Understanding Code Smell Detection via Code Review: A Study of the OpenStack Community

Xiaofeng Han¹, Amjed Tahir², Peng Liang^{1*}, Steve Counsell³, Yajing Luo¹

¹ School of Computer Science, Wuhan University, Wuhan, China

² School of Fundamental Sciences, Massey University, Palmerston North, New Zealand

³ Department of Computer Science, Brunel University London, London, United Kingdom

hanxiaofeng@whu.edu.cn, a.tahir@massey.ac.nz, liangp@whu.edu.cn, steve.counsell@brunel.ac.uk, luoyajing@whu.edu.cn

Abstract-Code review plays an important role in software quality control. A typical review process would involve a careful check of a piece of code in an attempt to find defects and other quality issues/violations. One type of issues that may impact the quality of the software is code smells - i.e., bad programming practices that may lead to defects or maintenance issues. Yet, little is known about the extent to which code smells are identified during code reviews. To investigate the concept behind code smells identified in code reviews and what actions reviewers suggest and developers take in response to the identified smells, we conducted an empirical study of code smells in code reviews using the two most active OpenStack projects (Nova and Neutron). We manually checked 19,146 review comments obtained by keywords search and random selection, and got 1,190 smell-related reviews to study the causes of code smells and actions taken against the identified smells. Our analysis found that 1) code smells were not commonly identified in code reviews, 2) smells were usually caused by violation of coding conventions, 3) reviewers usually provided constructive feedback, including fixing (refactoring) recommendations to help developers remove smells, and 4) developers generally followed those recommendations and actioned the changes. Our results suggest that 1) developers should closely follow coding conventions in their projects to avoid introducing code smells, and 2) review-based detection of code smells is perceived to be a trustworthy approach by developers, mainly because reviews are context-sensitive (as reviewers are more aware of the context of the code given that they are part of the project's development team).

Index Terms—Code Review, Code Smell, Mining Software Repositories, Empirical Study

I. INTRODUCTION

Code smells are identified as symptoms of possible code or design problems [1], which may potentially have a negative impact on software quality, such as maintainability [2], code readability [3], testability [4], and defect-proneness [5].

A large number of studies have focused on smell detection and removal techniques (e.g., [6], [7]) with many static analysis tools for smell detection; these include tools such as PMD¹, SonarQube², and Designite³. However, previous work [8], [9] indicated that the program context and domain are important in identifying smells. This makes it difficult for program analysis tools to correctly identify smells since contextual information

This work was partially funded by the National Key R&D Program of China with Grant No. 2018YFB1402800.

³ https://www.designite-tools.com

is rarely taken into account. Existing smell detection tools are also known to produce false positives [10], [11]; therefore, manual detection of smells could be considered more valuable than automatic approaches.

Code review is a process which aims to verify the quality of the software by detecting defects and other issues in the code, and to ensure that the code is readable, understandable, and maintainable. It has been linked to improved quality [12], reduced defects [13], reduced anti-patterns [14], and the identification of vulnerabilities [15]. Compared to smell detection static analysis tools, code reviews are usually performed by developers belonging to the same project [16], so it is possible that reviewers will take full account of contextual information and thus better identify code smells in the code.

However, little is known about the extent to which code smells are identified during code reviews, and whether developers (the code authors) take any action when a piece of code is deemed "smelly" by reviewers. Therefore, we set out to study the concept behind code smells identified in code reviews and track down actions taken after reviews were carried out. To this end, we mined code review discussions from the two most active OpenStack⁴ projects: Nova⁵ and Neutron⁶. We then conducted a quantitative and qualitative analysis to study how common it was for reviewers to identified code smells during code review, why the code smells were introduced, what actions they recommended for those smells, and how developers proceeded with those recommendations. In total, we analyzed 1,190 smell-related reviews got by manually checking 19,146 review comments to achieve our goal.

Our results suggest that: 1) code smells are not widely identified in modern code reviews, 2) following coding conventions can help reducing the introduction of code smells, 3) reviewers usually provide useful suggestions to help developers better fix the identified smells; while developers commonly accept reviewers' recommendations regarding the identified smells and they tend to refactor their code based on those recommendations, and 4) *review-based* detection of code smells is seen as a trustworthy mechanism by developers.

The paper is structured as follows: related work is presented

⁵ https://wiki.openstack.org/wiki/Nova

¹ https://pmd.github.io

² https://www.sonarqube.org

⁴ https://www.openstack.org

⁶ https://wiki.openstack.org/wiki/Neutron

in Section II, the study design and data extraction methods are explained in Section III, the results are presented in Section IV, followed by a discussion of the results in Section V, and the threats to the validity of the results are covered in Section VI, followed by conclusions and future work in Section VII.

II. RELATED WORK

A. Studies on Code Smells

A growing number of studies have investigated the impact of code smells on software quality, including defects [5], [17], maintenance [18], and program comprehension [3]. Other studies have looked at the impact of code smells on software quality using a group of developers working on a specific project [18]–[20].

Tufano *et al.* [21] mined version histories of 200 open source projects to study when code smells were introduced and the main reason behind their interaction. It was found that smells appeared in general as a result of maintenance and evolution activities. Sjôberg *et al.* [18] investigated the relationship between the presence of code smells and maintenance effort through a set of control experiments. Their study did not find significant evidence that the presence of smells led to increased maintenance effort. Previous studies also include work investigating the impact of different forms of smells on software quality, such as architectural smells [22], [23], test smells [4], [24], and spreadsheet smells [25].

A number of previous studies have investigated developer's perception of code smells and their impact in practice. A survey on developers' perception of code smells conducted by Palomba et al. [26] found that developer experience and system knowledge are critical factors for the identification of code smells. Yamashita and Moonen [8] reported that developers are moderately concerned about code smells in their code. A more recent study by Taibi et al. [27] replicated the two previous studies [8], [26] and found that the majority of developers always considered smells to be harmful; however, it was found that developers perceived smells as critical in theory, but not as much in practice. Tahir et al. [9] mined posts from Stack Exchange sites to explore how the topics of code smells and anti-patterns were discussed amongst developers. Their study found that developers widely used online forums to ask for general assessments of code smells or anti-patterns instead of asking for particular refactoring solutions.

B. Code Reviews in Software Development

Code review is an integral part in modern software development. In recent years, empirical studies on code reviews have investigated the potential code review factors that affect software quality. For example, McIntosh *et al.* [28] investigated the impact of code review coverage and participation on software quality in the Qt, VTK, and ITK projects. The authors used the incidence rates of post-release defects as an indicator and found that poorly reviewed code (e.g. with low review coverage and participation) had a negative impact on software quality. A study by Kemerer *et al.* [29] investigated the impact of review rate on software quality. The authors found that the *Personal Software Process* review rate was a significant factor affecting defect removal effectiveness, even after accounting for developer ability and other significant process variables.

Several studies [13], [28], [30] have investigated the impact of modern code review on software quality. Other studies have also investigated the impact of code reviews on different aspects of software quality, such as vulnerabilities [31], design decisions [32], anti-patterns [14], and code smells [33], [34]. Aziz and Apatta [33] examined review comments from code reviewers and described the need for an empirical analysis of the relationship between code smells and peer code review. Their preliminary analysis of review comments from OpenStack and WikiMedia projects indicated that code review processes identified a number of code smells. However, the study only provided preliminary results and did not investigate the causes or resolution strategies of these smells. A more recent study by Pascarella et al. [34] found that code reviews helped in reducing the severity of code smells in source code, but this was mainly a side effect to other changes unrelated to the smells themselves.

III. METHODOLOGY

A. Research Questions

The goal of this study is to investigate the concept behind code smells identified during code reviews and what actions are suggested by those reviewers and performed by developers in response to the identified smells. To achieve this goal, we formulated the following three research questions (RQs).

RQ1: Which code smells are the most frequently identified by code reviewers?

Rationale: This question aims to find out how frequent smells are identified by code reviewers and what particular code smells are repeatedly detected by reviewers. Such information can help in improving developers' awareness of these frequently identified code smells.

RQ2: What are the common causes for code smells that are identified during code reviews?

Rationale: This question investigates the main reasons behind the identified smells as explained by the reviewers or developers. When reviewing code, reviewers can express why they think the code under review may contain a smell. Developers can also reply to reviewers and explain how they introduced the smells. Understanding the common causes of smells identified manually by reviewers will shed some light on the effectiveness of manual detection of smells and help developers better understand the nature of identified smells and reduce such smells in the future.

RQ3: How do reviewers and developers treat the identified code smells?

Rationale: This question investigates the actions suggested by reviewers and those taken by developers on the identified smells. When a smell is identified, reviewers can provide suggestions to resolve the smell, and developers can then



Fig. 1: An overview of our data mining and analyzing process

decide on whether to fix or ignore it. This question can be further decomposed into three sub-research questions from the perspective of reviewers, developers, and the relationship between their actions:

RQ3.1: What actions do reviewers *suggest* to deal with the identified smells?

RQ3.2: What actions do developers *take* to resolve the identified smells?

RQ3.3: What is the *relationship* between the actions suggested by reviewers and those taken by developers?

B. OpenStack Projects and Gerrit Review Workflow

OpenStack is a set of software tools for building and managing cloud computing platforms. It is one of the largest open source communities. Based on most recent data, OpenStack projects contain around 13 million lines of code, contributed to by around 12k developers⁷. We deemed the platform to be appropriate for our analysis, since the community has long invested in its code review process⁸.

We then selected two of the most active OpenStack projects as our subject projects: Nova (a fabric controller) and Neutron (a network connectivity platform) - Table I provides an overview of the data obtained from the two projects. Both projects are written in Python, and use Gerrit⁹, a web-based code review platform built on top of Git. The Gerrit review workflow is explained next.

TABLE I: An overview of the subject projects (Nova and Neutron)

Project	Review Period	#Code Changes	#Comments
Nova	Jan 14 - Dec 18	22,762	156,882
Neutron	Jan 14 - Dec 18	15,256	152,429
Total		38,018	309,311

Gerrit provides a detailed code review workflow. First, a developer (author) makes a change to the code and submits

the code (patch) to the Gerrit server so that it can be reviewed. Then, verification bots check the code using static analysers and run automated tests. A reviewer (usually other developers that have not been involved in writing the code under review) will then conduct a formal review of the code and provide comments. The original author can reply to the reviewer's comments and action the required changes by producing a new revision of the patch. This process is repeated until the change is merged to the code base or abandoned by the author.

C. Mining Code Review Repositories

Fig. 1 outlines our data extraction and mining process. We mined code review data *via* the RESTful API provided by Gerrit, which returns the results in a JSON format. We used a Python script to automatically mine the review data in the studied period and store the data in a local database. In total, we mined 38,018 code changes and 309,311 review comments between Jan 2014 and Dec 2018 from the two projects.

D. Building the Keyword Set

To locate code review comments that include code smell discussions, we used several variations of terms referring to code smells or anti-patterns, including "code smell", "bad smell", "bad pattern", "anti-pattern", and "technical debt". In addition, considering that reviewers may point out the specific code smell by its name (e.g., dead code) rather than using generic terms, we included a list of code smell terms obtained from Tahir *et al.* [35], that extracted these smell terms from several relevant studies on this topic, including the first work on code smells by Fowler [1] and the systematic review by Zhang *et al.* [36]. The list of smell terms used in our study are shown in Table II.

Since the effectiveness of keyword-based mining approach relies on the set of keywords that are used in the search, we followed the systematic approach used by Bosu *et al.* [31] to identify the keywords included in our search. This includes the following steps¹⁰:

1) Build an initial keyword set.

10 implemented using the NLTK package: http://www.nltk.org

⁷ As of October 2020: https://www.openhub.net/p/openstack

⁸ https://docs.opendev.org/opendev/infra-manual/latest/developers.html

⁹ https://www.gerritcodereview.com

Code Smell Terms				
Accidental Complexity	Anti Singleton	Bad Naming	Blob Class	
Circular Dependency	Coding by Exception	Complex Class	Complex Conditionals	
Data Class	Data Clumps	Dead Code	Divergent Change	
Duplicated Code	Error Hiding	Feature Envy	Functional Decomposition	
God Class	God Method	Inappropriate Intimacy	Incomplete Library Class	
ISP Violation	Large Class	Lazy Class	Long Method	
Long Parameter List	Message Chain	Middle Man	Misplaced Class	
Parallel Inheritance Hierarchies	Primitive Obsession	Refused Bequest	Shotgun Surgery	
Similar Subclasses	Softcode	Spaghetti Code	Speculative Generality	
Suboptimal Information Hiding	Swiss Army Knife	Temporary Field	Use Deprecated Components	

TABLE II: Code smell terms included in our mining

- 2) Build a corpus by searching for review comments that contain at least one keyword of our initial keyword set (e.g., "dead" or "duplicated") in the code review data we collected in Section III-C.
- 3) Process the identified review comments which contain at least one keyword of our initial keyword set, and then apply the identifier splitting rules (i.e., "isDone" becomes "is Done" or "is_done" becomes "is done").
- 4) Create a list of tokens for each document in the corpus.
- 5) Clean the corpus by removing stopwords, punctuation, and numbers, and then convert all the words to lowercase.
- 6) Apply the Porter stemming algorithm [37] to obtain the stem of each token.
- 7) Create a Document-Term matrix [38] from the corpus.
- 8) Find the additional words co-occurred frequently with each of our initial keywords (co-occurrence probability of 0.05 in the same document).

After performing these eight steps, we found that no additional keywords co-occurred with each of our initial keywords, based on the co-occurrence probability of 0.05 in the same document. Therefore, we believe that our initial keyword set is sufficient to support the keyword-based mining method. Due to space constraints, we provide the initial set of keywords (which are the same as the final set of keywords) associated with code smells, in our replication package [39].

E. Identifying Smell-related Reviews in Keywords-searched **Review** Comments

We followed the following four steps to identify smellrelated reviews:

In step one, we developed a Python script to search for review comments that contained at least one of the keywords identified in Section III-D. The search returned a total of 18,082 review comments from the two projects.

In step two, to increase our verification process, two of the authors independently and manually analyzed the review comments obtained in step one to exclude comments *clearly* unrelated to code smells. If a review comment was deemed by both coders to be unrelated to code smells, it was then excluded. As a result of this step, the number of review comments was reduced to 3,666.

To illustrate this process, consider the following two re- 13 https://review.opendev.org/c/openstack/nova/+/100097 view comments that contain the keyword "dead". In the first ¹⁴ https://review.opendev.org/c/openstack/nova/+/91092

example, the reviewer commented that "why not to put the port on dead vlan first?"¹¹. Although this comment contains the keyword "dead", both coders thought that it was unrelated to code smells, and the comment was therefore excluded. In the second example, the reviewer commented "remove dead code"¹², which was regarded as related to "dead code" by the two coders and thus was included in the analysis.

In step three, two of the authors worked together to further manually analyze the remaining review comments. The same two authors carefully analyzed the contextual information of each review comment, including the code review discussions and associated source code to determine whether the code reviewers identified any smells in the review comments. We considered a comment to be related to code smell only when both coders agreed. The agreement between the two authors was calculated using Cohen's Kappa coefficient [40], which was 0.85. When the coders were unsure or disagreed, a third author was then involved in the discussion until an agreement was reached. This resulted in a reduction in the number of review comments to 1.235.

To better explain our selection process, consider the two examples in Fig. 2. In the top example¹³, the reviewer suggested adding another argument to the method to eliminate code duplication. Then the developer replied "Done", which implies an acknowledgment of the code duplication. We considered this as a clear smell-related review, and the review comment was retained for further analysis. In contrast, in the bottom example¹⁴, we observed that the comment was just used to explain the meaning of the "DRY" principle, but did not indicate that the code contained duplication according to the context. Thus, this comment was excluded from analysis.

Finally, in step four, we recorded the contextual information of each review comment in an external text file for further analysis, which contained: 1) a URL to the code change, 2) the type of the identified code smell, 3) the discussion between reviewers and developers, and 4) a URL to the source code. We ended up with a total of 1,174 smell-related reviews (we note that several review comments appearing in the same discussion were merged). An example of an

- 12 https://review.opendev.org/c/openstack/neutron/+/196893

¹¹ https://review.opendev.org/c/openstack/neutron/+/179314



Fig. 2: Review comments related to 'duplicated code': the top review is smell-related, while the bottom one is not.

extracted source file is shown below:

Code Change URL: http://alturl.com/2ne85
Code Smell: *Dead Code*Code Smell Discussions:

Reviewer: "Looks like copy-paste of above and, more importantly, dead code."

2) Developer: "yes, sorry for that."
Source Code URL: http://alturl.com/yai68

F. Identifying Smell-related Reviews in Randomly-selected Review Comments

Knowing that reviewers and developers may not use the same keywords as we used in Section III-E when detecting and discussing code smells during code review, we supplemented our keyword-based mining approach by including a randomly selected set of review comments from the rest of the review comments (291,229) that did not contain any of the keywords used in Section III-D. Based on 95% confidence level and 3% margin of error [41], we ended up with an additional 1,064 review comments. We then followed the same process of manual analysis (i.e., from **step two** to **step four** as described in Section III-E) to identify smell-related reviews in these randomly selected review comments. Finally, we identified a total of 16 smell-related reviews.

In addition to the reviews obtained by keywords search in Section III-E, we finally obtained a total of 1,190 smell-related reviews for further analysis. We provided a full replication package containing all the data, scripts, and results online [39].

G. Manual Analysis and Classification

For RQ1, in Sections III-E and III-F, we identified and recorded the smell type of each review when analyzing the review comments. When a reviewer used general terms (e.g., "smelly" or "anti-pattern") to describe the identified smell, we classified the type in these reviews as "general". The others were classified as specific smell (e.g., "duplicated code"). For RQ2, we adopted Thematic Analysis [42] to find the causes for the identified code smells in Sections III-E and III-F. We used MAXQDA¹⁵ - a software package for qualitative research - to code the contextual information of the identified code smells. Firstly, we coded the collected smellrelated reviews by highlighting sections of the text related to the causes of the code smell in the review. When no cause was found, we used "cause not provided/unknown". Next, we looked over all the code that we created to identify common patterns among them and generated themes. We then reviewed the generated themes by returning to the dataset and comparing our themes against it. Finally, we named and defined each theme. This process was performed by the same two coders in Sections III-E and III-F. A third author was involved in cases of disagreement by the two coders.

For RQ3, we decided to manually check the code reviews obtained in Sections III-E and III-F to identify the actions suggested by reviewers and taken by developers.

For RQ3.1, we categorized the actions recommended by reviewers into three categories, which are proposed in [35]:

- 1) **Fix**: recommendations are made to refactor the code smell.
- 2) **Capture**: detect that there may be a code smell, but no direct refactoring recommendations are given.
- 3) **Ignore**: recommendations are to ignore the identified smells.

For RQ3.2, we investigated how developers responded to reviewers that identified code smells in their code. We conducted this analysis in three steps: We first checked the developer's response to the reviewer in the discussion (Gerrit provides a discussion platform for both reviewers and developers). Second, we investigated the associated source code file(s) of the patch before the review was conducted, and the changes in the source code made after the review. Finally, if the developers neither responded to reviewers nor modified source code, we then checked the status (i.e., merged or abandoned) of the corresponding code change.

We considered the identified code smells to be solved in these three cases: 1) the original developer *self-admitted* a refactoring (as part of the review discussion), 2) changes were made in the source code file(s), and 3) the corresponding code change was then abandoned.

For RQ3.3, based on the results of RQ3.1 and RQ3.2, we categorized the relationship between the actions recommended by reviewers and those taken by developers into the following three categories:

- 1) A developer *agreed* with the reviewer's recommendations.
- 2) A developer *disagreed* with the reviewer's recommendations, or
- 3) A developer *did not respond* to the reviewer's comments.

These three categories were then mapped into two actions: 1) *fixed the smell* (i.e., refactoring was done) or 2) *ignored the*

¹⁵ https://www.maxqda.com/

change (i.e., no changes were performed to the source code with regard to the smell).

This process was conducted by the first author and the result of each step was cross-validated by another author. Again, a third author was involved in case of disagreement. In total, the manual analysis process took around thirty days of fulltime work of the coders. We also provided a full replication package containing all the data, scripts, and results from the manual analysis online [39].

IV. RESULTS

In this section, we present the results of our three RQs. We note that, due to space constraints, detailed results of our analysis are provided externally [39].

RQ1: Which code smells are the most frequently identified by code reviewers?

Figure 3 shows the distribution of code smells identified in the code reviews obtained in Sections III-E and III-F. In general, we identified 1,190 smell-related reviews. Compared to the number of all the review comments we obtained, we found that code smells are not commonly identified in code reviews. In addition, of all the code smells we identified, duplicated code is by far the most frequently identified smell by name, with exactly 620 instances. The smells of bad naming and dead code were also frequently identified, as they were discussed in 304 and 221 code reviews, respectively. There were 30 code reviews which identified long method, while other smells such as circular dependency and swiss army knife were discussed in only 4 code reviews. The rest of code reviews (11) used general terms (e.g., code smell) to describe the identified smells.



Fig. 3: Number of reviews for the identified code smells

RQ1: the most frequently identified smells in code reviews. Code smells are not widely identified in code reviews. Of the identified smells, *duplicated code*, *bad naming*, and *dead code* are the most frequently identified smells in code reviews.

RQ2: What are the common causes for code smells that are identified during code reviews?

For RQ2, we used Thematic Analysis to identify the com- ¹⁸ https://review.opendev.org/#/c/125384/ mon causes for the identified code smells as noted by code 19 https://review.opendev.org/#/c/181674/ reviewers or developers. We then identified five causes:

- Violation of coding conventions: certain violations of coding conventions (e.g. naming convention) cause the smell. (Example: "moreThanOneIp (CamelCase) is not our naming convention" ¹⁶).
- Lack of familiarity with existing code: developers introduced the smell due to unfamiliarity with the functionality or structure of the existing code. (Example: "this useless line because None will be returned by default" ¹⁷).
- Unintentional mistakes of developers: the developer forgets to fix the smell or introduces the smell by mistake. (Example: "You can see I renamed all of the other test methods and forgot about this one" ¹⁸).
- Improper design: the smell is identified to be related to • improper design of the code. (Example: "...If that's the case something is smelly (too coupled)..." ¹⁹).
- Detection by code analysis tools: the reviewer points out that the smell was detected by code analysis tools. (Example: "pass is considered as dead code by python coverage tool" ²⁰).



Fig. 4: Reasons for the identified smells

As demonstrated in Fig. 4, we found that the majority of reviews (70%) did not provide any explanation for the identified smells - in most cases, the reviewer(s) simply pointed out the problem, but did not provide any further reasoning for their decisions. 276 (23%) of the reviews indicate that violation of coding conventions is the main reason for the smell. For example, a reviewer suggested that the developer should adhere to the naming standard of 'test_[method under test]_[detail of what is being tested]', as shown below:

Link: http://alturl.com/urn8k

Reviewer: "Please adhere to the naming standard of 'test [method under test] [detail of what is being tested]' to ensure that future maintainers will have an easier time associating tests and the methods they target."

In addition, 40 (3%) of the reviews indicate that the smells were caused by developers' lack of familiarity with existing

16 https://review.opendev.org/#/c/147739/

- 17 https://review.opendev.org/#/c/147042/

20 https://review.opendev.org/#/c/143709/

code. An example of such a case is shown below. In this case, the reviewer pointed out that the exception handling should be removed. It could imply that the developer was not aware that the specific exception is not raised.

Link: http://alturl.com/ccjy3

Reviewer: "on block_device.BlockDeviceDict.from_api(), exception.InvalidBDMVolumeNotBootable does not raise. so it is necessary to remove the exception here."

Nineteen reviews attributed *unintentional mistakes of developers* (such as copy and paste) to be the cause of the smell, similar to the example shown below:

Link: http://alturl.com/zwz2x Reviewer: "I think you forgot to remove this." Developer: "Darn, yes bad copy / paste. Will fix it."

Eighteen reviews indicate that *improper design* was the cause for the identified smell. In the rest (6) of the reviews, reviewers would note that the smell was detected by code analysis tools. For example, a reviewer pointed out that the code 'pass' would be regarded as dead code by coverage tool.

Link: http://alturl.com/azt42 Reviewer: "you can remove 'pass', it's commonly considered as dead code by coverage tool"

RQ2: common causes for smells as identified during code reviews.

Taken overall, over half of the reviews did not provide an explanation of the cause of the smells. In terms of the formulated causes, *violation of coding conventions* is the main cause for the smells as noted by reviewers and developers.

RQ3: How do reviewers and developers treat the identified code smells?

RQ3.1: What actions do reviewers suggest to deal with the identified smells?

The results of this research question are shown in Table III. In the majority of reviews (870, representing 73% of the reviews), reviewers recommended *fix* for resolving the identified code smells. These fixes include either general directions (such as the name of a refactoring technique to be used) or specific actions (points to specific changes to the code base that could remove the smell). 303 (35%) of these fixes provided example code snippets to help developers better refactor the smells. Below is an example of a review that suggested a *fix* recommendation. In this example, the reviewer suggested removing duplicated code from a test case, and also provided a working example of how to apply "extract method" refactoring to define a new test method, so that it could be referenced from multiple methods to remove code duplication.

Link: http://alturl.com/c3g69

Reviewer: "I think you can do function that remove duplicated code, something like that following..."

```
def _compare(self, exp_real):
    for exp, real in exp_real:
        self.assertEqual(exp['count'],
        real.count)
        self.assertEqual(exp['alias_name
            '], real.alias_name)
        self.assertEqual(exp['spec'],
            real.spec)
```

272 reviews (23%) fell under the *capture* category. In those reviews, the reviewers just pointed to the presence of the smells, but did not provide any refactoring suggestions. In a small number of reviews (48, 4%), reviewers suggested ignoring the code smell found in the code review.

TABLE III: Actions recommended by reviewers to resolve smells in the code

Reviewer's recommendation		
Fix (without recommending any specific implementation)	567	
Fix (provided specific implementation)	303	
Capture (just noted the smell)	272	
Ignore (no side effects)	48	

RQ3.2: What actions do developers take to resolve the identified smells?

Table IV provides details of the number of reviews that identified code smells versus the number of fixes of the identified code smells. Of the 1,190 code smells identified in the reviews, the majority (1,029, representing 86%) were refactored by the developers after the review (i.e., changes were made to the patch). The remainder did not result in any changes in the code, indicating that the developers chose to ignore such recommendations. This could be a case where developers thought that those smells were not as harmful as suggested by the reviewers, or that there were other issues requiring more urgent attention, resulting in those smells being counted as technical debt in the code.

As per the results of RQ1, *duplicated code*, *bad naming*, and *dead code* were the most frequently identified smells by reviewers. Those smells were also widely resolved by developers. Over 508 (82%) *duplicated code*, 276 (91%) *bad naming*, and 210 (95%) *dead code* smell instances were refactored by developers after they were identified in the reviews. The proportion of other smells being fixed was nearly 78% (35/45). However, the sample size for these smells (35 instances) is still too small to make any generalisations.

Below is an example of a review with a recommendation by the reviewer to remove *dead code* in Line 132 of the original file (i.e., remove the pass statement); the developer then agreed to the reviewer's recommendation and deleted the unused code. Fig. 5 shows the code before review (5a) and after the action taken by the developer (5b).

Link: http://alturl.com/szswu Reviewer: "you can remove 'pass', it's commonly considered as dead code by coverage tool" Developer: "Done"

TABLE IV: Developers' actions to code smells identified during reviews

Code smell	#Reviews	#Fixed by developers	% of fixes
Duplicated Code	620	508	82%
Bad Naming	304	276	91%
Dead Code	221	210	95%
Long Method	30	25	83%
Circular Dependency	3	2	67%
Swiss Army Knife	1	1	100%
General Smell	11	7	64%
Total	1190	1029	86%

RQ3.3: What is the relationship between the actions suggested by reviewers and those taken by developers?

For answering this RQ, a map of reviewer recommendations and resulting developer actions is shown in Fig. 6. In 775 (65%) of the obtained reviews, developers agreed with the reviewers' suggestions and took exactly the same actions (either *fix* or *ignore*) as suggested by reviewers. Of those cases, there are 20 cases where developers agreed with reviewers on ignoring the smell (i.e., a smell has been identified, but the reviewer may think that the impact of the smell is minor). The example below shows a case where a reviewer pointed out that they could accept duplicated code if there was a reasonable justification and the developer gave their explanation and ignored the smell.

Link: http://alturl.com/s59so

Reviewer: "...I just don't like duplicated code but if there is a reasonable justification for this I can be sold cheaply and easily."

Developer: "we need create_vm here to support a lot of the other testing in this method. I agree it's duplicate code, but it's needed here too and this one is more complex that (sic) the test_config one...."

In 274 (23%) reviews, even when developers did not respond to reviewers directly in the review system, they still made the required changes to the source code files. We noted another 66 (5%) reviews where developers had different opinions from reviewers and decided to ignore the recommendations to refactor the code and remove the smell. In those cases, the developers themselves decided that the smells were either not as critical as perceived by the reviewers, or there are time or project constraints preventing them from implementing the changes, which is typically self-admitted technical debt [43]. An example review is shown below: **Link:** http://alturl.com/pzmzz

Reviewer: "This method has a lot duplicated code of '_apply_instance_name_template'. The differ in the use of 'index' and the CONF parameters. With a bit refactoring only one method would be necessary I guess."

Developer: "I thought to make / leave this separate in case one wants to configure the multi_instance_name_template different to that of single instance."

Similarly, there were also 75 (6%) reviews in which developers neither replied to reviewers nor modified the source code. For those cases, we assume that developers did not find the recommendations regarding how to deal with the specific smells in the code helpful, and therefore decided not to perform any changes. In all of those cases, no further explanation/reasons were provided by the developers on why they ignored these recommended changes.

RQ3: reviewers' recommendations and developers' actions.

In most reviews, reviewers provided fixing (refactoring) recommendations (e.g., in the form of code snippets) to help developers remove the identified smells. Developers generally followed those recommendations and performed the suggested refactoring operations, which then appeared in the patches committed after the review.

V. DISCUSSION

A. RQ1: The most frequently identified smells

In general, code smells are not commonly identified during code reviews. The results of RQ1 imply that *duplicated code*, *bad naming*, and *dead code* were, by far, the most frequently identified code smells in code reviews. The results regarding *duplicated code* are in line with previous findings which indicate that the smell is also frequently discussed among developers in online forums [9], and is also the smell that developers are most concerned about [8]. However, *dead code* and *bad naming* were not found to be ranked highly in previous studies [8].

The different results are due to the different context and domain, which are critical in identifying smells, as shown by previous studies [8], [9]. The results reported in these two previous studies [8], [9] are based on more generic investigation of code smells among online Q&A forums' users and developers. The context of some of these code smells was not fully taken into account, even if the developers may provide some specific scenario to explain their views. In contrast, our study is project-centric, and the context of the identified code smells during code reviews is known to reviewers and developers involved in identification and removal of the smells.

B. RQ2: The causes for identified smells

We identified five types of common causes for code smells in code reviews (RQ2). Among these, *violation of coding*

124 125 126 127 128 129 130 131	<pre>class AvailabilityZonePluginBase(object):</pre>		124 125 126 127 128 129 130 131	<pre>class AvailabilityZonePluginBase(object):</pre>
132	pass Cedric Brandily you can remove "pass", it's commonly considered as dead code by coverage tool Hirofumi Ichihara	Nov 24, 2015 ^ Nov 24, 2015 ^		
	Done			

(a) method before review.

(b) after change made by the developer.

Fig. 5: An example of a remove dead code operation after review (the change is highlighted in Line 132 (a))

Agreed with the Reviewer	No Response
Fixed the code 755	Fixed the code 274
	lgnored the change 75
Ignored the change	Disagreed with the Reviewer
20	Ignored the change
	66

Fig. 6: A treemap of the relationship between developers' actions in response to reviewers' recommendations regarding code smells identified in the code

conventions is the major cause of code smells identified in reviews. Coding conventions are important in reducing the cost of software maintenance while the existence of smells can increase this cost. We conjecture that this is because that developers may not be familiar with the coding conventions of their community and the system they implemented. For example, *duplicated code* and *dead code* may occur because developers are not aware of existing functionality, while *bad naming* may occur because developers are not familiar with the naming conventions. This reason can imply that developers can inadvertently violate coding conventions in their company or community, leading to smells or other problems. This may have a negative impact on software quality.

Another main observations is that more than half of reviewers (in review comments where they indicated that there was a code smell) simply pointed out the smell in the code, but did not provide any further explanation of why they considered that as a smell. One explanation for this is that the identified smells were simple or self-explanatory (e.g., *duplicated code*, *dead code*). Therefore, it is not expected that reviewers need to provide further explanation for these smells. Although the point of code review is to identify shortcomings (e.g., code smells) in the contributed code, understanding the causes of code smells can help practitioners understand how the code smell is introduced, and then take corresponding measures.

C. RQ3: The relationship between what reviewers suggest and the actions taken by developers

The results of RQ3 show that reviewers usually provide useful recommendations (sometimes in the form of code snippets) when they identify smells in the code and developers usually follow these suggestions. Given the constructive nature of most reviews, developers tend to agree with the review-based smell detection mechanism (i.e., where a reviewer detects and reports a smell) and in most cases they perform the recommended actions (i.e., refactoring their code) to remove the smell. We believe that this is because reviewers can take the contextual information into full account as the program context and domain are important in identifying smells [8], [9], [44].

Although not as frequent, there are cases where changes recommended by reviewers were ignored (see Figure 6). This situation is partially due to the different understanding of reviewers and developers about the severity of identified code smells, i.e., when a reviewer identifies a code smell to be resolved, a developer may not agree that this code smell must be fixed, such as technical debt [45].

D. Implications

First, although we built the initial set of keywords with 5 general code smell terms and 40 specific code smell terms, most of the smells were not identified in code reviews, such as *long parameter list, temporary field*, and *lazy class*. One potential reason is that code smells which are considered as problematic in academic research may not be considered as a pressing problem in industry. More research should be conducted with practitioners to explore existing code smells and to understand the driving force behind industry efforts on code smell detection and elimination. This will further help guide the design of next-generation code smell detection tools.

Second, *violation of coding conventions* is the main cause of code smells identified in code reviews. It implies that developers' lack of familiarity with the coding conventions in their company or organization can have a significantly negative impact on software quality. To reduce code smells, project leaders not only need to adopt code analysis tools, but also need to help and educate their developers to become familiar with the coding conventions adopted in the system.

Third, in smell-related reviews, reviewers usually give useful suggestions to help developers better fix the identified code smells and developers generally tend to accept those suggestions. It implies that review-based detection of smells is seen as a trustworthy mechanism by developers. In general, code reviews are useful for finding defects and locating code smells. Although code analysis tools (both static analyzers and dynamic (coverage-based) tools) are able to find some of those smells, their large outputs restrict their usefulness. Most tools are context and domain-insensitive, making their results less useful due to the potential false positives [10]. Context seems to matter in deciding whether a smell is bad or not [9], [11]. There have been some recent attempts to develop smelldetection tools that take developers-context into account [44], [46]. Still, other contextual factors such as project structure and developer experience are much harder to capture with tools. Code reviewers are much better positioned to understand and account for those contextual factors (as they are involved in the project) and therefore their assessment of smells might be trusted more by developers than automated detection tools. To increase reliability, it may be that we need a two-step detection mechanism; static analysis tools to identify smells (as they are faster than human assessment and also scalable) and then for reviewers to go through those smell instances. They should decide, based on the additional contextual factors, which of those smells should be removed and at what cost. The problem with such an approach is that most tools would probably produce large sets of outputs, making it impractical for reviewers working on a large code base.

VI. THREATS TO VALIDITY

External Validity: our study considered two major projects from the OpenStack community (Nova and Neutron), since those projects have invested a significant effort in their code review process (see Section III-B). Due to our sample size, our study may not be generalizable to other systems. However, we believe that our findings could help researchers and developers understand the importance of the manual detection of code smells better. Including code review discussions from other communities will supplement our findings, and this may lead to more general conclusions.

Internal Validity: the main threat to internal validity is related to the quality of the selected projects. It is possible that the projects we included do not provide a good representation of the types of code smells we included in our study. While we only selected two projects from the OpenStack community with Gerrit as their code review tool, OpenStack investment in code review processes and commitment to perform code review to their entire code base and following coding best practices make it a good candidate for our analysis.

Construct Validity: a large part of the study depends on manual analysis of the data, which could affect the construct validity due to personal oversight and bias. In order to reduce its impact, each step in the manual analysis (i.e., identifying smell-related reviews and classification) was conducted by at least two authors, and results were always cross-validated. The selection of the keywords used to identify the reviews which contain smell discussions is another threat to construct validity since reviewers and developers may use terms other than those that we used in our mining query. To minimize the impact of this threat, we first combined a list of code smell terms that developers and researchers frequently used, as reported in several previous studies. Then, we identified the keywords by following the systematic approach used by Bosu *et al.* [31] to minimize the impact of missing keywords due to misspelling or other textual issues. Moreover, we randomly selected a collection of review comments that did not contain any of our keywords to supplement our approach, reducing the threat to the construct validity.

Reliability: before starting our full scale study, we conducted a pilot run to check the suitability of the data source. Besides, the execution of all the steps in our study, including the mining process, data filtering, and manual analysis, was discussed and confirmed by at least two of the authors.

VII. CONCLUSIONS

Code review is a common software quality assurance practice. One of the issues that may impact software quality is the presence of code smells. Yet, little is known about the extent to which code smells are identified and resolved during code reviews. To this end, we performed an empirical study of code smell discussions in code reviews by collecting and analyzing code review comments from the two most active OpenStack projects (Nova and Neutron). Our results show that: 1) code smells are not commonly identified in code reviews, and when identified, duplicated code, bad naming, and dead code are, by far, the most frequently identified smells; 2) violation of coding conventions is the most common cause for smells as identified during code reviews; 3) when smells are identified, most reviewers provide recommendations to help developers fix the code and remove the smells (via specific refactoring operations or through an example code snippet); and 4) developers mostly agree with reviewers and remove the identified smells through the suggested refactoring operations.

Our results suggest that: 1) developers should follow the coding conventions in their projects to reduce code smell incidents; and 2) code smell detection via code reviews is seen as a trustworthy approach by developers (given their constructive nature) and smell-removal recommendations made by reviewers appear more actionable by developers. We found that the majority of smell-related recommendations were accepted by developers. We believe this is mainly due to the context-sensitivity of the reviewer-centric smell detection approach.

We plan to extend this work by studying code reviews in a larger set of projects from different communities. We also plan to explore, in more detail, the refactoring actions taken by developers when removing certain smells, and the reasons (e.g., trade-off in managing technical debt [45]) why developers disagreed with reviewers' recommendations or ignored the recommended changes by e.g., code smell discussions [47].

REFERENCES

- M. Fowler, *Refactoring: Improving the Design of Existing Code*, 2nd ed. Addison-Wesley Professional, 2018.
- [2] F. Palomba, G. Bavota, M. Di Penta, F. Fasano, R. Oliveto, and A. De Lucia, "On the diffuseness and the impact on maintainability of code smells: A large scale empirical investigation," in *Proceedings* of the 40th International Conference on Software Engineering (ICSE). ACM, 2018, pp. 1188–1221.
- [3] M. Abbes, F. Khomh, Y.-G. Gueheneuc, and G. Antoniol, "An empirical study of the impact of two antipatterns blob and spaghetti code on program comprehension," in *Proceedings of the 15th European Conference* on Software Maintenance and Reengineering (CSMR). IEEE, 2011, pp. 181–190.
- [4] A. Tahir, S. Counsell, and S. G. MacDonell, "An empirical study into the relationship between class features and test smells," in *Proceedings of the* 23rd Asia-Pacific Software Engineering Conference (APSEC). IEEE, 2016, pp. 137–144.
- [5] F. Khomh, M. Di Penta, and Y.-G. Gueheneuc, "An exploratory study of the impact of code smells on software change-proneness," in *Proceedings of the 16th Working Conference on Reverse Engineering (WCRE)*. IEEE, 2009, pp. 75–84.
- [6] N. Tsantalis and A. Chatzigeorgiou, "Identification of move method refactoring opportunities," *IEEE Transactions on Software Engineering*, vol. 35, no. 3, pp. 347–367, 2009.
- [7] N. Moha, Y.-G. Gueheneuc, L. Duchien, and A.-F. Le Meur, "Decor: A method for the specification and detection of code and design smells," *IEEE Transactions on Software Engineering*, vol. 36, no. 1, pp. 20–36, 2009.
- [8] A. Yamashita and L. Moonen, "Do developers care about code smells? an exploratory survey," in *Proceedings of the 20th Working Conference* on Reverse Engineering (WCRE). IEEE, 2013, pp. 242–251.
- [9] A. Tahir, J. Dietrich, S. Counsell, S. Licorish, and A. Yamashita, "A large scale study on how developers discuss code smells and anti-pattern in stack exchange sites," *Information and Software Technology*, vol. 125, 2020.
- [10] F. A. Fontana, J. Dietrich, B. Walter, A. Yamashita, and M. Zanoni, "Anti-pattern and code smell false positives: Preliminary conceptualisation and classification," in *Proceedings of the 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. IEEE, 2016, pp. 609–613.
- [11] T. Sharma and D. Spinellis, "A survey on software smells," *Journal of Systems and Software*, vol. 138, pp. 158–173, 2018.
- [12] R. A. Baker Jr, "Code reviews enhance software quality," in *Proceedings* of the 19th International Conference on Software Engineering (ICSE). ACM, 1997, pp. 570–571.
- [13] S. McIntosh, Y. Kamei, B. Adams, and A. E. Hassan, "An empirical study of the impact of modern code review practices on software quality," *Empirical Software Engineering*, vol. 21, no. 5, pp. 2146–2189, 2016.
- [14] R. Morales, S. McIntosh, and F. Khomh, "Do code review practices impact design quality? a case study of the Qt, VTK, and ITK projects," in *Proceedings of the 22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering (SANER).* IEEE, 2015, pp. 171–180.
- [15] A. Meneely, A. C. R. Tejeda, B. Spates, S. Trudeau, D. Neuberger, K. Whitlock, C. Ketant, and K. Davis, "An empirical investigation of socio-technical code review metrics and security vulnerabilities," in *Proceedings of the 6th International Workshop on Social Software Engineering (SSE)*. ACM, 2014, pp. 37–44.
- [16] S. McConnell, Code Complete. Pearson Education, 2004.
- [17] T. Hall, M. Zhang, D. Bowes, and Y. Sun, "Some code smells have a significant but small effect on faults," ACM Transactions on Software Engineering and Methodology, vol. 23, no. 4, pp. 1–39, 2014.
- [18] D. I. Sjøberg, A. Yamashita, B. C. Anda, A. Mockus, and T. Dybå, "Quantifying the effect of code smells on maintenance effort," *IEEE Transactions on Software Engineering*, vol. 39, no. 8, pp. 1144–1156, 2013.
- [19] F. Palomba, G. Bavota, M. Di Penta, R. Oliveto, D. Poshyvanyk, and A. De Lucia, "Mining version histories for detecting code smells," *IEEE Transactions on Software Engineering*, vol. 41, no. 5, pp. 462–489, 2015.
- [20] Z. Soh, A. Yamashita, F. Khomh, and Y.-G. Guéhéneuc, "Do code smells impact the effort of different maintenance programming activities?" in *Proceedings of the 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER).* IEEE, 2016, pp. 393–402.

- [21] 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 *Proceedings of the IEEE/ACM 37th IEEE International Conference on Software Engineering (ICSE)*, vol. 1. IEEE, 2015, pp. 403–414.
- [22] J. Garcia, D. Popescu, G. Edwards, and N. Medvidović, "Identifying architectural bad smells," in *Proceedings of the 13th European Conference* on Software Maintenance and Reengineering (CSMR). IEEE, 2009, pp. 255–258.
- [23] A. Martini, F. A. Fontana, A. Biaggi, and R. Roveda, "Identifying and prioritizing architectural debt through architectural smells: A case study in a large software company," in *Proceedings of the 12th European Conference on Software Architecture (ECSA)*. Springer, 2018, pp. 320– 335.
- [24] G. Bavota, A. Qusef, R. Oliveto, A. De Lucia, and D. Binkley, "Are test smells really harmful? an empirical study," *Empirical Software Engineering*, vol. 20, no. 4, pp. 1052–1094, 2015.
- [25] W. Dou, S.-C. Cheung, and J. Wei, "Is spreadsheet ambiguity harmful? detecting and repairing spreadsheet smells due to ambiguous computation," in *Proceedings of the 36th International Conference on Software Engineering (ICSE)*. ACM, 2014, pp. 848–858.
- [26] F. Palomba, G. Bavota, M. Di Penta, R. Oliveto, and A. De Lucia, "Do they really smell bad? a study on developers' perception of bad code smells," in *Proceedings of the 30th International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2014, pp. 101–110.
- [27] D. Taibi, A. Janes, and V. Lenarduzzi, "How developers perceive smells in source code: A replicated study," *Information and Software Technology*, vol. 92, pp. 223–235, 2017.
- [28] S. McIntosh, Y. Kamei, B. Adams, and A. E. Hassan, "The impact of code review coverage and code review participation on software quality: A case study of the Qt, VTK, and ITK projects," in *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR)*. ACM, 2014, p. 192–201.
- [29] C. F. Kemerer and M. C. Paulk, "The impact of design and code reviews on software quality: An empirical study based on psp data," *IEEE Transactions on Software Engineering*, vol. 35, no. 4, pp. 534– 550, 2009.
- [30] O. Kononenko, O. Baysal, L. Guerrouj, Y. Cao, and M. W. Godfrey, "Investigating code review quality: Do people and participation matter?" in 2015 IEEE International Conference on Software Maintenance and Evolution (ICSME), 2015, pp. 111–120.
- [31] A. Bosu, J. C. Carver, M. Hafiz, P. Hilley, and D. Janni, "Identifying the characteristics of vulnerable code changes: An empirical study," in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE)*. ACM, 2014, p. 257–268.
- [32] F. E. Zanaty, T. Hirao, S. McIntosh, A. Ihara, and K. Matsumoto, "An empirical study of design discussions in code review," in *Proceedings* of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM). ACM, 2018, pp. 1–10.
- [33] A. Nanthaamornphong and A. Chaisutanon, "Empirical evaluation of code smells in open source projects: preliminary results," in *Proceedings* of the 1st International Workshop on Software Refactoring (IWoR). ACM, 2016, pp. 5–8.
- [34] L. Pascarella, D. Spadini, F. Palomba, and A. Bacchelli, "On the effect of code review on code smells," in *Proceedings of the 27th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER).* IEEE, 2020.
- [35] A. Tahir, A. Yamashita, S. Licorish, J. Dietrich, and S. Counsell, "Can you tell me if it smells? a study on how developers discuss code smells and anti-patterns in stack overflow," in *Proceedings of the 22nd International Conference on Evaluation and Assessment in Software Engineering (EASE)*. ACM, 2018, pp. 68–78.
- [36] M. Zhang, T. Hall, and N. Baddoo, "Code bad smells: a review of current knowledge," *Journal of Software Maintenance and Evolution: research and practice*, vol. 23, no. 3, pp. 179–202, 2011.
- [37] M. F. Porter, "Snowball: A language for stemming algorithms," Open Source Initiative Osi, 2001.
- [38] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to data mining*. Pearson Education India, 2016.
- [39] X. Han, A. Tahir, P. Liang, S. Counsell, and Y. Luo, "Replication package for the paper understanding code smell detection via code review: A study of the openstack community," Jan. 2021. [Online]. Available: https://doi.org/10.5281/zenodo.4468035
- [40] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and Psychological Measurement*, vol. 20, no. 1, pp. 37–46, 1960.

- [41] G. D. Israel, "Determining sample size," Florida Cooperative Extension Service, Institute of Food and Agricultural Sciences, University of Florida, Florida, U.S.A, Fact Sheet PEOD-6, November 1992.
- [42] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006.
- [43] A. Potdar and E. Shihab, "An exploratory study on self-admitted technical debt," in 2014 IEEE International Conference on Software Maintenance and Evolution. IEEE, 2014, pp. 91–100.
- [44] N. Sae-Lim, S. Hayashi, and M. Saeki, "Context-based approach to prioritize code smells for prefactoring," *Journal of Software: Evolution* and Process, vol. 30, no. 6, pp. 1–24, 2018.
- [45] Z. Li, P. Avgeriou, and P. Liang, "A systematic mapping study on

technical debt and its management," Journal of Systems and Software, vol. 101, pp. 193–220, 2015.

- [46] F. Pecorelli, F. Palomba, F. Khomh, and A. De Lucia, "Developerdriven code smell prioritization," in *Proceedings of the 17th Working Conference on Mining Software Repositories (MSR)*. ACM, 2020, pp. 220–231.
- [47] S. Shcherban, P. Liang, A. Tahir, and X. Li, "Automatic identification of code smell discussions on stack overflow: A preliminary investigation," in *Proceedings of the 14th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*. ACM, 2020, pp. 1–6.