University of Groningen

Do practitioners intentionally repay their own Technical Debt and why?

Tan, Jie; Feitosa, Daniel; Avgeriou, Paris

Published in:
2021 IEEE International Conference on Software Maintenance and Evolution

DOI:
[10.1109/ICSME52107.2021.00029](https://doi.org/10.1109/ICSME52107.2021.00029)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Early version, also known as pre-print

Publication date:
2021

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Tan, J., Feitosa, D., & Avgeriou, P. (2021). Do practitioners intentionally repay their own Technical Debt and why? In *2021 IEEE International Conference on Software Maintenance and Evolution* (pp. 251-262) <https://doi.org/10.1109/ICSME52107.2021.00029>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Do practitioners intentionally self-fix Technical Debt and why?

Jie Tan, Daniel Feitosa, Paris Avgeriou
University of Groningen, The Netherlands
{j.tan, d.feitosa, p.avgeriou}@rug.nl

Abstract—The impact of Technical Debt (TD) on software maintenance and evolution is of great concern, but recent evidence shows that a considerable amount of TD is fixed by the same developers who introduced it; this is termed *self-fixed TD*. This characteristic of TD management can potentially impact team dynamics and practices in managing TD. However, the initial evidence is based on low-level source code analysis; this casts some doubt whether practitioners repay their own debt intentionally and under what circumstances. To address this gap, we conducted an online survey on 17 well-known Java and Python open-source software communities to investigate practitioners’ intent and rationale for self-fixing technical debt. We also investigate the relationship between human-related factors (e.g., experience) and self-fixing. The results, derived from the responses of 181 participants, show that a majority addresses their own debt consciously and often. Moreover, those with a higher level of involvement (e.g., more experience in the project and number of contributions) tend to be more concerned about self-fixing TD. We also learned that the sense of responsibility is a common self-fixing driver and that decisions to fix TD are not superficial but consider balancing costs and benefits, among other factors. The findings in this paper can lead to improving TD prevention and management strategies.

Index Terms—Self-fixed technical debt, Survey, Human factors

I. INTRODUCTION

The term *technical debt* (TD) is commonly used to express the adoption of suboptimal solutions in software development, that impact maintainability and evolvability [1], [2]. To help practitioners in managing TD, researchers have investigated potential causes and effects of TD [3]–[5], its various types (e.g., Code [6], [7], Design [8], [9], and Test Debt [10], [11]) and the role of tools [12], [13]. The community has also invested effort into studying specific forms of TD such as *self-admitted TD*, i.e., when developers acknowledge introducing TD themselves (e.g., in source code comments) [14], [15]. More recently, there is work on *self-fixed TD*, i.e., when developers repay TD that they introduced themselves [16].

Although investigating self-fixed TD is still rather preliminary, initial findings are noteworthy. For example, Tan et al. [16] found that more than two-thirds of TD items in source code are self-fixed, and that Defect Debt tends to be self-fixed frequently, while Test Debt and Design Debt are likely to be repaid by other developers [16]. However, current research on this topic, similarly to most work on TD, is based on source code analysis and not on the opinion of software developers. Consequently, there is no evidence of whether developers consider the differences between various debt types when

repaying their own debt and what else they take into account. We advocate that, in order to comprehend the phenomenon of self-fixed TD more thoroughly, it is paramount to consider the practitioners’ intent and experiences.

To address this shortcoming, this paper reports on a study aiming at examining the phenomenon of self-fixed TD from the perspective of practitioners. To this end, we conducted a survey with 17 well-known Java and Python open-source software communities, receiving 181 valid responses (89 for Java and 92 for Python). The survey included questions on the practitioners’ attitude towards introducing and repaying TD in general and regarding five types of debt: Code Debt, Defect Debt, Design Debt, Documentation Debt and Test Debt. These five types have been extensively investigated in other studies [17]–[19], and were also the focus of a related work on self-fixed TD [16]. The survey also looks into practitioners’ reasons to introduce and self-fix TD, as well as human factors that may influence their responses (e.g., experience and contribution levels).

This study offers a number of contributions. First, it provides insights on the relevance of self-fixed TD to practitioners and the recurrence with which they address their own debt. Second, it elaborates on the circumstances that lead to self-fixing, also providing a list of reasons (both technical and non-technical) and highlighting how these associate with reasons to introduce TD. These contributions can benefit both researchers and practitioners. On the one hand, researchers can build on this evidence to fine-tune TD-related tools and further study the phenomenon of self-fixed TD. On the other hand, practitioners can gain a broader understanding of the intent for introducing and self-fixing TD, which can help them devise or refine TD prevention and management strategies.

The remaining sections are organized as follows. The paper starts by summarizing related work and contextualizing our study in Section II. Next, Section III reports on the goal, research questions and methodology of the study. Then, the results are presented in Section IV, while Section V elaborates on their interpretation and implications to researchers and practitioners. Section VI reports on the threats to the validity of our study, leading to Section VII where we conclude the paper and outline directions for future work.

II. RELATED WORK

Recent studies have shown that a large percentage of TD is paid back during software evolution [17], [18], [20], [21]. Tan

et al. [16], established that a considerable amount of TD items are repaid by the same developers who introduced them. More recently, a number of studies focused on self-admitted technical debt (SATD): TD items that are introduced intentionally and explicitly documented in code comments [19], [22], [23]. Some related findings indicate that the majority of SATD is removed by the same developer who introduced it [14], [19], [24]. However, all these previous studies drew conclusions only through source code analysis (including source code [17] and code comments [24]) and focused only on one language (i.e., either Python [16] or Java [14]). In our study, we conducted a survey to investigate the practitioners' perspective on how they self-fix TD for both Java and Python.

Other studies have focused on the rationale in TD management, e.g., the reasons to introduce or (not) repay TD. For example, Rios et al. [3] reported that non-technical reasons, e.g., *deadline* and *inappropriate planning*, play a significant role in the occurrence of TD items. Some studies also pointed out that the leading causes of technical debt are architectural choices [25] and that developers are frequently forced to introduce new TD due to already existing TD [26]. In addition, Maldonado et al. [14] found that developers self-admit TD to track potential future bugs, code that needs improvements or areas to implement new features, while they mostly remove SATD when they are fixing bugs or adding new features. Some SATD may also be accidentally removed when entire classes or methods are dropped [27]. Researchers have also investigated the reasons for not paying TD back, e.g., *lack of organizational interest*, *low priority on the debt*, *focus on short-term goals*, *cost*, and *lack of time* [28]. In contrast to these studies, we focus on the reasons influencing developers' decisions to deal with their own debt, i.e., why they introduce and self-fix TD.

There is also work on the influence of human factors on incurring and managing TD. Some findings indicate that most developers that incur TD have low project-related experience [29] while developers' seniority and commit frequency are negatively correlated with the amount of introduced TD [30]. In addition, Besker et al. [5] found that some organizational factors, such as developers' experience and startup founders' knowledge, can influence the intentional accumulation of TD. However, another study found that the developers' participation level (i.e., the total lines of code edited by a developer) and their experience in the project positively correlate with the amount of TD introduced by them, while communication skills have barely any impact on TD [31]. In contrast to these studies, we focused on remediation of TD rather than its introduction, and in the context of self-fixed TD.

III. RESEARCH APPROACH

The goal of our study, described according to the Goal-Question-Metrics (GQM) approach [32], is to “*analyze practitioners' attitudes for the purpose of investigating self-fixed TD with respect to the intent and rationale for remediating it and the influence of human factors in doing so from the point*

of view of software practitioners in the context of open-source software development”. In the following subsections, we describe the research questions, the survey design, participants and projects selection, as well as data analysis in detail.

A. Research questions

We further refine the goal of this study into the following research questions:

RQ₁: Do practitioners intentionally self-fix TD?

This research question aims at investigating to what extent developers repay their own debt intentionally. More specifically, we analyze practitioners' opinions to examine which types of debt they self-fix intentionally. The results can shed further light on how developers prioritize the remediation of technical debt. For example, developers may introduce specific types of debt with the intention of repaying it soon thereafter, thus prioritizing this item above others.

RQ₂: What reasons motivate practitioners to introduce and self-fix TD?

This research question aims at identifying practitioners' rationale for introducing and subsequently repaying their own debt. Analyzing the possible reasons can help to manage debt (in the case that TD is unavoidable) or prevent it (if possible). An example of management is to build tools that account for and highlight the conditions under which developers are more likely to repay a debt item. An example of prevention is to identify those reasons for introducing debt, that can be eliminated depending on a particular team's culture and composition, e.g., by making a more fine-grained planning or avoiding overloading the team.

RQ₃: How do human factors influence practitioners in self-fixing TD?

Practitioners in different roles, with different experience and contribution levels may have different attitudes towards self-fixing TD. For example, developers with a leading role tend to monitor and manage the development tasks more often than others [33], [34], which may reflect in a stronger attitude towards repaying their own debt. Also, practitioners with more experience are more likely to repay technical debt [19]. Thus, investigating the relationship between characteristics of practitioners and how they self-fix TD can help comprehend and subsequently support their decision-making processes when managing TD. For example, team composition can be adjusted to foster TD remediation by taking into account the diversity of roles, experiences, and cultures. In addition, researchers can also use the results to offer different tool options for practitioners with different characteristics.

B. Survey Design

According to our goal and the guidelines for selecting empirical methods in software engineering research [35], we decided to use a survey. A survey is a comprehensive research method for describing, comparing, or explaining knowledge and behavior [36]; it thus suits our focus on collecting and explaining practitioners' perceptions on self-fixed TD. More specifically, we decided to use a web-based questionnaire

to conduct the survey because it is time- and cost-effective, and also suitable for collecting data from a large number of developers in geographically diverse locations [37].

The flow of the survey design is presented in Fig. 1. The black box represents the introductory welcome page, and the numbered white boxes depict steps in which participants answer questions. The data collection was planned to last two months and the questionnaire was made accessible online between February and March 2021. To mitigate privacy concerns, all invited participants were informed about the survey with an invitation text describing the target open-source project and programming language (see Section III-C), together with the motivation and goal of this study. Moreover, we did not collect any identifiable information in the survey. The templates used for communication are available in the replication package¹.

Fig. 1. Design of the survey in our study.

As shown in Fig. 1, the survey started with a welcome page explaining the structure of the questionnaire, and that the survey is anonymous. On step ①, the survey starts with a form with demographic questions. In particular, we inquired about the participants' role in the target software project (e.g., developer, tester), their educational background, the number of years they have been working in software development, the number of years they have been working on the project, and the number of commits they have contributed to the project. Moreover, since this survey targeted a broad audience, we wanted to assess how balanced the population is and, thus, we also asked for the participants' sex and country (as optional questions).

Subsequently, the survey included five general questions (see Table I) related to technical debt in step ②. But before asking any questions, we start by describing three terms to set a baseline for common understanding: what *technical debt* is, and what it means to *introduce* it and *repay* it. We then ask GQ1 to identify whether the participants are familiar with the aforementioned descriptions of these terms. Next, to investigate whether practitioners are aware of introducing and repaying their TD, we record the frequency (from *always* to *never*), with which practitioners introduced (GQ2) and self-fixed (GQ4) TD. We also ask participants the reasons why they have introduced (GQ3) and self-fixed (GQ5) TD. The options for GQ3 stem from the results of a previous study [38], which reported the top ten most cited TD causes as informed by 107 participants. GQ5 is an open-ended question as there was no established knowledge in the literature on reasons for self-fixing TD. Finally, for GQ3, in addition to the listed options, we add the option *Other* with a text field for participants to provide more details.

¹<https://github.com/sftdsoftware/icsme>

TABLE I
TD GENERAL QUESTIONS OF THE SURVEY

No.	TD General Question Description
GQ1	Were you familiar with the presented description of technical debt before this questionnaire?
GQ2	How often do you introduce technical debt?
GQ3	For what reasons have you introduced technical debt?
GQ4	How often do you repay your own technical debt?
GQ5	For what reasons do you repay your own technical debt?

The questions in step ③ covers all five types of debt. For each debt type, we provide the definition and ask the participants about their recollection of (DQ1) *noticing* and (DQ2) *self-fixing* TD items. For that, we use a Likert scale with the options *definitely not*, *probably not*, *I am not sure*, *probably yes*, and *definitely yes*.

C. Projects and Participants Selection

For this study, we selected projects that are written in either Java or Python, which are ranked as the top two most popular object-oriented programming languages². Although both of them are strongly typed languages, Java is statically typed while Python is dynamically typed. Because of the dynamic features, Python projects are more change-prone, resulting in developers spending more effort on software maintenance [39] and software quality improvement [40]. Thus, selecting systems written in these two programming languages, helps to increase the external validity of our study. Moreover, investigating both languages can help to compare the findings to the previous results that focus on source code analysis [16], [18]. In the following, we describe the process of selecting projects and participants for the survey.

We started by compiling a list of systems that have been widely used in the studies that investigated technical debt through source code analysis (see Section II). From this list, we filtered out projects with less than 5,000 commits or 200 contributors, to ensure that a sufficient number of practitioners are familiar with the maintenance of the projects for this survey. Ultimately, the following eleven projects were selected: Camel [14], [15], Cassandra [41], Chromium [41], [42], Elasticsearch [43], Gerrit [14], [15], Hadoop [44], [45], HBase [46], Guava [47], [48], Jenkins [47], [48], RxJava [15], [47], TensorFlow [24]. Considering that, except for TensorFlow, these are Java projects, we still needed to supplement the list with pertinent Python projects. In addition to complying with the same thresholds for the number of commits and contributions, we chose the Python projects with the most stars on GitHub: Ansible, Django, Flask, Pandas, Scikit-Learn, and Superset. The number of stars of a GitHub repository can be seen as a proxy of its popularity [49] as more stars indicate that more users are interested in the project [50].

²According to the Tiobe Index, which is one of the best-known indexes of programming languages popularity: <https://www.tiobe.com/tiobe-index/>, visited March 2021

TABLE II
MAPPING BETWEEN SURVEY QUESTIONS, VARIABLES AND RESEARCH QUESTIONS

No.	Variables	RQs
BQ1 – BQ7	demographic information, i.e., the practitioners’ role, educational background, sex, country, contribution, experience in the project and software development	RQ ₃
GQ1	practitioner’s familiarity with technical debt	RQ ₁
GQ2	the frequency of introducing technical debt	RQ ₂
GQ3	the reasons why developers have introduced TD	RQ ₂
GQ4	the frequency of self-fixing technical debt	RQ _{2,3}
GQ5	the reasons why practitioners have self-fixed TD	RQ ₂
DQ1	the types of debt that practitioners have noticed	RQ ₁
DQ2	the types of debt that practitioners have self-fixed	RQ _{1,3}

To obtain the target sample of participants, we extracted each contributor’s name, email address, and the time of each commit from GitHub. Next, for each project, we consider the contributors who had submitted at least two commits within the last three years; this helps to target contributors who still have a recollection of their work in the project and whose email address is likely to still be in use, thus improving the response rate and the response quality. To mitigate privacy concerns, we did not retain any information after sending invitations. In total, we sent out 7638 emails.

D. Data Analysis

Table II presents the mapping between the questions in the survey, the collected variables and research questions. The three sets of survey questions correspond to the three steps in the survey design (i.e., step ① – ③ in Fig. 1): background (demographic) questions (BQ), general questions about technical debt (GQ), and detailed questions about TD types (DQ). In the following, we describe in detail how each of the research questions makes use of the responses.

To answer RQ₁, we first investigate to what extent developers are familiar with the description of technical debt from the responses of GQ1, using descriptive statistics. Next, we answer to what extent developers self-fix technical debt intentionally per TD type; for that, we use the responses of DQ2 from step ③ of the survey (i.e., *Do you recall self-fixing technical debt* (per type)?). Since the responses to this question use a five-point Likert scale (i.e., from *definitely not* to *definitely yes*—see Section III-B), the results are presented in the form of stacked bar charts and comparisons are assessed using statistical tests. We also make a supplementary discussion in conjunction with the types of debt that participants have noticed in source code (DQ1), since identifying TD items may contribute to one’s knowledge of TD.

To answer RQ₂, i.e., the reasons why practitioners introduce and self-fix TD, we first investigate the extent to which they are aware of introducing and self-fixing debt (GQ2 and GQ4, respectively). For that, we create and analyze a heat map comparing various frequencies of TD introduction to various frequencies of self-fixing it (see Section III-B) and evaluate the association between them using several statistical tests. Then, we qualitatively discuss the reasons provided by

participants to explain why they decide to introduce (GQ3) and self-fix technical debt (GQ5). Since GQ3 provides the option *Other* (with a text field) and GQ5 is an open-ended question, we followed an open-coding approach [51], [52] to analyze the answers collected from the survey respondents; a similar approach was followed in many recent software engineering studies, e.g., [53], [54]. To conduct open coding, the first two authors independently assigned one or more tags to each answer; each tag is meant to summarize a reason. The level of inter-rater agreement between the classifications of the two authors was measured and conflicts were solved by a discussion with the third author. Since participants can provide multiple reasons to introduce and self-fix TD, the associations between reasons are presented in the form of a chord diagram and investigated using descriptive statistics.

For RQ₃, to explore how the human factors (characteristics of the practitioners such as roles, experience and contribution) affect their self-fixing behavior, we investigate the relationship between the demographic information of practitioners (step ① in Fig. 1) and their perspective of self-fixing TD (step ②③ in Fig. 1). For that, we first calculate, for each demographic group, the proportions of different ratings for the frequency (GQ4) and certainty (DQ2) of self-fixing TD. Then, we conducted several statistical tests to explore whether one demographic group tends to rate the frequency or certainty higher or lower than other groups.

IV. RESULTS

Overall, we obtained 184 responses to our survey and the response rate is 2.41%. Among them, 181 responses were valid: 89 for Java and 92 for Python. The three invalid responses lack specific information as, e.g., respondents choose “Other” and answered “NA” for several questions. The distribution of the population is presented in Table III. In the following, we present the results of the study organized by research question.

A. Do practitioners intentionally self-fix technical debt? (RQ₁)

To answer this research question, we first investigated to what extent developers are familiar with the concept of technical debt (GQ1 in Table I). Fig. 2 depicts the answers through a diverging stacked bar chart to visualize the percentages that describe the distribution of responses. The left side shows the answers of the 89 Java participants (49%), while the answers of the 92 Python participants (51%) are presented on the right.

Fig. 2. Respondents’ knowledge level of technical debt

According to Fig. 2, the vast majority (i.e., 92%) of the participants were familiar or somewhat familiar with technical debt before this survey. This increases our confidence that the data we collect is (mostly) derived from knowledgeable

TABLE III
CHARACTERISTICS OF THE RESPONDENTS

Characteristics	Java	Python	Characteristics	Java	Python
Main role in the project			Experience in the project		
Developer*	76.40%	69.56%	< 1 year	8.99%	33.70%
Software Architect	8.99%	14.13%	1–2 years	26.97%	36.96%
Team Leader	6.74%	4.35%	3–5 years	32.58%	15.22%
Product Owner	3.37%	5.43%	> 5 years	31.46%	14.13%
Tester**	1.12%	2.17%	Contribution to the project		
System Administrator	1.12%	1.09%	< 10 commits	11.24%	43.48%
DevOps Engineer	1.12%	1.09%	10–50 commits	12.36%	14.13%
Engineering Manager	1.12%		50–100 commits	10.11%	14.13%
Project Manager		2.17%	100–200 commits	12.36%	6.52%
Educational background			200–500 commits	14.61%	7.61%
Master degree	48.31%	30.43%	> 500 commits	39.33%	14.13%
Bachelor degree	39.33%	36.96%	Software development experience		
Ph.D. degree	8.99%	20.65%	< 5 year	12.36%	21.74%
High school	1.12%	4.35%	5–10 years	16.85%	26.09%
Other / not shared	2.25%	7.61%	10–20 years	44.94%	28.26%
Geographical location			20–35 years	23.60%	18.48%
Europe	52.81%	29.35%	> 35 years	1.12%	4.35%
North America	21.35%	39.13%	Not shared	1.12%	1.09%
Asia	10.11%	23.91%	Respondent's sex		
South America	5.62%	2.17%	Male	87.64%	90.22%
Oceania	2.25%	1.09%	Female	3.37%	6.52%
Not shared	7.87%	4.35%	Other	1.12%	
			Not answered	7.87%	3.26%

*Developer, incl. Software Engineer

**Tester, incl. Quality Assurance

individuals (a thorough discussion of threats to validity is found in Section VI). In addition, Java participants seem slightly more familiar than Python participants. To further evaluate the significance of such difference, we calculated the Wilcoxon Rank Sum test [55] (p -value = 0.02) and Cliff's Delta Effect Size [56] (δ = 0.1468) on the ratings of the Java and Python respondents, since we cannot assume the populations are normally distributed. The results reveal that the difference between the familiarity of Java and Python participants with TD is significant (i.e., p -value < 0.05) but the effect size is negligible (i.e., $|\delta|$ < 0.147). Thus, different levels of familiarity with the concept of TD between Java and Python participants have almost no effect on the disparity between the populations in subsequent findings.

Fig. 3. Distribution of self-fixed TD types

Delving into the essence of the research question, we looked at whether participants recalled self-fixing each type of debt (DQ2). While incurring TD can be both intentional and unintentional, self-fixing in this case is intentional: if practitioners recall paying back their own TD items, it is safe

to assume they did it on purpose (at least for those items). The answers are summarized in Fig. 3. This figure shows that, for each type of TD, more than half of the participants mentioned that they *definitely* or *probably* have self-fixed items. To investigate the significance of this observation, we first combined the answers into three categories: (1) positive (*definitely* or *probably* self-fixed TD); (2) neutral (*I am not sure*); and (3) negative (*definitely not* or *probably not* self-fixed TD). Next, we conducted a Scott-Knott Effect Size Difference (ESD) test [57] to group the three categories of answers into statistically distinct ranks based on the combined proportions of all debt types. As a variant of the Scott-Knott test [58], the Scott-Knott ESD test evaluates a non-normally distributed dataset and merges any two statistically distinct groups that have a negligible effect size into a single group [57]. The result shows that the three categories belong to different groups and, thus, the percentage differences observed are significant. Since the percentage of positive answers (category #1) is considerably greater than the other two, it implies that **developers often display a positive attitude towards self-fixing TD**.

Looking into individual deb types, participants seem to be slightly more concerned about self-fixing Code Debt issues, less about Test Debt, and mostly uncertain about Design Debt. To assess the significance of these differences, we conducted a Scott-Knott ESD test to group TD types into statistically distinct ranks based on the Likert scale. The result shows that with the exception of Code Debt, there is no significant distinction between debt types. Consequently, **as Code Debt displays the highest positive score, it seems to be the main concern of the participants**.

Comparing the languages, we observed from Fig. 3 that Java participants seem to be more likely to acknowledge that they have *definitely* or *probably* self-fixed TD, as more than half of them mentioned so for each debt type. To assess this particular observation, we conducted two Scott-Knott ESD tests (for Java and Python populations respectively) to group TD types into statistically distinct ranks based on their Likert scores. The result shows that there are three types of TD (i.e., Code Debt, Test Debt, and Defect Debt) belonging to the same group in Java, while Code Debt alone makes up one group in Python. **Although Code Debt is reaffirmed as an important concern in both languages, Test Debt and Defect Debt are as relevant as Code Debt only to the Java participants**.

B. What reasons motivate practitioners to introduce and self-fix TD? (RQ2)

To answer this research question, we first investigated the extent to which practitioners are introducing and self-fixing TD. Fig. 4 shows a heat map depicting the percentages of participants who self-fixed TD at different frequencies for a given frequency of TD introduction in Java and Python, respectively. A Wilcoxon Signed Rank test [59] reveals no significant difference between Java and Python (p -value = 0.29), suggesting that Java and Python participants introduce and self-fix TD in a similar fashion.

Fig. 4. Percentage of participants indicating the frequency of self-fixing TD (column) for a given frequency of TD introduction (row).

Fig. 5. Reasons to introduce technical debt

Analyzing the population as a whole, almost half of the participants (48%) have acknowledged that they self-fix TD *on a regular basis* and a half of those (27% of the general population) acknowledge also introducing debt sometimes only. This combination seems more recurrent than any other. To assess the significance of this observation, we conducted a Scott-Knott ESD test to group the frequencies of self-fixing into statistically distinct ranks based on introduction frequencies. The result shows participants that self-fix TD *on a regular basis* belong in a distinct group, and participants that do it *sometimes* in another distinct group. The results indicate that **although participants sometimes have to introduce TD, when they do, they also tend to pay it back.**

In the following, we examine the reasons that lead practitioners to introduce and self-fix TD. Regarding the former, 37 of all the 181 participants also selected the option *Other* and provided custom opinions. We applied an open-coding approach and extracted one or more reasons from each custom answer that could either fit into one of the ten pre-defined reasons or become a new reason. By the end of the coding, five new reasons had been created. Fig. 5 shows the 15 reasons

that can lead to TD introduction from the participants' perspective. The overall frequency of each reason is the number of participants who mentioned it. Apart from the reasons listed in the Figure, 20 participants selected that they *are not aware of introducing TD*. Another four participants indicated that *TD introduction is inevitable*, e.g., "all code is technical debt" and "TD introduction is a part of software development"; one of them mentioned also the reason *non-adoption of good practice*.

To facilitate the interpretation of the reasons to incur TD, we grouped them into more general categories, which we adopted from Rios et al. [3]. The categories are also shown in Fig. 5, from which we observe that *deadline* is cited by more than half of the participants, indicating that it is a factor that normally contributes to the appearance of TD. During the open-coding, some custom answers were coded as *deadline*, e.g., "it is evaluated and decided to be worth incurring this technical debt to finish the task earlier". Furthermore, *lack of experience*, *inappropriate planning*, and *team overload* are other reasons cited by at least 30% of the participants. Among them, *lack of experience* also includes custom answers such as "lack of learning" and "not fully understanding how the project will evolve over time".

The five new reasons that appeared during the open-coding are: *workload balance* referring to the trade-off between the cost and benefit of introducing TD, e.g., "prioritize mainline tasks and de-prioritize secondary tasks", and "increasing tech debt is desirable"; *not well-understood technology or feature*, e.g., "experiments to introduce new technologies" and "a feature might not be well understood and it is more cost effective to do it in a suboptimal way"; *debt inheritance* referring to new debt introduced due to a pre-existing TD item, e.g., "working in an area that is already burdened with technical debt"; *backward compatibility*, e.g., "backward compact and minor version release rules" and "backportability"; and *lack of refactoring*, e.g., "ongoing major refactoring".

Regarding the reasons to self-fix TD, we summarized 25 reasons from the responses (all open-ended in this case), as shown in Fig. 6. Similarly to Fig. 5, we also mapped the reasons into more general categories. For that, we adopted the categories provided by Rios et al. [3] again, but renamed some categories to fit the context. During the open-coding, answers were often mapped into multiple reasons. For example, one participant answered: "improving code quality, maintainability, making new features easier to implement". This answer was coded as *improve quality*, *improve maintainability*, and *prepare for new features*.

After examining the reasons for self-fixing TD, we noticed that *sense of responsibility* and *avoid higher cost* are the most frequently mentioned, each being cited by at least 37 participants (approx. 20%). For example, one participant mentioned "I feel ownership over the code, also, better to get rid of it before it bites you", which was coded as both *sense of responsibility* and *avoid higher cost*.

Most participants mentioned multiple reasons when answering why they introduce and self-fix TD. Thus, to gain a better understanding on the relationship between these rea-

Fig. 6. Reasons to self-fix technical debt

Fig. 7. Association between reasons

sons, we calculated the number of times that every pair of categories was mentioned together (i.e., totally 1536 relations), and plotted a chord diagram [60] illustrating the relations between categories of reasons (cf. Fig. 7). The categories for introducing and self-fixing TD are colored red and blue, respectively, and the width of the links conveys the strength of the relation.

From Fig. 7, we can observe that among all the relations, the reasons to introduce TD (in red) seem to display the strongest relationships with themselves. In particular, the reasons related to *planning and management* and *methodology* were most often mentioned simultaneously (127 times), followed by *planning and management* and *people* (121 times). Regarding

the reasons for self-fixing TD, it is noticeable that the links between the self-fixing and introduction categories are wider than those between self-fixing reasons themselves. In other words, the reasons for self-fixing TD were mentioned more often together with those to introduce TD than with themselves (629 vs. 112 times). Especially the reasons to self-fix TD related to *development* were often mentioned when the reasons to introduce TD concerned *planning and management* (108 times).

Altogether, we observe both proactive reasons (e.g., *sense of responsibility*, *avoid higher cost*) and reactive reasons (e.g., *TD overload*, *Repayment request*) to self-fix TD. More importantly, the evidence suggests that **regardless of the reason that lead a practitioner to introduce TD, she is keen to repay the debt proactively** not only due to an internal driver (e.g., *sense of responsibility* or *improve collaboration*) but also for strategic reasons (e.g., *avoid higher debt* or *improve maintainability*). Further evidence that repayment is thoughtfully considered, is that most reasons to self-fix TD co-occur with reasons to introduce TD related to *planning and management*.

C. How do human factors influence practitioners in self-fixing TD? (RQ₃)

To answer this research question, we followed a process similar to that used by Wan et al. [61] and Zou et al. [62] to analyze 11 demographic groups derived from the *team roles*, *development experience*, *project experience* and *contribution to project*. For comparison reasons, we also consider an additional group (All) including all participants. We explain the demographic groups in the following paragraphs.

Regarding the *team roles*, one group corresponds to developers (Dev), including software engineers. Another group consists of respondents who are involved in the project or team management and support (PM), i.e., software architect, team leader, product owner, system administrator, DevOps engineer, engineering manager and project manager. In addition to these two groups, we also received three responses from testers. However, we did not consider testers as a separate group because the sample size does not fit the threshold for statistical analysis [63], which would also increase the randomness of the results [64].

Regarding the levels of *software development experience*, we distinguish between respondents with: (a) high experience (ExpSdHigh), as the 25% with the most experience in years [62], i.e., ≥ 18 years in this survey; low experience (ExpSdLow), as the 25% with the least experience in years, i.e., ≤ 5 years; and medium experience (ExpSdMed), i.e., the remaining respondents. For *project experience*, since the survey offered discrete options, we selected the maximum and minimum interval that approximates the top and bottom 25%. In particular, we defined a high experience group (ExpProHigh) with those that have five or more years in a project (approx. 23%), a low experience group (ExpProLow) with those that joined a project less than one year ago (approx. 22%), and a medium experience group

(ExpProMed) with the remaining respondents. The *contributions* to projects are grouped in a similar fashion, as follows: high contribution level (ContribHigh), with participants that submitted 500 commits or more (approx. 27%); low contribution level (ContribLow), with those that submitted up to 10 commits (approx. 28%); and medium contribution level (ContribMed), with the remaining respondents.

Fig. 8 presents the frequencies of self-fixing TD as perceived by respondents in the eleven demographic groups. As shown in the figure, the majority of participants in all of the 12 groups (79%–97%) mentioned they repay their TD *always*, *on a regular basis*, or *sometimes* (when it is absolutely essential). In particular, about half of the participants in the various groups (44%–54%) considered they self-fix TD *on a regular basis*.

Fig. 8. The percentages of how often participants in each demographic group have self-fixed TD

Fig. 9 depicts the certainty of respondents in each demographic group about having self-fixed all five TD types (combined); this is in contrast to Fig. 8 which reflects frequency, independently of TD types. For example, in the group All, out of a total of 905 ratings, 36.5%, 21.9%, 15.0%, 11.4%, and 15.2% respectively were *definitely yes*, *probably yes*, *I am not sure*, *probably not*, and *definitely not*. As shown in the figure, more than half of the participants in most groups mentioned that they have self-fixed TD (43%–80%). However, these percentages are slightly lower than the percentages of the participants who mentioned that they have self-fixed TD *always*, *on a regular basis*, or *sometimes* (i.e., 79%–97% in Fig. 8). This discrepancy is due to the fact that we focused on individual types of TD when asking about the certainty level; thus, the participants might only self-fix specific types.

To investigate whether one demographic group tends to self-fix TD more or less frequently than other groups, we conducted pairwise Fisher's exact tests [65] with Bonferroni-Holm corrections [66] on the proportions of respondents from the different demographic groups. In detail, those statistical tests were conducted within four sets of demographic groups, i.e., a set of different roles (Dev versus PM), a set of different levels of software development experience (ExpSdHigh versus ExpSdMed versus ExpSdLow), a set of different contribution levels (ContribHigh versus ContribMed versus ContribLow) and a set of different levels of project experience (ExpProHigh versus ExpProMed ver-

Fig. 9. The percentages of self-fixing certainty levels for participants in each demographic group

sus ExpProLow). Moreover, for the demographic groups in each group set, we conducted eight times pairwise Fisher's exact tests with Bonferroni-Holm corrections to validate the significance of rating differences among these groups. They are the tests on the proportions of different frequencies ratings, together with the attitudes towards self-fixing TD (i.e., *always*, *on a regular basis*, and *sometimes*, when it is absolutely essential) and the negative ones (i.e., *almost never* and *never*). Furthermore, similar to RQ₁, we also conducted the tests twice: one to compare *definitely yes* and *probably yes* between the demographic groups; and another to compare *definitely not* and *probably not*. The results per set, are as follows:

1) *Roles*: We found that 85% of the participants in the Dev group and 96% of those in the PM group demonstrate a rather proactive stance towards self-fixing TD, i.e., they repay their own TD at least when necessary. The result of a Fisher's exact test shows that the difference observed between the two groups is significant (p -value = 0.01). Although a similar disparity can be observed in Fig. 9, a Fisher's exact test shows it is not substantial. Altogether, the findings suggest that **people in both types of roles engage in self-fixing TD, although those in project management and support seem to do it more often.**

2) *Software development experience*: After applying pairwise Fisher's exact tests with Bonferroni-Holm correction, we found no statistically significant difference between the three groups, i.e., **the level of development experience seems to have no effect on how participants repay their own TD.**

3) *Contributions to the projects*: For the groups with various levels of contribution, we found that 98% of the participants with high contribution level (ContribHigh) have a positive attitude towards self-fixing TD, and that 90% of them mentioned that they have *definitely* or *probably* self-fixed TD. These percentages were higher than the proportions for practitioners with medium (ContribMed) and low (ContribLow) contribution levels (87%/55% and 80%/43% respectively). Similarly, we conducted Fisher's exact tests for the three possible pairwise comparisons (ContribHigh versus ContribMed, ContribHigh versus ContribLow, and ContribMed versus ContribLow). After applying the Bonferroni-Holm correction, we found that participants with high contribution level (ContribHigh) differ significantly

from the other two groups (ContribMed and ContribLow) in all of the comparisons (p -values < 0.05).

These findings suggest that **practitioners with high contribution level display a significantly more positive attitude towards self-fixing TD**. In addition, more than two-fifths of the participants with low contribution level (42%) and about one-third of the participants with medium contribution level (27%) mentioned that they have not self-fixed TD in the projects. As this difference is significant (p -value = 0.04), we identified further evidence that **the level of contribution reflects the self-fixing attitude, both positively and negatively**.

4) *Project experience*: For groups with different levels of experience in the projects, we found that 93% of participants with higher experience (ExpProHigh) consider fixing their own debt at least when necessary. The percentage for those with medium and lower experience (ExpProMed and ExpProLow) is 89% and 79%, respectively. The results of the Fisher’s exact tests with Bonferroni-Holm correction suggest a significant difference between participants with high and low experience (p -value = 0.02), and no significance between these two groups and ExpProMed (medium experience). This finding is further supported by the differences observed in Fig. 9, which are all statistically significant. Altogether, these results suggest that **participants who have been involved in the projects for a longer period also seem to be more keen to repay their own TD**.

V. DISCUSSION

A. Interpretation of results

To the best of our knowledge, this is the first study to inquire practitioners on how they repay their own debt. Moreover, we argue that this work is complementary to previous studies that used source code analysis. Thus, we start our discussion with the results that can be compared more directly.

One important finding is that, although most Java and Python practitioners acknowledge that they have self-fixed TD, Java participants were slightly more certain. This may be due to the fact that, among the five debt types, 76%–84% of Java participants acknowledged that they have noticed similar debt, while the proportions are 53%–71% for Python participants. We also found that although all participants are likely to acknowledge that they have self-fixed Code Debt, there are some differences between Java and Python practitioners. For example, the former are similarly concerned about Test Debt, while Python participants have self-fixed Test Debt the least. The results are consistent with the findings obtained from source code level, i.e., Test Debt is commonly self-fixed in Java [24], while Test Debt appears to show the lowest self-fixing rate among the five debt types in Python [16]. We also found that only 4% of Java participants mentioned that they have not noticed Test Debt, while the percentage of Python participants is over 26%. One possible explanation may be related to the population of the study: the percentage of Java participants with higher experience and contribution level is about twice that of Python participants; since repaying test

smells is time-consuming [67], it may be more commonly performed by experienced developers.

Regarding the reasons to introduce and self-fix TD, one can notice that non-technical reasons seem to be the most cited ones. Almost three-quarters of the reasons for introducing TD are related to *planning and management*, *lack of knowledge*, *people* and *organizational*, while most reasons to self-fix TD are related to *planning and management*, *people* and *knowledge*. The observations related to introducing TD are consistent with the literature [3]. There is no literature to compare the self-fixing results, but we do observe that these reasons are quite similar to those for introducing TD.

Despite the dominance of non-technical reasons, we found that technical reasons were frequently mentioned together with non-technical ones. For example, the most recurrent relations involve reasons to introduce TD related to *methodology* and *planning and management*. The latter (non-technical) reason is also the most often associated with reasons to self-fix TD related to *development*. This intertwining between technical and non-technical reasons reflects the multi-faceted decision-making process related to repaying one’s own debt. Moreover, we found that practitioners might introduce TD because of *debt inheritance*, which is consistent with a related finding that developers are frequently forced to introduce new TD due to already existing TD [26] (“the poor are getting poorer”).

Furthermore, the most cited reason to self-fix TD is the practitioners’ *sense of responsibility*, which indicates the relevance of further investigating the impact of human factors on self-fixing TD. The other two most cited reasons for self-fixing TD are to *avoid higher cost* and *improve maintainability*, which indicate a consideration of the trade-off between cost and benefit in order to find the appropriate moment to self-fix TD. This result aligns with related findings that mentioned *focus on short-term goals* and *cost* as reasons when participants are considering whether or not to repay TD [28].

Finally, the results also suggest that the practitioners’ level of involvement (i.e., experience in the project and contribution level) can be related to how they deal with their own TD. This observation may be at least partially explained by a sense of code or project ownership: those with higher levels of involvement were significantly more positive towards self-fixing. Furthermore, practitioners in a management role were also found to display a more positive attitude. To assess the impact of practitioners’ responsibility on repaying their own TD, we divided the participants into two groups based on whether or not they mentioned *sense of responsibility* as a reason to self-fix TD. The results show that 95% and 86% of the participants that did and did not mention responsibility, respectively, have self-fixed TD at least when necessary, although this difference is not statistically significant. In addition, all of the participants that mentioned *sense of responsibility* acknowledged that they have *definitely* or *probably* self-fixed TD, compared to only half of the other participants. We assessed this disparity with a Fisher’s exact test and confirmed it to be significant. Thus, the sense of responsibility is an important driver to self-fix TD and should be properly motivated and rewarded.

B. Implications to researchers and practitioners

The results suggest that Java and Python practitioners have different attitudes towards self-fixing different types of TD. Thus, researchers can develop tools to prioritize TD remediation by assigning different weights to different TD types and giving suggestions to developers. For example, Test Debt and Defect Debt can be prioritized along with Code Debt for Java projects. In addition, our findings advance the state of the art as they expand the currently-known list of reasons to introduce TD and provide a list of reasons to repay one's own TD. Broader knowledge of such reasons can lead to more informed decisions when designing future studies on TD. For example, the detailed reasons to introduce and self-fix TD can help researchers in devising and evaluating prevention strategies for the most cited reasons for introducing TD and more timely management strategies for the most cited reasons to self-fix TD. One concrete suggestion is to encourage admission of TD (especially incurred under pressure) and explicitly allocate time to assess the current risk of existing TD.

We learned that practitioners are often willing to address their debt, displaying a sense of responsibility and collaboration. Also, it seems that repayment decisions are not made light-heartedly (e.g., due to pressure or at every single chance) but by weighing various factors, such as time, expected quality and confidence of efficacy. Thus, team leaders should encourage their team to admit TD, discuss it and document the corresponding items in their backlog. We also hope that practitioners gain a broader view of possible reasons to introduce and self-fix TD and how common they are; this may help them to make more informed decisions when managing TD and improving software quality. For example, team leaders can strengthen team communication and promote collaboration, which can boost self-fixing as we observed that “*improving collaboration*” is a recurrent reason to self-fix TD. Finally, software engineers should be given incentives to fix their own TD as a means of increasing their technical knowledge, which was also observed as a reason to self-fix TD (*experience and knowledge increased*).

VI. THREATS TO VALIDITY

Construct validity regards the connection between the research questions and the study objects. Although the survey participants are at the center of the study design, the questionnaire responses are the medium through which we capture their perception. Thus, the study's findings depend on how the participants interpret the questions and how we interpret the responses. To mitigate potential bias in interpretation, we provided definitions of the necessary TD-related concepts. Also, the level of clarity in answers to open-ended questions (e.g., reasons to self-fix TD) confirms to some extent that the participants understood the underlying concepts.

Furthermore, we used log-transformations in Scott-Knott Effect Size Difference (ESD) test to mitigate the skewness of data distribution; as pointed by Tantithamthavorn et al. [57], different results might be obtained if other transformations were used (e.g., Blom [68] or Box-Cox [69] transformations).

This limitation is to some extent addressed by repeating each test multiple times to assess if the distribution of bias and variance are approximately normally distributed (as per the Central Limit Theorem) [70].

External validity concerns threats to the generalizability of our findings. Given the population of our study, one cannot expect the results to apply to every development team. However, we believe that our findings can be extrapolated to a considerable portion of the Java and Python open-source communities, since our study included participants from popular open-source projects in various domains and with diverse team sizes. While Python and Java are the two most popular programming languages, we cannot generalize our results beyond these languages.

Reliability considers the bias from the researchers in data collection or data analysis. To mitigate bias from the interpretation of open-ended responses (i.e., reasons), we applied open coding in two steps to categorize the replies. Specifically, the first two authors manually classified different reasons of introducing and self-fixing TD. To measure the level of agreement between the classifications of independent researchers, we estimated the inter-rater agreement using Krippendorff's alpha [71] to be 0.88³, which was calculated based on the classification of all reasons. In the conflicting cases, the first and second authors discussed with the third author until they achieved consensus. Finally, to support the replication and reproduction of this study, we created a replication package⁴ with all the necessary data and scripts to run the analyses.

VII. CONCLUSIONS

This paper reports an empirical study investigating the phenomenon of self-fixed technical debt, i.e., when developers repay TD items that they have introduced themselves. In particular, we surveyed 181 practitioners spread among 17 projects (written in Java or Python) to explore participants' perception of this phenomenon and under what circumstances they self-fix TD.

The results show that most practitioners display a positive attitude towards self-fixing TD, and seem mostly attentive to Code Debt. Moreover, we found that non-technical reasons (e.g., related to planning and management) are the most considered both when introducing and self-fixing TD, and are often mentioned together with technical reasons (e.g., related to methodology). In particular, many developers mention a sense of responsibility as a factor for self-fixing, and that decisions to repay are not made easily but by balancing costs and benefits, among other factors.

In the future, we plan to investigate how the reasons portrayed by practitioners can be (semi-)automatically identified through software artefacts such as code, documentation and messages in issue trackers. We also plan to develop a process and associated tooling to incorporate TD self-fixing in TD-related decisions.

³Krippendorff's inspection of the tradeoffs between statistical techniques establishes that it is customary to require $\alpha \geq 0.80$ [72].

⁴<https://github.com/sftdsoftware/icsme>

REFERENCES

- [1] Z. Li, P. Avgeriou, and P. Liang, "A systematic mapping study on technical debt and its management," *Journal of Systems and Software*, vol. 101, no. C, pp. 193–220, Mar 2015.
- [2] P. Kruchten, R. L. Nord, and I. Ozkaya, "Technical debt: From metaphor to theory and practice," *Ieee software*, vol. 29, no. 6, pp. 18–21, 2012.
- [3] N. Rios, R. O. Spínola, M. Mendonça, and C. Seaman, "The practitioners' point of view on the concept of technical debt and its causes and consequences: a design for a global family of industrial surveys and its first results from brazil," *Empirical Software Engineering*, vol. 25, no. 5, pp. 3216–3287, 2020.
- [4] W. N. Behutiye, P. Rodríguez, M. Oivo, and A. Tosun, "Analyzing the concept of technical debt in the context of agile software development: A systematic literature review," *Information and Software Technology*, vol. 82, pp. 139–158, 2017.
- [5] T. Besker, A. Martini, R. E. Lokuge, K. Blincoe, and J. Bosch, "Embracing technical debt, from a startup company perspective," in *Proceedings of the 34th IEEE International Conference on Software Maintenance and Evolution (ICSME '18)*. Madrid, Spain: IEEE, 2018, pp. 415–425.
- [6] 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," *Empirical Software Engineering*, vol. 23, no. 3, pp. 1188–1221, 2018.
- [7] D. Di Nucci, F. Palomba, D. A. Tamburri, A. Serebrenik, and A. De Lucia, "Detecting code smells using machine learning techniques: are we there yet?" in *Proceedings of the IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER '18)*. Campobasso, Italy: IEEE, 2018, pp. 612–621.
- [8] M. Abidi, M. S. Rahman, M. Openja, and F. Khomh, "Are multi-language design smells fault-prone? an empirical study," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 30, no. 3, pp. 1–56, 2021.
- [9] T. Sharma, P. Singh, and D. Spinellis, "An empirical investigation on the relationship between design and architecture smells," *Empirical Software Engineering*, vol. 25, no. 5, pp. 4020–4068, 2020.
- [10] D. Spadini, F. Palomba, A. Zaidman, M. Bruntink, and A. Bacchelli, "On the relation of test smells to software code quality," in *Proceedings of the 34th IEEE International Conference on Software Maintenance and Evolution (ICSME '18)*. Madrid, Spain: IEEE, 2018, pp. 1–12.
- [11] J. De Bleser, D. Di Nucci, and C. De Roover, "Assessing diffusion and perception of test smells in scala projects," in *Proceedings of the IEEE/ACM 16th International Conference on Mining Software Repositories (MSR '19)*. Montreal, QC, Canada: IEEE, 2019, pp. 457–467.
- [12] P. Avgeriou, D. Taibi, A. Ampatzoglou, F. A. Fontana, T. Besker, A. Chatzigeorgiou, V. Lenarduzzi, A. Martini, N. Moschou, I. Pigazzini *et al.*, "An overview and comparison of technical debt measurement tools," *IEEE Software*, 2020.
- [13] V. Lenarduzzi, T. Besker, D. Taibi, A. Martini, and F. A. Fontana, "A systematic literature review on technical debt prioritization: Strategies, processes, factors, and tools," *Journal of Systems and Software*, vol. 171, p. 110827, 2021.
- [14] E. D. S. Maldonado, R. Abdalkareem, E. Shihab, and A. Serebrenik, "An empirical study on the removal of self-admitted technical debt," in *Proceedings of the 33th IEEE International Conference on Software Maintenance and Evolution (ICSME '17)*. Shanghai, China: IEEE, Sep. 2017, pp. 238–248.
- [15] M. Iammarino, F. Zampetti, L. Aversano, and M. Di Penta, "Self-admitted technical debt removal and refactoring actions: Co-occurrence or more?" in *Proceedings of the 35th IEEE International Conference on Software Maintenance and Evolution (ICSME '19)*. Cleveland, OH, USA: IEEE, 2019, pp. 186–190.
- [16] J. Tan, D. Feitosa, and P. Avgeriou, "An empirical study on self-fixed technical debt," in *Proceedings of the 3rd International Conference on Technical Debt (TechDebt '20)*. Seoul, Republic of Korea: ACM, 2020, pp. 11–20.
- [17] M. T. Baldassarre, V. Lenarduzzi, S. Romano, and N. Saarimäki, "On the diffuseness of technical debt items and accuracy of remediation time when using sonarqube," *Information and Software Technology*, vol. 128, 2020.
- [18] G. Digkas, M. Lungu, P. Avgeriou, A. Chatzigeorgiou, and A. Ampatzoglou, "How do developers fix issues and pay back technical debt in the apache ecosystem?" in *Proceedings of the IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER '18)*. Campobasso, Italy: IEEE, 2018, pp. 153–163.
- [19] G. Bavota and B. Russo, "A large-scale empirical study on self-admitted technical debt," in *Proceedings of the IEEE/ACM 13th Working Conference on Mining Software Repositories (MSR '16)*. Austin, TX, USA: IEEE, May 2016, pp. 315–326.
- [20] J. Tan, D. Feitosa, P. Avgeriou, and M. Lungu, "Evolution of technical debt remediation in python: A case study on the apache software ecosystem," *Journal of software-Evolution and process*, 2020.
- [21] V. Lenarduzzi, N. Saarimäki, and D. Taibi, "On the diffuseness of code technical debt in java projects of the apache ecosystem," in *Proceedings of the 2nd IEEE/ACM International Conference on Technical Debt (TechDebt '19)*. Montreal, QC, Canada: IEEE, 2019, pp. 98–107.
- [22] A. Potdar and E. Shihab, "An exploratory study on self-admitted technical debt," in *Proceedings of the 30th IEEE International Conference on Software Maintenance and Evolution (ICSME '14)*. Victoria, British Columbia, Canada: IEEE, 2014, pp. 91–100.
- [23] G. Fucci, F. Zampetti, A. Serebrenik, and M. Di Penta, "Who (self) admits technical debt?" in *Proceedings of the 36th IEEE International Conference on Software Maintenance and Evolution (ICSME '20)*. Adelaide, Australia: IEEE, 2020, pp. 672–676.
- [24] J. Liu, Q. Huang, X. Xia, E. Shihab, D. Lo, and S. Li, "An exploratory study on the introduction and removal of different types of technical debt in deep learning frameworks," *Empirical Software Engineering*, vol. 26, no. 2, pp. 1–36, 2021.
- [25] N. A. Ernst, S. Bellomo, I. Ozkaya, R. L. Nord, and I. Gorton, "Measure it? manage it? ignore it? software practitioners and technical debt," in *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering (FSE '15)*. Bergamo, Italy: ACM, 2015, pp. 50–60.
- [26] T. Besker, A. Martini, and J. Bosch, "Technical debt cripples software developer productivity: a longitudinal study on developers' daily software development work," in *Proceedings of the First International Conference on Technical Debt (TechDebt '18)*. Gothenburg, Sweden: ACM, 2018, pp. 105–114.
- [27] F. Zampetti, A. Serebrenik, and M. Di Penta, "Was self-admitted technical debt removal a real removal? an in-depth perspective," in *Proceedings of the 15th International Conference on Mining Software Repositories (MSR '18)*. Gothenburg, Sweden: ACM, 2018, p. 526–536.
- [28] S. Freire, N. Rios, B. Gutierrez, D. Torres, M. Mendonça, C. Izurieta, C. Seaman, and R. O. Spínola, "Surveying software practitioners on technical debt payment practices and reasons for not paying off debt items," in *Proceedings of the 24th Evaluation and Assessment in Software Engineering (EASE '20)*. Trondheim, Norway: ACM, 2020, pp. 210–219.
- [29] T. Amanatidis, A. Chatzigeorgiou, A. Ampatzoglou, and I. Stamelos, "Who is producing more technical debt? a personalized assessment of td principal," in *Proceedings of the XP2017 Scientific Workshops*. Cologne, Germany: ACM, 2017, pp. 1–8.
- [30] R. Alfayez, P. Behnamghader, K. Srisopha, and B. Boehm, "An exploratory study on the influence of developers in technical debt," in *Proceedings of the First IEEE/ACM International Conference on Technical Debt (TechDebt '18)*. Gothenburg, Sweden: IEEE, 2018, pp. 1–10.
- [31] M. J. Salamea and C. Farré, "Influence of developer factors on code quality: a data study," in *Proceedings of the 19th IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C '19)*. Sofia, Bulgaria: IEEE, 2019, pp. 120–125.
- [32] R. Van Solingen, V. Basili, G. Caldiera, and H. D. Rombach, "Goal Question Metric (GQM) approach," in *Encyclopedia of Software Engineering*. Hoboken, NJ, USA: John Wiley & Sons, Inc., Jan 2002, pp. 528–532.
- [33] E. Oliveira, E. Fernandes, I. Steinmacher, M. Cristo, T. Conte, and A. Garcia, "Code and commit metrics of developer productivity: a study on team leaders perceptions," *Empirical Software Engineering*, vol. 25, no. 4, pp. 2519–2549, 2020.
- [34] A. Trendowicz and J. Münch, "Factors influencing software development productivity—state-of-the-art and industrial experiences," *Advances in computers*, vol. 77, pp. 185–241, 2009.
- [35] S. Easterbrook, J. Singer, M.-A. Storey, and D. Damian, "Selecting empirical methods for software engineering research," in *Guide to advanced empirical software engineering*. Springer, 2008, pp. 285–311.
- [36] A. Fink, *The survey handbook*. SAGE Publications, 2003.

- [37] T. C. Lethbridge, S. E. Sim, and J. Singer, "Studying software engineers: Data collection techniques for software field studies," *Empirical Software Engineering*, vol. 10, no. 3, p. 311–341, July 2005.
- [38] N. Rios, R. O. Spínola, M. Mendonça, and C. Seaman, "The most common causes and effects of technical debt: First results from a global family of industrial surveys," in *Proceedings of the ACM/IEEE 12th International Symposium on Empirical Software Engineering and Measurement (ESEM '18)*. Oulu, Finland: ACM, 2018, pp. 1–10.
- [39] B. Wang, L. Chen, W. Ma, Z. Chen, and B. Xu, "An empirical study on the impact of python dynamic features on change-proneness," in *Proceedings of the 27th International Conference on Software Engineering and Knowledge Engineering (SEKE '15)*. Pittsburgh, USA: KSI Research Inc., 2015, pp. 134–139.
- [40] B. Ray, D. Posnett, V. Filkov, and P. Devanbu, "A large scale study of programming languages and code quality in github," in *Proceedings of the ACM SIGSOFT 22nd International Symposium on Foundations of Software Engineering (FSE '14)*. Hong Kong, China: ACM, 11 2014, pp. 155–165.
- [41] S. Wehaibi, E. Shihab, and L. Guerrouj, "Examining the impact of self-admitted technical debt on software quality," in *Proceedings of the 23rd IEEE international conference on software analysis, evolution, and reengineering (SANER '16)*, vol. 1. Osaka, Japan: IEEE, 2016, pp. 179–188.
- [42] S. Bellomo, R. L. Nord, I. Ozkaya, and M. Popeck, "Got technical debt? surfacing elusive technical debt in issue trackers," in *Proceedings of the 13th IEEE/ACM Working Conference on Mining Software Repositories (MSR '16)*. Austin, TX, USA: IEEE, 2016, pp. 327–338.
- [43] J. Flisar and V. Podgorelec, "Enhanced feature selection using word embeddings for self-admitted technical debt identification," in *Proceedings of the 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA '18)*. Prague, Czech Republic: IEEE, 2018, pp. 230–233.
- [44] M. Yan, X. Xia, E. Shihab, D. Lo, J. Yin, and X. Yang, "Automating change-level self-admitted technical debt determination," *IEEE Transactions on Software Engineering*, vol. 45, no. 12, pp. 1211–1229, 2018.
- [45] F. Zampetti, A. Serebrenik, and M. Di Penta, "Automatically learning patterns for self-admitted technical debt removal," in *Proceedings of the 27th IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER '20)*. London, ON, Canada: IEEE, 2020, pp. 355–366.
- [46] F. Palomba, G. Bavota, M. Di Penta, F. Fasano, R. Oliveto, and A. De Lucia, "A large-scale empirical study on the lifecycle of code smell co-occurrences," *Information and Software Technology*, vol. 99, pp. 1–10, 2018.
- [47] A. Lerina and L. Nardi, "Investigating on the impact of software clones on technical debt," in *Proceedings of the 2nd IEEE/ACM International Conference on Technical Debt (TechDebt '19)*. Montreal, QC, Canada: IEEE, 2019, pp. 108–112.
- [48] D. Tsoukalas, D. Kehagias, M. Siavvas, and A. Chatzigeorgiou, "Technical debt forecasting: an empirical study on open-source repositories," *Journal of Systems and Software*, vol. 170, p. 110777, 2020.
- [49] H. Borges, A. Hora, and M. T. Valente, "Understanding the factors that impact the popularity of github repositories," in *Proceedings of the 32th IEEE International Conference on Software Maintenance and Evolution (ICSME '16)*. Raleigh, NC, USA: IEEE, 2016, pp. 334–344.
- [50] J. Tsay, L. Dabbish, and J. Herbsleb, "Influence of social and technical factors for evaluating contribution in github," in *Proceedings of the 36th International Conference on Software Engineering (ICSE '14)*. Hyderabad, India: ACM, 2014, pp. 356–366.
- [51] C. Seaman, "Qualitative methods in empirical studies of software engineering," *IEEE Transactions on software engineering*, vol. 25, no. 4, pp. 557–572, 1999.
- [52] M. B. Miles and A. M. Huberman, *Qualitative data analysis: An expanded sourcebook*. SAGE Publications, 1994.
- [53] E. Aghajani, C. Nagy, M. Linares-Vásquez, L. Moreno, G. Bavota, M. Lanza, and D. C. Shepherd, "Software documentation: the practitioners' perspective," in *Proceedings of the 42nd IEEE/ACM International Conference on Software Engineering (ICSE '20)*. Seoul, Republic of Korea: IEEE, 2020, pp. 590–601.
- [54] S. L. Vadlamani and O. Baysal, "Studying software developer expertise and contributions in stack overflow and github," in *Proceedings of the 36th IEEE International Conference on Software Maintenance and Evolution (ICSME '20)*. Adelaide, Australia: IEEE, 2020, pp. 312–323.
- [55] R. A. Armstrong, "When to use the bonferroni correction," *Ophthalmic and Physiological Optics*, vol. 34, no. 5, pp. 502–508, 2014.
- [56] N. Cliff, *Ordinal methods for behavioral data analysis*. Psychology Press, 2014.
- [57] C. Tantithamthavorn, S. McIntosh, A. E. Hassan, and K. Matsumoto, "An empirical comparison of model validation techniques for defect prediction models," *IEEE Transactions on Software Engineering*, vol. 43, no. 1, pp. 1–18, 2016.
- [58] A. J. Scott and M. Knott, "A cluster analysis method for grouping means in the analysis of variance," *Biometrics*, pp. 507–512, 1974.
- [59] B. Rosner, R. J. Glynn, and M.-L. T. Lee, "The wilcoxon signed rank test for paired comparisons of clustered data," *Biometrics*, vol. 62, no. 1, pp. 185–192, 2006.
- [60] Z. Gu, L. Gu, R. Eils, M. Schlesner, and B. Brors, "circlize implements and enhances circular visualization in r," *Bioinformatics*, vol. 30, no. 19, pp. 2811–2812, 2014.
- [61] Z. Wan, X. Xia, A. E. Hassan, D. Lo, J. Yin, and X. Yang, "Perceptions, expectations, and challenges in defect prediction," *IEEE Transactions on Software Engineering*, vol. 46, no. 11, pp. 1241–1266, 2018.
- [62] W. Zou, D. Lo, Z. Chen, X. Xia, Y. Feng, and B. Xu, "How practitioners perceive automated bug report management techniques," *IEEE Transactions on Software Engineering*, vol. 46, no. 8, pp. 836–862, 2018.
- [63] S. Siegel and N. Castellan, *Nonparametric statistics for the behavioral sciences*, 2nd ed. McGraw-hill, 1988.
- [64] R. V. Lenth, "Some practical guidelines for effective sample size determination," *The American Statistician*, vol. 55, no. 3, pp. 187–193, 2001.
- [65] R. A. Fisher, "On the interpretation of χ^2 from contingency tables, and the calculation of p," *Journal of the Royal Statistical Society*, vol. 85, no. 1, pp. 87–94, 1922.
- [66] M. Aickin and H. Gensler, "Adjusting for multiple testing when reporting research results: the bonferroni vs holm methods," *American journal of public health*, vol. 86, no. 5, pp. 726–728, 1996.
- [67] 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.
- [68] N. Mittas and L. Angelis, "Ranking and clustering software cost estimation models through a multiple comparisons algorithm," *IEEE Transactions on Software Engineering*, vol. 39, no. 4, pp. 537–551, 2013.
- [69] M. Azzeh, A. B. Nassif, and L. L. Minku, "An empirical evaluation of ensemble adjustment methods for analogy-based effort estimation," *Journal of Systems and Software*, vol. 103, pp. 36–52, 2015.
- [70] D. S. Moore, *The Basic Practice of Statistics*, 8th ed. W.H. Freeman & Company, 2018.
- [71] K. Krippendorff, "Computing krippendorff's alpha-reliability," *Departmental Papers (ASC)*, pp. 1–12, Jan 2011.
- [72] —, *Content analysis: An introduction to its methodology*. Sage publications, 2018.