Characterizing and Detecting Duplicate Logging Code Smells

Zhenhao Li

A Thesis
in
The Department
of
Computer Science and Software Engineering

Presented in Partial Fulfillment of the Requirements
For the Degree of
Master of Applied Science (Software Engineering) at
Concordia University
Montréal, Québec, Canada

August 2019

CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the thesis prepared

By: **Zhenhao Li**

Entitled: **Characterizing and Detecting Duplicate Logging Code Smells**

and submitted in partial fulfillment of the requirements for the degree of

Master of Applied Science (Software Engineering)

complies with the regulations of this University and meets the accepted standards with respect to originality and quality.

Signed by the final examining commitee:

______________________________ Chair
Dr. Nikolaos Tsantalis

______________________________ Examiner
Dr. Olga Ormandjieva

______________________________ Examiner
Dr. Juergen Rilling

______________________________ Supervisor
Dr. Tse-Hsun Chen, Dr. Weiyi Shang

Approved by ______________________________
Dr. Leila Kosseim, Graduate Program Director

1 August 2019 ______________________________
Dr Amir Asif, Dean
Faculty of Engineering and Computer Science

Abstract

Characterizing and Detecting Duplicate Logging Code Smells

Zhenhao Li

Developers rely on software logs for a wide variety of tasks, such as debugging, testing, program comprehension, verification, and performance analysis. Despite the importance of logs, prior studies show that there is no industrial standard on how to write logging statements. Recent research on logs often only considers the appropriateness of a log as an individual item (e.g., one single logging statement); while logs are typically analyzed in tandem. In this thesis, we focus on studying duplicate logging statements, which are logging statements that have the same static text message. Such duplications in the text message are potential indications of logging code smells, which may affect developers' understanding of the dynamic view of the system. We manually studied over 3K duplicate logging statements and their surrounding code in four large-scale open source systems: Hadoop, CloudStack, ElasticSearch, and Cassandra. We uncovered five patterns of duplicate logging code smells. For each instance of the code smell, we further manually identify the problematic (i.e., require fixes) and justifiable (i.e., do not require fixes) cases. Then, we contact developers in order to verify our manual study result. We integrated our manual study result and developers' feedback into our automated static analysis tool, DLFinder, which automatically detects problematic duplicate logging code smells. We evaluated DLFinder on the four manually studied systems and four additional systems: Kafka, Flink, Camel and Wicket. In total, combining the results of DLFinder and our manual analysis, we reported 91 problematic code smell instances to developers and all of them have been fixed. This thesis provides an initial step on creating a logging guideline for developers to improve the quality of logging code. DLFinder is also able to detect duplicate logging code smells with high precision and recall.

Statement of Originality

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners. I understand that my thesis may be made electronically available to the public.

Acknowledgement

Foremost, I would like to express my greatest gratitude to my supervisors Dr. Tse-Hsun Chen and Dr. Weiyi Shang for your patient guidance and encouragement on my research and life. Without your supervision and invaluable support, nothing of this would have been possible.

Apart from my supervisors, I would like to sincerely thank my thesis examiners, Dr. Ormandjieva and Dr. Rilling, for their extremely valuable and constructive suggestions. Furthermore, I appreciate Dr. Jinqiu Yang for her valuable contribution and guidance in my research.

I am very lucky to have lively communications and fruitful discussions with all the members of SPEAR and SENSE. I learned so much from all of you, it is my honor and pleasure to work with you all.

Last but not least, I would like to express my special thanks to my parents. Words can hardly express my gratitude and feelings towards you. Your unconditional support sustains me thus far and keeps me going.

Contribution of Authors

DLFinder: Characterizing and Detecting Duplicate Logging Code Smells

Zhenhao Li: quantitative analysis, manual analysis, program implementation, writing, editing and proofing

Tse-Hsun (Peter) Chen: research supervisor, funding, experimental guidance, manual analysis, writing, editing and proofing

Jinqiu Yang: writing, editing and proofing

Weiyi Shang: research supervisor, funding, writing, editing and proofing

A Comprehensive Study on Duplicate Logging Statements

Zhenhao Li: quantitative analysis, manual analysis, program implementation, writing, editing and proofing

Tse-Hsun (Peter) Chen: research supervisor, funding, experimental guidance, manual analysis, writing, editing and proofing

Jinqiu Yang: writing, editing and proofing

Weiyi Shang: research supervisor, funding, writing, editing and proofing

Characterizing and Detecting Duplicate Logging Code Smells

Zhenhao Li: data analysis, program implementation, writing and proofing

Related publication

Zhenhao Li, Tse-Hsun (Peter) Chen, Jinqiu Yang, and Weiyi Shang, “DLFinder: Characterizing and Detecting Duplicate Logging Code Smells”, in Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019, pp. 152–163.

Zhenhao Li, Tse-Hsun (Peter) Chen, Jinqiu Yang, and Weiyi Shang, “A Comprehensive Study on Duplicate Logging Statements”, IEEE Transactions on Software Engineering, TSE, *under review.*

Zhenhao Li, "Characterizing and Detecting Duplicate Logging Code Smells", in Proceedings of the 41st International Conference on Software Engineering: Companion Proceedings, ICSE-SRC 2019, Montreal, QC, Canada, May 25-31, 2019, pp. 147–149.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Software logs are widely used in software systems to record system execution behaviors. Developers use the generated logs to assist in various tasks, such as debugging [73, 71, 20], testing [15, 28, 13], program comprehension [23, 59], system verification [9, 6], and performance analysis [14, 69]. A logging statement (i.e., code that generates a log) contains a static message, to-be-recorded variables, and log verbosity level. As an example, a logging statement may be written as *logger.error(“Interrupted while waiting for fencing command: ” + cmd);*. In this example, the static text message is *“Interrupted while waiting for fencing command:”*, and the dynamic message is from the variable *cmd*, which records the command that is being executed. The logging statement is at the *error* level, which is the level for recording failed operations [2].

Even though developers have been analyzing logs for decades [31], there exists no industrial standard on how to write logging statements [20, 49]. Prior studies often focus on recommending where logging statements should be added into the code (i.e., *where-to-log*) [76, 75], and what information should be added in logging statements (i.e., *what-to-log*) [59, 73, 51]. A few recent studies [12, 24] aim to detect potential problems in logging statements. However, these studies often only consider the appropriateness of one single logging statement as an individual item; while logs are typically analyzed in tandem [73, 14]. In other words, we consider that the appropriateness of a log is also influenced by other logs that are generated in system execution.

In particular, an intuitive case of such influence is duplicate logs, i.e., multiple logs that have the same text message. Even though each log itself may be impeccable, duplicate logs may affect developers' understanding of the dynamic view of the system. For example, as shown in Figure 1, there are two logging statements in two different *catch* blocks, which are associated with the same *try* block. These two logging statements have the same static text message and do not include any other error-diagnostic information. Thus, developers cannot easily distinguish what is the occurred

```
...
} catch (AlreadyClosedException closedException) {
      s_logger.warn("Connection to AMQP service is
          lost.");
} catch (ConnectException connectException) {
      s_logger.warn("Connection to AMQP service is
          lost.");
}
...
```

Figure 1: An example of duplicate logging code smell that we detected in CloudStack. The duplicate logging statements in the two *catch* blocks contain insufficient information (e.g., no exception type or stack trace) to distinguish what may be the occurred exception.

exception when analyzing the produced logs. Since developers rely on logs for debugging and program comprehension [59], such duplicate logging statements might be confusing and negatively affect developers' activities in maintenance and quality assurance.

Hence, in this thesis, we focus on studying duplicate logging statements in the source code to help developers mitigate the impact as discussed in the example above and further improve logging practices. We conducted a manual study on four large-scale open source systems, namely Hadoop, CloudStack, ElasticSearch, and Cassandra. We first used static analysis to identify all duplicate logging statements, which are defined as two or more logging statements that have the same static text message. We then manually study all the (over 3K) identified duplicate logging statements and uncovered five patterns of *duplicate logging code smells*. We follow prior code smell studies [8, 19], and consider duplicate logging code smell as a "surface indication that usually corresponds to a deeper problem in the system". However, *not* all of the duplicate logging code smell are problematic and require fixes (i.e., *problematic duplicate logging code smells*). In particular, context (e.g., surrounding code and usage scenario of logging) may play an important role in identifying fixing opportunities. Hence, we further categorized duplicate logging code smells into *problematic* or *justifiable* cases. In addition to our manual analysis, we sought confirmation from developers on the manual analysis result: For the problematic duplicate logging code smells, we reported them to developers for fixing. For the justifiable ones, we communicated with developers for discussion (e.g., emails or posts on developers' forums).

We implemented a static analysis tool, DLFinder, to automatically detect *problematic* duplicate logging code smells. DLFinder leverages the findings from our manual study, including the uncovered

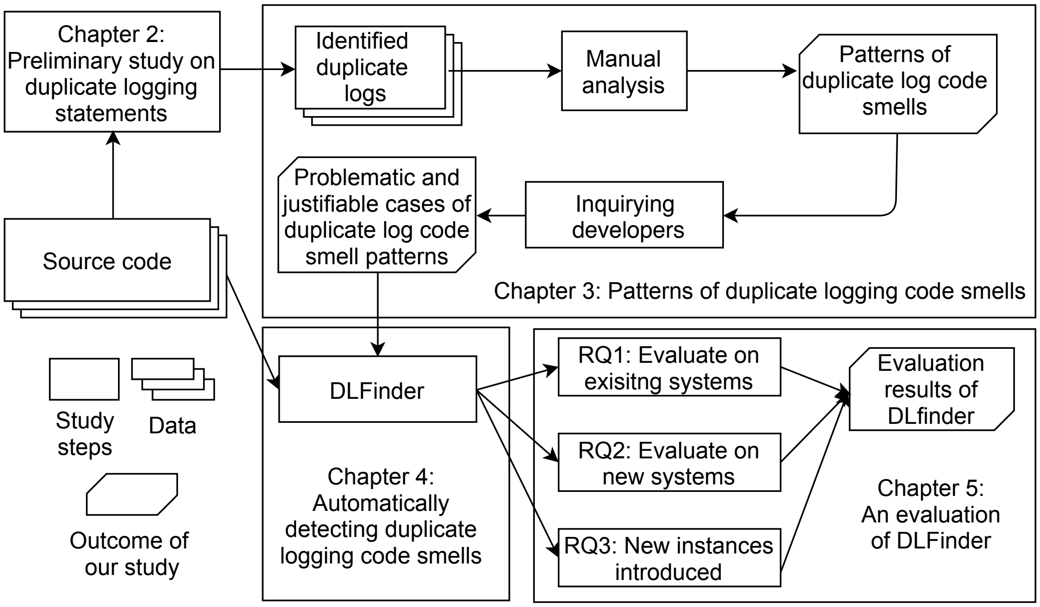
Chapter 2: Preliminary study on duplicate logging statements
Identified duplicate logs
Manual analysis
Patterns of duplicate log code smells
Problematic and justifiable cases of duplicate log code smell patterns
Inquirying developers
Source code
Chapter 3: Patterns of duplicate logging code smells
DLFinder
RQ1: Evaluate on exisitng systems
RQ2: Evaluate on new systems
RQ3: New instances introduced
Evaluation results of DLfinder
Chapter 4: Automatically detecting duplicate logging code smells
Chapter 5: An evaluation of DLFinder
Study steps
Data
Outcome of our study

Figure 2: The overall process of the thesis. The term "duplicate logging statements" is referred as "duplicate logs" for simplification.

patterns of duplicate logging code smells and the categorization on problematic and justifiable cases. We evaluated DLFinder on eight systems: four are from the manual study and four are new systems (Kafka, Flink, Camel and Wicket). We also applied DLFinder on the updated versions of the four manually studied systems. The evaluation shows that the uncovered patterns of the duplicate logging code smells also exist in the additional systems, and duplicate logging code smells may be introduced over time. An automated approach such as DLFinder can help developers avoid duplicate logging code smells as systems evolve. In total, we reported 91 instances of duplicate logging code smell to developers and all the reported instances are fixed.

Figure 2 shows the overall process of the thesis.

In summary, this thesis makes **the following contributions:**

- We uncovered five patterns of duplicate logging code smells through an extensive manual study on over 3K duplicate logging statements.
- We presented a categorization of duplicate logging code smells (i.e., problematic or justifiable), based on both our manual assessment (i.e., studying the logging statement and its surrounding code) and developers' feedback.
- We proposed DLFinder, a static analysis tool that integrates our manual study result and developers' feedback to detect problematic duplicate logging code smells. We evaluated DLFinder for both the accuracy and generalization (i.e., on new systems and on the newer versions as systems evolve).
- We reported 91 instances of problematic duplicate logging code smells to developers (DLFinder is able to detect 81 of them), and all of the reported instances are fixed.

Our study provides an initial step on creating a logging guideline for developers to improve the quality of logging code. DLFinder is also able to detect duplicate logging code smells with high precision and recall.

Thesis organization. The rest of this thesis is organized as follows: Chapter 2 surveys related work and describes the preliminary study on duplicate logging statements including how we prepare the data for manual study and the studied systems. Chapter 3 discusses the process and the results of our manual study, and also developers' feedback on our results. Chapter 4 discusses the implementation details of DLFinder. Chapter 5 evaluates DLFinder for both the accuracy and generalization and the threats to validity of our study. Finally, Chapter 6 concludes the thesis.

Chapter 2

Related Works and Preliminary Study on Duplicate Logging Statements

In this chapter, we survey the related works of this thesis and discuss our preliminary study on duplicate logging statements.

2.1 Related Works

We discuss three areas of related research: studies on logging practices, improving logging practices, and code smells.

2.1.1 Studies on logging practices

Previous studies show that software logs are extensively used and analyzed for various tasks, such as error diagnosis [73, 71], deployment verification [57], load testing [15, 28], understanding code quality [58], security monitoring[44], program comprehension [23, 59], and performance analysis[14, 69, 45]. Comparing to their studies which provide logging understanding or solve particular problems by analyzing logs, we aim to improve the the logging practice. Thus our work can potentially benefit those various tasks related to logging.

There are several studies on characterizing the logging practices in software systems [72, 11, 20]. Yuan et al. [72] conducted a quantitative characteristics study on log messages for large-scale open source C/C++ systems. Chen et al. [11] replicated the study by Yuan et al. [72] on Java open-source projects. Both of their studies found that log message is crucial for system understanding and maintenance. Fu et al. [20] studied where developers in Microsoft add logging statements in the code and summarized several typical logging strategies. They found that developers often add logs to check the returned value of a method. Prior studies focus on studying where do people add logging code and how often do developers modify logging code. However, these studies do not analyze the quality of the log lines. Different from prior studies, in this thesis, we focus on manually understanding duplicate logging code smells. We also discuss potential approaches to detect and fix these code smells based on different contexts (i.e., surrounding code).

2.1.2 Improving logging practices

Zhao et al. [75] proposed a tool that determines how to optimally place logging statements given a performance overhead threshold. Zhu et al. [76] provided a tool for suggesting log placement using machine learning techniques. Yuan et al. [73] proposed an approach that can automatically insert additional variables into logging statements to enhance the error diagnostic information. Chen et al. [12] concluded five categories of logging anti-patterns from code changes, and implemented a tool to detect the anti-patterns. Hassani et al. [24] identified seven root-causes of the log-related issues from log-related bug reports.

Compared to prior studies, we study logging code smells that may be caused by duplicate logs, with a goal to help developers improve logging code. The logging problems that we uncovered in this study are not discovered by prior work. We conducted an extensive manual study through obtaining a deep understanding on not only the logging statements but also the surrounding code, whereas prior studies usually only look at the problems that are related to the logging statement itself. The context of the logging code (i.e., the semantic and structure of the surrounding code) affects the potential impact of logging code smells. Hence, code context of logging statements should also be considered to precisely uncover the potential problems in log. Compared to prior studies [12, 24] that focus on logging statement itself, we also analyze the surrounding code when we are detecting patterns of duplicate logging code smells, such as IC and IE.

2.1.3 Code smells

Code smells can be indications of bad design and implementation choices, which may affect software systems' maintainability [63, 5, 61, 41], understandability [10, 4], and performance [67]. To mitigate the impact of code smells, studies have been proposed to detect code smells [47, 46, 48, 56, 25].

Duplicate code (or code clones) is a kind of code smells which may be caused by developers copying and pasting a piece of code from one place to another [53, 74]. Such code clones may indicate quality problems. There are many studies that focus on studying the impact of code clones [30, 33, 22], and detecting them [32, 39, 55].

In this thesis, we study duplicate logging code smells, which are not studied in prior duplicate code studies. A number of studies also investigate the prevalence and characteristics of micro-clones (i.e., code clones of 4 or fewer lines of code) [43, 26, 27]. Some instances of the problematic duplicate logging code smells in our study might also be micro-clones (e.g., IC, IE, and LM), however, the impact of micro-clones on logging practice is not considered in these works. Our study might provide insights for future studies on the relationship between micro-clones and logging practice, and the detection of duplicate logging code smells might also help with identifying micro-clones and further alleviate the impact of micro-clones on maintenance.

2.2 Preliminary study on duplicate logging statements

We first introduce our studied systems and then discuss how we define and identify duplicate logging statements. Finally, we conduct a preliminary quantitative study on the prevalence of duplicate logging statements.

2.2.1 Definition and how to identify duplicate logging statements

We define duplicate logging statements as logging statements that have identical static text messages. We focus on studying the log message because such semantic information is crucial for log understanding and system maintenance [59, 72]. As an example, the two following logging statements are considered duplicate: "*Unable to create a new ApplicationId in SubCluster*" + ***subClusterId.getId()***, and "*Unable to create a new ApplicationId in SubCluster*" + ***id***.

To prepare for a manual study, we identify duplicate logging statements by analyzing the source code using static analysis. In particular, the static text message of each logging statement is built by concatenating all the strings (i.e., constants and values of string variables) and abstractions of the non-string variables. We also extract information to support the manual analysis, such as the types of variables that are logged, and the log level (i.e., *fatal*, *error*, *warn*, *info*, *debug*, or *trace*). Log levels represent the verbosity level of the log and can be used to reduce logging overheads in production (e.g., only logging *info* level or above) [37, 72]. If two or more logging statements have the same static text message, they are identified as duplicate logging statements. We exclude logging statements with only one word in the static text message since those logging statements usually do not contain much static information, and are usually used to record the value of a dynamic variable

Table 1: An overview of the studied systems. The top four systems are used for manual study in Chapter 3. The bottom four systems are used for evaluating our code smell detection tool in Chapter 5.

System	Version	Release date	LOC	Num. of logs	Num. of dup. logs	Num. of dup. log sets	Med. words in dup. logs	Med. words in non-dup. logs
Cassandra	3.11.1	Oct. 2017	358K	1.6K	113 (7%)	46	7	7
CloudStack	4.9.3	Aug. 2017	1.18M	11.7K	2.3K (20%)	865	8	8
ElasticSearch	6.0.0	Nov. 2017	2.12M	1.7K	94 (6%)	40	6	7
Hadoop	3.0.0	Nov. 2017	2.69M	5.3K	496 (9%)	217	6	6
Camel	2.21.1	Apr. 2018	1.68M	7.3K	2.3K (32%)	886	6	6
Flink	1.7.1	Dec. 2018	177K	2.5K	1.4K (56%)	203	6	7
Kafka	2.1.0	Nov. 2018	542K	1.5K	406 (27%)	104	5	7
Wicket	8.0.0	May. 2018	381K	0.4K	45 (11%)	21	6	8

during system execution.

2.2.2 Studied Systems

Table 1 shows the statistics of the studied systems. We identify duplicate logging statements from the top four large-scale open source Java systems in the table for our manual analysis: Hadoop, CloudStack, ElasticSearch, and Cassandra, which are commonly used in prior studies for log-related research [36, 12, 24]. The studied systems also use popular Java logging libraries (e.g., Log4j [2] and SLF4J [3]). Hadoop is a distributed computing framework, which is composed of four subsystems: Hadoop Common, Hadoop Distributed File System, YARN, and MapReduce. CloudStack is a cloud computing platform, ElasticSearch is a distributed search engine, and Cassandra is a NoSQL database system. These systems belong to different domains and are well maintained. In our study, we study all Java source code files in the main branch of each system and exclude test files, since we are more interested in studying duplicate logging statements that may affect log understanding in production.

2.2.3 Prevalence of duplicate logging statements

Before we conduct a more detailed manual analysis on duplicate logging statements, we first study how prevalent these duplicate logging statements are in the studied systems. Table 1 shows that there is a non-negligible number of duplicate logs. In particular, for between 6% to 20% of all the logging statements (94 to 2.3K duplicate logs) in each studied system, there exists at least another logging statement that is a duplicate. We group the duplicate logging statements into duplicate log sets, each set containing logging statements with the same message (see Table 1). In general, there are 40–865 log sets in the studied systems, with an average of 2.29 to 2.66 duplicate logs per set.

Intuitively, a short log is more likely to have a duplicate. For example, short logs like *"Exception Occurred: + e"* may exist in different places in the code. To study whether the duplicate logs are merely because they are such short logs, we compare the number of words in duplicate and non-duplicate logs (see Table 1). We find that the median number of words in duplicate logs is similar to that of non-duplicate logs (both range from 6 to 8 words). We further conduct a Mann-Whitney U test to determine whether the difference is statistically significant, and compute the Cliff's delta effect size to quantify the difference [7, 18]. We choose the Mann-Whitney U test and Cliff's delta because they do not have any assumptions of the sample population. We find that the number of words in non-duplicate logs is larger than that of duplicate logs (with a p-value < 0.01 for all studied systems). However, such differences only have a *negligible* [18] effect in three studied systems (Cliff's delta is 0.08, 0.045, 0.146 for Hadoop, Cassandra, and CloudStack, respectively), and a *small* [18] effect in one studied system (Cliff's delta is 0.27 for ElasticSearch).

There is a non-negligible number of duplicate logging statements in the studied systems (6% to 20% of all the logs). Duplicate logging statements have a similar level of semantics (i.e., number of words) compared to non-duplicate logging statements.

Chapter 3

Patterns for Duplicate Logging Code Smells

In this chapter, we conduct a manual study to uncover patterns of potential code smells that may be associated with duplicate logging statements (i.e., *duplicate logging code smells*). Similar to prior code smell studies, we consider duplicate logging code smells as a *"surface indication that usually corresponds to a deeper problem in the system"* [8, 19]. Such duplicate logging code smells may be indications of logging problems that require fixes.

Furthermore, we categorize each code smell instance as either problematic (i.e., require fixes) or justifiable (i.e., do not require fixes), by understanding the surrounding code. Not every duplicate logging code smell is problematic. Intuitively, one needs to consider the code context to decide whether a code smell instance is problematic and requires fixes. As shown in prior studies [76, 20, 36], logging decisions, such as log messages and log levels, are often associated with the structure and semantics of the surrounding code. In addition to the manual analysis, we also ask for developers' feedback regarding both the problematic and justifiable cases. By providing a more detailed understanding of code smells, we may better assist developers to improve logging practices and inspire future research.

3.1 Manual study process

We conduct a manual study by analyzing all the duplicate logging statements identified from the studied systems. In total, we studied 1,168 sets of duplicate logging statements in the four studied systems (more than 3K logging statements in total; each set contains two or more logging statements with the same static message).

The process of our manual study involves five phases:

- Phase I: We manually studied 289 randomly sampled (based on 95% confidence level and 5% confidence interval [7]) sets of duplicate logging statements and the surrounding code to derive an initial list of duplicate logging code smell patterns. All disagreements were discussed until a consensus was reached.
- Phase II: We *independently* categorized *all* of the 1,168 sets of duplicate logging statements to the derived patterns in Phase I. We did not find any new patterns in this phase. The results of this phase have a Cohen's kappa of 0.806, which is a substantial-level of agreement [42].
- Phase III: We discussed the categorization results obtained in Phase II. All disagreements were discussed until a consensus was reached.
- Phase IV: We further studied all logging code smell instances that belong to each pattern in order to identify justifiable cases of the logging code smell that may not need fixes. The instances that do not belong to the category of justifiable are considered potentially problematic and may require fixes.
- Phase V: We verified both the problematic instances of logging code smells and the justifiable ones with developers by creating issue reports and pull requests, sending emails, or posting our findings on developers' forums such as Stack Overflow. In particular, we reported every instance that we believe to be problematic (i.e., require fixes). We also reported a number of instances for each justifiable category.

3.2 Patterns of duplicate logging code smells

In total, we uncovered five patterns of duplicate logging code smells. Table 2 lists the uncovered code smell patterns and the corresponding examples. Table 3 shows the number of problematic code smell instances for each pattern. Below, we discuss each pattern according to the following template:
Description. A description of the pattern of duplicate logging code smell.
Example. An example of the pattern.
Code smell instances. Discussions on the manually-uncovered code smell instances. We also discuss the justifiable cases if we found any.
Developers feedback. A summary of developers' feedback on both the problematic and justifiable cases.

Table 2: Patterns of duplicate logging code smells and corresponding examples.

Name	Example

Inadequate information in catch blocks (IC)	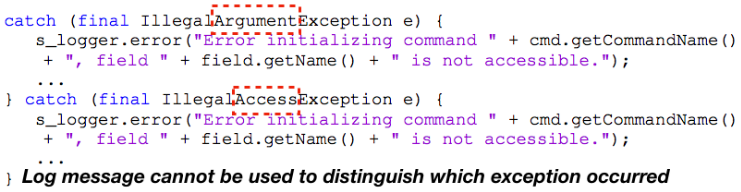
	catch (final IllegalArgumentException e) {
	s_logger.error("Error initializing command " + cmd.getCommandName()
	+ ", field " + field.getName() + " is not accessible.");
	...
	} catch (final IllegalAccessException e) {
	s_logger.error("Error initializing command " + cmd.getCommandName()
	+ ", field " + field.getName() + " is not accessible.");
	...
	} ***Log message cannot be used to distinguish which exception occurred***

Inconsistent error-diagnostic information (IE)	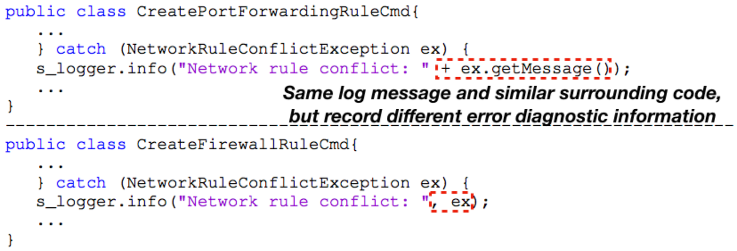
	public class CreatePortForwardingRuleCmd{
	...
	} catch (NetworkRuleConflictException ex) {
	s_logger.info("Network rule conflict: " + ex.getMessage());
	...
	}
	Same log message and similar surrounding code, but record different error diagnostic information
	public class CreateFirewallRuleCmd{
	...
	} catch (NetworkRuleConflictException ex) {
	s_logger.info("Network rule conflict: ", ex);
	...
	}

Log message mismatch (LM)	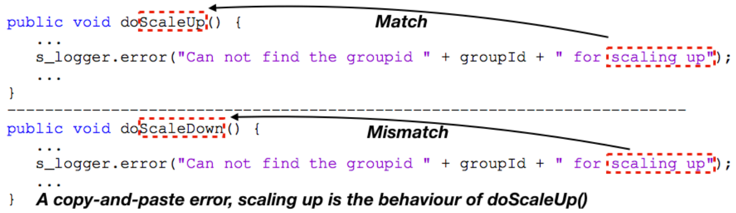
	public void doScaleUp() {
	Match
	...
	s_logger.error("Can not find the groupid " + groupId + " for scaling up");
	...
	}
	public void doScaleDown() {
	Mismatch
	...
	s_logger.error("Can not find the groupid " + groupId + " for scaling up");
	...
	} ***A copy-and-paste error, scaling up is the behaviour of doScaleUp()***

Inconsistent log level (IL)	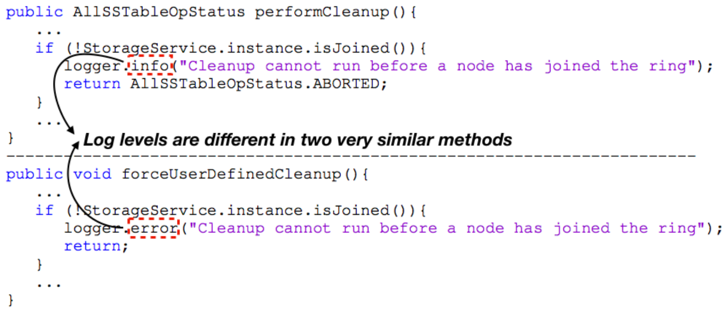
	public AllSSTableOpStatus performCleanup(){
	...
	if (!StorageService.instance.isJoined()){
	logger.info("Cleanup cannot run before a node has joined the ring");
	return AllSSTableOpStatus.ABORTED;
	}
	...
	}
	Log levels are different in two very similar methods
	public void forceUserDefinedCleanup(){
	...
	if (!StorageService.instance.isJoined()){
	logger.error("Cleanup cannot run before a node has joined the ring");
	return;
	}
	...
	}

Duplicate log in polymorphism (DP)	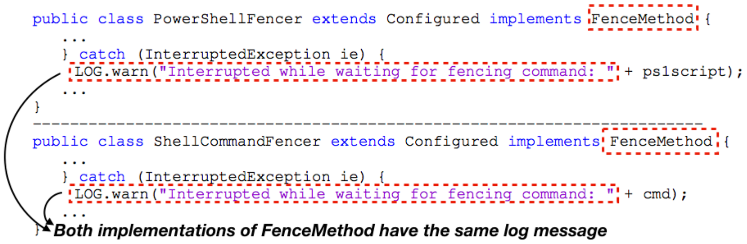
	public class PowerShellFencer extends Configured implements FenceMethod{
	...
	} catch (InterruptedException ie) {
	LOG.warn("Interrupted while waiting for fencing command: " + ps1script);
	...
	}
	public class ShellCommandFencer extends Configured implements FenceMethod{
	...
	} catch (InterruptedException ie) {
	LOG.warn("Interrupted while waiting for fencing command: " + cmd);
	...
	} ***Both implementations of FenceMethod have the same log message***

Table 3: Number of problematic instances (Prob.) verified by our manual study and developers' feedback, and total number of instances (Total) including non-problematic instances.

	IC		**IE**		**LM**		**IL**		**DP**	
	Prob.	Total	Prob.	Total	Prob.	Total	Prob.	Total	Prob.	Total
Cassandra	1	1	0	1	0	0	0	3	2	2
CloudStack	8	8	4	14	27	27	0	47	107	107
ElasticSearch	1	1	0	5	1	1	0	9	3	3
Hadoop	5	5	0	0	9	9	0	17	27	27
Total	15	15	4	20	37	37	0	76	139[1]	139

[1] Developers acknowledged the problem but we did not report all the instances, because systematic refactoring of DP would require supports from logging libraries.

3.2.1 Pattern 1: Inadequate information in catch blocks (IC)

Description. Developers usually rely on logs for error diagnostics when exceptions occur [70]. However, we find that sometimes, duplicate logging statements in different *catch* blocks of the same *try* block may cause debugging difficulties since the logs fail to tell which exception occurred.

Example. As shown in Table 2, in the `ParamProcessWorker` class in CloudStack, the *try* block contains two *catch* blocks; however, the log messages in these two *catch* blocks are identical. Since both the exception message and stack trace are not logged, once one of the two exceptions occurs, developers may encounter difficulties in finding the root causes and determining the occurred exception.

Code smell instances. After examining all the instances of IC, we find that all of them are potentially problematic and require fixes. For all the instances of IC, none of the exception type, exception message, and stack trace are logged.

Developers' feedback. We reported all the problematic instances of IC (15 instances) by using pull requests. All the pull requests were accepted by developers and the fixes were integrated to the studied systems. Developers agree that IC will cause confusion and insufficient information in the logs, which may increase the difficulties of error diagnostics.

3.2.2 Pattern 2: Inconsistent error-diagnostic information (IE)

Description. We find that sometimes duplicate logging statements for recording exceptions may contain inconsistent error-diagnostic information (e.g., one logging statement records the stack trace and the other does not), even though the surrounding code is similar.

Example. As shown in Table 2, the two classes `CreatePortForwardingRuleCmd` and `CreateFirewallRuleCmd` in CloudStack have similar functionalities. The two logging statements have the same static text message and are in methods with identical names (i.e., *create()*, not shown due to space restriction). The *create()* method in `CreatePortForwardingRuleCmd` is about creating rules for port forwarding, and the method in `CreateFirewallRuleCmd` is about creating rules for firewalls. These two methods have very similar code structure and business logic. However, the two logging statements record different information: One records the stack trace information and the other one only records the exception message (i.e., *ex.getMessage()*). Since the two logging statements have similar context, the error-diagnostic information recorded by the logs may need to be consistent for the ease of debugging. We reported this example, which is now fixed to have consistent error-diagnostic information.

Code smell instances. As shown in Table 3, we find 20 instances of IE, and four of them are considered problematic based on our understanding. From the remaining instances of IE, we find three justifiable cases that may not require fixes.

Justifiable case IE.1: Duplicate logging statements record general and specific exceptions. For 11/20 instances of IE, we find that the duplicate logging statements are in the *catch* blocks of different types of exception. In particular, one duplicate logging statement is in the *catch* block of a generic exception (i.e., the `Exception` class in Java) and the other one is in the *catch* block of a more specific exception (e.g., application-specific exceptions such as `CloudRuntimeException`). In all of the 11 cases, we find that one log would record the stack trace for `Exception`, and the duplicate log would only record the type of the occurred exception (e.g., by calling *e.getMessage()*) for a more specific exception. The rationale may be that generic exceptions, once occurred, are often not expected by developers [70], so it is important that developers record more error-diagnostic information.

Justifiable case IE.2: Duplicate logging statements are in the same catch block for debugging purposes. For 3/20 instances of IE, we find that the duplicate logging statements are *in the same catch* block and developers' intention is to use a duplicate logging statement at debug level to record rich error-diagnostic information such as stack trace (and the log level of the other logging statement could be error). The extra logging statements at debug level help developers debug the occurred exception and reduce logging overhead in production [37] (i.e., logging statements at debug level are turned off).

Justifiable case IE.3: Having separate error-handling classes. For 2/20 instances, we find that the error-diagnostic information is handled by creating an object of an error-handling class. As an example from CloudStack:

```
public final class LibvirtCreateCommandWrapper {
    ...
        } catch (final CloudRuntimeException e) {
            s_logger.debug("Failed to create volume: "
                + e.toString());
            return new
                CreateAnswerErrorHandler(command, e);
        }
    ...
}

public class KVMStorageProcessor {
    ...
        } catch (final CloudRuntimeException e) {
            s_logger.debug("Failed to create volume: ",
                e);
            return new
                CopyCmdAnswerErrorHandler(e.toString());
        }
    ...

}
```

In this example, extra logging is added by using error-handling classes (i.e., `CreateAnswerErrorHandler` and `CopyCmdAnswerErrorHandler`) to complement the logging statements. As a consequence, the *actual* logged information is consistent in these two methods: One method records *e.toString()* in the logging statement and records the exception variable *e* through an error-handling class; the other method records *e* in the logging statement and records *e.toString()* through an error-handling class.

Developers' feedback. We reported all the instances of IE (four in total) that we consider problematic to developers as pull requests, all of which are accepted by developers. Moreover, we ask developers whether our conjecture was correct for each of the justifiable cases of IE. We received positive feedback that confirms our manual analysis on the justifiable cases.

3.2.3 Pattern 3: Log message mismatch (LM)

Description. We find that sometimes after developers copy and paste a piece of code to another method or class, they may forget to change the log message. Hence, this results in having duplicate logging statements that record inaccurate system behaviors.

Example. As an example, in Table 2, the method *doScaleDown()* is a code clone of *doScaleUp()* with very similar code structure and minor syntactical differences. However, developers forgot to change the log message in *doScaleDown()*, after the code was copied from *doScaleUp()* (i.e., both log messages contain *scaling up*). Such instances of LM may cause confusion when developers analyze the logs.

Code smell instances. We find that there are 37 instances of LM that are caused by copying-and-pasting the logging statement to new locations without proper modifications. For all the 37 instances, the log message contains words of incorrect class or method name that may cause confusion when analyzing logs.

Developers' feedback. Developers agree that the log messages in LM should be changed in order to correctly record the execution behavior (i.e., update the copy-and-pasted log message to contain the correct class/method name). We reported all the 37 instances of LM that we found through pull requests, and all of the reported instances are now fixed.

3.2.4 Pattern 4: Inconsistent log level (IL)

Description. Log levels (e.g., *fatal*, *error*, *info*, *debug*, or *trace*) allow developers to specify the verbosity of the log message and to reduce logging overhead when needed (e.g., *debug* is usually disabled in production) [37]. A prior study [72] shows that log level is frequently modified by developers in order to find the most adequate level. We find that there are duplicate logging statements that, even though the log messages are exactly the same, the log levels are different.

Example. In the IL example shown in Table 2, the two methods, which are from the same class `CompactionManager`, have very similar functionality (i.e., both try to perform cleanup after compaction), but we find that the log levels are different in these two methods.

Code smell instances. We find three justifiable cases in IL that may be developers' intended behavior. We do not find problematic instances of IL after communicating with developers – Developers think the problematic instances identified by our manual analysis may not be problems.

Justifiable case IL.1: Duplicate logging statements are in the catch blocks of different types of exception. Similar to what we observed in IE, we find that for 8/76 instances, the log level for a more generic exception is usually more severe (e.g., *error* level for the generic Java `Exception` and *info* level for an application-specific exception). Generic exceptions may be more unexpected to developers [70], so developers may use a log level of higher verbosity (e.g., *error* level) to record exception messages.

Justifiable case IL.2: Duplicate logging statements are in different branches of the same method. There are 35/76 instances belong to this case. Below is an example from ElasticSearch, where a set of duplicate logging statements may occur in the same method but in different branches.

```
if (lifecycle.stoppedOrClosed()) {
    logger.trace("failed to send ping transport
        message", e);
} else {
    logger.warn("failed to send ping transport
        message", e);
}
```

In this case, developers already know the desired log level and intend to use different log levels due to the difference in execution (i.e., in the if-else block).

Justifiable case IL.3: *Duplicate logging statements are followed by error-handling code.* There are 18/76 instances that are observed to have such characteristics: In a set of duplicate logging statements, some statements have log levels of higher verbosity, and others have log levels of lower verbosity. However, the duplicate logging statement with lower verbosity log level is followed by additional error handling code (e.g., *throw a new Exception(e);*). Therefore, the error is handled elsewhere (i.e., the exception is re-thrown), and may be recorded at a higher-verbosity log level.

Developers' feedback. In all the instances of IL that we found, developers think that IL may not be a problem. In particular, developers agreed with our analysis on the justifiable cases. However, developers think the problematic instances of IL from our manual analysis may also not be problems. We concluded the following two types of feedback from developers on the "suspect" instances of IL (i.e., 15 problematic ones from our manual analysis out of the 76 instances of IL). The first type of developers' feedback argues the importance of semantics and usage scenario of logging in deciding the log level. A prior study [72] suggests that logging statements that appear in syntactically similar code, but with inconsistent log levels, are likely problematic. However, based on developers' feedback that we received, IL still may not be a concern, even if the duplicate logging statements reside in very similar code. A developer indicated that "conditions and messages are important but the *context* is even more important". As an example, both of the two methods may display messages to users. One method may be displaying the message to *local* users with a *debug* logging statement to record failure messages. The other method may be displaying the message to *remote* users with an *error* logging statement to record failure messages (problems related to remote procedure calls may be *more severe* in distributed systems). Hence, even if the code is syntactically similar, the log level has a reason to be different due to the different semantics and purposes of the code (i.e., referred to as different *contexts* in developers' responses). Our findings show that `future studies should consider both the syntactic structure and semantics of the code when suggesting log levels.`

The second type of developers' feedback acknowledges the inconsistency. However, developers are reluctant to fix such inconsistencies since developers do not view them as concerns. For example, we reported the instance of IL that we discussed in Table 2 to the developer. The developer replied: "I think it should probably be an *ERROR* level, and I missed it in the review (could make an argument either way, I do not feel strongly that it should be *ERROR* level vs *INFO* level." Our opinions (i.e., from us and prior studies [72, 37]) differ from that of developers' regarding whether such inconsistencies are problematic. On one hand, whether an instance of IL is problematic or not can be subjective. This shows the importance of including perspectives from multiple parties (e.g., user studies, discussions with developers) in future studies of software logging practice. On the other hand, the discrepancy also indicates the need of establishing a guidance for logging practice and further even enforcing such standard.

3.2.5 Pattern 5: Duplicate logging statements in polymorphism (DP)

Description. Classes in object-oriented languages are expected to share similar functionality if they inherit the same parent class or if they implement the same interface (i.e., polymorphism). Since log messages record a higher level abstraction of the program [59], we find that even though there are no clones among a parent method and its overridden methods, such methods may still contain duplicate logging statements. Such duplicate logging statements may cause maintenance overhead. For example, when developers update one log message, he/she may forget to update the log message in all the other sibling classes. Inconsistent log messages may cause problems during log analysis [24, 1].

Example. As shown in Table 2, the two classes (`PowerShellFencer` and `ShellCommandFencer`) in Hadoop both extend the same parent class and implement the same interface. These two classes share similar behaviors. The inherited methods in the two classes have identical log message. However, as the system evolves, developers may not always remember to keep the log messages consistent in the two inherited methods, which may cause problems during system debugging, understanding, and analysis.

Code smell instances. We find that all the 139 instances of DP are potentially problematic that may be fixed by refactoring. In most of the instances, the parent class is an abstract class, and the duplicate logging statements exist in the overridden methods of the subclasses. We also find that in most cases, the overridden methods in the subclasses are very similar with minor differences (e.g., to provide some specialized functionality), which may be the reason that developers use duplicate logging statements.

Developers' feedback. Developers generally agree that DP is associated with logging code smells and specific refactoring techniques are needed. One developer comments that:

"You want to care about the logging part of your code base in the same way as you do for business-logic code (one can argue it is part of it), so salute DRY (do-not-repeat-yourself)."

Resolving DP often requires systematic refactoring. However, to the best of our knowledge, current Java logging frameworks, such as SLF4J and Log4j 2, do not support refactoring logging statements. The way to resolve DP is to ensure that the log message of the parent class can be reused by the subclasses, e.g., storing the log message in a static constant variable. We received similar suggestions from developers on how to refactor DP, such as *"adding a method in the parent class that generates the error text for that case:* logger.error(notAccessible(field.getName()));*"*, or *"creat[ing] your own Exception classes and put message details in them".* However, we find that without supports from logging frameworks, even though developers acknowledged the issue of DP, they do not want to *manually* fix the code smells. Similar to some code smells studied in prior research [29, 60], developers may be reluctant to fix DP due to additional maintenance overheads but limited supports (i.e., need to manually fix hundreds of DP instances). *In short, logging frameworks should provide better support to developers in creating log "templates" that can be reused in different places in the code.*

We manually uncovered five patterns of duplicate logging code smells and six justifiable cases (for the IE and IL pattern) where the code smell instances may not need fixes. In total, our study helped developers fix 56 problematic duplicate logging code smells in the studied systems.

Chapter 4

DLFinder: Automatically Detecting Problematic Duplicate Logging Code Smells

The manual study in Chapter 3 uncovers five patterns of duplicate logging code smells, and provides guidance in identifying *problematic* logging code smells. To help developers automatically detect such problematic code smells and improve logging practices, we propose an automated approach, specifically a static analysis tool, called DLFinder. DLFinder uses abstract syntax tree (AST) analysis, data flow analysis, and text analysis. In this chapter, we discuss how DLFinder detects each pattern of duplicate logging code smell.

4.1 Detecting inadequate information in catch blocks (IC)

DLFinder first locates the *try-catch* blocks that contain duplicate logging statements. Specifically, DLFinder finds the *catch* blocks of the same *try* block that catch different types of exceptions, and these catch blocks contain the same set of duplicate logging statements. Then, DLFinder uses data flow analysis to analyze whether the handled exceptions in the *catch* blocks are logged (e.g., record the exception message). DLFinder detects an instance of IC if none of the logging statements in the *catch* blocks record either the stack trace or the exception message.

4.2 Detecting inconsistent error-diagnostic information (IE)

DLFinder first identifies all the *catch* blocks that contain duplicate logging statements. Then, for each *catch* block, DLFinder uses data flow analysis to determine how the exception is logged by analyzing the usage of the exception variable in the logging statement. Namely, the logging statement records 1) the entire stack trace, 2) only the exception message, or 3) nothing at all. Then, DLFinder compares how the exception variable is used/recorded in each of the duplicate logging statements. DLFinder detects an instance of IE if a set of duplicate logging statements that appear in *catch* blocks has an inconsistent way of recording the exception variables (e.g., the log in one *catch* block records the entire stack trace, and the log in another *catch* block records only the exception message, while the two catch blocks handle the same type of exception). Note that for each instance of IE, the multiple *catch* blocks with duplicate logging statements in the same set may belong to different *try* blocks. In addition, DLFinder decides if an instance of IE belongs to one of the three justifiable cases (IE.1–IE.3). If so, the instance is marked as potentially justifiable, and thus, excluded by DLFinder.

4.3 Detecting log message mismatch (LM)

LM is about having an incorrect method or class name in the log message (e.g., due to copy-and-paste). Hence, DLFinder analyzes the text in both the log message and the class-method name (i.e., concatenation of class name and method name) to detect LM by applying commonly used text analysis approaches [16]. DLFinder detects instances of LM using four steps: 1) For each logging statement, DLFinder splits class-method name into a set of words (i.e., *name set*) and splits log message into a set of words (i.e., *log set*) by leveraging naming conventions (e.g., camel cases) and converting the words to lower cases. 2) DLFinder applies stemming on all the words using Porter Stemmer [52]. 3) DLFinder removes stop words in the log message for each system. We find that there is a significant number of words that are generic across the log messages in a system (e.g., on, with, and process). Hence, we obtain the stop words by finding the top 50 most frequent words (each of our four studied systems has an average of 3,178 unique words in the static text messages) across all log messages in a system [68]. 4) For every logging statement, between the name set (i.e., from the class-method name) and its associated log set, DLFinder counts the number of common words shared by both sets. Afterward, DLFinder detects an instance of LM if the number of common words is inconsistent among the duplicate logging statements in one set.

For the LM example shown in Table 2, the common words shared by the first pair (i.e., method *doScaleUp()* and its log) are “scale, up”, while the common word shared by the second pair is “scale”. Hence, DLFinder detects an LM instance due to this inconsistency. The rationale is that the number of common words between the class-method name and the associated logging statement is subject to change if developers make copy-and-paste errors on logging statements (e.g., copy the logging statement in *doScaleUp()* to method *doScaleDown()*), but forget to update the log message to match with the new method name "doScaleDown". However, the number of common words will remain unchanged (i.e., no inconsistency) if the logging statement (after being pasted at a new location) is updated respectively.

4.4 Detecting inconsistent log level (IL)

DLFinder detects an instance of IL if duplicate logging statements in one set (i.e., have the same static text message) have inconsistent log level. Furthermore, DLFinder checks whether an instance of IL belongs to one of the three justifiable cases (IL.1–IL.3). If so, the instance is marked as justifiable and DLFinder excludes this instance in the detection result.

4.5 Detecting duplicate logs in polymorphism (DP)

DLFinder generates an object inheritance graph when statically analyzing the Java code. For each overridden method, DLFinder checks if there exist any duplicate logging statements in the corresponding method of the sibling and the parent class. If there exist such duplicate logging statements, DLFinder detects an instance of DP. Note that, based on the feedback that we received from developers (Chapter 3), we do not expect developers to fix instances of DP. DP instances can be viewed more as technical debts [35] and our goal is to propose an approach to detect DP instances to raise the awareness from the research community and developers regarding this issue.

Chapter 5

An Evaluation of DLFinder and Threats to Validity

In this chapter, we evaluate our tool by answering three research questions.

5.1 RQ1: What is the accuracy of DLFinder detecting duplicate logging code smells in the four manually studied systems?

5.1.1 Motivation

DLFinder was implemented based on the duplicate logging code smells uncovered from the four manually studied systems. Since we obtain the ground truth (i.e., problematic code smells) in these four systems from our manual study, the goal of this RQ is to evaluate the detection accuracy of DLFinder.

5.1.2 Approach

We applied DLFinder on the same versions of the systems that we used in our manual study (Chapter 3). We calculated the precision and recall of DLFinder in detecting problematic duplicate logging code smells. Precision is the percentage of problematic code smell instances among all the detected instances, and recall is the percentage of problematic code smell instances that DLFinder is able to detect.

Table 4: The results of DLFinder in RQ1 and RQ2. In each pattern, Pro. is the number of problematic instances as the ground-truth, Det. is the number of instances detected by DLFinder, and T.Det. is the number of true problematic instances detected by DLFinder.

Research questions		**IC**			**IE**			**LM**			**IL**			**DP**		
		Pro.	T.Det.	Det.	Pro.	T.Det.	Det.	Pro.	T.Det.	Det.	Pro.	T.Det.	Det.	Pro.	T.Det.	Det.
RQ1: applying DLFinder on the same software versions as the manual study	**Cassandra**	1	1	1	0	0	0	0	0	4	0	0	2	2	2	2
	CloudStack	8	8	8	4	4	4	27	24	186	0	0	12	107	107	107
	ElasticSearch	1	1	1	0	0	0	1	0	15	0	0	0	3	3	3
	Hadoop	5	5	5	0	0	0	9	7	44	0	0	1	27	27	27
	Precision / Recall	100% / 100%			100% / 100%			12.4% / 83.8%			N/A			100% / 100%		
RQ2: applying DLFinder on additional systems	**Camel**	1	1	1	0	0	0	14	10	95	0	0	3	29	29	29
	Flink	0	0	0	2	2	2	4	4	41	0	0	0	24	24	24
	Kafka	0	0	0	0	0	0	3	3	15	0	0	0	14	14	14
	Wicket	1	1	1	0	0	0	1	1	4	0	0	0	1	1	1
	Precision / Recall	100% / 100%			100% / 100%			11.6% / 81.8%			N/A			100% / 100%		
	Total	17	17	17	6	6	6	59	49	404	0	0	18	207	207	207

5.1.3 Results and discussion

The first five rows of Table 4 show the results of RQ1. Note that all the numbers in Table 4 show instances of problematic code smells which require fixes since DLFinder focuses on detecting problematic ones and excludes the justifiable cases. For the patterns of IC, IE, and DP, DLFinder detects all the problematic instances of duplicate logging code smells (100% in recall) with a precision of 100%. For the IL pattern, since we do not find any problematic instances (as discussed in Chapter 3), both of the columns of problematic instances in ground truth (*Pro.*) and detected (*T.Det.*) in Table 4 are 0.

For the LM pattern, DLFinder achieves a recall of 83.8% (i.e., DLFinder detects 31/37 problematic LM instances). We manually investigate the six instances of LM that DLFinder cannot detect. We find that the problem is related to typos in the log message. For example, developers may write "mlockall" instead of "mLockAll". Hence, the text in the log message cannot be matched with the method name when we split the word using camel cases. The precision of detecting problematic LM instances is modest because, in many false positive cases, the log messages and class-method names are at different levels of abstraction: The log message describes a local code block while the class-method name describes the functionality of the entire method. For example, *encodePublicKey()* and *encodePrivateKey()* both contain the duplicate logging statement *"Unable to create KeyFactory"*. The duplicate logging statement describes a local code block that is related to the usage of the *KeyFactory* class, which is different from the major functionalities of the two methods (i.e., as expressed by their class-method names). Nevertheless, DLFinder detects the LM instances with a high recall, and developers may quickly go through the results to identify the true positives (it took the author of this thesis less than 10 minutes on average to go through the LM result of each system to identify true positives).

To further evaluate our detection approach for LM, we compare our detection results with a baseline. We use random prediction algorithm as our baseline, which is commonly used as the baseline in prior studies [66, 65, 64]. The random prediction algorithm predicts the label of an item (i.e., whether a set of duplicate logging statements belong to LM) based on the distribution of the training data. For each system, we use our manually labeled results (which are discussed and verified in the previous chapters) as the training data. Note that we only compare the detection results of LM with the baseline. The reason is that pattern IC, IE, IL, and DP are relatively independent and well-defined, unlike LM which depends on the semantics of the logging statement and its surrounding code. We repeat the random prediction 30 times (as suggested by previous studies [21, 17]) for each system to reduce the biases. Finally, we report the average precision and recall that are computed based on the 30 times of iterations. Figure 3 and Figure 4 shows how the precision and recall of our approach compared to that of the baseline. The average precision and recall for the baseline are 3.47% and 4.03%, respectively, for the four studied systems. Our detection approach achieves a precision and recall of 12.4% and 83.8%, respectively. In short, our approach is significantly better than the baseline and is able to have a very high recall in the four manually studied systems.

5.2 RQ2: What is the accuracy of DLFinder detecting duplicate logging code smells in additional studied systems?

5.2.1 Motivation

The goal of this RQ is to study whether the uncovered patterns of duplicate logging code smells are generalizable to other systems.

5.2.2 Approach

We applied DLFinder to four additional systems that are not included in the manual study in chapter 3: Camel, Flink, Kafka, and Wicket, which are all large-scale open source Java systems. Details of the systems are presented in Table 1. Similar to our manual study, we manually collect the problematic duplicate logging code smells in the additional systems, i.e., the ground-truth used for calculating the precision and recall of DLFinder. Note that the collected ground-truth of the additional systems is only used in this evaluation, but not in designing the patterns in DLFinder. (There are also no new patterns found in this process.)

5.2.3 Results and discussion

The second half of Table 4 shows the results of the additional systems. In total, we found 26 problematic code smell instances (DLFinder detects 22) in these systems and all of them are reported and fixed. Compared to the four systems in RQ1, DLFinder has similar precision and recall values in the additional systems. DLFinder detected three instances of IL in Camel; however, based on the manual investigation and developers' feedback, these IL instances are not problematic. Similar to what we discuss in Chapter 3, the differences in the log level are related to having different semantics in the code. Different from a prior study [72], we found that all IL instances are not problematic in the eight evaluated systems. Future studies are needed to investigate the effect of IL. DLFinder detects DP instances with 100% in recall and precision; however, developers are reluctant to fix them due to limited support from logging frameworks. Nevertheless, the patterns of duplicate logging code smells that we uncovered can still be found in other systems.

DLFinder achieves good precision and recall in the additional systems. Similar to our observation in RQ1, we find that DLFinder cannot detect some LM instances due to typos in log message. We also compare our LM detection results with the baseline mentioned in RQ1 using the same approach. The average precision and recall for DLFinder are 11.6% and 81.8%, respectively, which are considerably better than the precision (1.75%) and recall (1.55%) of the baseline. In summary, apart from the manually studied systems in RQ1, DLFinder also achieves noticeably better precision and recall than the baseline and is able to have a reasonably high recall in the additional systems.

5.3 RQ3: Are new code smell instances introduced over time?

5.3.1 Motivation

The purpose of this RQ is to investigate whether new duplicate logging code smells instances are introduced during the evolution of systems.

5.3.2 Approach

We applied DLFinder on the latest versions of the four studied systems, i.e., Hadoop, CloudStack, ElasticSearch, and Cassandra, and compare the results with the ones on previous versions. The gaps of days between the manually studied versions and the new versions vary from 77 days to 297 days.

5.3.3 Results and discussion

Table 5 shows that new instances of code smells are introduced during software evolution. These detected code smell instances are all problematic and are all reported and fixed (except for DP). As mentioned in Chapter 3 and 4, our goal of detecting DP is to show developers the logging technical debt in their systems. In short, we found that duplicate logging code smells are introduced over time, and an automated approach such as DLFinder can help developers avoid duplicate logging code smells as the system evolves.

The duplicate logging code smells exist in both manually studied and additional systems. In total, DLFinder is able to detect 81 out of 91 problematic code smell instances. We also find that new instances of logging code smells are introduced as systems evolve.

5.4 Threats to Validity

In this chapter, we discuss the threats to validity of this thesis.

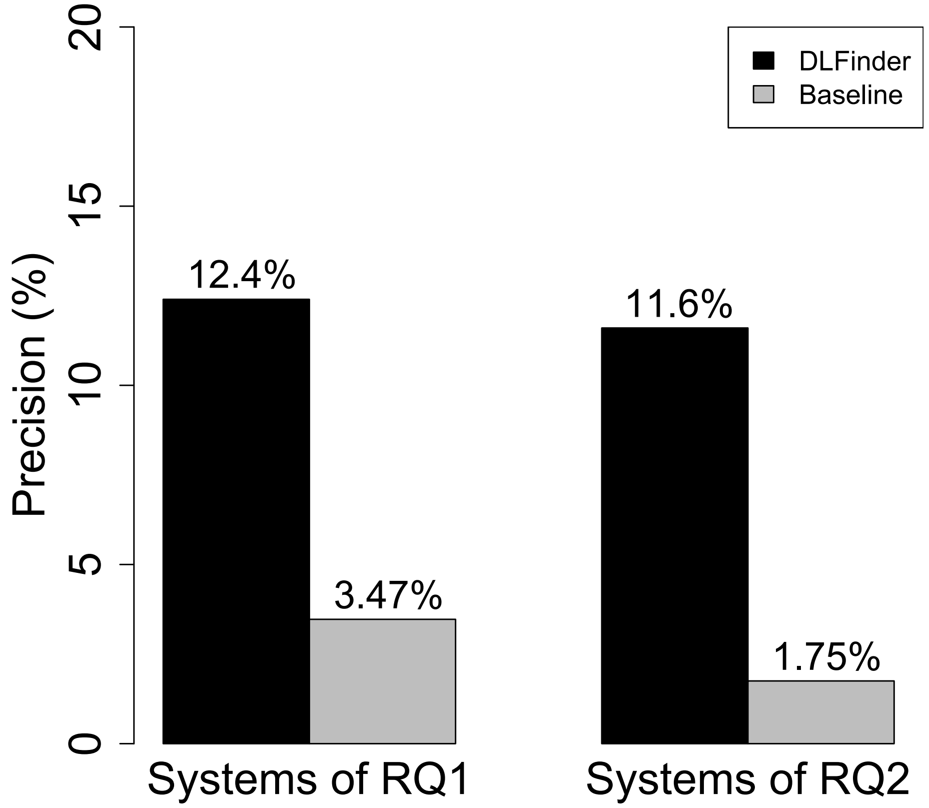

Figure 3: The precision of DLFinder detecting LM on the systems of RQ1 and RQ2 respectively, compared with the baseline (random prediction).

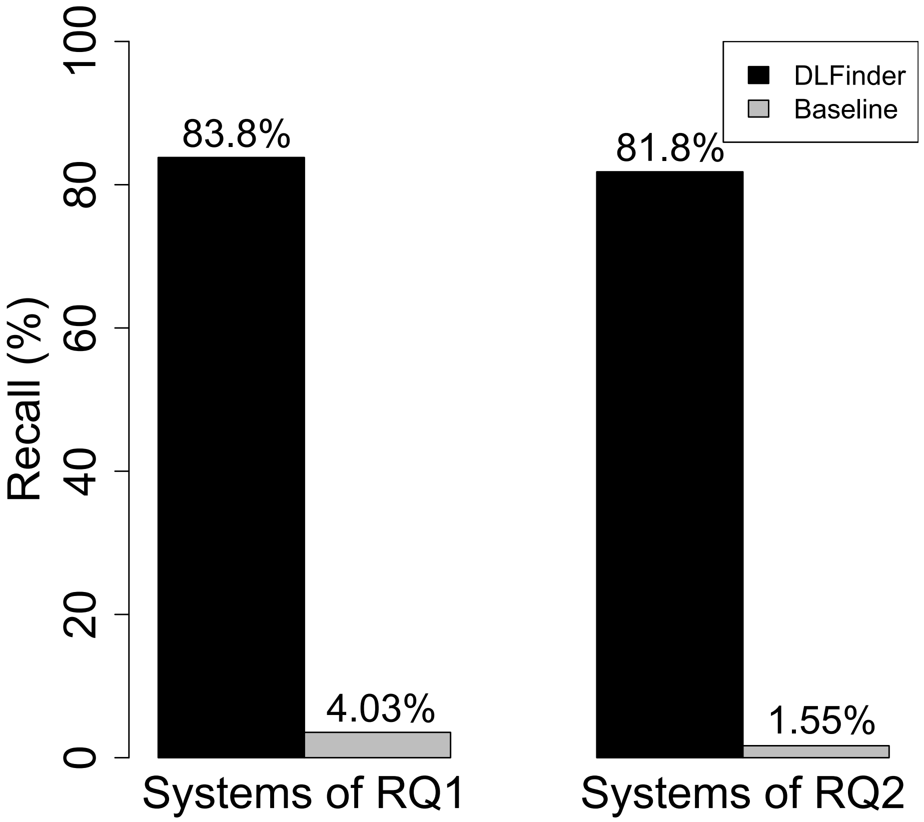

Figure 4: The recall of DLFinder detecting LM on the systems of RQ1 and RQ2 respectively, compared with the baseline (random prediction).

Table 5: The result of RQ3: applying DLFinder to the newer versions of the studied systems. Gap. shows the duration of time in days between the original (Org.) and the newer release (New.)

	Releases		**IC**	**IE**	**LM**	**IL**	**DP**
	Org., New.	Gap.					
Cassandra	3.11.1, 3.11.3	294	0	0	0	0	1
CloudStack	4.9.3, 4.11.1	297	5	0	2	0	0
ElasticSearch	6.0.0, 6.1.3	77	0	0	0	0	0
Hadoop	3.0.0, 3.0.3	208	0	0	2	0	21
Total	-	-	5	0	4	0	22

5.4.1 Construct Validity

A potentially more intuitive approach for detecting duplicate logging statements may rely on code clone detection tools such as CCFinder [32] and CCLearner [38]. However, it is not clear if duplicate logging statements are indeed associated with code clones. Moreover, the performance of clone detection tools is dependent on the thresholds [54, 34, 50, 40]. Choosing the optimal thresholds is a non-trivial task and the value may differ across systems [62]. Future works are needed for investigating the relationship between duplicate logging statements and code clones.

5.4.2 Internal validity

We define duplicate logging statements as two or more logging statements that have the same static text message. We were able to uncover five patterns of duplicate logging code smells and detect many code smell instances. However, logging statements with non-identical but similar static texts may also cause problems to developers. Future studies should consider different types of duplicate logging statements (e.g., logs with similar text messages). We conducted manual studies to uncover the patterns of code smells, study their potential impact and examine duplicate logging statements that are not classified by the automated clone detection tool as clones. To avoid biases, we examine the data independently. For most of the cases we reach an agreement. Any disagreement is discussed until a consensus is reached. In order to reduce the subjective bias, we have contacted the developers to confirm the uncovered patterns and their impact. We only do exact match for log message when we are identifying duplicate logging statements, and we are able to get a large set of results. However logs with slight differences may still cause confusion to operators. Future works might apply fuzzy match to possibly uncover more potential problems under logging practices.

5.4.3 External validity

We conducted our study on four large-scale open source systems in different domains. We found that our uncovered patterns and the corresponding problematic and justifiable cases are common among the studied systems. However, our finding may not be generalizable to other systems. Hence, we studied whether the uncovered patterns exist in four other systems. We found that the patterns of code smells also exist in these four systems and we did not find any new code smell patterns in our manual verification. Our studied systems are all implemented in Java, so the results may not be generalizable to systems in other programming languages. Future studies should validate the generalizability of our findings in systems in other programming languages.

Chapter 6

Conclusion and Future Work

In this chapter, we summarize process and contributions discussed in this thesis. In addition, we propose potential future works that might be complementary to this thesis for better understanding duplicate logging statements and improving logging practices.

6.1 Summary of the thesis

Duplicate logging statements may affect developers' understanding of the system execution. In this thesis, we study over 3K duplicate logging statements in four large-scale open source systems (Hadoop, CloudStack, ElasticSearch, and Cassandra). We uncover five patterns of duplicate logging code smells. Further, we assess the impact of each code smell and find not all are problematic and need fixes. In particular, we find six justifiable cases where the uncovered patterns of duplicate logging code smells may not be problematic. We received confirmation from developers on both the problematic and justifiable cases. Combining our manual analysis and developers' feedback, we developed a static analysis tool, DLFinder, which automatically detects problematic duplicate logging code smells. We applied DLFinder on the four manually studied systems and four additional systems. In total, we reported 91 problematic code smell instances in the eight studied systems to developers and all of them are fixed. DLFinder successfully detects 81 out of the 91 instances.

This thesis highlights the importance of the context of the logging code, i.e., the nature of logging code is highly associated with both the structure and the functionality of the surrounding code. Future studies should consider the code context when providing guidance to logging practices, more advanced logging libraries are needed to help developers improve logging practice and to avoid logging code smells.

6.2 Future work

This thesis makes a major contribution towards the goal of improving logging practices. However, there are many open challenges and research opportunities that may complement this work and further provide a logging guideline for developers to improve the quality of logging code. We now highlight some aspects for future work.

6.2.1 Investigating the relationship between duplicate logging statements and code clones

During the process of our manual analysis in this thesis, we noticed that some duplicate logging statements or duplicate logging code smells may be related to code clones. Code clone or duplicate code is considered a bad programming practice and an indication of deeper maintenance problems. Prior studies focus on the detection of code clones by applying text-based approaches or abstract syntax tree analysis on program source code. However, these studies did not study code clones from the perspective of logging code. Logging statements might also be copied along with other code since cloning is often performed hastily without much attention on the context. Some cloned logging statements (e.g., the LM logging code smell discussed in this thesis) may fail to correctly record the runtime behaviors; and thus, increase maintenance difficulties. Hence, a formal investigation of the relationship between duplicate logging statements and code clones might be needed for improving logging practices, and may inspire future code clone studies.

6.2.2 Examining the impact of similar logging statements on logging practice

In this thesis, we apply exact matching on log messages to identify duplicate logging statements. However, even though each log itself may be impeccable, similar logs may also affect the analysis of logs (e.g., log searching). Similar logs might be generated from logging statements that are highly similar, but with minor differences in static messages or dynamic variables. Future work may apply some specific techniques (e.g., clustering) to characterize similar logging statements and provide an approach to better distinguish them then further help improve logging practices.

6.2.3 Better leveraging the semantic and structural information of logging statements

As we discussed in this thesis, the semantic and structural information of the code is crucial for making logging decisions. Bad logging decision that does not properly leverage those information may cause confusing logs that affect maintenance (e.g., the LM logging code smell). Hence, code context of logging statements should also be considered for future work to precisely uncover potential problems in logs and improve logging practices.

Bibliography

[1] Changes to JobHistory makes it backward incompatible. `https://issues.apache.org/jira/browse/HADOOP-4190`. Last checked April 4th 2018.

[2] Log4j. http://logging.apache.org/log4j/2.x/.

[3] Simple logging facade for Java (SLF4J). `http://www.slf4j.org`. Last checked Feb. 2018.

[4] S. L. Abebe, S. Haiduc, P. Tonella, and A. Marcus. The effect of lexicon bad smells on concept location in source code. In *2011 IEEE 11th International Working Conference on Source Code Analysis and Manipulation*, pages 125–134, Sept 2011.

[5] Iftekhar Ahmed, Caius Brindescu, Umme Ayda Mannan, Carlos Jensen, and Anita Sarma. An empirical examination of the relationship between code smells and merge conflicts. In *Proceedings of the 11th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM '17, pages 58–67, 2017.

[6] Howard Barringer, Alex Groce, Klaus Havelund, and Margaret H. Smith. Formal analysis of log files. *JACIC*, 7(11):365–390, 2010.

[7] S. Boslaugh and P.A. Watters. *Statistics in a Nutshell: A Desktop Quick Reference*. In a Nutshell (O'Reilly). O'Reilly Media, 2008.

[8] D. Budgen. *Software Design*. Addison-Wesley, 2003.

[9] Nimrod Busany and Shahar Maoz. Behavioral log analysis with statistical guarantees. In *Proceedings of the 38th International Conference on Software Engineering*, ICSE '16, pages 877–887, 2016.

[10] C. Chapman, P. Wang, and K. T. Stolee. Exploring regular expression comprehension. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 405–416, Oct 2017.

[11] Boyuan Chen and Zhen Ming (Jack) Jiang. Characterizing logging practices in java-based open source software projects – a replication study in apache software foundation. *Empirical Software Engineering*, 22(1):330–374, Feb 2017.

[12] Boyuan Chen and Zhen Ming (Jack) Jiang. Characterizing and detecting anti-patterns in the logging code. In *Proceedings of the 39th International Conference on Software Engineering*, ICSE '17, pages 71–81, 2017.

[13] Boyuan Chen, Jian Song, Peng Xu, Xing Hu, and Zhen Ming (Jack) Jiang. An automated approach to estimating code coverage measures via execution logs. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018*, pages 305–316, 2018.

[14] Tse-Hsun Chen, Weiyi Shang, Ahmed E. Hassan, Mohamed Nasser, and Parminder Flora. Cacheoptimizer: Helping developers configure caching frameworks for hibernate-based database-centric web applications. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2016, pages 666–677, 2016.

[15] Tse-Hsun Chen, Mark D. Syer, Weiyi Shang, Zhen Ming Jiang, Ahmed E. Hassan, Mohamed Nasser, and Parminder Flora. Analytics-driven load testing: An industrial experience report on load testing of large-scale systems. In *Proceedings of the 39th International Conference on Software Engineering: Software Engineering in Practice Track*, ICSE-SEIP '17, pages 243–252, 2017.

[16] Tse-Hsun Chen, Stephen W. Thomas, and Ahmed E. Hassan. A survey on the use of topic models when mining software repositories. *Empirical Software Engineering*, 21(5):1843–1919, 2016.

[17] Tse-Hsun Chen, Shang Weiyi, Zhen Ming Jiang, Ahmed E. Hassan, Mohamed Nasser, and Parminder Flora. Detecting performance anti-patterns for applications developed using object-relational mapping. In *Proceedings of the 36th International Conference on Software Engineering (ICSE)*, pages 1001–1012, 2014.

[18] Norman Cliff. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin*, 114(3):494–509, November 1993.

[19] M. Fowler and K. Beck. *Refactoring: Improving the Design of Existing Code.* Addison-Wesley object technology series. 1999.

[20] 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 *Proceedings of the 36th International Conference on Software Engineering*, ICSE-SEIP '14, pages 24–33, 2014.

[21] Andy Georges, Dries Buytaert, and Lieven Eeckhout. Statistically rigorous java performance evaluation. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2007, October 21-25, 2007, Montreal, Quebec, Canada*, pages 57–76, 2007.

[22] Nils Göde and Rainer Koschke. Frequency and risks of changes to clones. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu , HI, USA, May 21-28, 2011*, pages 311–320, 2011.

[23] Ahmed E. Hassan, Daryl J. Martin, Parminder Flora, Paul Mansfield, and Dave Dietz. An Industrial Case Study of Customizing Operational Profiles Using Log Compression. In *ICSE '08: Proceedings of the 30th international conference on Software engineering*, pages 713–723, Leipzig, Germany, 2008. ACM.

[24] Mehran Hassani, Weiyi Shang, Emad Shihab, and Nikolaos Tsantalis. Studying and detecting log-related issues. *Empirical Software Engineering*, 2018.

[25] Felienne Hermans, Martin Pinzger, and Arie van Deursen. Detecting and refactoring code smells in spreadsheet formulas. *Empirical Software Engineering*, 20(2):549–575, 2015.

[26] Judith F. Islam, Manishankar Mondal, and Chanchal K. Roy. A comparative study of software bugs in micro-clones and regular code clones. In *26th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2019, Hangzhou, China, February 24-27, 2019*, pages 73–83, 2019.

[27] Judith F. Islam, Manishankar Mondal, Chanchal K. Roy, and Kevin A. Schneider. Comparing bug replication in regular and micro code clones. In *Proceedings of the 27th International Conference on Program Comprehension*, ICPC '19, pages 81–92, 2019.

[28] Zhen Ming Jiang, Ahmed E. Hassan, Gilbert Hamann, and Parminder Flora. Automatic identification of load testing problems. In *Proceedings of 24th International Conference on Software Maintenance (ICSM)*, pages 307–316, 2008.

[29] Brittany Johnson, Yoonki Song, Emerson Murphy-Hill, and Robert Bowdidge. Why don't software developers use static analysis tools to find bugs? In *Proceedings of the 2013 International Conference on Software Engineering*, ICSE '13, pages 672–681, 2013.

[30] Elmar Jürgens, Florian Deissenboeck, Benjamin Hummel, and Stefan Wagner. Do code clones matter? In *31st International Conference on Software Engineering, ICSE 2009, May 16-24, 2009, Vancouver, Canada, Proceedings*, pages 485–495, 2009.

[31] Suhas Kabinna, Cor-Paul Bezemer, Weiyi Shang, and Ahmed E. Hassan. Logging library migrations: A case study for the apache software foundation projects. In *Proceedings of the 13th International Conference on Mining Software Repositories*, MSR '16, pages 154–164, 2016.

[32] Toshihiro Kamiya, Shinji Kusumoto, and Katsuro Inoue. Ccfinder: A multilinguistic token-based code clone detection system for large scale source code. *IEEE Transactions on Software Engineering*, 28(7):654–670, 2002.

[33] Cory Kapser and Michael W. Godfrey. Cloning considered harmful. *Reverse Engineering, Working Conference on*, 0:19–28, 2006.

[34] I. Keivanloo, F. Zhang, and Y. Zou. Threshold-free code clone detection for a large-scale heterogeneous java repository. In *Proceedings of the 22nd International Conference on Software Analysis, Evolution, and Reengineering*, SANER '15, pages 201–210, 2015.

[35] Philippe Kruchten, Robert L. Nord, and Ipek Ozkaya. Technical debt: From metaphor to theory and practice. *IEEE Softw.*, 29(6):18–21, 2012.

[36] Heng Li, Tse-Hsun (Peter) Chen, Weiyi Shang, and Ahmed E. Hassan. Studying software logging using topic models. *Empirical Software Engineering*, Jan 2018.

[37] Heng Li, Weiyi Shang, and Ahmed E. Hassan. Which log level should developers choose for a new logging statement? *Empirical Software Engineering*, 22(4):1684–1716, Aug 2017.

[38] L. Li, H. Feng, W. Zhuang, N. Meng, and B. Ryder. Cclearner: A deep learning-based clone detection approach. In *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 249–260, 2017.

[39] Zhenmin Li, Shan Lu, Suvda Myagmar, and Yuanyuan Zhou. Cp-miner: Finding copy-paste and related bugs in large-scale software code. *IEEE Trans. Software Eng.*, 32(3):176–192, 2006.

[40] Yun Lin, Zhenchang Xing, Yinxing Xue, Yang Liu, Xin Peng, Jun Sun, and Wenyun Zhao. Detecting differences across multiple instances of code clones. In *Proceedings of the 36th International Conference on Software Engineering*, ICSE 2014, pages 164–174, 2014.

[41] Umme Ayda Mannan, Iftekhar Ahmed, Rana Abdullah M. Almurshed, Danny Dig, and Carlos Jensen. Understanding code smells in android applications. In *Proceedings of the International Conference on Mobile Software Engineering and Systems*, pages 225–234, 2016.

[42] Mary L. McHugh. Interrater reliability: the kappa statistic. *Biochemia Medica*, 22(3):276–282, 2012.

[43] Manishankar Mondal, Chanchal K. Roy, and Kevin A. Schneider. Micro-clones in evolving software. In *Proceedings of the 25th International Conference on Software Analysis, Evolution and Reengineering*, pages 50–60, 2018.

[44] Mirko Montanari, Jun Ho Huh, Derek Dagit, Rakesh Bobba, and Roy H. Campbell. Evidence of log integrity in policy-based security monitoring. In *International Conference on Dependable Systems and Networks Workshops*, DSN '12, pages 1–6, 2012.

[45] Karthik Nagaraj, Charles Edwin Killian, and Jennifer Neville. Structured comparative analysis of systems logs to diagnose performance problems. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation*, NSDI '12, pages 353–366, 2012.

[46] Hung Viet Nguyen, Hoan Anh Nguyen, Tung Thanh Nguyen, Anh Tuan Nguyen, and Tien N. Nguyen. Detection of embedded code smells in dynamic web applications. In *Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*, ASE 2012, pages 282–285, 2012.

[47] F. Palomba, G. Bavota, M. Di Penta, R. Oliveto, A. De Lucia, and D. Poshyvanyk. Detecting bad smells in source code using change history information. In *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 268–278, Nov 2013.

[48] Chris Parnin, Carsten Görg, and Ogechi Nnadi. A catalogue of lightweight visualizations to support code smell inspection. In *Proceedings of the 4th ACM Symposium on Software Visualization*, SoftVis '08, pages 77–86, 2008.

[49] A. Pecchia, M. Cinque, G. Carrozza, and D. Cotroneo. Industry practices and event logging: Assessment of a critical software development process. In *Proceedings of th 37th International Conference on Software Engineering*, ICSE '15, pages 169–178, 2015.

[50] N. H. Pham, H. A. Nguyen, T. T. Nguyen, J. M. Al-Kofahi, and T. N. Nguyen. Complete and accurate clone detection in graph-based models. In *2009 IEEE 31st International Conference on Software Engineering*, pages 276–286, 2009.

[51] He Pinjia, Zhuangbin Chen, Shilin He, and Michael R. Lyu. Characterizing the natural language descriptions in software logging statements. In *Proceedings of the 33rd IEEE international conference on Automated software engineering*, pages 1–11, 2018.

[52] Martin F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.

[53] F. Rahman, C. Bird, and P. Devanbu. Clones: What is that smell? In *2010 7th IEEE Working Conference on Mining Software Repositories (MSR 2010)*, pages 72–81, May 2010.

[54] Chanchal K. Roy, James R. Cordy, and Rainer Koschke. Comparison and evaluation of code clone detection techniques and tools: A qualitative approach. *Sci. Comput. Program.*, 74(7):470–495, 2009.

[55] Chanchal Kumar Roy and James R. Cordy. NICAD: accurate detection of near-miss intentional clones using flexible pretty-printing and code normalization. In *The 16th IEEE International Conference on Program Comprehension*, ICPC '08, pages 172–181, 2008.

[56] Jan Schumacher, Nico Zazworka, Forrest Shull, Carolyn Seaman, and Michele Shaw. Building empirical support for automated code smell detection. In *Proceedings of the 2010 ACM-IEEE International Symposium on Empirical Software Engineering and Measurement*, ESEM '10, pages 8:1–8:10, 2010.

[57] Weiyi Shang, Zhen Ming Jiang, Hadi Hemmati, Bram Adams, Ahmed E. Hassan, and Patrick Martin. Assisting developers of big data analytics applications when deploying on hadoop clouds. In *Proceedings of the 2013 International Conference on Software Engineering*, ICSE '13, pages 402–411, 2013.

[58] Weiyi Shang, Meiyappan Nagappan, and Ahmed E. Hassan. Studying the relationship between logging characteristics and the code quality of platform software. *Empirical Software Engineering*, 20(1):1–27, Feb 2015.

[59] Weiyi Shang, Meiyappan Nagappan, Ahmed E. Hassan, and Zhen Ming Jiang. Understanding log lines using development knowledge. In *Proceedings of the 2014 IEEE International Conference on Software Maintenance and Evolution*, ICSME '14, pages 21–30, 2014.

[60] Danilo Silva, Nikolaos Tsantalis, and Marco Tulio Valente. Why we refactor? confessions of github contributors. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE '16, pages 858–870, 2016.

[61] D. I. K. Sjøberg, A. Yamashita, B. C. D. Anda, A. Mockus, and T. Dybå. Quantifying the effect of code smells on maintenance effort. *IEEE Transactions on Software Engineering*, 39(8):1144–1156, Aug 2013.

[62] Nikolaos Tsantalis, Laleh M. Eshkevari, Davood Mazinanian, and Danny Dig. Accurate and efficient refactoring detection in commit history. In *Proceedings of the 40th International Conference on Software Engineering*, ICSE '18, 2018.

[63] M. Tufano, F. Palomba, G. Bavota, R. Oliveto, M. Di Penta, A. De Lucia, and D. Poshyvanyk. When and why your code starts to smell bad. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, volume 1, pages 403–414, May 2015.

[64] Harold Valdivia Garcia and Emad Shihab. Characterizing and predicting blocking bugs in open source projects. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, MSR 2014, pages 72–81, 2014.

[65] Xin Xia, David Lo, Emad Shihab, Xinyu Wang, and Xiaohu Yang. Elblocker: Predicting blocking bugs with ensemble imbalance learning. *Information & Software Technology*, 61:93–106, 2015.

[66] Xin Xia, Emad Shihab, Yasutaka Kamei, David Lo, and Xinyu Wang. Predicting crashing releases of mobile applications. In *Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM 2016, Ciudad Real, Spain, September 8-9, 2016*, pages 29:1–29:10, 2016.

[67] Xiao Xiao, Shi Han, Charles Zhang, and Dongmei Zhang. Uncovering javascript performance code smells relevant to type mutations. In Xinyu Feng and Sungwoo Park, editors, *Programming Languages and Systems*, pages 335–355, 2015.

[68] Jinqiu Yang and Lin Tan. SWordNet: Inferring semantically related words from software context. *Empirical Software Engineering*, 19(6):1856–1886, 2014.

[69] Kundi Yao, Guilherme B. de Pádua, Weiyi Shang, Steve Sporea, Andrei Toma, and Sarah Sajedi. Log4perf: Suggesting logging locations for web-based systems' performance monitoring. In *Proceedings of the 2018 ACM/SPEC International Conference on Performance Engineering*, ICPE '18, pages 21–30, 2018.

[70] Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, Yongle Zhang, Pranay U. Jain, and Michael Stumm. Simple testing can prevent most critical failures: An analysis of production failures in distributed data-intensive systems. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*, OSDI'14, pages 249–265, 2014.

[71] Ding Yuan, Haohui Mai, Weiwei Xiong, Lin Tan, Yuanyuan Zhou, and Shankar Pasupathy. Sherlog: Error diagnosis by connecting clues from run-time logs. In *Proceedings of the 15th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 143–154, 2010.

[72] Ding Yuan, Soyeon Park, and Yuanyuan Zhou. Characterizing logging practices in open-source software. In *ICSE 2012: Proceedings of the 2012 International Conference on Software Engineering*, pages 102–112, Piscataway, NJ, USA, 2012. IEEE Press.

[73] Ding Yuan, Jing Zheng, Soyeon Park, Yuanyuan Zhou, and Stefan Savage. Improving software diagnosability via log enhancement. In *ASPLOS '11: Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems*, pages 3–14, Newport Beach, California, USA, 2011. ACM.

[74] Min Zhang, Tracy Hall, and Nathan Baddoo. Code bad smells: a review of current knowledge. *Journal of Software Maintenance*, 23(3):179–202, 2011.

[75] Xu Zhao, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan, and Yuanyuan Zhou. Log20: Fully automated optimal placement of log printing statements under specified overhead threshold. In *Proceedings of the 26th Symposium on Operating Systems Principles*, SOSP '17, pages 565–581, 2017.

[76] Jieming Zhu, Pinjia He, Qiang Fu, Hongyu Zhang, Michael R. Lyu, and Dongmei Zhang. Learning to log: Helping developers make informed logging decisions. In *Proceedings of the 37th International Conference on Software Engineering*, ICSE '15, pages 415–425, 2015.