERE` as categories rather than levels. In such cases, specific partitions will dynamically not cause mismatches to be detected, potentially affecting the transformations performed. We utilize equal range size of DOI partitions as the initial step of research because it is the most straightforward way to deal with the matching. In the future, we will consider Machine Learning approaches to improve our algorithms.

The above scheme creates equivalently-sized partitions. Elements tend not to change often over time; therefore, a naïve partitioning may result in uneven distribution. Luckily, Mylyn supports customizable element decay rates (cf. § 3.3.1), where—as other elements become more interesting—less and less frequently edited elements lose their “interest” at a specific rate. The default decay rate does not suffice because Mylyn was not originally designed to process the sheer number of modifications typical found in Git repositories contributed by multiple developers working on many tasks. Instead, it was designed to record *IDE interactions* made by a *single* developer working on a *single* task. Hence, the default decay rate causes elements to decrease in DOI rapidly; thus, we decreased the rate significantly. The decay rate was calculated by analyzing two open-source projects (i.e., `selenium` and `IRCT`). It was later tweaked to fit into all of the projects we studied. Although several partitioning schemes were attempted, we found that combining equivalently-sized partitions with a reduced decay rate worked the best.

Once partitions are formed, potential mismatches are discovered (step 7) by comparing partitions with the current statement levels from step 2. If a mismatch is found, the statement is marked for potential transformation barring heuristics.

### 3.5. Feature Logging Statement Classification & Heuristics

Distinguishing between feature logging statements and other kinds of logging—and transformation guidance—is accomplished via heuristics (step 8). We conducted a pilot study to assist in the design of heuristics through trial and error. We used an early version of the tool and sent pull requests to large and popular open-source projects with a small subset of the tool’s output. After that, we obtained feedback from the developers. We investigated the code surrounding logging statements based on their feedback to improve our approach. Heuristics are used to avoid undesirably transforming logging statements, e.g., lines 13, 18, and 25 of `lst. 4`. Here is a list of heuristics, with their abbreviations in bold font:

1. **WS**: Treat particular levels as log *categories*.
2. **LOW**: Never *lower* the level of logging statements:
  - (a) **CTCH**: appearing within catch blocks.
  - (b) **IFS**: immediately following branches (e.g., `if`, `else`, `switch`).
  - (c) **KEYL**: *having* particular keywords in their log messages.
3. **CNDS**: Never *change* the level of logging statements immediately following branches whose condition contains a log level.
4. **KEYR**: Never *raise* the level of logging statements *without* particular keywords in their log messages. This only applies for target levels **WARNING** or **SEVERE** (**WARN** and **ERROR** in SLF4J) —typically used in more critical situations—and does not apply when **WS** is enabled.
5. **INH**: Only *consistently* transform the level of logging statements appearing in overriding methods.
6. **TDIST**: Only transform the level of logging statements up to a transformation distance threshold.

*Logging Categories.* For **WS**, a developer may choose to treat **WARNING** and **SEVERE** as logging statement categories rather than traditional levels. Such log levels are more likely to be associated with failures or potential bugs rather than the feature implementation surrounding the logs. This way, developers can denote that logging statements with such levels have special semantics and are not on a “sliding scale.” Denoting **WARNING** and **SEVERE** as categories can be used to avoid transforming lines 13 and 25 of lst. 4. In SLF4J, developers could use level **WARN** and **ERROR** as logging statement categories.

*Catch Blocks & Branches.* Logging statements appearing in catch blocks may serve as error notifications [21]. This critical logging information for bug detection in catch blocks should not be overlooked, so we proposed **CTCH** that ensures that the level of these statements is never reduced. For example, line 25, lst. 4 did not have its level lowered despite the lowering of a level at line 22 due to recent non-local edits. **IFS** is similar to **CTCH**, with an example at line 13, lst. 4. Below is an abbreviated example from `blueocean-plugin` [22], where two similar logging statements appear in each of the different contexts:

```
try {
    node = execution.getNode(action.getUpstreamNodeId());
} catch (IOException e) {
    LOGGER.warning("Couldn't retrieve upstream node: " + e);}
if (node == null) {
    LOGGER.warning("Couldn't retrieve upstream node (null)");}
```

*Log Wrapping.* **CNDS** prevents logging semantic violations when run-time level checks redundantly guard logging statements. Consider the below example from `guava` [23]:

```
1 if (logger.isLoggable(Level.SEVERE))
2  logger.log(Level.SEVERE, message(ctxt), exception);
```

Altering the level at line 2 without also changing the level at line 1 would be counterproductive. More sophisticated analysis is necessary to handle such cases, which are not typical. Moreover, other approaches [24] do not deal with this.

*Keywords.* **KEYL** and **KEYR** help distinguish feature logging statements using keywords, which originated from our evaluation and developer feedback. For the former, we manually assessed the transformations made by earlier versions of our approach, comparing them to the surrounding context. For the latter, developers commented on the transformations made by earlier versions of our tool. In both cases, we noted common keywords that appeared in logging statement messages. These words elicit strong emotions. Rather than the feature implementation surrounding the logs, they are more likely to have an association with probable issues. Stopgap words are used to maximize coverage, and keywords must appear in the literal parts of the log message construction. Keywords for **KEYL** include “fail,” “disabl,” “error,” and “exception;” keywords for **KEYR** include “stop,” “shut,” “kill,” “dead,” and “not alive.”

For example, in lst. 4, the levels at lines 8–9 are allowed to *increase* from **FINER** to **FINE** because the target levels are neither **WARNING** nor **SEVERE**, passing **KEYR**. The level at line 22 was allowed to *decrease* from **FINE** to **FINEST** as there are no “anti-lowering” keywords (**KEYL**). In contrast, the levels at lines 18 without local edits and 25 were not lowered, partly due to having anti-lowering keywords.

While the keywords are not exhaustive, they were derived via an extensive study and assistance of open-source developers. Nevertheless—in the future—we will explore using machine learning (ML) for broader classification, as well as adding more keywords related to security and privacy.

*Subtyping.* **INH**—formulated using developer feedback—applies when mismatches are found in methods involved with inheritance. Specifically, if method  $M'$  overrides method  $M$  and both  $M$  and  $M'$  include level mismatches, the target levels must be consistent as to preserve a behavioral subtyping-like [25, 26] relationship w.r.t. logging.

*Transformation Distance.* Because logging is pervasive, our approach may suggest widespread modifications. As such, **TDIST**—also from developer feedback—curtails the degree of level transformations using a threshold, which is a setting



in our tool. Finally, all mismatches that pass the enabled heuristics are transformed via AST node replacements (step 9).

## 4. Evaluation

### 4.1. Implementation

Our approach is implemented as an open-source Eclipse IDE [14] plug-in [27] and built upon JGit [16], for Git extraction, and Mylyn [13], for DOI manipulation. Eclipse is leveraged for its extensive source-to-source transformation support [28] and that it is entirely open-source for all Java development. It also supports widely-used build systems, so that projects using different IDEs may also use our plug-in.

Mylyn provides many DOI facilities, including model tuning, interaction event modeling, and interaction prediction. Integrating with Mylyn also makes it possible to—in the future—combine historical edit events from Git with IDE interactions—potentially leading to a richer DOI model. It may also be possible to extend Mylyn to populate DOI models with version control events [29], solving an open Eclipse bug [30].

Eclipse ASTs with source symbol bindings are used as an intermediate representation. Two popular logging frameworks, namely, JUL [18] and SLF4J [31], are currently supported. Heuristics (§ 3.5) are presented as tool options.

Although we depend on Mylyn for DOI and Eclipse for code analysis, transformation, and preview-panes, it may be possible to convert our tool to a GitHub App [32] that would monitor Git commits, periodically update an isolated DOI model, and generate recommended modifications as pull requests. This future work may follow recent work on refactoring bots [33].

### 4.2. Experimental Evaluation

#### 4.2.1. Research Questions

We answer the following questions:

**RQ1.** How applicable is our tool to and how does it behave with real-world open-source software?

**RQ2.** Can our tool help bring focus to buggy code?

**RQ3.** Are our tool’s results acceptable? What is its impact?

**RQ1** answers whether the proposed approach scales, in terms of SLOC, number and usages of logging statements, and revision history length, to real-world projects. It also provides insight on how logging statements are used and the contexts for which they appear by (i) assessing heuristics applicability, (ii) studying the degree and directionality of mismatch between edit frequency and logging levels, and (iii) measuring level distribution before and after our tool’s application for fuller logging level spectrum usage by developers.

**RQ2** inquires about our tool’s ability to increase buggy code focus by altering feature logging statement levels. To help potentially expose bugs, **RQ2** assesses whether our tool increases levels of feature logging statements in the context of bug fixes (i.e., “buggy” contexts) and likewise decreases levels in non-“buggy” contexts. Lastly, **RQ3** gauges whether the transformations are acceptable and their impact on developer communities. The former effectively evaluates the heuristics, while the latter assesses the extent to which our tool’s transformations affect developers as a whole.

To answer **RQ1**, quantitative (§ 4.2.2) and qualitative (§ 4.2.3) analyses are performed. To answer **RQ2**, we applied our tool to software versions leading up to buggy and non-buggy feature log contexts mined from software repositories. Finally, a pull request study is issued to answer **RQ3** (§ 4.2.5). Our dataset [34] is available.

#### 4.2.2. Quantitative Analysis

We ran our tool on a large corpus.

*Subject Selection & Details.* This phase involves 18 open-source Java applications and libraries of varying sizes and domains (tab. 1). Subjects were also

subject	HEAD	KLOC	Kcms	$\delta$ KLOC	fw	logs	fails	trns	ctch	ifs	cncls	keyl	keyr	inh	$\overline{dist}$ ( $\sigma_{dist}$ )	$\sigma_{pre}$	$\sigma_{post}$	low	ise	t (m)
bc-java	a66e904	701.57	6.00	1,213	1	56	4	11	5	13	6	16	0	0	2.64 (1.30)	1.49	2.01	11	0	62.78
blueocean*	c3bfac5	49.32	4.06	138	2	109	0	13	11	6	5	57	9	0	1.54 (0.63)	1.65	1.74	11	2	2.93
californium	5b026ab	79.35	1.00	144	1	986	0	141	23	123	95	317	176	2	1.52 (0.73)	N/A	N/A	118	22	7.51
CoreNLP	d7c356db	535.07	6.00	43,686	2	358	0	146	0	87	21	55	0	0	3.10 (1.43)	1.84	2.16	146	0	444.99
deviation*	88751d6	6.52	0.08	25	1	91	0	11	1	7	8	48	3	0	2.64 (1.23)	1.69	2.07	11	0	0.06
errorprone	2118ba3	164.65	3.99	517	1	12	0	1	1	0	1	9	0	0	2.00 (0.00)	1.61	1.86	1	0	17.36
guava	71de406	393.69	1.00	297	1	36	0	0	0	0	0	33	0	0	0.00 (N/A)	1.52	1.52	0	0	29.63
hollow	ff635ee	68.60	0.90	159	1	31	0	6	0	3	0	22	0	0	3.50 (1.12)	1.02	1.98	5	1	2.86
IRCT	d67f539	42.29	0.89	194	1	13	0	6	0	0	0	2	1	0	2.00 (1.41)	1.22	1.51	6	0	2.07
JaepFX	14c2a4c	24.06	0.37	121	1	21	0	2	4	0	7	8	0	0	2.50 (1.50)	1.47	1.69	2	0	0.77
jdp*	515b7e7	35.86	1.00	100	2	799	0	180	17	15	4	84	399	0	1.70 (0.63)	0.00	1.79	139	41	3.18
jenkins	4f0436a	160.60	1.00	25	1	503	22	60	29	52	54	244	22	0	2.17 (1.25)	1.73	1.84	58	2	3.81
jsilhouette	8de4e64	2.35	0.03	6	1	22	0	0	0	8	0	14	0	0	0.00 (N/A)	0.00	0.00	0	0	0.04
junit5	3969585	73.59	5.98	381	1	9	2	4	0	0	0	2	0	0	3.00 (1.00)	0.00	1.67	4	0	75.25
OpenGrok	0f3b14d	71.65	4.93	377	1	592	3	61	44	49	41	353	13	0	2.70 (1.47)	1.81	2.07	53	8	15.20
selenium	a49fb60	93.61	1.00	76	1	94	0	28	5	6	0	48	4	0	2.96 (0.94)	1.00	1.75	27	1	1.45
SpotBugs	ecc7994	187.78	1.00	94	2	96	0	44	2	15	6	26	0	0	1.57 (0.96)	1.74	2.30	44	0	26.59
WALA	b73a0b2	202.45	1.00	236	2	145	0	39	1	38	0	22	16	1	1.41 (0.54)	0.94	0.94	38	1	9.20
Total		2,893.01	40.22	47,787	1.28	3,973	31	753	143	422	248	1,360	643	3	2.13 (1.23) <sup>†</sup>	1.75	2.05	674	78	705.68

\* blueocean is blueocean-plugin, deviation is deviation-upload, and jdp is java-design-patterns.

<sup>†</sup> Overall average and standard deviation.

Table 1: Quantitative analysis results. Column **Kcms** is the thousands of commits analyzed. Column  $\delta$ **KLOC** is the thousands of source lines of code. Column **fw** is the number of logging frameworks detected. Column **logs** is the sum of candidate logging statements. Column **fails** is the number of statements where the current level could not be extracted. Column **trns** is the number of transformed logging statements. Column **ctch** is the sum of log levels not lowered due to **CTCH**. Column **ifs** is the sum of log levels not lowered due to **IFS**. Column **keyl** is the sum of log levels not lowered due to **KEYL**. Column **inh** is the sum of levels not modified due to **INH**. Column  $\overline{dist}$  is the average level transformation distance and corresponding standard deviation ( $\sigma_{dist}$ ). Columns  $\sigma_{pre}$  and  $\sigma_{post}$  are the stdev of log levels *before* and *after* transformation for JUL, respectively. Column **lower** is the number of lowered log levels. Column **raise** is the number of raised log levels. Column **t (m)** is the total running time in minutes.

chosen to include logging statements using either JUL or SLF4J and have at least one source code change to a method containing a logging statement (i.e., a candidate statement). Column **KLOC** ranges from  $\sim 6\text{K}$  for **deviation-upload** to  $\sim 535\text{K}$  for **CoreNLP**. Column **Kcms** can significantly affect the number of candidate statements. For subjects with long Git histories, we choose a large number of commits to obtain more candidates. For others, a relatively small number of commits was adequate to obtain all logging statements as candidates.

*Execution & Run Time.* The analysis was executed on an Intel Xeon-E3 with four cores, 31 GB RAM, and a 20 GB maximum heap size. Column **t (m)** averages 10.66 secs per candidate statement and 0.89 secs per KLOC changed ( $\delta\text{KLOC}$ ). The running time is highly related to analyzed source lines per commit ( $\delta\text{KLOC}/\text{Kcms}$ ), with a 0.98 Pearson correlation coefficient that ranges from  $-1$  to  $1$  and  $1$  is an exact linear correlation. **CoreNLP**, having a particularly long and intricate Git history, is an outlier, taking over half of the running time.

**Takeaway 1:** Logging statements process in 10.66 secs and one  $\delta\text{KLOC}$  in 0.89 secs, with the processing time highly connected to LOC/commit (0.98 Pearson correlation).

*Log Level Reinvigoration.* We successfully analyzed 99.22% of  $\sim 4\text{K}$  candidate statements (column **logs**) across 18 subjects.<sup>4</sup> Failures in column **fails** include when the logging level is stored in a variable.

**Takeaway 2:** 99.22% of  $\sim 4\text{K}$  logging statements were successfully analyzed.

Column **trns** denotes 19% of candidate logging statements are transformed. This metric is not a success rate as it is dependent on the mismatches found and the enabled heuristics. Instead, it demonstrates that the transformations made by our tool—when applied to real-world projects—are subtle enough that they

---

<sup>4</sup>**junit5** only has 9 candidates as this version mainly uses custom logging. It is included since its previous versions use JUL and are subsequently studied.

may be appealing to more risk-averse projects. It also shows that manual effort can be labor-intensive, involving multi-developer, historical analysis and transformation, as manually modifying 753 logging statements is non-trivial.

**Takeaway 3:** The transformations were subtle (19%) yet labor-alleviating (753 transformations).

*Heuristics.* To more fully understand the effects, all heuristics were enabled except log categories and transformation distance (**WS** and **TDIST**, respectively, in § 3.5) in this phase. Columns **ctch**, **ifs**, and **keyl** are the sums of levels not lowered due to **CTCH** (3.60% of candidates), **IFS** (10.62%), and **KEYL** (34.23%), respectively. Columns **cnnds** and **keyr** are the sums of levels not transformed due to **CNDS** (6.24%) and not raised due to **KEYR** (16.18%), respectively.

The discrepancy between **keyl** (34.23%) and **keyr** (16.18%) suggests that our tool is more frequently attempting to lower levels than raise them. This tendency may be due to few program elements changing at a specific time. As such, it is expected that levels would be lowered more often than they are raised, but **KEYL** is curbing the lowering. Thus, contrary to previous studies [9], log messages may play a more significant role in determining log types than their placement.

In summary, we filtered out 2,819 non-feature logs that are calculated by adding the total number of logs in columns **ctch**, **ifs**, **cnnds**, **keyl**, **keyr**, and **inh**, and evaluated 1,154 feature logs, which account for 29.05% of all candidate logging statements.

*Transformation.* To ascertain transformation grade, we consider the distance between two adjacent levels to be 1.

**Takeaway 4:** Transformations (avg. 2.13 levels) were not overly drastic yet far enough to be noticeable.

To discover transformation directionality, columns **low** and **rse** depict the number of lowered (89.51%) and raised (10.36%) log levels, respectively, stipulating that our tool more frequently lowers levels. As with columns **keyl** and

**keyr**, this is most likely indicative of the relatively small number of features that developers focus on at a particular time.

**Takeaway 5:** Levels are typically lowered (89.51%), potentially facilitating focus on fewer features of interest.

*Log Level Diversity.* Evident from column  $\sigma_{pre}$ , which is the stdev of log levels *before* transformation,<sup>5</sup>—averaging only 1.75—is that the full spectrum of available log levels is not always utilized. Column  $\sigma_{post}$ , on the other hand, is stdev *after* transformation—averaging 2.05.

**Takeaway 6:** Our tool increased logging level distribution by 17.14%, thereby utilizing more level spectrum.

#### 4.2.3. Qualitative Analysis

We discuss several instances where our tool did and did not work as intended. While evaluating **guava**, one manual modification to unit test code where a log level was being tested (i.e., `assertEquals(Level.INFO, record.getLevel())`) was necessary. Our tool transformed this tested log level to **FINEST**, which failed the test suite. While more sophisticated analysis is needed to handle such cases, we conjecture they are rare; we only found one.

The following transformation occurred in **selenium**. The color **red** represents lines removed and **green** lines added:

```
1 if (!check.isAlive()) {
2 - LOG.info("Node is not alive: " + check.getMessage());
3 + LOG.severe("Node is not alive: " + check.getMessage());
4   // Throw an exception to force another check sooner.
5   throw new UnsupportedOperationException("Node can't..");}
```

The logging statement at line 2 indicates a failure before an exception throw on line 5. The level is erroneously **INFO**. Because this code area is a “hot spot,” i.e., being either frequently or recently edited, fortunately, our tool fixed this level by transforming it to **SEVERE** (line 3). This fix was later incorporated into **selenium**’s mainline [35].

---

<sup>5</sup>Only includes JUL. N/A indicates that JUL was not used.

The following transformation occurred in **CoreNLP**:

```
1 if (o == null) {
2   logger.severe("READER ERROR: Failed to find...");
3 - logger.severe("This happens because a few relation...");
4 + logger.fine("This happens because a few relation...");
```

In place of using string concatenation in a single statement, line 2 logs an event overview, while line 3 logs the details. Although the heuristics worked for line 2, they failed for line 3, resulting in an incorrect transformation at line 4. This pattern was not observed in other subjects.

The following transformation took place in **Jenkins**:

```
1 - LOGGER.log(INFO, "{0} main build action completed:{1}");
2 + LOGGER.log(FINEST, "{0} main build action completed:{1}");
```

As the feature associated with the log at line 1 was not of recent developer interest, our tool correctly lowered its level on line 2. This transformation, with developers expressing that, “[it is p]robably a good idea: [i]t’s time we started removing this from the general system log [36],” was also accepted into **Jenkins**’ mainline. A subsequent comment for a similar transformation further motivated our approach by stating that, “I [ha]ve [grew] so used to these messages over the years.”

#### 4.2.4. Bug Study

To answer **RQ2**, we mined software repositories for bug and non-bug fixes with nearby feature logs and applied our tool to versions leading up to the changesets.

*Methodology.* To discover changesets having buggy and non-buggy logging contexts, we used `gitcproc` [37], a tool for processing and classifying Git commits that has been used previously [38–40] to automatically detect and analyze bug fixes in GitHub. Natural language processing (NLP) is used on commit log messages to identify bug fixes, and changesets can be queried for specific keywords. As we focus on lines with logging statements *surrounding* modified lines, e.g.,

	orig. level	bug	LOW	KEYR	dir
1	INFO	<i>T</i>	N/A	N/A	NONE
2	FINEST/FINER/FINE	<i>T</i>	N/A	<i>T</i>	RAISE
3	FINEST/FINER/FINE	<i>T</i>	N/A	<i>F</i>	NONE
4	FINER/FINE/INFO	<i>F</i>	<i>T</i>	N/A	NONE
5	FINER/FINE/INFO	<i>F</i>	<i>F</i>	N/A	LOWER
6	FINEST	<i>F</i>	N/A	N/A	NONE

Table 2: “Ideal” level transformation directions for feature logs. Columns **orig. level** and **bug** denote the original log level and whether the log is in a buggy context, respectively. Columns **LOW** and **KEYR** represent whether the statement passes any of **CTCH**, **IFS**, **KEYL** and **KEYR**, respectively. These columns help to hone in the analysis on feature logs. Finally, column **dir** portrays the ideal level direction. N/A is either *T* or *F*.

line 4, lst. 2, we altered `gitcproc` to search *outside* changesets, set the query keywords to those corresponding to logging APIs,<sup>6</sup> and manually examined the results.

Once commits containing buggy and non-buggy feature logging statement contexts were identified, we ran our tool on the project version of the *immediately preceding* commits. The tool examines up to 1000 commits in each project version’s git history. All heuristics in § 3.5 except for **TDIST** were enabled to distinguish feature logging statements in contexts, **TDIST** was set to `INT_MAX`, and  $\leq 1\text{K}$  commits leading up to the preceding commits were processed. For each subject,  $\leq \sim 20$  of each context type was chosen, such that the number of project versions considered was minimized. This increased the likelihood of successful project building as the selected versions may have been intermediate. Otherwise, contexts were chosen randomly. Only 12 subjects were used as some had long running times by `gitcproc`, no logging contexts, or only unbuildable versions.

To assess our tool’s transformations when applied to project versions immediately preceding commits containing buggy and non-buggy feature logging contexts, we define an “ideal” level direction to be used as an oracle (tab. 2).

<sup>6</sup>For succinctness, only JUL was used.



subject	vers	ctxts		rse		low		none		i=a
		bug	¬bug	idl	act	idl	act	idl	act	
bc-java	8	0	10	0	0	2	2	8	8	10
blueocean	6	7	3	4	2	0	0	6	8	6
errorprone	2	0	2	0	0	0	0	2	2	2
guava	11	1	14	0	0	1	0	14	15	14
hollow	10	1	9	0	0	2	0	8	10	8
IRCT	10	11	20	0	0	8	10	23	21	29
JacpFX	4	3	1	0	0	0	1	4	3	3
jenkins	5	9	22	4	0	3	2	24	29	24
OpenGrok	6	28	20	0	1	6	12	42	35	35
selenium	17	20	20	0	0	6	8	34	32	34
SpotBugs	18	20	20	0	0	6	6	34	34	36
WALA	4	1	3	1	0	0	0	3	4	3
Total	101	101	144	9	3	34	41	202	201	204
		245								

Table 3: Feature logging statements in change contexts. Column **vers** depicts the number of subject versions, column **ctxts** the buggy (**bug**) and non-buggy (**¬bug**) logging contexts extracted, and columns **rse**, **low**, and **none** the feature log levels that are ideally (**idl**) and actual were (**act**) raised, lowered, and not altered, respectively. Column **i=a** is the number of matching ideal and actual level directions.

For example, lst. 2 fixes a bug accepting invalid temperatures. Here, **FINER** logs at lines 4–5 are in a *buggy* context. Higher levels, e.g., **FINE** (lines 8–9, lst. 4), may have helped bring attention to this bug *earlier*, i.e., by documenting invalid temperature values *more* prominently (row 2, tab. 2). Conversely, a **FINE** log, e.g., line 19, lst. 1, that is *not* in a buggy context may have its level lowered (e.g., line 22, lst. 4) so that *other* logs *in* buggy contexts are more noticeable (row 5). Rows 1 and 6 are boundaries—logs in buggy contexts at the highest level cannot be raised and vice-versa. Logs in row 3 are in buggy contexts but do not have their levels raised due to failing **KEYL**, while logs in row 4 are in non-buggy contexts but are not lowered due to failing **LOW**.

*Results.* In tab. 3, column **ctxts** the buggy averages 1 per vers and non-buggy **¬bug** averages 1.43 per vers.

**Takeaway 7:** Levels of feature logs in change contexts were transformed in ideal directions  $\sim 83\%$  ( $204/245$ ) of the time, potentially bringing problematic feature implementations into higher focus and exposing bugs.

*Discussion.* Higher/lower levels of feature logs in buggy/non-buggy contexts may have helped reveal/highlight problematic feature implementations, especially considering that such logs, e.g., lines 4–5 of lst. 2, typically include critical variables. Having these logs appear more prominently may induce fixes, e.g., lines 3 and 6. Unnecessarily altering levels in equally essential to avoid false positives that introduce noise. As buggy code tends to be more frequently and recently edited [41], our approach is well-suited to ideally adjust—if necessary—feature logging statement levels in directions that may divulge bugs and avoid time-consuming fixes.

#### 4.2.5. Pull Request Study

To answer **RQ3**, we submitted pull requests containing our tool’s transformations.

*Results.* As of this writing, requests have been made to 16 projects— 2 pull requests have been accepted, 5 have been refused, and 11 have been pending. Projects that merged transformations include **Jenkins**, a well-known continuous integration (CI) engine, having  $\sim 15\text{K}$  stars and  $\sim 6.1\text{K}$  forks on GitHub, and **selenium**, a prominent web application testing and automation tool, having  $\sim 17\text{K}$  stars and  $\sim 5.6\text{K}$  forks. As suggested by these statistics, although only two requests been accepted so far, the merged transformations have far-reaching impact as the projects include libraries and frameworks that are widely used in diverse circles—**selenium**, for example, is used by 55,931 other projects [42]. Furthermore, the acceptance and rejection rates are comparable to that of previous work [43].

**Takeaway 8:** 2 projects, both widely-used, having  $\geq 56\text{K}$  integrations—thus ensuring developer impact—accepted pull requests at rates compara-

ble to previous work [43].

*Discussion.* Our pull requests were rejected by five different projects of varying sizes. Developers from these projects, with the exception of the `guava` developer, are more indifferent than developers from projects with accepted pull requests. We got a response that said, “We are not interested in your tool. Please do not send any further PRs from this address.” Furthermore, PRs with fewer changes are more likely to be accepted, whereas those with more changes are prone to be rejected.

Apparent during the study was that our approach encourages developers to actively consider how their logging statement levels evolve alongside their core software. Feedback from rejected requests includes questions on whether or not our approach applies to very mature and largely stable projects that are in “maintenance mode” [44, 45]. In this scenario, developers respond to bug reports, resulting in *consistent* modifications to *diverse* system modules. In such cases, our tool will never pick more “interesting” parts of the system as they are all equally or not equally “interesting;” application/system code that is under active development may be more amenable.

A question [46] was also raised regarding the approach’s applicability to library vs. application code. Notably, parts of a library that are important to library developers may not match the interest level of application developers using the library. This problem—which may be more prevalent with public API implementations—touches on a broader issue of log composition, i.e., application developers’ log intentions may not coincide with the application’s dependencies in general.

#### 4.3. Threats to Validity

Subjects may not be representative of real-world log usage. For mitigation, subjects were chosen from diverse domains and sizes, as well as those used in previous studies [47, 48].

Git repositories with very large commits may be too coarse-grained to detect program element modifications accurately. This limitation is standard among

approaches that mine software repositories. Furthermore, our keyword related heuristics cannot work with typos. There were still  $\sim 2\text{K}$  logging statements in our study whose levels were not altered due to keywords, suggesting that typos are not pervasive. Adding support for the use of misspelled keywords is straight-forward.

Moreover, during the pull request study (§ 4.2.5), there was at least one instance where developers desired to modify logging statement levels eventually but had not yet done so.

We may miss some of these “wrapped” logs that are excluded by the **CNDS** heuristic. However, we did not find a significant number of such logs during our evaluation (only  $\sim 6.3\%$ ). Moreover, such logs are likely to reside in developer-defined library code that is used by many features, making it difficult to correlate DOI. Lastly, we understand that the pattern involving **CNDS** is used to improve performance. However, with the relatively recent introduction of logging API accepting  $\lambda$ -expression, we expect that developers will increasingly use those instead of the pattern.

## 5. Related Work

Logging frameworks may include logger hierarchies [8, 12], where log objects can be designated to log certain features implementations. Individual loggers can thus be enabled or disabled to facilitate focus on particular feature implementations. However, because they are typically many features (i) whose interests change over time [15] and (ii) whose implementations are not localized [49], developers are still burdened with manually maintaining logger hierarchies.

Li et al. [9] determine log levels for *new* logging statements to be consistent with other logging practices within the same project. On the other hand, we focus on log level *evolution* and how levels relate to feature interest. While their work can be retargeted for evolution by treating the logging statement under consideration as “new” and subsequently predicting a possibly new level, the goals of their approach are quite different from ours. Firstly, they focus on

*general* logging statements, whereas we focus on feature logs. Secondly, they aim to predict a level that fits well with the current *project* trends, including log placement, existing logs in the same file, and message content. On the contrary, our goal is to better align feature logging statement levels with current *feature* interest, which may have little bearing on placement. In other words, feature interest at a particular point in time may be well-independent of the logging practices previously employed.

Chen and Jiang [7] and Hassani et al. [10] detect and correct mismatches between log messages and levels. However, they do not consider how varying developer interests in particular features affect logging levels over time. While these approaches are useful in discovering error logging statements that use lower-than-normal logging levels, they may not be as useful for event-type logs, which may be more tied to features.

Li et al. [43] predict log revisions by mining the correlation between logging context and code modifications with the premise that logging code in similar contexts deserves similar modifications. As far as we can tell, however, they do not explicitly deal with logging levels. Kabinna et al. [3] and Li et al. [24] determine the likelihood of log change but do not suggest a specific modification. Li et al. [50] predict—using topic modeling—the likelihood that a particular code snippet should include logging. Shang et al. [51] examine log lines using development knowledge to resolve log inquiries, Yuan et al. [52] add appropriate logging statements to enhance failure diagnosis, and Zhu et al. [53] evaluate log parsers.

Several approaches combat information overload. Haas et al. [54] use static analysis to detect unnecessary source code. Fowkes et al. [55] introduce auto-folding for source code summarization. Information overload is also an issue for logging, and many approaches enhance logging statements by, e.g., determining log placement [56, 57], optimizing the number of logging statements [58], and enriching log messages [11, 59, 60]. Xu et al. [61] mine log output to detect problems.

Khatchadourian et al. [62] also integrate the Mylyn DOI model but for

Aspect-Oriented Programming (AOP). To the best of our knowledge, we are the first to manipulate a DOI model using software repository mining programmatically. Other work mines software repositories for evolution [63], detecting refactorings [19], and design flaw detection [64]. Additionally, approaches support software evolution more generally, e.g., by refactoring programs to use enumerated types [65], default methods [66], and lambda expressions [48]. Bhat-tacharya et al. [67] also use graphs for software evolution.

## 6. Conclusion & Future Work

Our automated approach “reinvigorates” feature logging statement levels based on “interest” of surrounding code as determined by software repository mining. Distinguishing feature logs is performed via introduced heuristics. The approach is implemented as an Eclipse IDE plug-in, using JGit and Mylyn, and evaluated on 18 projects with  $\sim 3$  MLOC and  $\sim 4$ K logging statements. Our tool successfully analyzes 99.22% of logging statements, increases log level distributions by  $\sim 20\%$ , increases the focus of logs within bug fix contexts  $\sim 83\%$  of the time, and integrated transformations into several projects. In the future, we will explore leveraging existing Mylyn task contexts, expanding the heuristics, and surveying developers. In addition, we will enhance our algorithm to partition DOI ranges by using Machine Learning approaches.

## Acknowledgments

We would particularly like to thank all the developers that offered extremely helpful feedback through our pull requests. These include Manu Sridharan, Daniel Beck, and David P. Baker. We would also like to thank Oren Friedman and Walee Ahmed for their help with experiments and infrastructure, respectively, as well as Shigeru Chiba and Hidehiko Masuhara for their thoughtful insight and feedback. Support for this project was provided by PSC-CUNY Awards #617930049 and #638010051, jointly funded by The Professional Staff Congress and The City University of New York.

## References

- [1] Mik Kersten and Gail C. Murphy. “Mylar: a degree-of-interest model for IDEs”. In: *International Conference on Aspect-Oriented Software Development*. Chicago, Illinois: ACM, 2005, pp. 159–168. ISBN: 1-59593-042-6. DOI: 10.1145/1052898.1052912.
- [2] A. Rozinat and W. M. P. van der Aalst. “Conformance Testing: Measuring the Fit and Appropriateness of Event Logs and Process Models”. In: *International Conference on Business Process Management*. BPM ’05. Nancy, France: Springer-Verlag, 2005, pp. 163–176. ISBN: 3540325956. DOI: 10.1007/11678564\_15.
- [3] Suhas Kabinna, Cor-Paul Bezemer, Weiyi Shang, Mark D. Syer, and Ahmed E. Hassan. “Examining the Stability of Logging Statements”. In: *Empirical Softw. Engg.* 23.1 (Feb. 2018), pp. 290–333. ISSN: 1382-3256. DOI: 10.1007/s10664-017-9518-0.
- [4] Jiaqi Tan, Xinghao Pan, Soila Kavulya, Rajeev Gandhi, and Priya Narasimhan. “SALSA: Analyzing Logs as State Machines”. In: *USENIX Conference on Analysis of System Logs*. WASL ’08. San Diego, California: USENIX Association, 2008, p. 6.
- [5] Yi Zeng, Jinfu Chen, Weiyi Shang, and Tse-Hsun Chen. “Studying the characteristics of logging practices in mobile apps: a case study on F-Droid”. In: *Empirical Softw. Engg.* 24 (Feb. 2019), pp. 3394–3438. ISSN: 1573-7616. DOI: 10.1007/s10664-019-09687-9.
- [6] M. D. Syer, Z. M. Jiang, M. Nagappan, A. E. Hassan, M. Nasser, and P. Flora. “Leveraging Performance Counters and Execution Logs to Diagnose Memory-Related Performance Issues”. In: *International Conference on Software Maintenance*. Sept. 2013, pp. 110–119. DOI: 10.1109/ICSM.2013.22.
- [7] Boyuan Chen and Zhen Ming Jiang. “Characterizing and Detecting Anti-Patterns in the Logging Code”. In: *International Conference on Software Engineering*. ICSE ’17. Buenos Aires, Argentina: IEEE Press, 2017, pp. 71–81. ISBN: 9781538638682. DOI: 10.1109/ICSE.2017.15.
- [8] Oracle. *Logger (Java SE 10 & JDK 10)*. 2018. URL: <http://docs.oracle.com/javase/10/docs/api/java/util/logging/Logger.html> (visited on 02/29/2020).
- [9] Heng Li, Weiyi Shang, and Ahmed E. Hassan. “Which Log Level Should Developers Choose for a New Logging Statement?” In: *Empirical Softw. Engg.* 22.4 (Aug. 2017), pp. 1684–1716. ISSN: 1382-3256. DOI: 10.1007/s10664-016-9456-2.
- [10] Mehran Hassani, Weiyi Shang, Emad Shihab, and Nikolaos Tsantalis. “Studying and detecting log-related issues”. In: *Empirical Software Engineering* (Mar. 2018). ISSN: 1573-7616. DOI: 10.1007/s10664-018-9603-z.

- [11] Pinjia He, Zhuangbin Chen, Shilin He, and Michael R. Lyu. “Characterizing the Natural Language Descriptions in Software Logging Statements”. In: *International Conference on Automated Software Engineering*. ASE ’18. Montpellier, France: ACM, 2018, pp. 178–189. ISBN: 9781450359375. DOI: 10.1145/3238147.3238193.
- [12] Apache Software Foundation. *Log4j*. Log4j 2 Architecture. 2020. URL: [http://logging.apache.org/log4j/2.x/manual/architecture.html#Logger\\_Hierarchy](http://logging.apache.org/log4j/2.x/manual/architecture.html#Logger_Hierarchy) (visited on 06/12/2020).
- [13] Eclipse Foundation, Inc. *Eclipse Mylyn Open-Source Project*. 2020. URL: <http://eclipse/gC> (visited on 06/12/2020).
- [14] Eclipse Foundation, Inc. *Eclipse IDEs*. Aug. 2020. URL: <http://eclipse/gD> (visited on 03/02/2020).
- [15] Mik Kersten and Gail C. Murphy. “Using Task Context to Improve Programmer Productivity”. In: *ACM Symposium on the Foundations of Software Engineering*. SIGSOFT ’06/FSE-14. Portland, Oregon, USA: ACM, 2006, pp. 1–11. ISBN: 1-59593-468-5. DOI: 10.1145/1181775.1181777.
- [16] Eclipse Foundation, Inc. *JGit*. 2020. URL: <http://eclipse/gF> (visited on 03/02/2020).
- [17] QOS.ch. *SLF4J Manual. Typical usage pattern*. 2019. URL: [http://www.slf4j.org/manual.html#typical\\_usage](http://www.slf4j.org/manual.html#typical_usage) (visited on 03/03/2020).
- [18] Oracle. *java.util.logging (Java SE 10 & JDK 10)*. 2018. URL: <http://docs.oracle.com/javase/10/docs/api/java/util/logging/package-summary.html> (visited on 03/02/2020).
- [19] Nikolaos Tsantalis, Matin Mansouri, Laleh M. Eshkevari, Davood Mazinianian, and Danny Dig. “Accurate and Efficient Refactoring Detection in Commit History”. In: *International Conference on Software Engineering*. ICSE ’18. Gothenburg, Sweden: ACM, 2018, pp. 483–494. ISBN: 978-1-4503-5638-1. DOI: 10.1145/3180155.3180206.
- [20] Hung-Fu Chang and Audris Mockus. “Evaluation of Source Code Copy Detection Methods on FreeBSD”. In: *International Working Conference on Mining Software Repositories*. MSR ’08. Leipzig, Germany: ACM, 2008, pp. 61–66. ISBN: 9781605580241. DOI: 10.1145/1370750.1370766.
- [21] Oracle. *The catch Blocks*. 2021. URL: <https://docs.oracle.com/javase/tutorial/essential/exceptions/catch.html> (visited on 06/15/2021).
- [22] Josh McDonald. *blueocean-plugin/DownstreamJobListener.java at 522fd48 · jenkinsci/blueocean-plugin*. Jenkins. Jan. 24, 2018. URL: <http://git.io/JvVQ0> (visited on 03/05/2020).
- [23] Colin Decker. *guava/EventBus.java at 0cd4e9f · google/guava*. Google. Aug. 5, 2014. URL: <http://git.io/JvV7P> (visited on 03/05/2020).



- [24] Heng Li, Weiyi Shang, Ying Zou, and Ahmed Hassan. “Towards Just-in-Time Suggestions for Log Changes”. In: *Empirical Softw. Engg.* 22.4 (Aug. 2017), pp. 1831–1865. ISSN: 1382-3256. DOI: 10.1007/s10664-016-9467-z.
- [25] Gary T. Leavens and William E. Weihl. “Reasoning about Object-Oriented Programs That Use Subtypes”. In: *European Conference on Object-Oriented Programming*. OOP-SLA/ECOOP ’90. ACM, 1990, pp. 212–223. ISBN: 0897914112. DOI: 10.1145/97945.97970.
- [26] Barbara H. Liskov and Jeannette M. Wing. “A Behavioral Notion of Subtyping”. In: *ACM Transactions on Programming Languages and Systems* 16.6 (1994), pp. 1811–1841. ISSN: 0164-0925. DOI: 10.1145/197320.197383.
- [27] Anonymous. *Reinvigorate Logging Levels*. Version v1.9.0. Anonymized repository. 2020. URL: <https://anonymous.4open.science/r/d301551f-765a-4825-9f01-79d570e0cc94> (visited on 06/17/2020).
- [28] Dirk Bäumer, Erich Gamma, and Adam Kiezun. “Integrating refactoring support into a Java development tool”. Oct. 2001. URL: <http://people.csail.mit.edu/akiezun/companion.pdf>.
- [29] Jens von Pilgrim. *Create context from git diff. Mylyn*. Eclipse Community Forums. Eclipse Foundation. Apr. 17, 2014. URL: <http://eclip.se/gJ> (visited on 06/08/2020).
- [30] Sam Davis. *433030 – Add to Task Context does not work*. Eclipse Foundation. Apr. 17, 2014. URL: <http://eclip.se/gI> (visited on 06/08/2020).
- [31] QOS.ch. *SLF4J. Simple Logging Facade for Java*. 2019. URL: <http://www.slf4j.org> (visited on 03/02/2020).
- [32] GitHub, Inc. *Building Apps*. GitHub Developer Guide. 2020. URL: <http://git.io/JfS9E> (visited on 06/08/2020).
- [33] Vahid Alizadeh, Mohamed Amine Ouali, Marouane Kessentini, and Meriem Chater. “RefBot: Intelligent Software Refactoring Bot”. In: *International Conference on Automated Software Engineering*. ASE ’19. San Diego, California: IEEE Press, 2019, pp. 823–834. ISBN: 9781728125084. DOI: 10.1109/ASE.2019.00081.
- [34] Anonymous. *Automated Evolution of Feature Logging Statement Levels Using Git Histories and Degree of Interest*. Zenodo, Aug. 2020. DOI: 10.5281/zenodo.3698983.
- [35] Anonymous. *Rejuvenate log levels by anonymous • Pull Request #7737 • SeleniumHQ/selenium*. Following this URL may reveal author identity. Selenium. Feb. 18, 2020. URL: <http://git.io/Jfd2W> (visited on 06/17/2020).

- [36] Daniel Beck. *Reduce log levels for successful run completion and update center polling events by anonymous* · Pull Request #4345 · *jenkinsci/jenkins*. Following this URL may reveal author identity. Jenkins. Nov. 5, 2019. URL: <https://git.io/JJWzA> (visited on 07/21/2020).
- [37] Casey Casalnuovo, Yagnik Suchak, Baishakhi Ray, and Cindy Rubio-González. “GitProc: A Tool for Processing and Classifying GitHub Commits”. In: *International Symposium on Software Testing and Analysis*. ISSTA ’17. Santa Barbara, CA, USA: ACM, 2017, pp. 396–399. ISBN: 978-1-4503-5076-1. DOI: 10.1145/3092703.3098230.
- [38] Sirine Gharbi, Mohamed Wiem Mkaouer, Ilyes Jenhani, and Montassar Ben Messaoud. “On the Classification of Software Change Messages Using Multi-label Active Learning”. In: *Symposium on Applied Computing*. SAC ’19. ACM/SIGAPP. Limassol, Cyprus: ACM, 2019, pp. 1760–1767. ISBN: 978-1-4503-5933-7. DOI: 10.1145/3297280.3297452.
- [39] Yuchi Tian and Baishakhi Ray. “Automatically Diagnosing and Repairing Error Handling Bugs in C”. In: *ACM Symposium on the Foundations of Software Engineering*. ESEC/FSE 2017. Paderborn, Germany: ACM, 2017, pp. 752–762. ISBN: 978-1-4503-5105-8. DOI: 10.1145/3106237.3106300.
- [40] Raffi Khatchadourian, Yiming Tang, Mehdi Bagherzadeh, and Baishakhi Ray. “An Empirical Study on the Use and Misuse of Java 8 Streams”. In: *International Conference on Fundamental Approaches to Software Engineering*. Ed. by Heike Wehrheim and Jordi Cabot. FASE ’20. ETAPS. Cham: Springer International Publishing, Apr. 2020, pp. 97–118. ISBN: 978-3-030-45234-6. DOI: 10.1007/978-3-030-45234-6\_5.
- [41] Manishankar Mondal, Chanchal K. Roy, and Kevin A. Schneider. “Identifying Code Clones Having High Possibilities of Containing Bugs”. In: *International Conference on Program Comprehension*. ICPC ’17. Buenos Aires, Argentina: IEEE Press, 2017, pp. 99–109. ISBN: 9781538605356. DOI: 10.1109/ICPC.2017.31.
- [42] GitHub, Inc. *Network Dependents* • *SeleniumHQ/selenium*. 2020. URL: <http://git.io/JfS90> (visited on 06/10/2020).
- [43] Shanshan Li, Xu Niu, Zhouyang Jia, Ji Wang, Haochen He, and Teng Wang. “Log-tracker: Learning Log Revision Behaviors Proactively from Software Evolution History”. In: *International Conference on Program Comprehension*. ICPC ’18. Gothenburg, Sweden: ACM, 2018, pp. 178–188. ISBN: 978-1-4503-5714-2. DOI: 10.1145/3196321.3196328.
- [44] David P. Baker. *Rejuvenate logging statement levels by anonymous* • Pull Request #3435 • *google/guava*. Following this URL may reveal author identity. Google. Apr. 5, 2019. URL: <http://git.io/Jfdal> (visited on 06/17/2020).
- [45] Simon Stewart. *Rejuvenate log levels by anonymous* • Pull Request #7170 • *SeleniumHQ/selenium*. Following this URL may reveal author identity. Selenium. May 7, 2019. URL: <http://git.io/JfdaA> (visited on 06/17/2020).

- [46] David P. Baker. *Rejuvenate log levels by anonymous • Pull Request #3713 • google/guava*. Following this URL may reveal author identity. Google. Nov. 20, 2019. URL: <http://git.io/Jfdae> (visited on 07/27/2020).
- [47] Raffi Khatchadourian, Yiming Tang, Mehdi Bagherzadeh, and Syed Ahmed. “Safe Automated Refactoring for Intelligent Parallelization of Java 8 Streams”. In: *International Conference on Software Engineering*. ICSE ’19. ACM/IEEE. Montréal, QC, Canada: IEEE Press, May 2019, pp. 619–630. DOI: 10.1109/ICSE.2019.00072.
- [48] Ameya Ketkar, Ali Mesbah, Davood Mazinanian, Danny Dig, and Edward Aftandilian. “Type Migration in Ultra-large-scale Codebases”. In: *International Conference on Software Engineering*. ICSE ’19. Montreal, Quebec, Canada: IEEE Press, 2019, pp. 1142–1153. DOI: 10.1109/ICSE.2019.00117.
- [49] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean-Marc Loingtier, and John Irwin. “Aspect-oriented programming”. In: *European Conference on Object-Oriented Programming*. Springer, 1997, pp. 220–242. ISBN: 978-3-540-69127-3. DOI: 10.1007/bfb0053381.
- [50] Heng Li, Tse-Hsun Chen, Weiyi Shang, and Ahmed E. Hassan. “Studying software logging using topic models”. In: *Empirical Software Engineering* 23.5 (Oct. 2018), pp. 2655–2694. ISSN: 1573-7616. DOI: 10.1007/s10664-018-9595-8.
- [51] W. Shang, M. Nagappan, A. E. Hassan, and Z. M. Jiang. “Understanding Log Lines Using Development Knowledge”. In: *International Conference on Software Maintenance and Evolution*. Sept. 2014, pp. 21–30. DOI: 10.1109/ICSME.2014.24.
- [52] Ding Yuan, Soyeon Park, Peng Huang, Yang Liu, Michael M. Lee, Xiaoming Tang, Yuanyuan Zhou, and Stefan Savage. “Be Conservative: Enhancing Failure Diagnosis with Proactive Logging”. In: *Operating Systems Design and Implementation*. OSDI ’12. USENIX. Hollywood, CA, USA: USENIX Association, 2012, pp. 293–306. ISBN: 9781931971966.
- [53] Jieming Zhu, Shilin He, Jinyang Liu, Pinjia He, Qi Xie, Zibin Zheng, and Michael R. Lyu. “Tools and Benchmarks for Automated Log Parsing”. In: *International Conference on Software Engineering: Software Engineering in Practice*. ICSE-SEIP ’19. Montreal, Quebec, Canada: IEEE Press, 2019, pp. 121–130. DOI: 10.1109/ICSE-SEIP.2019.00021.
- [54] Roman Haas, Rainer Niedermayr, Tobias Roehm, and Sven Apel. “Is Static Analysis Able to Identify Unnecessary Source Code?” In: *ACM Trans. Softw. Eng. Methodol.* 29.1 (Jan. 2020). ISSN: 1049-331X. DOI: 10.1145/3368267.
- [55] Jaroslav Fowkes, Pankajan Chanthirasegaran, Razvan Ranca, Miltiadis Allamanis, Mirella Lapata, and Charles Sutton. “Autofolding for Source Code Summarization”. In: *IEEE Transactions on Software Engineering* 43.12 (Dec. 2017), pp. 1095–1109. ISSN: 0098-5589. DOI: 10.1109/TSE.2017.2664836.

- [56] Qiang Fu, Jieming Zhu, Wenlu Hu, Jian-Guang Lou, Rui Ding, Qingwei Lin, Dongmei Zhang, and Tao Xie. “Where Do Developers Log? An Empirical Study on Logging Practices in Industry”. In: *Companion Proceedings of the 36<sup>th</sup> International Conference on Software Engineering*. ICSE Companion 2014. Hyderabad, India: ACM, 2014, pp. 24–33. ISBN: 978-1-4503-2768-8. DOI: 10.1145/2591062.2591175.
- [57] Jieming Zhu, Pinjia He, Qiang Fu, Hongyu Zhang, Michael R. Lyu, and Dongmei Zhang. “Learning to Log: Helping Developers Make Informed Logging Decisions”. In: *International Conference on Software Engineering*. ICSE ’15. Florence, Italy: IEEE Press, 2015, pp. 415–425. ISBN: 978-1-4799-1934-5. URL: <http://dl.acm.org/citation.cfm?id=2818754.2818807>.
- [58] Sangeeta Lal, Neetu Sardana, and Ashish Sureka. “Analysis and Prediction of Log Statement in Open Source Java Projects”. In: *Doctoral Consortium at the International Conference on Open Source Systems*. Buenos Aires, Argentina, May 2017, pp. 65–80.
- [59] Ding Yuan, Jing Zheng, Soyeon Park, Yuanyuan Zhou, and Stefan Savage. “Improving Software Diagnosability via Log Enhancement”. In: *ACM Trans. Comput. Syst.* 30.1 (Feb. 2012), 4:1–4:28. ISSN: 0734-2071. DOI: 10.1145/2110356.2110360.
- [60] Zhenhao Li, Tse-Hsun Chen, Jinqiu Yang, and Weiyi Shang. “DLFinder: Characterizing and Detecting Duplicate Logging Code Smells”. In: *International Conference on Software Engineering*. ICSE ’19. Montreal, Quebec, Canada: IEEE Press, 2019, pp. 152–163. DOI: 10.1109/ICSE.2019.00032.
- [61] Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael I. Jordan. “Detecting Large-scale System Problems by Mining Console Logs”. In: *Symposium on Operating Systems Principles*. SOSP ’09. SIGOPS. Big Sky, Montana, USA: ACM, 2009, pp. 117–132. ISBN: 978-1-60558-752-3. DOI: 10.1145/1629575.1629587.
- [62] Raffi Khatchadourian, Awais Rashid, Hidehiko Masuhara, and Takuya Watanabe. “Detecting Broken Pointcuts Using Structural Commonality and Degree of Interest”. In: *Sci. Comput. Program.* 150.C (Dec. 2017), pp. 56–74. ISSN: 0167-6423. DOI: 10.1016/j.scico.2017.06.011.
- [63] Huzefa Kagdi, Michael L. Collard, and Jonathan I. Maletic. “A Survey and Taxonomy of Approaches for Mining Software Repositories in the Context of Software Evolution”. In: *J. Softw. Maint. Evol.* 19.2 (Mar. 2007), pp. 77–131. ISSN: 1532-060X. DOI: 10.1002/smr.344.
- [64] Daniel Ratiu, Stéphane Ducasse, Tudor Gundefinedrba, and Radu Marinescu. “Using History Information to Improve Design Flaws Detection”. In: *European Conference on Software Maintenance and Reengineering*. CSMR ’04. IEEE, 2004, pp. 223–232. ISBN: 076952107X. DOI: 10.1109/csmr.2004.1281423.

- [65] Raffi Khatchadourian. “Automated refactoring of legacy Java software to enumerated types”. In: *Automated Software Engineering* 24.4 (Dec. 2017), pp. 757–787. ISSN: 0928-8910. DOI: 10.1007/s10515-016-0208-8.
- [66] Raffi Khatchadourian and Hidehiko Masuhara. “Automated Refactoring of Legacy Java Software to Default Methods”. In: *International Conference on Software Engineering*. ICSE '17. Buenos Aires, Argentina: IEEE Press, 2017, pp. 82–93. ISBN: 978-1-5386-3868-2. DOI: 10.1109/ICSE.2017.16.
- [67] Pamela Bhattacharya, Marios Iliofotou, Iulian Neamtiu, and Michalis Faloutsos. “Graph-based Analysis and Prediction for Software Evolution”. In: *International Conference on Software Engineering*. ICSE '12. Zurich, Switzerland: IEEE Press, 2012, pp. 419–429. ISBN: 978-1-4673-1067-3. URL: <http://dl.acm.org/citation.cfm?id=2337223.2337273>.