

Toward Better Understanding and Documentation of Rationale for Code Changes

Khadijah Al Safwan

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Science and Applications

Francisco Javier Servant, Chair

Eli Tilevich

Na Meng

Muhammad Ali Gulzar

Thomas D LaToza

July 28, 2023

Blacksburg, Virginia

Keywords: Software Engineering, Software Evolution and Maintenance, Revision Control
Systems, Software Changes Rationale, Applied Machine Learning.

Copyright 2023, Khadijah Al Safwan

Toward Better Understanding and Documentation of Rationale for Code Changes

Khadijah Al Safwan

(ABSTRACT)

Software development is driven by the development team's decisions. Communicating the rationale behind these decisions is essential for the projects success. Although the software engineering community recognizes the need and importance of rationale, there has been a lack of in-depth study of rationale for code changes. To bridge this gap, this dissertation examines the rationale behind code changes in-depth and breadth. This work includes two studies and an experiment. The first study aims to understand software developers' need. It finds that software developers need to investigate code changes to understand their rationale when working on diverse tasks. The study also reveals that software developers decompose the rationale of code commits into 15 separate components that they could seek when searching for rationale. The second study surveys software developers' experiences with rationale. It uncovers issues and challenges that software developers encounter while searching for and recording rationale for code changes. The study highlights rationale components that are needed and hard to find. Additionally, it discusses factors leading software developers to give up their search for the rationale of code changes. Finally, the experiment predicts the documentation of rationale components in pull request templates. Multiple statistical models are built to predict if rationale components' headers will not be filled. The trained models are effective in achieving high accuracy and recall. Overall, this work's findings shed light on the need for rationale and offer deep insights for fulfilling this important information need.

Toward Better Understanding and Documentation of Rationale for Code Changes

Khadijah Al Safwan

(GENERAL AUDIENCE ABSTRACT)

Software developers build software by creating and changing the software's code. In this process, developers make decisions and other developers need to understand these decisions. The rationale behind code changes is an important piece of information that leads to development success if well explained and understood. In this work, we study the developers' need for rationale by conducting two studies and an experiment. In the first study, we found that software developers often need to look into the rationale behind code changes to understand them better while working on different tasks. We identified 15 different parts of rationale that developers seek when searching for rationale for code changes. The second study focused on the experiences of software developers when looking for and recording rationale. We discovered some challenges that developers face, like difficulty in finding specific rationale parts and the factors that make developers give up searching for rationale. The experiment predicts if developers would document rationale in specific templates. We built models to predict if certain parts of rationale would be left empty, and the models were effective. Overall, this research provides a better understanding of software developers' need, and it provides valuable insights to help fulfill this important information need.

Dedication

To my husband, for always believing in me and supporting me.

Acknowledgments

First and foremost, my deepest gratitude is to my advisor, Dr. Francisco Servant, whose guidance has been instrumental in my success throughout my Ph.D. journey. His support and mentorship have been invaluable, and I am deeply indebted to him for the knowledge and skills I have gained under his supervision.

I extend my appreciation to my Ph.D. committee members, Dr. Eli Tilevich, Dr. Na Meng, Dr. Muhammad Ali Gulzar, and Dr. Thomas LaToza, for generously volunteering their time and expertise to review my work and provide feedback. Their thoughtful participation in my milestones presentation and my documents review have contributed to my work quality. I would like to express my gratitude to my fellow lab mates, both past and present, for creating an inspiring and collaborative environment that has fostered the exchange of innovative ideas and exciting projects. The challenges posed by COVID-19 have only reinforced my appreciation for the intellectual conversations and knowledge I have gained from being a part of this remarkable research community.

I am deeply appreciative of the distinguished professors at Virginia Tech, who have played a pivotal role in shaping my professional development. Their dedication to teaching and mentorship has been a privilege to experience, and the transition from student to teaching assistant to instructor has been a transformative journey enriched by their support.

I must acknowledge that my Ph.D. journey was marked by its share of difficulties, but I am immensely grateful for the transformative growth it has brought into my life. It is with profound appreciation that I extend my thanks to my loving husband, Dr. Mohammed Alaboalirat, whose unwavering love and support have been my constant pillar of strength throughout this academic pursuit. His belief in me has been a driving force behind my achievements, and I am truly fortunate to have him by my side.

Last but not least, I would like to extend my heartfelt gratitude to my parents, siblings, family, and friends for being my supporters from the outset. Their love and encouragement have been a constant source of motivation and inspiration. My accomplishments, including the completion of my Ph.D., are a testament to their belief in me, and I am forever grateful for their presence in my life.

Once again, I express my sincere appreciation to everyone who has played a part in my journey, as your support and encouragement have been indispensable to my success.

Contents

- List of Figures** **xii**

- List of Tables** **xiii**

- 1 Introduction** **1**
 - 1.1 Thesis Statement 1
 - 1.2 Understanding developers' need 2
 - 1.3 Surveying developers' experiences 3
 - 1.4 Predicting rationale components' documentation 3
 - 1.5 Results and Contributions 4
 - 1.6 Outline 6

- 2 Understanding developers' need for rationale of code changes** **7**
 - 2.1 Background and Motivation 7
 - 2.2 Research Questions 8
 - 2.3 Approach 9
 - 2.3.1 Developer Interviews 10
 - 2.3.2 Survey I 13
 - 2.3.3 Survey II 13

2.3.4	Participants' Recruitment	13
2.4	Tasks for which rationale is needed	16
2.4.1	Research Method	16
2.4.2	Results	18
2.5	Seeking rationale of internal code vs. external code	25
2.5.1	Research Method	25
2.5.2	Results	25
2.6	Components of rationale	30
2.6.1	Research Method	30
2.6.2	Results	31
2.7	Threats to validity	35
2.8	Summary	37
3	Surveying developers' experiences with rationale of code changes	38
3.1	Background and Motivation	38
3.2	Research Questions	39
3.3	Approach	40
3.4	Experience with rationale	41
3.4.1	Research Method	41
3.4.2	Results	41

3.5	Experience with individual components	45
3.5.1	Research Method	45
3.5.2	Results	46
3.6	Factors leading to give up rationale search	49
3.6.1	Research Method	49
3.6.2	Results	50
3.7	Experience comparison	58
3.7.1	Research Method	58
3.7.2	Results	58
3.8	Threats to validity	60
3.9	Conclusion	61
4	Predicting Rationale Components Documentation in PRT	62
4.1	Background and Motivation	63
4.2	Research Questions	64
4.3	Approach	65
4.3.1	Data Preparation	65
4.3.2	Correlation analysis	67
4.3.3	Statistical Modeling	70
4.3.4	Models Evaluation & Analysis	71

4.4	PRT rationale header correlation analysis	72
4.5	Predicting PRT rationale header fill	74
4.6	Comparing features groups performance	75
4.7	Threats to validity	76
4.8	Summary	77
5	Discussion and Implications	78
5.1	Rationale of code commits vs. rationale in other contexts	78
5.2	Providing a better understanding of rationale need	79
5.3	Our vision: practitioners documentation of rationale	80
5.4	Our vision: tools support in documenting rationale	82
5.5	Our vision: rationale documentation automation	85
5.6	Our vision: benefits of good documentation	86
6	Related Work	88
6.1	Rationale Management in Software Requirements, Design, and Architecture	89
6.2	Rationale in Software Evolution and Maintenance	90
6.2.1	Components of the Rationale of Code Changes	90
6.2.2	Tasks that need the Rationale of Code Changes	90
6.2.3	Experience with the Rationale of Code Changes	91
6.3	Rationale Documentation Support	92

6.3.1	Rationale documentation in code review context	93
6.3.2	Tools to support software changes documentation	93
7	Conclusions	95
	Bibliography	97

List of Figures

2.1	Demographics of our interview and survey participants	15
2.2	Tasks for which rationale for code commits is needed	27
3.1	Experience of developers needing rationale	42
3.2	Experience of developers finding rationale	43
3.3	Experience of developers recording rationale	44
3.4	Experience of developers with individual components of the rationale	47
3.4	Experience of developers with individual components of the rationale (cont.)	48
3.5	Give up factor	52
3.6	Cross-dimensional analysis of developers' experience with the individual components of rationale of code commits	59
4.1	Classification results of classifiers trained using 10-fold cross-validation	74
4.2	Classification results of <i>Random Forest</i> classifier trained with different set/-groups of features	75

List of Tables

2.1	Preliminary model of rationale of code commits	12
2.2	Tasks for which rationale for code commits is needed	18
2.3	Resulting model of the rationale of code commits	32
3.1	Factors leading developers to give up the search for rationale of code changes	51
4.1	Features extracted for each pull request	68
4.2	Correlation between features and PRT rationale header fill	73

Chapter 1

Introduction

Software development is driven by decisions from various stakeholders, following some *rationale*, at every stage of the software development life-cycle [1]. Given the complexity of software and its development teams, the effective communication of development decisions' rationale is expected to play an important role in software project success [2]. Multiple studies involving software history and developers' information needs establish a strong demand for rationale [1, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]. Our work is motivated by these empirical studies that highlight the importance of rationale.

Although the software engineering community recognizes rationale need and importance, there is a lack of in-depth study of rationale for changes. To bridge this gap, this thesis aims to better understand and document the rationale for code changes.

1.1 Thesis Statement

To achieve a better understanding and documentation of rationale for code changes, we devise the following hypothesis:

Thesis Statement: *The rationale of code changes can be needed in many software development tasks and involve many components, which can often be hard to find, and automatic classification could be utilized to assist developers in documenting rationale.*

To address our hypothesis, we perform a series of studies and experiments that we describe in detail in the following chapters. Each chapter addresses one of the three hypotheses we divided the thesis statement into. The following sub-sections enumerate those studies, drawing up the specific hypothesis that needs to be validated for each research study alongside the corresponding research questions.

1.2 Understanding software developers' need for rationale of code changes

Rationale of code changes is informally defined as the answer to the question: “*why was this code implemented this way?*” [3, 7]. However, this informal question could be easily interpreted in many different ways. We aim to understand *what developers mean* by the rationale of code changes and the *context* of their need in practice.

For this study, we apply a developer-centric mixed-methods approach. We interviewed software developers and distributed two surveys to study software developers' perspectives on rationale for code changes. This study investigates the following hypothesis:

Sub-Hypothesis H_1 : *The rationale of code changes can be needed in many software development tasks and involve many components.*

To verify our hypothesis, we investigate the following questions:

- What are the tasks in which developers need to find the rationale for code commits?
- How often do developers seek rationale of their team's internal code vs. others' code?
- Which components do software developers decompose the rationale of code commits into?

1.3 Surveying software developers' experiences with rationale of code changes

Many research studies about information needs convey the importance of understanding the rationale of code changes. Rationale is the most common [3] and important [5] information need to understand from code-change history. Unfortunately, it can also be quite challenging to find an answer for rationale of code changes [5, 7]. We aim to understand the developers' efforts and issues associated with rationale of code changes needing, finding, and recording. For this study, we use the same developer-centric mixed-methods approach as the first study. We interviewed software developers and distributed two surveys to study software developers' experiences. This study investigates the following hypothesis:

Sub-Hypothesis H_2 : *The rationale of code changes can often be hard to find.*

To verify our hypothesis, we investigate the following questions:

- What is the experience of developers needing, finding, and recording the rationale of code commits?
- What is the experience of developers needing, finding, and recording the individual components of the rationale of code commits?
- Would comparing the experience of developers needing, finding, and recording the individual components of the rationale of code commits with each other reveal areas for improvement?
- What makes software developers give up their search for rationale of code commits?

1.4 Predicting the documentation of rationale components in pull request templates

Rationale of code changes is frequently sought in code review [9, 10, 13, 14]. Open Source GitHub repositories prescribe *Pull Request Templates (PRT)* [16] to guide software developers

in describing their code changes for code review. Multiple *rationale components* were present in recent empirical studies of pull request templates' content [17, 18]. Although rationale components are requested in prescribed PRT headers, the components may not be suited for every Pull Request (PR). We aim to streamline the documentation of rationale for code change components by predicting if a PRT rationale header will not be filled in individual PRs.

For this study, we take a data-centric analysis approach. First, we mine software repositories for PRs information. Then, we perform correlation analysis between PRT rationale components headers being filled and PRs characteristics. Finally, we build statistical models to predict if a PRT header of rationale components will not be filled in individual PRs. This study investigates the following hypothesis:

Sub-Hypothesis H_3 : *An automatic classification could be utilized to assist developers in documenting rationale.*

To verify our hypothesis, we investigate the following questions:

- How are different pull request characteristics correlated with filling a pull request template rationale header?
- To what extent can classification algorithms predict if a pull request's template rationale header will be filled?
- How effective are different feature groups in predicting if a pull request's template rationale header will be filled?

1.5 Results and Contributions

The first two studies of this thesis provide scientific contributions, they advance the knowledge about software developers need of rationale for code changes. The first study establish-

ing a *common understanding* of software developers *need* of rationale for code changes. The results of this study highlight eight tasks for which rationale for code commits is needed: programming, working on bugs, communication, tools, documentation, project management, testing, and specifications. The software engineering community can benefit from knowing the context of rationale need. For example, open-source tools or libraries could consider that some users might need to find the rationale of code changes before installing or using their tool. In addition, the results decompose rationale for code commits into 15 different components: goal, need, benefits, constraints, alternatives, selected alternatives, dependency, committer, time, location, modifications, explanation of modifications, validation, maturity stage, and side effects. The decomposition of rationale into components is a step toward improved communication of rationale for code changes.

The second study results bring awareness to developers' *efforts* and *issues* associated with rationale for code commits need, finding, and recording. For example, the results highlight components of rationale that are needed but really hard to find: alternatives, constrains, and side effects. Those components are also rarely recorded and are the ones that our study participants struggled most to find. Additionally, the study results present eight factors leading software developers to give up their search for rationale of code commits: codebase state, documentation, effort management, developer knowledge, interpersonal (emotions), impact on productivity, personnel, and time management.

The third study serves as a demonstration to how possible it is to assist software developers in documenting the rationale of code changes. In this study we built a machine learning approach to predict if a rationale header in pull request templates is likely to be empty. The recommended approach for predicting if a rationale component will not be filled could help save developers' time and effort.

1.6 Outline

The body of this dissertation, Chapters 2-4, presents experiments' methods and results to provide evidence to support the thesis statement. Chapter 5 discuss the dissertation findings and their implications. Chapter 6 describes this dissertation's related work. Chapter 7 outline future work directions and concluding remarks.

Chapter 2

Understanding developers’ need for rationale of code changes

In this chapter, we aim to understand *what developers mean* by the rationale of code changes and their *need* for it in practice. First, we study the context in which software developers need the rationale for code changes. We set out to discover the *tasks* and the *target code* for which developers need the rationale for code changes. Second, we study the specific *pieces of information* in which developers decompose the rationale of code changes.

2.1 Background and Motivation

The rationale of code commits is informally defined as the answer to the question: “*why was this code implemented this way?*” [3, 7]. However, software developers could easily interpret this informal question in many different ways, potentially as disparate as: “*what is the purpose of this code?*” [6]; “*why where [these changes] introduced?*” [4]; or “*why was it done this way?*” [7] — all of which request different answers. In past studies, software developers mentioned all these different interpretations when asked about rationale. Thus, we formed our intuition that it could be decomposed into multiple components, each addressing different aspects of the question.

Efforts to study rationale in-depth have been carried out in the design context, decomposing it into various more-specific components [19]. In the context of software maintenance, Burge et al. [1] *prescriptively* propose some questions that may answer rationale [1]. We, instead, take a *descriptive* approach, *i.e.*, we aim to discover *how developers decompose* the rationale of code changes — as opposed to conceptually and rigorously decomposing the concept.

We take a developer-centric approach to discover *decomposition of the rationale of code changes* by the rationale of code changes and their *need* for it in practice. We used a mixed-methods approach in our study, involving interviews and surveys of software developers. The decomposition of the rationale of code changes provides: (1) a *common language* to use when discussing rationale, which practitioners can use to (2) assess and (3) strengthen the quality of their rationale sharing and documentation processes.

2.2 Research Questions

To understand some of the contexts under which the rationale of code commits is needed, we ask the following two research questions:

- ***RQ: What tasks do software developers need to find the rationale for code commits?*** Discovering the more extensive set of *tasks* in which *rationale for code commits is needed* will enable researchers to tailor their support to the context of each task. While the rationale for code commits has been identified as a need in previous task-specific studies (*e.g.*, during code review [13]), an exhaustive list of tasks in which developers needed it is still unknown.

- **RQ: How often do developers seek rationale of their team’s internal code vs. others’ code?** Our intuition is that developers may seek the rationale for code commits of internal and external code outside their team. If our intuition is validated, future research efforts (e.g., rationale retrieval support) should accommodate for the fact that developers may or may not own the code for which they seek rationale.

To discover an extensive set of rationale for code commits components that developers believe would compose a high-quality detailed description, we ask the following research question:

- **RQ: Which components do software developers decompose the rationale of code commits into?** A model of rationale for code commits will inform developers wanting to improve their documentation of rationale of code commits — whether they aim to document it fully or more thoroughly.

2.3 Approach

Our study uses a mixed-methods approach, combining developer interviews and a survey. Mixed methods have been successfully employed in other studies of software developers, *e.g.*, [3, 5, 20, 21]. The developer interviews allowed us to *qualitatively* discover details about multiple aspects of the rationale of code commits through rich one-on-one conversations with developers. The surveys enabled us to reach more participants and extend our *quantitative* findings. In the following text, we present the design of our interview, surveys, and participant recruitment. This chapter and the following chapter’s studies share the interview and surveys.

2.3.1 Developer Interviews

We designed and refined our interview script through five pilot sessions. We ran a first pilot interview at the early stages of designing our interview script, in which we asked general open questions about the rationale of code changes. After the first pilot, we improved the interview script, making it more structured, making the questions more specific, and adding a preliminary model of rationale (Table 2.1) to the script. Then, we ran a second pilot to test these improvements. After it, we changed some questions' wording to make them explicit. We ran the third and fourth pilots with experienced and beginner developers. We aimed to test the new improvements and check the time required for our interview. After the fourth pilot, we finalized the interview kit (available in the research artifacts package) by adding introductions to the different sections of the interview script. Finally, we ran a fifth pilot to check for the entire interview process. The process included: advertisement, screening survey, scheduling, interview session, and analysis of responses. After this pilot, we were ready to recruit and interview participants.

Our interview consisted of three parts. The first part focused on finding the tasks in which rationale of code commits is needed. The second part focused on discovering the components that form the rationale of code commits. The third part aimed to understand the experiences of developers needing, finding, and recording rationale of code commit and its components, and to understand when developers give up searching for it.

Part 1: Tasks with Rationale Need. We started our interviews by giving our participants the definition of *rationale of code* that is most common in the research literature, *i.e.*, the answer to “*why is the code this way?*” [3, 7]. We did this to ensure that all participants had a uniform definition of the concept we discussed. Next, we asked them to describe real situations in which they investigated a code commit to understand its rationale. This way,

they grounded their answers in real experiences. We used their answers to identify the tasks in which they needed rationale. We provide more details about this analysis in Section 2.4.2.

Part 2: Components of Rationale. Right after giving our participants the definition of rationale of code commits and asking them to describe real situations in which they needed it, we asked them to decompose the rationale of code commits into components. By asking this question after they had been thinking about their own experiences searching for the rationale of code commits, we intended to stimulate the participants' memories and set them in the proper context, as well as to maximize the number of components that they would report. Then, we showed them a *preliminary model* (see Table 2.1) of components of the rationale of code commits that we created by studying the research literature — including components to which researchers have referred as *rationale* [1, 4, 6, 7, 22]. We used this preliminary model as a probe to prime our participants and get them in the right frame of reference. We asked participants to critique and extend the preliminary model — considering their previous decomposition — to the extent that they believed necessary to build a *final model* of all the components of the rationale of code commits. For any component added by participants, we asked them to describe it with a name, question, and example answer. We presented the same preliminary model to all participants — *i.e.*, we did not show the modified models to other interview participants.

Using the preliminary model as a probe served multiple purposes. The preliminary model clarified the scope of our study. It also allowed developers to discuss an extensive set of components. We also believe that it allowed us to reach saturation of answers much faster (interviewing fewer participants) than if we had only relied on our participants' experiences and decompositions — since situations in which many components are needed simultaneously may be rare or because people's memory is generally unreliable.

Table 2.1: Preliminary model of rationale of code commits

Component	Component Expressed as a Question	Literature References
Goal	What do you want to achieve?	[1, 6]
Need	Why do you need to achieve that?	[4, 22]
Location	What artifacts were changed?	[1]
Modifications	What specific changes were performed in the artifacts?	[6, 22]
Alternatives	What other alternatives did you have?	[7]
Selected alternative	Why did you make those specific changes and not others?	[6, 7]
Validation	How do those specific changes achieve the goal?	[1, 7]
Benefits	What is the benefit of what you want to achieve?	[1, 6]
Costs	What risks could come from these changes?	[1, 6]

However, using a preliminary model, we had the risk of introducing confirmation bias [23]. We took multiple measures to reduce this potential bias. First, we presented the preliminary model neutrally, as “*this model*” — avoiding potentially-biasing adjectives, such as “ours” or “preliminary”. Second, we built it from the research literature, reducing the risk of inserting our own opinions. Third, we presented the preliminary model to participants only after they had produced their own decomposition without having seen it. Fourth, we asked participants to consider their own decomposition when critiquing and extending the preliminary model. We believe that we were successful with these efforts since our final model of rationale of code commits (see Table 2.3) is much more extensive than the preliminary model (see Table 2.1). The preliminary model had only nine components, whereas the final model have 15.

Part 3: Experience with Rationale. Next, we asked participants to rate their experiences needing, finding, and recording rationale of code commits and its components in Likert-scale-style questions. We also asked an open question about what makes developers give up their search for rationale.

2.3.2 Survey I

Once we had identified the components of the rationale of code commits through our interviews, we used a survey to obtain more answers about developers' experiences needing, finding, and recording rationale of code commits and its components. We refined our survey through four pilot versions, improving its clarity and the time required to complete it. Our survey included the same Likert-scale-style questions we asked our interview participants about their experiences. However, the reference model of rationale of code commits we gave survey respondents was the final model resulting from our analysis of the rationale components. Our experiences' results include the answers we obtained from our interviews and the survey.

2.3.3 Survey II

Once we had analyzed our interviews to identify the tasks for which rationale is needed and the factors leading our participants to give up their search for the rationale of code commits, we created a short survey to obtain more responses. We ran two pilots to test the new survey. The pilots provided positive feedback and an understanding of the time required to complete the survey. Our survey included an introduction and the same open questions we asked our interview participants about the tasks and factors. We obtained results for the tasks and factors by analyzing the answers from our interviews and Survey II.

2.3.4 Participants' Recruitment

We used snowball sampling [24] to recruit participants for our study, *i.e.*, we asked them to refer our study to their contacts. We advertised our study in mailing lists in our university

that covered software developers of diverse experience, *e.g.*, developing various university software systems, and graduate students with professional software development experience. We also advertised it through public channels and social media, *e.g.*, developers' communities on Slack. We compensated interview participants with a \$20 Amazon gift card and encouraged survey participation by raffling a \$50 gift card. figure 2.1 represents the demographics of our interview, survey I, and survey II participants.

We interviewed 20 participants after having discarded three other interviews for various reasons: one participant could not describe an example of seeking rationale of code commits, another voluntarily expressed lack of experience throughout the interview, and we found out that the last one knew our interview materials.

We analyzed 26 survey I responses after having discarded two responses. We discarded two surveys we deemed done carelessly, taking less than 10 minutes. We determined this cut-off point through our pilot surveys; we asked one pilot participant to complete the survey carelessly, which took 10 minutes. We specifically asked the interview participants not to fill out the survey, eliminating the chance of including duplicates in our results.

For this survey, we did not have a target number of responses. We kept the survey open for a month, resulting in 28 total responses (before filtration).

We analyzed 26 survey II responses after having discarded 60 spam responses. We identified those spam responses since the spammers did not address the asked open questions. Instead, the spammers submitted the same spam text answer with different contact information for each response. We believe a large number of spam responses resulted from the compensation's raffle odds of winning (1 in 50 chance). For this second survey, the advertisement emphasized: "not to fill the survey if you have participated in the study before". Our target was to reach a similar number to the previous survey I. We also kept the survey open for a month, resulting

in 86 total responses (before filtration). Although we received/analyzed the same number of responses as survey I, the participants are different, which can be observed from the demographics figure (figure 2.1).

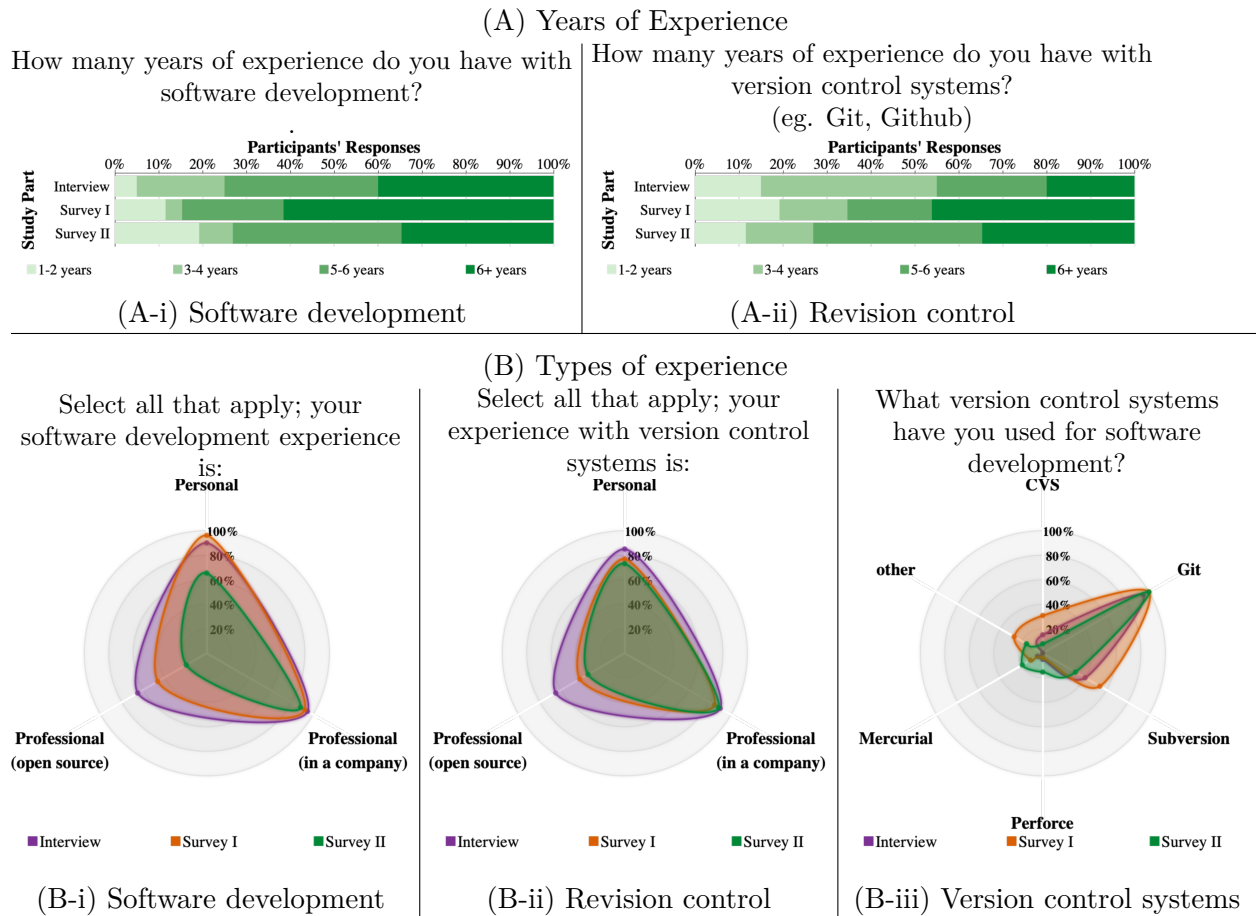


Figure 2.1: Demographics of our interview and survey participants

2.4 RQ: What are the tasks in which software developers need to find the rationale for code commits?

2.4.1 Research Method

We asked our (20) interview and (26) survey participants to describe a situation in which they needed to understand the rationale of a code commit. We analyzed their responses qualitatively, using closed coding to extract the task they performed when they needed rationale for code changes. Three coders were involved in the coding. We reached saturation in our observed tasks in interview 6 (of 20) and in our observed subtasks in survey 14 (of 26).

Analysis data: To base our observations on real developer experiences, we asked our participants (with an open-ended question) to describe an actual situation in which they needed to understand the rationale of code commits. Then, we analyzed their descriptions of these real-world experiences to identify the tasks they were performing, in which they experienced the need to know the rationale of code commits. We specifically asked our participants: *“Tell me about one time in which you investigated a code commit to understand its rationale. Why did you need to find the rationale for that specific code commit? What was the rationale for that specific code commit?”*

Analysis method: We analyzed our participants' answers using closed coding [25] (also used in *grounded theory* [26]) — *i.e.*, we coded our participants' answers, labeling them with categories according to a pre-existing set of codes (a *codebook*). Studies of the tasks developers perform at their job already exist in the research literature. So, we used a closed codebook containing the list of developer tasks that Begel and Simon [27] captured when observing software developers at work [27]. We analyzed our participants' responses and

labeled them with the task(s) from our codebook that we identified they were performing. We allowed multiple labels for each response, since our participants sometimes mentioned being involved in multiple tasks when needing the rationale of code changes.

First, two coders held a code discussion session to reach a common understanding of the scope of each developer task and subtask in the codebook. In this session, the two coders provided a detailed description for each task to better delineate their difference. Then, the two coders performed their coding. They labeled participant responses with multiple labels in two scenarios: if multiple tasks were mentioned (*e.g.*, “debugging” during “code review”), or when the response could fit multiple tasks (*i.e.*, if it was not clear which one of the many was being described). In both situations, they had discussions to agree on the labels that better fit each response, *i.e.*, the superset of tasks that were either mentioned or would easily fit the description in the text. They also allowed the addition of new tasks and subtasks if they still needed to be added to the codebook. The two coders coded about 52% of the tasks similarly, and they coded about 59% of the subtasks similarly. Next, the two coders held a joint focused-coding session and resolved disagreements. After this joint coding session, the two coders resolved most disagreements but still disagreed on coding two participants’ responses. Therefore, a third coder provided additional independent coding for these two responses. Finally, after a final discussion that reviewed all three codings, all coders agreed on how to label these two responses.

Analysis evaluation: We reached saturation [28, 29] in our identified set of tasks; our participants mentioned no new tasks after the sixth interview (out of 20) and in our subtasks after the 14th survey (out of 26).

2.4.2 Results

We show the list of tasks (and subtasks) during which our participants needed to understand the rationale of code commits in Table 2.2. Table 2.2 also contains a description and example (abstracted from our participants' responses) that we created to clarify the scope of each task during our coding sessions. Finally, we report in figure 2.2A the relative frequency with which each task and subtask was mentioned.

Table 2.2: Tasks for which rationale for code commits is needed

Task	Subtask	Description	Example (abstracted from participants' responses)
Programming	Reading	Reading source code to understand various aspects of it, like design and features.	A Developer navigates through the commits of a project to understand how a specific feature was implemented. The feature spans multiple classes, and the developer finds the commit that introduced the feature. The developer reads and investigates the rationale of the changes.
	Writing	Writing source code to implement, refactor, improve, and maintain the codebase.	A developer tries to break down a big function in a previous commit to improve code reuse and testing. This leads the developer to ask questions about the rationale of the changes.
	Proofreading	Reading source code to look for and solve issues before submitting the code for review or boarding.	A developer proofreads code to ensure that variable names are clear and informative before committing. This verification of variable names leads the developer to ask questions about the rationale of the changes.
	Code Reviewing	Reviewing source code as part of the code review process.	A reviewer assesses the correctness and quality of the commit under review, which requires understanding the rationale of the changes.
Specifications	Writing	Writing specifications for functional or non-functional requirements.	A development team wants to write a new specification document to improve the functionality of a system. The team decides to assign the task to a new team member, to obtain a fresh perspective. While considering possible improvements, the developer also studies old code commits to understand how the system evolved. This effort takes the developer to ask questions about the rationale of the changes.

Table 2.2: Tasks for which rationale for code commits is needed (cont.)

Task	Subtask	Description	Example (abstracted from participants' responses)
Working on Bugs	Reproducing	Reproducing the situation in which a previously reported bug was observed.	A developer tries to reproduce a race condition that happens occasionally. The developer investigates the edge cases introduced in a previous commit to guess when the race condition emerges, which leads the developer to ask questions about the rationale of the changes.
	Reporting	Reporting the existence of a bug through formal/informal communication methods to inform the team members.	A developer discovers a bug in the codebase. To write a bug report, the developer wants to refer to a particular commit as a suspect of introducing the bug. In the effort to find the suspect commit, the developer needs to understand the rationale of the investigated commits.
	Triaging	Evaluating a reported bug in terms of validity, severity, urgency, and needed work, before assigning a developer and a due fix time.	A Project Manager (PM) uses git blame to figure out who introduced code associated with a newly reported bug. In this effort, the PM needs to understand the rationale of the multiple code changes that introduced the buggy code.
	Debugging	finding the source code that contains the bug and fixing it.	A developer investigates his/her assigned bug, returns to the previous commits, reads their code, reflects on their rationale, understands how the bug was introduced, and writes a fix.
	Postmortem Analysis *	Exploring and analyzing bugs, which might have already been resolved, for research or to improve productivity.	A PM studies the history of bugs to take preventive measures in the future. In this effort, the PM needs to understand how these bugs were introduced and the rationale behind the commits introducing them.
Testing	Writing	Writing test code that checks the behavior of the software against its specifications.	A developer wants to write several tests for new functionality. The developer studies the latest code commits to design his/her test strategy, which also involves understanding the rationale of those changes.
	Running	Running tests to compare the behavior of the software with its specifications.	Before a developer runs a test suite for a system, the developer looks at the code commits to design his/her test strategy. In this effort, the developer tries to understand the rationale behind the changes performed to the code under test.

Table 2.2: Tasks for which rationale for code commits is needed (cont.)

Task	Subtask	Description	Example (abstracted from participants' responses)
Communication	Learning	Learning software best practices, conventions, technologies, skills, and tools.	A researcher looks at the commits in the repository of an open-source ML library to learn the undocumented mathematics of the ML approach. This leads the developer to ask questions about the rationale of the changes in the commits..
	Coordinating	Coordinating with members of the same or another team to achieve a smooth integration.	Team A and Team B are developing intersecting functionality. A developer in Team A looks at the commits of Team B to avoid redundancy, sees how Team A can leverage parts of Team B's code, and builds a shared vision of the two teams. In this process, developer A requires understanding the rationale for some of the code changes by team B.
	Mentoring	Advising, guiding, and one-to-one teaching another software developer.	A mentor looks at the mentee's code and commits to checking how well the mentee practiced the design process, which leads the mentor to ask questions about the rationale of the mentee's changes.
Documentation	Searching	Searching for reference documentation, tutorials, commit messages, or any other sources of knowledge that can help accomplish a task.	A developer looks for the documentation of a particular piece of code. Because inline comments in the code are unclear, the developer runs git blame to look for related commits. The developer reads the related commit messages and asks about the rationale of the code changes.
	Writing	Writing documentation of software artifacts (<i>e.g.</i> , tutorials, inline code comments, or user-generated JavaDoc) or the software process (<i>e.g.</i> , commit messages, process documents, or meeting minutes).	A developer works on writing a tutorial on using a library that the team newly migrated to. Since the library documentation is incomplete, the developer looks at commits from the library repository to understand the rationale behind some of its functionality.
	Reading	Reading documentation of software artifacts or software processes.	A developer reads the project changelog and figures out where the support for a library was added. The developer reads the commit to understand the rationale of this change.

Table 2.2: Tasks for which rationale for code commits is needed (cont.)

Task	Subtask	Description	Example (abstracted from participants' responses)
Tools	Discovering	Discovering various aspects of a tool or library before adoption, such as its supported features, popularity, and version history.	A software team is assessing whether to replace an existing component with a new tool that is not fully documented yet. A developer goes through the commits in the repository of the new tool to discover its capabilities and to look for what the team needs. In this effort, the developer needs to understand the rationale of the historical changes to the tool.
	Installing	Installing tools for developers to use for their tasks.	A developer tries to install Docker Composer, and finds it incompatible with other installed tools. The developer finds another repository that was previously encountered and solves the same problem. The developer then tries to understand the rationale behind the changes that solved the compatibility issue.
	Using	Using external tools that are not part of the default API libraries, either by calling from code or by impacting the execution environment.	A developer wants to use some GUI controls for their website. The developer looks at a similar code repository that uses these controls and goes through its commits to learn how to use the GUI controls. In this effort, the developer needs to understand the rationale of code commits.
	Building	Compiling the source code for the software to launch and execute correctly.	A developer's code is not building successfully. The developer goes back to an older code version to determine why the build was successful then. The developer reads commits related to the problematic code and reflects on the rationale behind the changes before the build starts failing.
Project Management	Checking out	Checking out a software artifact to a separate machine to work on it independently.	A developer pulls the changes made to an existing branch. The developer looks at the pull commits to understand their impact on his/her tasks and eventually needs to understand the rationale behind these changes.
	Reverting	Reverting to an older version of the software repository.	A developer inserts an incorrect commit on the project. The team does not know exactly which commit is the problematic one. A developer goes through the changed files for each commit to backtrack the problem. The developer needs to understand the rationale of various code changes in the history to revert those inserted after the incorrect one.

Table 2.2: Tasks for which rationale for code commits is needed (cont.)

Task	Subtask	Description	Example (abstracted from participants' responses)
	Deploying *	Moving software from one controlled environment to another, <i>e.g.</i> , merging the development branch into the production branch, releasing the software to users, moving the software to a different environment on the customer end, or staging to test the software using real data.	Some tests fail after moving the project from the development branch to the production branch. A developer investigates the commits in the repository and finds a configuration change that might affect the deployment. The developer reads the commit thoughtfully to find out the goal of each configuration flag and why they were introduced.

* Subtasks observed in this study that were not originally in the codebook (*i.e.*, that were not observed in Begel and Simon's study of developer tasks [27]).

Our participants reported needing to understand the rationale of code changes in a wide diversity of software development tasks: *programming, working on bugs, communication, tools, documentation, project management, testing, and specifications*. Their responses covered all (8) tasks in the codebook and the majority of subtasks (23 out of 34). Another signal of the wide diversity of tasks requiring understanding the rationale of code changes is the fact that a good portion of our participant responses (35%) described multiple tasks.

For many of the reported tasks and subtasks, it is intuitive that developers would need to understand the rationale of code changes for them, *e.g.*, in *code review*, or when *reading* documentation. Others are initially surprising (*e.g.*, when *installing* or *discovering* tools) but become less surprising after hearing our participants' examples (column 4 in Table 2.2). For example, when a tool is poorly documented, some practitioners resort to understanding its code changes to understand its behavior. Another possibility is that the tool is well documented, but software developers are discovering the tool through unofficial search engines. Liu, Li, and Tilevich [30] recent study of the practice of code search discovered that developers use various tools rather than specialized code search engines.

Our main takeaway from these observations is that when the rationale of code changes is well documented, it could help software developers in many different situations. The help could even be in less intuitive situations, such as when they need to understand the rationale of code changes in other codebases they do not own (*e.g.*, to understand the behavior of external tools). Good documentation of the rationale of code changes should consider the fact that developers seeking it may be involved in a wide diversity of tasks, which could cause them to search for it in different places and search for different dimensions of it under different contexts. Furthermore, it would also be beneficial if the documentation of the rationale of code changes could be understood even by people outside the development team of the software project (24% of responses mentioned seeking the rationale of code changes outside their project).

We hope this new understanding of how many tasks can benefit from a well-documented rationale of code changes encourages developers to document it well and appropriately for a wide diversity of tasks (*i.e.*, in adequate locations and levels of detail suitable for many contexts).

We also learned from the tasks and subtasks our participants reported but were not initially included in the codebook: *postmortem analysis* of bugs, and software *deployment*. We believe that these tasks were not observed initially by Begel and Simon [27] because they are intuitively performed less regularly and are therefore harder to observe. However, our participants reported these tasks in their experiences of needing to understand the rationale of code changes: they needed it to understand why an old bug was introduced in the first place (to avoid similar ones in the future), and they needed it to have a stronger understanding of the new changes that they were deploying. We learn from these observations that the rationale of code changes is needed for both frequent and infrequent tasks.

Similarly, our participants did not mention some developer tasks from the codebook. Some of

them were tasks that do not necessarily involve code changes: tools (*finding*), specifications (*reading*), and communication (*finding people*). Others are tasks that are more focused on *providing* information than *requesting* it: communication (*persuasion, meeting prep, interacting with managers, and teaching*). The remaining ones are tasks that we believe practitioners may be more likely to think of them as part of a broader task (and thus may not mention them specifically): programming (*commenting*), project management (*check in*), communication (*asking questions, email, and meetings*). Overall, there are more tasks and subtasks that require understanding the rationale of code changes than those that may not.

Finally, we also measured the ratio of times our participants mentioned any task or sub-task (figure 2.2A). The tasks that our participants most often reported were *programming* (32% of mentions), followed by *working on bugs* (22%), and *communication* (14%). Within them, the most popular subtasks were *debugging*, *code reading*, *code review*, *learning*, and *documentation search*.

Most often reported tasks teach us that the most typical scenarios for practitioners needing to understand the rationale of code changes are those that involve *debugging* code or *reading* it to *learn* something about it, often during *code review*, and by also *searching* for its documentation. Therefore, developers documenting the rationale of code commits to help other developers in the most typical scenarios should document the aspects of the code change that could be informative for these tasks.

However, in addition to those scenarios, there is a long tail of tasks that were much less often mentioned in our practitioners' scenarios but still required understanding the rationale of code commits. Therefore, developers aiming to provide extensive documentation of the rationale of code changes should document all the aspects of rationale that could be relevant for all those tasks and subtasks.

2.5 RQ: How often do developers seek the rationale of code commits within their team’s internal code vs. external code?

2.5.1 Research Method

We used the same method to answer this research question as we did for the previous tasks research question (see Section 2.4.1). We again analyzed our participants’ descriptions of their experiences investigating the rationale of code commits, this time looking for whether they were investigating internal or external code. As with the tasks research question, two coders performed open coding individually and resolved disagreements afterward in a joint focused-coding session.

2.5.2 Results

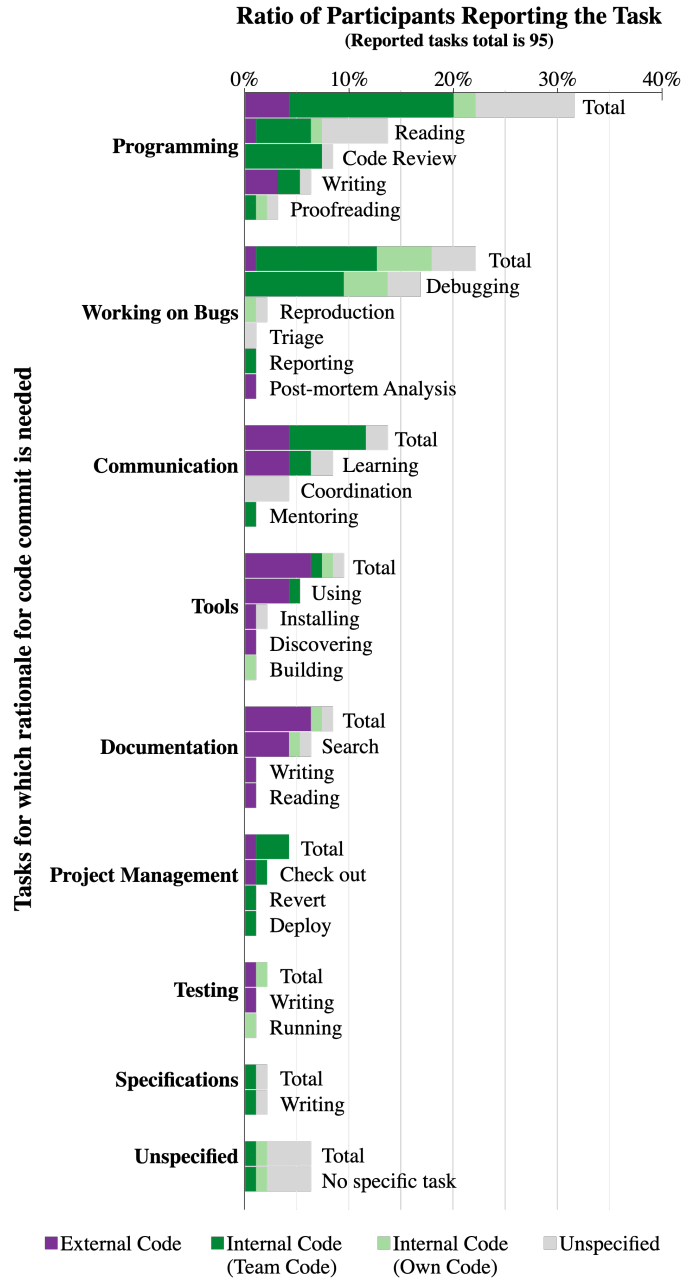
We display in figure 2.2B the breakdown of how often our participants reported investigating internal code vs. external code. Internal code refers to code changes owned by our participants or developers in the development team for which our participants are contributors to the team’s project. External code refers to code changes owned by developers from external projects’ teams to our participants. As we did for the previous research question, when it was unclear which source code was being investigated, we coded the answer as “unspecified”.

Our participants mostly investigated the rationale of internal code changes (53%), but they also investigated it for external code changes (24%). The external code could be code changes to open-source libraries or external projects in the same organization related to the participant’s current project. An example of external code changes, as one of our participants


mentioned, is “*I was trying to look at forks of that repository and how other people solved those issues.*” In most cases, our participants were investigating code changes that were written by other members of their team— internal code (team code) (41%); However, in other situations, the investigated code change was written by the same person investigating it — internal code (own code) (12%). Some of our participants had forgotten the rationale of their own code changes after some time had passed. One participant said, “*It was just a decision made by the person who wrote that who also happened to be me.*” We also connected our participants’ targeted code change to the software development task they mentioned performing when seeking the rationale. We report targeted code change per task in figure 2.2A.

In figure 2.2A, we observe three types of tasks. The first type includes tasks for which our participants reported needing to understand only the rationale of internal code. Examples of these are: *code review* and *debugging*. This observation supports an intuitive guess that this type of tasks would only need rationale of code commits in internal code, since they often only involve internal code.

Tell me about one time in which you investigated a code commit to understand its rationale. Why did you need to find the rationale for that specific code commit? What was the rationale for that code commit?



(A) Targeted code change per task



(B) Targeted code change

Figure 2.2: Tasks for which rationale for code commits is needed

Other types of tasks are those for which our participants reported needing to understand only the rationale of external code. Examples of these are several tasks under the *tools* and the *documentation* categories. We were less surprised that our participants would need to understand the rationale of code changes to external code when the developers investigate tools since tools can often be developed externally. One participant was reading the changelog of a tool he/she was using.

“ ...one of the software that I work with is called MXNet. It is a machine learning, deep learning library, and I saw in the changelog for MXNet that they added support for one of the libraries it depends on, open CV, which is a computer vision library. It changed how they loaded JPEG, and they included support in MXNet ... In the MXNet changelog, they made reference to the commit in open CV where this occurred. So I looked at that commit and in the detailed commit message... ”

we were more surprised that they also needed to understand the rationale of code changes to external code when writing documentation. However, our participants often created internal documentation about external code to better keep track of the external code's behavior. For example, one of the participants was trying to write a tutorial for his team about a specific tool, and so he had to discover the tool in various ways, which led to multiple questions about the rationale of code changes to the tool.

Finally, for other tasks, our participants mentioned the need to understand the rationale of both internal and external code changes. Examples of these tasks are *learning* and *code writing*. It is more intuitive to understand why *learning* could require understanding the rationale of external code changes to learn from external projects. However, we learned that our participants also often have that need when writing code. For example, when they are trying to port or incorporate external code into their internal project. One participant mentioned:

“ I worked on a derivative of the Bitcoin codebase for two years (named LBRYcrd). There were many occasions when I had to trace the history of some lines of suspicious code. Often I would find bugs in the LBRY-specific portion, but rarely in Bitcoin... Just recently, as I was working on the Golang version of that, I had to understand where the UPnP code portion came from to see if the author had created a newer version (that might have fixed some bugs I was seeing with it). ”

These observations may explain why other studies observed that finding the rationale of code changes was sometimes easy, sometimes hard [3, 5]. It would be harder to find when the developer is seeking rationale for *external* code changes, since they would have access to fewer documentation resources for external code changes than for internal code changes. The fact that our participants need rationale of external code could explain the challenge of finding rationale and its consistent occurrence as information need in information needs’ studies like [3, 5].

More generally, our participants needed to understand the rationale of external code changes most of our studied development tasks (all but “specification design” in Table 2.2). This observation motivates the creation of external-facing tools and practices to improve the documentation and explanation of the rationale of code changes, not only to internal stakeholders but also to external ones. We discuss ideas for how to achieve this in Section 5.4.

2.6 RQ: Which components do software developers decompose the rationale of code commits into?

2.6.1 Research Method

Data: The data for this research question is a large set of rationale for code commits' components. We proposed nine components based on the literature references to rationale. In addition, our participants proposed 18 different components of rationale for code commits.

Analysis method: Our goal is to create a mental model and derive a taxonomy of the components. Therefore, we used card sorting [31] to discover the components of rationale for code commits. Card sorting is a widely used inexpensive method with three phases: preparation (participants selection and cards creation), execution (cards sorting into meaningful groups), and analysis (hierarchies formation).

For the card sorting preparation phase, we prepared the cards of all the components from the preliminary model and the participants proposed final model. Then, for the execution phase, we (two coders) performed individual open card sorting. We individually sorted the cards without using predefined groups. During the individual sorting, we aggregated those cards that described similar components. For example, we aggregated into “*Side Effects*”: the preliminary component “*Costs*”, and the “*Merge Conflict/Success*”, “*Limitation*”, and “*Impact*” components that were mentioned by different participants. After that, both the two coders collaboratively consolidated the two sets of individually-aggregated components, comparing them and deciding on disagreements. Then, we characterized each of the resulting aggregated components with a name, a question, and an example answer to the question based on a hypothetical commit. Finally, we categorized the resulting components into themes for the analysis phase.

Analysis evaluation: Our discovered model of rationale for code commits (Table 2.3) is of 15 different components. The fact that many participants added and some removed components suggest that our participants were not strongly biased towards simply agreeing with the preliminary model (Table 2.1). Also, although we interviewed 20 software developers, we reached saturation in the 15th interview.

2.6.2 Results

We display in Table 2.3 the model of rationale of code commits that we discovered. It represents the union of all the models that our participants reported. As we discussed in Section 2.3.1, we obtained this model aggregating all the components that were mentioned by at least one participant in their final interview model of rationale of code commits. Each participant built their final model by adding and removing components to the preliminary model, while also considering their own rationale decomposition. Altogether, our participants reported 27 components of the rationale of code commits, adding 18 components to the nine components in the preliminary model. Since many of those components reported very similar concepts, we aggregated them using card sorting to obtain the final model that we show in Table 2.3. This resulting model of rationale of code commits includes 15 components into which developers decompose it. We categorized the resulting components into four themes.

Our goal with this rationale model was to gather an extensive set of specific components of the rationale of code commits, which developers may be looking for when they need it. For that reason, when a participant removed a component from the preliminary model, we still kept it in our resulting final model (in Table 2.3). Besides, only a few participants removed components.

Table 2.3: Resulting model of the rationale of code commits

Theme	Component	Component Expressed as Question	Example Answer
Change Objective	*Goal	What did the developer want to achieve?	The code is this way because the developer <i>wants to modify</i> the usage of try/catch blocks to account for unexpected Exceptions.
	*Need	Why did the developer need to achieve that?	The code is this way because the developer <i>needs</i> to improve Exception handling by June 1st as per the new <i>company demand</i> to eliminate exceptions before release.
	*Benefits	What is the benefit of what the developer wants to achieve?	The code is this way because handling exceptions that were not considered before will benefit in increasing the system's quality.
Change Design (pre-implementation assessment)	Constraints	What were the constraints limiting the developer implementation choice?	The code is this way because the developer choices are <i>limited</i> by the team development <i>guidelines</i> that <i>prohibit hard-coded String use</i> .
	Alternatives	What other alternatives did the developer have?	The code is this way because the <i>alternative bucket sort implementation</i> option is not feasible since the maximum value is unknown.
	*Selected Alternative	Why did the developer make those specific changes and not others?	The code is this way because heap sort has the <i>advantage</i> of being space <i>efficient</i> and has a <i>predictable speed</i> . Other sorting algorithms options are not as efficient and predictable.
	Dependency	What other changes does this change depend on?	The code is this way because it <i>depends</i> on the <i>API response format</i> , which needs to be updated to provide JSON format.
Change Execution	Committer	Who changed the code?	The code is this way because it was introduced by <i>Developer X</i> , who was our <i>short-term consultant hired to improve the security</i> of our software system.
	Time	Why were the changes made at that time?	The code is this way because it was introduced <i>four months ago</i> to meet <i>3.0 release cycle</i> .
	*Location	What artifacts were changed?	The code is this way because, in our MVC architecture, the <i>model</i> , <i>view</i> , and <i>controller</i> are updated together when introducing a new data field.
	*Modifications	What specific changes were performed in the artifacts?	The code is this way because the developer <i>altered</i> the user interfaces' look and feel, including <i>color</i> and <i>layout</i> .

Table 2.3 Resulting model of the rationale of code commits (cont.)

Theme	Component	Component Expressed as Question	Example Answer
	Explanation of Modifications	What are the details of the implementation?	The code is this way because the look and feel are <i>altered by</i> : (1) Changing all interfaces colors to match color palettes provided by web accessibility guidelines and (2) Changing all interfaces' layouts to be responsive to screen size.
Change Evaluation (post-implementation assessment)	*Validation	How did those specific changes achieve the goal?	The code is this way because the goal is to account for an edge case, which newly added <i>test cases</i> show the edge cases examples and verify code success in handling the edge case.
	Maturity Stage	How mature is this code?	The code is this way because it is an <i>experimental hack</i> created to explore ways to fix a persistent bug.
	*Side Effects	What are the side effects of the change?	The code is this way because of <i>side effect</i> mitigations; Temporary if statements are added to <i>avoid integration test failure</i> until the API update.

* Components that were included in the preliminary model of rationale of code commits. We extended the preliminary component *Costs* to *Side Effects* to include other side effects mentioned by participants *e.g.*, *Impact*.

When participants decided to remove components from the preliminary model, they mentioned two main reasons: overlap with other components, and the component being out of scope. In terms of overlap among components, one participant thought that *goal* and *need* could be the same most of the time and preferred to merge them, deleting the *goal* component. Another thought that *need* is included in *benefits* and *cost*, deleting the *need* component. Another participant deleted *benefits* because it is included in *goal*. Another one considered *location* as part of *modification*.

We believe it is possible that different components' answers can be the same in some cases. For a single code commit, components of the same theme (see Table 2.3) may have very similar answers to their expressive question. However, they will be different in many other cases, making it worthwhile to separate those components. We illustrate the differences between components in Table 2.3 by including the components expressed as questions and

different example answers for different ones.

Other participants removed components that they considered out of the scope of rationale. From our 20 interview participants: two participants removed *modifications* because they considered it too low-level; three participants removed *location* because it would not tell why the changes were made; three participants removed *alternatives*, e.g., “*alternatives is not something that you actually implement!*”; and one participant deleted *validation*, saying that “*validation answers why the code is correct, not the rationale*”. Despite these disagreements, the majority of our interview participants (18, 17, 17, and 19, respectively) considered that these components do belong in the rationale of code commits.

Furthermore, our participants generally provided positive comments about the preliminary model — describing it as e.g., “*a good model,*” “*detailed,*” “*thorough,*” “*comprehensive,*” “*holistic,*” or “*exhaustive.*” They thought that the model “*formally define[s] rationale*” and that “*the components seem to be related to each other, but classified differently to each other.*” One participant said that the model is a

“logical framework for thinking through rationale because [it is] a sort of wide-open concept. It is a little bit hard to know how to think about [rationale]. [The model] makes sense as a directed way to understand a specific commit or a series of commits. Why they are the way they are.”

Many participants added components to the preliminary model. As we mentioned earlier, we used card sorting to aggregate them into the preliminary model and with each other. The 18 components proposed by participants were: *technical requirement, timeliness, documentation, guidelines, non-feasible alternative, opinion selected alternative, constraints, dependency, committer, time/date, explanation of modifications, result, environment, scope for future development, quality, merge conflict/success, limitation, and impact*. For each added component, we also asked participants to describe them with a name, expressive question, and example answer to the question.

The fact that many participants added and some removed components suggest that our participants were not strongly biased towards simply agreeing with the preliminary model. More importantly, it also suggests that different developers seek different components at different times. Our study throws light into this phenomenon. Thus, we posit that the rationale of code commits would be much easier to comprehend, search for, and document when expressed as its components– not necessarily all of them at all times, but the ones relevant to each situation.

2.7 Threats to validity

To answer our research questions, we asked both open and quantitative questions. We scheduled the interview sessions to be relatively long (two hours), ensuring that we gave the participants enough time to express their ideas and share their thoughts. At the beginning of each interview section, we asked the participants to “*answer the questions in [their] own words and provide as much detail as [they] feel is relevant to address each question*”. We also placed an open question at the end of the interview to allow the participants to share any additional information about the topic.

The methods we used in our study, interviews and surveys, can be affected by bias and inaccurate responses. This effect could be intentional or unintentional. We gave gift cards to the interview participants and some survey participants, which could have biased our results. We indicated that the compensation is for the time spent and not the answers given to mitigate these concerns. We repeatedly and constantly used phrases to encourage the participants to provide their own honest opinions, using the phrase “*based on your experience*” in most of the questions. We also clearly indicated that the participants should “*feel free to change/add/delete components or not.*” Sometimes, we also indicated that “*there*

is no right or wrong answer; we are interested in what you think and your perspective.”

We also took multiple steps to reduce potential confirmation bias [23] resulting from using a preliminary model. We asked participants to describe their own examples and decomposition of rationale into components before they ever saw the preliminary model. We formed the preliminary model based on knowledge from the research literature and presented it neutrally. The fact that the preliminary was largely extended from 9 components into 15 validates that potential confirmation bias was minimal in our study.

Another threat to the validity of our study is drawing conclusions based on recollected memories [32]. We are interested in capturing developers' opinions about what components constitute rationale, independently of how accurate their recollection is. We encouraged participants to take their time to recall situations and to report the components that mattered in their experience.

Our studied developers may not fully represent the whole developer population. To mitigate this threat, we recruited a diverse population with diverse types and amounts of experience (figure 2.1). Our studied population was similar to the ones previously studied in the literature since we obtained similar answers for our two questions about rationale that were already studied by Tao et al. [5]. Furthermore, we have reached saturation for qualitative analysis of our open-ended interview questions.

2.8 Summary

Any efforts aiming to improve the process of discovering the rationale of code changes will require a good understanding of the software developers' need and its context. We applied a qualitative study to establish a common understanding of developers' need. We performed a series of interviews with software developers to discover the tasks for which rationale for code commits is needed, and the components into which developers decompose the rationale of code commits. We found that software developers decompose rationale of code commits into 15 components along four themes. We also found that they need rationale for code commits to complete various software development tasks.

Our decomposition of rationale for code commits is a descriptive representation that practitioners can use to improve their documentation and communication of rationale. Additionally, researchers and tool builders can support the management of the rationale of code commits using our discovered components of rationale and the experiences of software developers with them. While we do not expect practitioners to document all components in all situations, they now have an extensive list of components to judge which ones are relevant for each situation.

Chapter 3

Surveying developers' experiences with rationale of code changes

In this chapter, we aim to explore software developers' *experiences* needing, finding, and recording rationale for code changes and its individual components.

3.1 Background and Motivation

In the specific scope of code changes, rationale is a major information need. Many research studies about information needs convey the importance of understanding the rationale of code changes. Rationale is the most common [3] and important [5] information need to understand from code-change history, and it is very frequently sought in various software development tasks including code review [9, 10, 13, 14]. Unfortunately, it can also be quite difficult to find an answer for rationale of code changes [5, 7]. We build on those studies' observations, aiming to provide an in-depth survey of developers' experiences. We aim to understand developers' efforts and issues associated with rationale of code changes need, finding, and recording.

We take a developer-centric approach to study software developers' *experiences* when seeking rationale for code changes and its individual components. We used a mixed-methods approach in our study, involving interviews and surveys of software developers. We are set

out to discover imbalances in developers' need, finding, and recording experiences, and the factors leading developers to give up their search for rationale. Such understanding of developers' efforts and issues will reveal areas of improvement in developers' practices. Those areas could be targeted by researchers to: (1) better document, (2) develop support tools, and (3) automate documenting rationale for code changes.

3.2 Research Questions

To explore software developers' experience needing, finding, and recording the rationale of code commits, we ask the following research questions:

- ***RQ: What is the experience of developers needing, finding, and recording the rationale of code commits?*** We investigated this research question to understand the effort developers dedicate to seek and document the rationale of code commits. Tao et al. [5] discovered that finding the rationale of code commits is *very important*, and it is *easy* or *hard to find* depending on how well-documented it is [5]. We replicate those two questions of their study, and extend it by asking developers five additional questions in three different contexts: needing, finding, and recording rationale.
- ***RQ: What is the experience of developers needing, finding, and recording the individual components of the rationale of code commits?*** We studied how developers need, find, and record different components differently. This study will enable developers to improve their documentation of rationale efforts, *e.g.*, by focusing on documenting the most needed or most hard-to-find components.

- **RQ: What makes software developers give up their search for rationale of code commits?** We wanted to discover issues that may limit software developers from effectively finding rationale for code commits. Identifying these issues will inform future research efforts that are targeted to support information needs' finding
- **RQ: Would comparing the experience of developers needing, finding, and recording the individual components of the rationale of code commits with each other reveal areas for improvement?** We performed a cross-dimensional study (*i.e.*, comparing need vs. finding vs. recording components) to investigate areas for improvement in current recording and finding practices of rationale of code commits. Identifying gaps, *e.g.*, between needed and recorded components, will provide valuable recommendations for developers who want to improve their documentation of rationale for code commits.

3.3 Approach

To understand these aspects of the rationale of code changes in depth, we used the same mixed-methods approach of the previous study (Section 2.3), combining developer interviews and a survey. Other methods (*e.g.*, mining software repositories) could be applied to study different aspects (*e.g.*, the extent to which code commits fulfill the need for their rationale), but our goal was to understand the developers' need in depth from a developer perspective.

3.4 RQ: What is the experience of software developers needing, finding, and recording the rationale of code commits?

3.4.1 Research Method

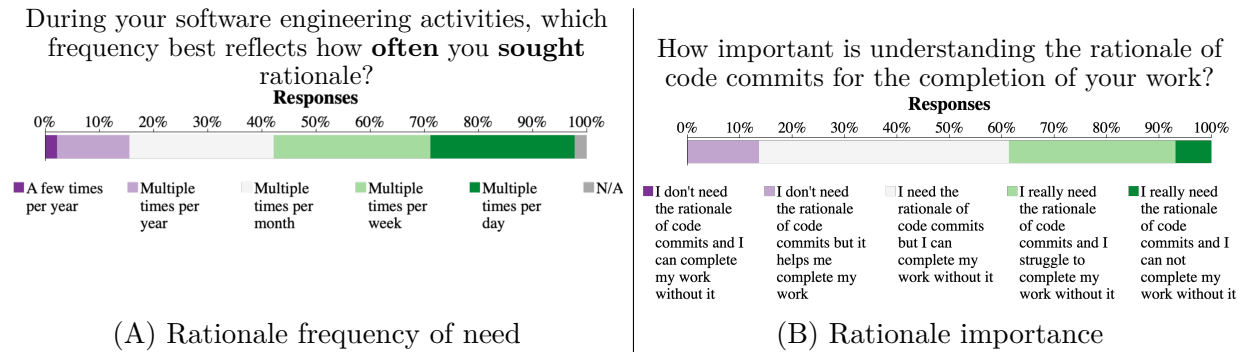
The *data* for this research question is quantitative. We asked both the interview and survey participants the likert-scale-style question shown in figures 3.1, 3.2, and 3.3. To *analyse* the participants' experiences responses, we compare and report the responses' statistics.

3.4.2 Results

figures 3.1, 3.2, and 3.3 show the interview/survey questions which we asked about the developers' needing, finding, and recording experiences along with the distribution of the participants' responses for each question.

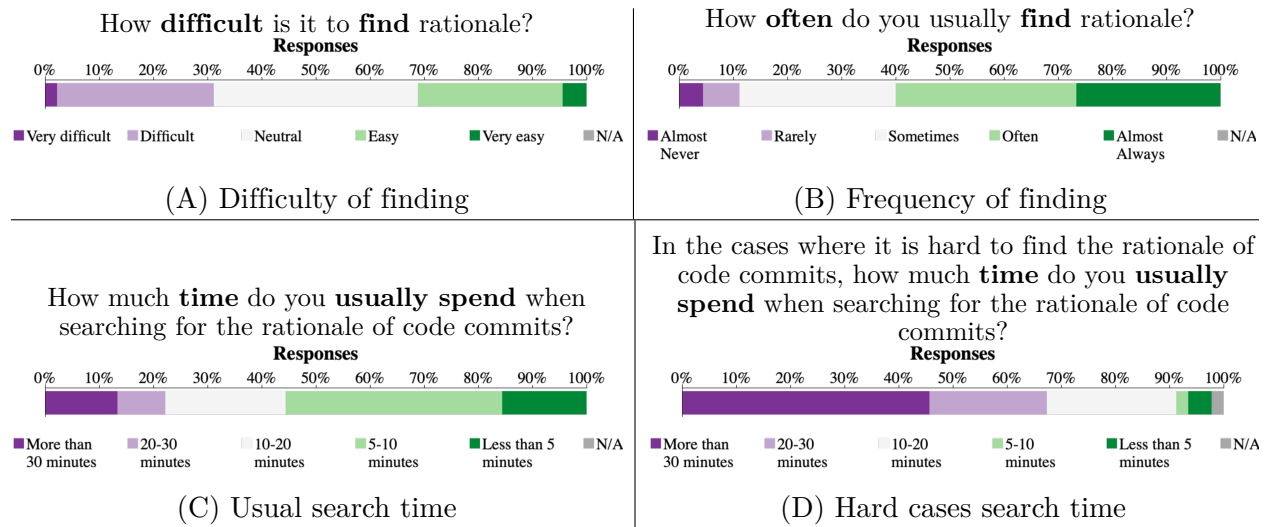
Need: The participants of our study reported needing to seek rationale with diverse frequencies (see figure 3.1A): multiple times per day (27%), multiple times per week (29%), multiple times per month (27%), multiple times per year (13%), and a few times per year (2%). Overall, the majority (56%) of our study participants need rationale relatively frequently: multiple times per week or more often. The inconsistency of the need for rationale for code commits is because of the diversity of the software developers' roles and their work activities. One participant said

“I do a lot more than just software engineering on a yearly basis. And so there are periods of time when I am doing primarily software engineering, and there are large periods of my work time that I am not.”

Figure 3.1: Experience of developers **needing** rationale

When asked about how important it is to understand the rationale of code commits, 86% of our participants reported needing the rationale of code commits (see figure 3.1B), from which: 7% cannot complete their work without understanding it, 30% struggle to complete their work without it, and 46% can complete their work without it but still need it. The remaining 17% do not need the rationale of code commits, but 15% of the 17% report that it still helps them complete their work. A very similar question was studied by Tao et al. [5] [5], whose participants “*generally considered knowing the rationale of a change as the top priority in change-understanding tasks*”. Our finding is aligned with theirs since a majority of our participants reported needing the rationale of code commits, which validates that we are studying a similar population of developers.

finding: Our participants’ responses in figure 3.2A indicate that the difficulty of finding the rationale of code commits, in general, is not easy nor difficult. Software developers (on average) selected neutral difficulty in finding the rationale of code commits. This finding also generally agrees with Tao et al. [5]’s, since their participants reported that the rationale of code commits was generally easy to find, but sometimes hard, depending on “*the availability and quality of the change description*” [5].

Figure 3.2: Experience of developers **finding** rationale

Regardless of how hard it is, we were also interested in how often developers end up finding the rationale of code commits altogether. For this aspect, our study participants' responses are positive (see figure 3.2B). Most software developers find the rationale of code commits often or almost always. Only a few participants (11%) rarely or almost never find the rationale of code commits.

In addition to studying whether software developers find the rationale of code commits, we also studied how much time they spend searching for it. figure 3.2C and 3.2D shows the times that our participants reported spending when searching for rationale. In the usual cases, slightly more than half of our participants (55%) spend less than 10 minutes. However, in the hard cases of searching for rationale, only slightly less than half of our participants (46%) spend more than 30 minutes searching for the rationale of code commits. One participant said about the time they spend searching for rationale in the hard cases that it “*depends how responsive the other person is.*” A considerable amount of time, 68% of the participants spend more than 20 minutes, is spent by software developers when it is hard to search for

the rationale of code commits. When considering the relatively high frequency with which developers search for rationale of code commits, it can be a rather time-consuming task.

Recording: Regarding the frequency of recording rationale in general, figure 3.3 shows our participants' responses. Our participants reported high involvement in recording the rationale for their code changes.

During your software engineering activities, how **often** do you **record** rationale?




Figure 3.3: Experience of developers **recording** rationale

The majority of them (71%) reported recording the rationale of code commits often (31%) or almost always (40%). These ratios are very similar to the frequencies with which they report needing it and finding it. However, there may be multiple explanations for why these two ratios are similar. It could be that the teams to which our participants belong are generally diligent about documenting the rationale for their code changes, and that is why they report finding it at similar frequencies when they need it. However, there could be more nuance to this observation: our participants may sometimes find the rationale that was documented, and other times find it after asking their colleagues (because it was not documented). Also, the rationale found by our participants is a reflection of the documentation habits of their teammates than theirs. Our observation motivates further study of the extent to which the specific rationale that developers document is what ends up helping their teammates later find it (*e.g.*, with observational studies or mining of software repositories). We take one further step in understanding the similarity between the frequency of recording and finding the rationale of code changes in more depth (see Section 3.7), in which we observed that such similarity is generally preserved for individual components of it.

3.5 RQ: What is the experience of developers needing, finding, and recording the individual components of the rationale of code commits?

3.5.1 Research Method

Just like the previous research question, the *data* for this research question is quantitative. We asked a likert-scale-style question shown in figure 3.4. To *analyze* the participants' experiences responses, we compare and report the responses' statistics. For the individual rationale components, the answers include only the interview participants that included the components in their final model. Whenever we aggregated components through card sorting, we also aggregated the responses about the experience. The reported experiences with rationale components combine the aggregated interview and survey responses.

In addition to the responses' statistics, we wanted to cluster components with similar experiences. Clustering the components made it simpler for us to compare the experiences. We used Scott-Knott [33] clustering algorithm to group the components which have similar software developers' experiences. Scott-Knott is a hierarchical clustering algorithm that serves in the Analysis of Variance (ANOVA) contexts. For the individual experiences (need, finding, and recording), the algorithm compares the means of all component's responses. This comparison results in a non-overlapping grouping of rationale components. We present the algorithm outputted groups by the red border lines in figure 3.4.

3.5.2 Results

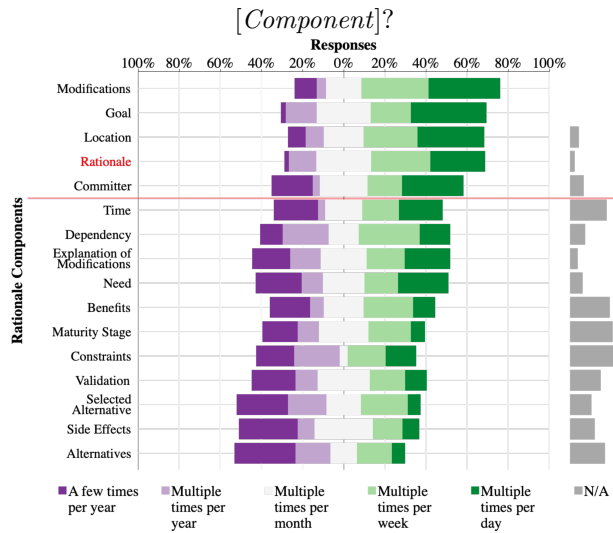
We plot the distribution of answers to our questions about individual components of the rationale for code commits in figures 3.4A–3.4C. We cluster components into similar groups according to the mean value of their responses using the Scott-Knott [33] algorithm. We sort the components in our figures by the mean value of their responses, and we use red horizontal lines to separate clusters. As a reference point, we also include in each figure the responses that our participants gave for rationale in general.

Need: figure 3.4A-i shows the distribution of responses for how frequently developers need each component of the rationale of code commits. Overall, the frequency with which developers need different components of the rationale of code commits is highly similar for all components and rationale itself in general. In this case, the Scott-Knott algorithm returns only two very-similar clusters. While some of the most often needed components (like *modifications*, *location* or *committer*) are normally automatically recorded by revision control systems, many other components are similarly often needed and are not recorded automatically (like *need* or *dependency*, or *constraints*). These results show that practitioners would benefit from regularly recording these frequently-needed components.

figure 3.4A-ii shows the relative importance of each component to understand the rationale of code commits reported by developers. These results show that most developers mentioned that most components are important enough to understand the rationale of code commits better if they knew that component. We also observe that developers wanting to document the most important component of rationale should focus on documenting the *goal* of their changes, since the other most-important components (*modifications*, *location*) are already recorded by revision control.


(A) Experience of developers **needing** individual components of rationale

During your software engineering activities, which frequency best reflects how **often** you sought



(A-i) Frequency of need


How **important** is finding each **component** (for understanding the rationale of code commits)?



(A-ii) Components importance


(B) Experience of developers **finding** individual components of rationale

For the components of rationale of code commits that you seek, how **often** do you usually find



(B-i) Rationale components frequency of finding

For the components of rationale of code commits that you seek, how **difficult** is it to find



(B-ii) Rationale components difficulty of finding

Figure 3.4: Experience of developers with individual components of the rationale


(C) Experience of developers **recording** individual components of rationale

Figure 3.4: Experience of developers with individual components of the rationale (cont.)

finding: figure 3.4B shows the relative frequency and difficulty of finding reported by developers for each component. Unsurprisingly, the most frequently found components (and also the easiest to find) are those automatically tracked by revision control (*committer*, *modifications*, and *location*), followed by *goal* and *time*. However, the frequency (and easiness) of finding drops quickly for all other components, bringing our attention to a clear problem in finding the remaining components. These results highlight the need to improve documentation for the other components since they are hard to find. This clear divide could also explain why developers talking about rationale (in general) say that sometimes it is much harder to find rationale than other times [5] and it takes longer (figure 3.2D).

Recording: figure 3.4C shows the relative frequency with which developers reported to record components of the rationale of code commits. Again unsurprisingly, the most frequently recorded components are those recorded automatically by revision control, but again the frequency of recording drops dramatically for the remaining components (which

probably explains why they are hard to find). These results show that, even if developers claim to frequently record rationale in general, there are many components that they are not recording frequently (even if they are relatively often needed).

3.6 RQ: What makes software developers give up their search for rationale of code commits?

3.6.1 Research Method

We asked our (20) interview and (26) survey participants an open question about when they give up their search for rationale of code changes. We analyzed their responses qualitatively, using open coding to extract the factors that our participants reported. Multiple coders were involved in the coding. We reached saturation in our observed tasks in response 20 of 46.

Data: For this research question, we asked our participants an open question: “*When do you give up the search for rationale of code changes?*”. We asked this question during our interviews and survey. We used a survey in addition to the interviews to reach saturation in our observations — an initial analysis of the interview responses showed that we were still making new observations in the last responses.

Analysis method: We analyzed our participants’ answers using open coding [25] (also used in *grounded theory* [26]). We decided to apply open coding for this research question (as opposed to closed coding) because it was not clear that any pre-existing list of “reasons to give up tasks” would fit this context well. We decided that it would be more appropriate to have our codebook emerge from our participants’ responses in this case. We labeled each participant’s answer with one or multiple codes that best described the reported factor(s).

First, two coders individually coded the interview responses. Then, they had a discussion about merging their codebooks and reaching a common understanding of the scope of each code. During this meeting, they also categorized the resulting codes under three broader themes: *project-centric*, *human-centric*, and *team-centric* factors. Then, they individually reviewed and updated the labels of each participant's response according to the merged codebook. After this step, the two coders reached an agreement ratio of about 85%. Then, they held a joint focused-coding session where they resolved all disagreements.

This first effort did not allow us to reach saturation — the last response mentioned a new factor that was not previously observed. Thus, we decided to obtain more responses to the same question, running an additional survey that we described in Section 2.3.3.

Two coders individually coded the survey responses using the codebook obtained from their analysis of the interview responses. They allowed the addition of new codes to the codebook if new factors were observed. The coders coded about 46% of the survey responses similarly. Then, they held a joint focused-coding session to compare their coding and resolved all disagreements. After coding all the survey responses, no new codes were added to the codebook.

Analysis evaluation: After coding all the interview and survey responses, we reached saturation in our observed give-up factors —we observed no new factors after the 20th response out of 46 (the last interview response).

3.6.2 Results

We present the results of this research question in Table 3.1 and figure 3.5. Table 3.1 shows the factors that our participants reported for giving up the search for the rationale of code changes. It also contains a description of each factor. figure 3.5 shows the ratio of how many

Table 3.1: Factors leading developers to give up the search for rationale of code changes

Factor Category	Give Up Factor	Give Up Factor Description
Project-centric Factors	Codebase state	The codebase state includes (1) the repository state (<i>e.g.</i> , the number of commits in a pull request under review) and (2) the code aspects (<i>e.g.</i> , the code is/is not running, the code has design/readability issues, the code is reviewed).
	Documentation	The documentation factor includes (1) artifacts documentation (<i>e.g.</i> , Javadocs, inline comments) and (2) development process documentation (<i>e.g.</i> , commit messages, log messages).
Human-centric Factors	Effort management	Effort management involves the developers' assessment of their work activities in terms of the required (1) effort and (2) time. This assessment is typically followed by a decision to possibly give more priority to some of these activities, discard or replace some others.
	Developer knowledge	Developer knowledge involves the developer possessing/awareness of (1) background knowledge to comprehend the code change under inspection (<i>e.g.</i> , programming language, mathematics, technical knowledge) or (2) project knowledge (<i>e.g.</i> , project design aspects, implementation aspects, and testing aspects).
	Interpersonal (Emotions)	Interpersonal factor involves human sentiment or attitude in the situation (<i>e.g.</i> , frustration, confusion, fear of obtrusion to the code review).
Team-centric Factors	Impact on productivity	Impact on productivity includes the developers' awareness of the impact of their actions on the project progress (<i>e.g.</i> , work progress stalls).
	Personnel	Personnel includes aspects related to the development team personnel (<i>e.g.</i> , developer is not around, project manager changed).
	Time management	Time management includes the team time-related aspects (<i>e.g.</i> , time crisis (deadline), team time-frame for the project).

times our participants reported each factor. Our eight observed factors were mentioned 69 times by our 46 participants since some of them (16) mentioned multiple factors.

figure 3.5 shows our three observed categories of factors being reported relatively similar rates: the most often reported factors were project-centric factors (39%), followed by human-centric (30%) and team-centric (19%) factors. The remaining 12% of responses form the *unspecified* category: these either did not provide a valid specific answer to the question (6%) or reported never giving up the search (6%). Next, we describe each category in more detail.




Figure 3.5: Give up factor

Unspecified factors: An example of an invalid response stated why it is important to understand rationale without mentioning giving up or giving up actions. Examples of participants reporting not giving up either stated it explicitly or reported giving up only temporarily. Sometimes, giving up temporarily meant *switching to another task*, leaving the rationale-demanding task for a later time:

“Essentially, if I cannot talk to the person who committed it, I will usually just postpone until they are back online.”

Other times, our participants paused to rest and returned to the task later with a fresh mind or after acquiring more background knowledge.

“I usually do not give up. I mean, I just go run or sleep, and then I try again the next day. With a clear mind, or something like that.”

“Well, sometimes, if I want to do something and implement something math-intense, perhaps I will go first and revise my math knowledge behind this. ”

Project-centric factors: Our participants most reported give up factor is *codebase state* (20%). When the code base state is problematic (*e.g.*, “commit is quite large”, “changes are too much to be able to track them”, “variables are not properly named”), our participants may give up their search for rationale of code changes. This observation shows that past efforts to aid in the comprehension of source code, *e.g.*, efforts to untangle large code changes [34, 35, 36], have additional benefits. Any effort that helps improve the comprehension of source code would also help developers not give up their search for the rationale of changes to that code.

Our participants also reported giving up due to poor/lack of *documentation* (19%). Lack of good documentation here refers to both the source code documentation, *e.g.*, JavaDocs, and the documentation of the development process, *e.g.*, commit messages.

Poor documentation in source code is a well-known problem [37], and we found that it affects the search for rationale too. One participant mentioned:

“In well-documented code it is very rare that doing this is necessary.”

Our participants also highlighted the importance of documenting the change itself and not only the code.

“Committing with a summary might not be enough sometimes. It is better to add a broader description to give a more detailed idea about the changes. ”

In general, our participants communicated giving up the search for rationale of code commits when the “commit lacks description”, or “old documentation is not found on GitHub”. They also give up when the commit message “is vague”, “not illustrative enough”, or “not descriptive enough to help me understand what is going on”.

Even when best practices are specified to document code well, they can be followed inconsistently or the granularity of the documentation can be too coarse.

“ I wish the commit messages were more granular so that they exist on each file level instead of the whole commit event. ”

The automatic generation of commit messages has been proposed in multiple research efforts [38, 39]. Our observed responses call for attention from the researchers working in the commit message auto-generation area. Our observed responses call for designing techniques that generate more granular commit messages, potentially documenting some of the rationale components that we discovered.

Human-centric factors: These factors refer to the internal state of the developer, more specifically to the developers' sense of effort (17%), knowledge (9%), and interpersonal emotions (4%).

Some of our participants evaluated the *effort* required to find the rationale of code commits. We believe that effort management is an intuitive give-up factor since developers can spend long hours or even days attempting to find the rationale for code changes. Developers often have several tasks to do, and they want to move on. They can get frustrated if they spend a long time and effort being blocked by a single task. Our participants commented on the time spent search for rationale of code changes:

“If it takes half an hour, it is not worth spending more time [on the search]. Then, I will ask others and interrupt their work”

“I would say at the hour mark, because at that point I would be like, “I just need to do something, and I will make a best-case judgment call.””

“If it is taking me more than 30 minutes [to find the rationale of a code commit], I will start trying and get in touch with the person who wrote the code, at which point it becomes an asynchronous process. So I am no longer sitting there trying to figure out the rationale; I will email/message the person, then I move on with whatever I was doing.”

Other participants mentioned *lack of knowledge* as a reason to give up searching for rationale. Their responses referred to background knowledge and skills, or knowledge about the project. Examples of the first case are when a developer is not fluent in a particular programming language or not familiar with the underlying mathematical foundations of the code. Examples of the second case are when a developer is responsible for certain project areas (*e.g.*, back/front-end) and lacks knowledge of other project areas. As a result, we see the existence of this factor as a normal part of software development. Educating developers could mitigate this factor, but in many cases, it is not an affordable option in terms of cost and developer time. In other words, this specific give-up factor might be one of the main prompts for documenting the code change rationale since it is inevitable on many occasions. Furthermore, this factor, perhaps, encourages documenting rationale in a way that other developers can understand regardless of their detailed knowledge of the project.

Finally, other participants expressed that *interpersonal (emotions)* play a role in giving up their search for rationale of code changes. One participant answered:

“ No specific time cut-off, I would say. I would say it is more...how frustrated I am ”

Other emotions that we observed were confusion, fear of obtrusiveness, and anger. Software developers' sentiments and emotions are important aspects that could impact their work [40, 41, 42]. Our findings suggest that negative emotions like frustration impact software developers while searching for rationale for code commits. Therefore, we encourage future studies of the relationship between sentiments and rationale. We also encourage future studies of how to mitigate the impact of these sentiments while still satisfying developer information needs.

Team-centric factors: Our participants expressed that changes in the project *personnel* (12%) introduce obstacles to communication and could prevent them from finding the change owners' rationale. One participant described a give-up situation caused by personnel that left

as

“[The change] was a decision made by [the previous project manager] that is not even here anymore. We do not really understand the reason, but we make a new decision from this point forward.”

Personnel changes are common in software development teams. In such situations, we think that standard documentation guidelines could make the search for rationale for code changes easier. One participant described his/her company guidelines:

“At my company, typically we link each commit, branch, and PR to a ticket number, which should include details and discussion about the change and why it is being made.”

Our participants also give up their search for rationale of code changes to accommodate project-related *time management*. For example, in a rush to fix a bug before approaching a software release, one participant said that he/she came up with some rationale for a buggy code change to provide a temporary fix to the bug. One participant answered:

“When I have to, it could be when I am not getting other things done because I am searching for this rationale. Working intimately with another partner, our progress was halted until I completed or gave up. In other contexts that I have somewhat been associated with a large team, I do not need to understand it as long as other people on the team do.”

Composite factors: Sixteen participants (34%) mentioned more than one give-up factor. Twelve of them mentioned two factors in conjunction. For example:

“I will look at the commit (code and description) as well as the pull request, including the commit and the ticket relating to that pull request. If the implementer can still be contacted, I will do so. If none of this leads to any results, I will usually give up. ”

The remaining 4 participants mentioned two or more factors in disjunction. For example:

“ Usually, there are certain scenarios where I give up finding rationale, if I cannot run the code, or if it is a language that I cannot understand, programming language. Or the code base is too big. ”

When participants report factors in conjunction, they show that many factors need to co-exist for them to give up. Most of the time, they mention the lack of documentation in addition to another factor. The second factor is sometimes the codebase state (*3 times*), showing that the commit is large or complicated, making the search for rationale a tedious process. Some other times, the second factor besides the absence of documentation is the absence of the author (*3 times*), which deprives the developer of seeking the rationale from, possibly, the two most helpful information sources.

When our participants report factors in disjunction, they show multiple individual reasons why they may give up (one of them is enough). For example, they may give up to save effort (effort management factor) given the large size of the commit (Codebase state), or because they were not sure they had the required project knowledge to understand the code change (Developer knowledge) given the poor code readability state (Codebase state).

Another case that we found interesting is when the participants use a combination of factors mixing negative and positive instances of them (*3 times*). For example, a participant's decision to give up seeking rationale was based on a close project deadline, but the decision was eased by the fact that the code was running correctly, in which case seeking rationale was not a pressing need.

“ ...Then, the main concern was to finish the project before the deadline. Sometimes in such a situation, if the codes work, then I do not try to check the rationale. ”

We believe that the search time required to find the rationale of code commits is likely connected with the difficulty of finding it (Section 3.4.2), *i.e.*, the rationale for code commits may be easy to find when it is available and well documented. When the components sought by the developer are documented where the developer expects, the rationale for code commit could be easy and fast to find.

The problems revealed by the team and project factors are the reliance on the availability of

the change owner, the trade-off between searching for rationale and productivity/time, and the quality of a project's codebase and documentation. We discuss in Section 5.4 how future tools could support developers in their search for the rationale of code commits, informed by our observations.

3.7 RQ: Would comparing the experience of developers needing, finding, and recording the individual components of the rationale of code commits with each other reveal areas for improvement?

3.7.1 Research Method

The data for this research question is the same as Section 3.5. We performed a cross-experience analysis to further investigate the developers' experiences with rationale components' need, finding, and recording. We paired every two experiences (need-finding, need-recording, and finding-recording) and then used the median responses of each component to discover needed areas of improvement.

3.7.2 Results

figures 3.6A and 3.6B show that software developers are most frequently finding and recording the most frequently needed components of rationale. Most of the components are in the middle frequency of need and finding. However, this result reveals that many components are not too frequently needed, but when they are needed, they are really hard to find. Devel-

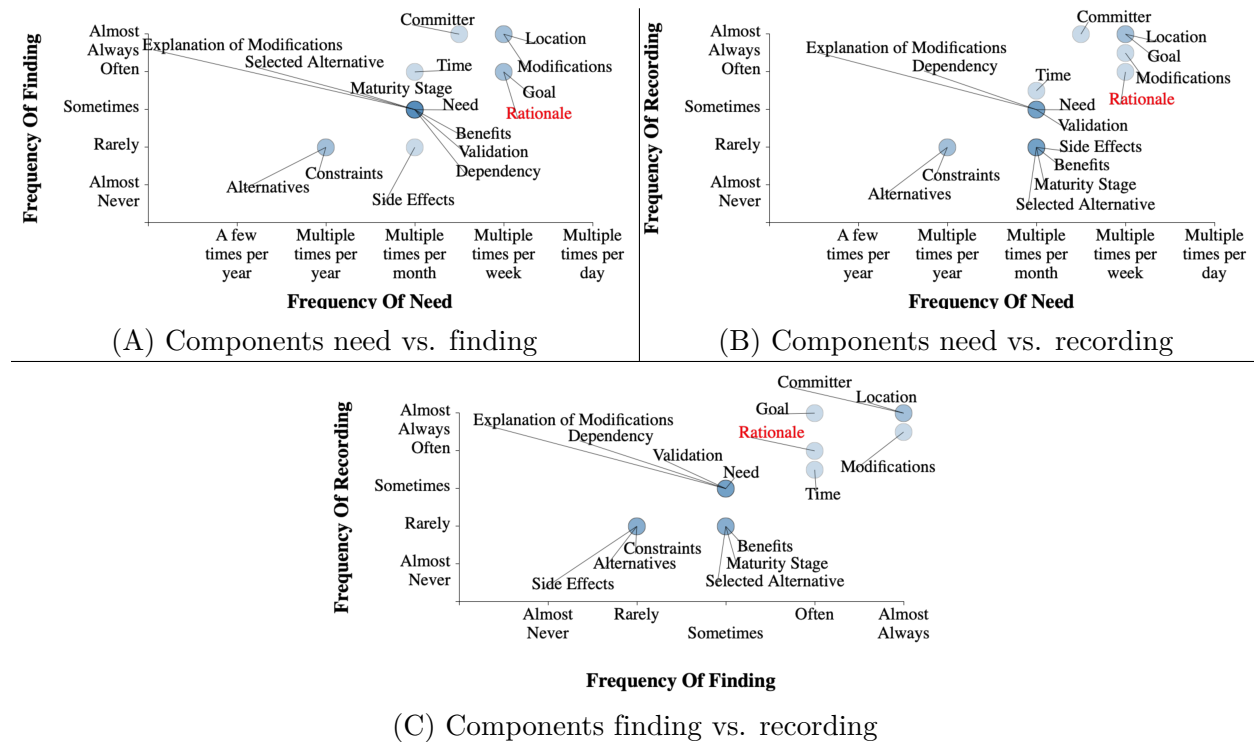


Figure 3.6: Cross-dimensional analysis of developers' experience with the individual components of rationale of code commits

opers most struggled to find *side effects*, *alternatives*, and *constraints*, even if they need to find them on average multiple times per month and per year. In these cases, the difficulty of finding these components may overcome their limited frequency of need. Thus, practitioners may want to pay more attention to documenting these not-so-frequently-needed components.

The difficulty of finding rationale depends on many factors, *e.g.*, the complexity of code commits, the developers' documentation of code changes, and the need to discover the rationale.

One of the participants said about giving up searching rationale

"I would completely give up if I could not find any record in our system and the author was someone who either is no longer at our company or is somebody who just does not write code anymore. Yeah. I give up when I have exhausted all the possibilities, but if I really need to know, I would keep trying until I figured it out."

For the components of rationale that are not easy to find, guidelines could be established,

and tools could be developed to simplify finding these components. One participant said about finding rationale:

“From my experience, the rationale, it is easier to figure out once your team kind of has standards or guidelines.”

The recording of rationale goes hand in hand with the finding of rationale (see figure 3.6C). Unsurprisingly, not recording some components makes it hard to find them later. The rarely recorded components were: *side effects*, *alternatives*, *constraints*, *selected alternative*, *maturity stage*, and *benefits* — even when developers need to find them on average multiple times per year (*alternatives* and *constraints*) and per month (remaining ones). Identifying this group of rarely recorded components should encourage researchers to develop tools specifically focused on recording or answering them. For example, a technique to evaluate the maturity stage of a commit will aid developers in seeking this component without the need for other developers to document it manually.

3.8 Threats to validity

Our studied developers may not fully represent the whole developer population. To mitigate this threat, we recruited a diverse population with diverse types and amounts of experience (figure 2.1). Our studied population was similar to the ones previously studied in the literature since we obtained similar answers for our two questions about rationale that were already studied by Tao et al. [5].

3.9 Conclusion

Developers invest valuable time and resources in the process of discovering the rationale of code commits, which they perform frequently and is difficult. We applied a mixed-methods approach in this study to survey the developers' experiences. We performed a series of interviews with software developers to discover software developers' experiences (needing, finding, and recording) rationale of code commits and its components, and the factors leading developers to give up their search for rationale. Then, we ran two online surveys to obtain more responses about developers' experiences with rationale of code commits. We discovered that software developers have different experiences with different components. We also found a set of *human*, *team*, and *project* factors leading software developers to give up their search for rationale.

Our findings suggest that there is a space for both researchers and practitioners to improve the practices of managing the rationale of code changes. For example, developers need to find most rationale components with similar frequency, but they mostly only record and find those components automatically recorded by revision control systems. Therefore, researchers and tool builders could help software developers automatically record rationale components that are not already recorded by revision control systems. Another way to improve the practices of managing the rationale of code changes is acting on our discovered factors leading software developers to give up their search for rationale of code commits. For example, we discovered that *changes in project personnel* make it hard for software developers to find the rationale behind code changes. Practitioners could mitigate this factor by establishing standard documentation guidelines, making the search for rationale for code changes easier.

Chapter 4

Predicting Rationale components

Documentation in Pull Request

Templates

In previous chapters 2 and 3, we found that rationale of code changes is needed for many software development tasks, including *code review*. We also found that the rationale of code changes is composed of many components that are needed, recorded, and found at different levels. Rationale components need, recording, and finding are likely to be different for different software development tasks.

This chapter aims to help software developers document rationale components in the code review context. We investigate prescribed pull request templates for requested rationale components. Then, we experiment with different machine learning classification algorithms to predict if a pull request template header of rationale will not be filled. Software developers time and effort could be saved if a pull request template header is predicted not be filled.

4.1 Background and Motivation

Pull-based software development model is a widely popular practice in Open Source Software (OSS) [43]. This model involves contributors forking the main branch of a repository, making independent code changes, and subsequently submitting a pull request (PR) to merge their contributions into the main repository. The core team, responsible for maintaining the repository, reviews these proposed contributions and determines whether to merge or reject the changes [44].

In the process of opening a PR, contributors provide a PR description that plays a pivotal role in enhancing collaboration efficiency with the maintainers. A *well-crafted PR description* can improve collaboration *efficiency* by providing *clear and concise information* to the maintainers, making it easier for them to review and understand the proposed changes [17].

To assist contributors in crafting PR descriptions, GitHub repositories prescribe *Pull Request Templates (PRT)* [16]. When a repository *prescribes a PRT*, project contributors automatically see the template’s contents when opening a pull request. Recent research work empirically studied these templates to explore their contents, impacts, and developers’ perceptions [17].

Some *rationale components* were present in Li *et al.* exploration of PRT contents. Their study also discovered that PRTs are perceived to be useful, but using PRT comes with its own challenges. “Templates require too many fields” was identified by Li *et al.* [17] survey as one of the top challenges from both contributors and maintainers.

Developers will likely not fill all PRT headers for every pull request, including rationale components headers. To assist the developers’ practice of filling PRT when opening a pull request, we present a novel application of machine learning classification to predict if a PRT header of rationale will not be filled. Predicting if a PRT header will not be filled could be

helpful for customizing PRT of individual PRs, *i.e.*, requiring fewer fields. Such predictor could assist in (1) filtering likely to be empty fields and/or (2) ranking/reordering template fields.

4.2 Research Questions

To predict if a PRT rationale header will not be filled in individual PRs, we ask the following research questions:

- ***RQ: How are different pull request characteristics correlated with filling a pull request template rationale header?*** We investigate this research question to discover features that may be useful for predicting if a pull request template is going to be filled or left empty. Statistically significant correlation between features and a template header being filled will help selecting features for a predictor.
- ***RQ: To what extent can classification algorithms predict if a pull request's template rationale header will be filled?*** We studied how different prediction algorithms are capable of predicting if a pull request's template rationale headers will be filled. This study recommend the best performing algorithms to help software developers save time and effort.
- ***RQ: How effective are different feature groups in predicting if a pull request's template rationale header will be filled?*** We investigate how different group of features perform in the prediction task. An understanding of the features contribution to the predictor success could lead to building a lightweight predictor with comparable prediction performance.

4.3 Approach

We use a data-centric analysis approach to build a technique capable of predicting if a PRT header of rationale will not be filled. We start by mining software repositories to extract pull requests' information (*e.g.*, PR patch characteristics). Then, we process pull requests' information by applying qualitative and automated labeling methods. Afterward, we study the correlation between mined features and the PRT header of rationale being filled. In the next step, we utilize machine learning classification algorithms to predict if a PRT header of rationale will not be filled. Finally, we test the predictors' performance using different groups of features. The following sections present an overview of each step of our analysis approach.

4.3.1 Data Preparation

In this section, we introduce our data collection and preprocessing processing steps.

Dataset: We started our data collection with GHArchive ¹ dataset. GHArchive, short for GitHub Archive, is a project that collects and provides public access to a comprehensive record of activities (events) on the GitHub platform. GHArchive “record the public GitHub timeline, archives it, and makes it easily accessible for further analysis”. We downloaded and extracted GHArchive data for the second half (Jun 01-Dec 31) of the 2021 year. We hosted this 1.8TB of data in a MongoDB ² database and indexed parts of it using Apache Solr ³ for quick querying.

For this study, we are only interested in Pull Request (PR) events out of the 20 different GitHub events. Furthermore, we are only interested in PR events at “open” time. Those

¹<https://www.gharchive.org/>

²<https://www.mongodb.com/>

³<https://solr.apache.org/>

PR events include general information about the PR like: PR creation time, PR title, PR description, PR head and base repositories, and PR creator.

Since our study is interested in the code changes for which the PRs are created, we extended the GHArchive PR events by extracting the PR patch using GitHub API ⁴. Our extracted PR patches are in raw Git patch format that we later transformed to extract the PR patch information like: PR patch commits, PR patch commits files, PR patch commit messages, etc.

In addition to the PR events, we needed to collect information about repositories' prescribed Pull Request Templates (PRT) ⁵. We used GitHub API to collect repositories PRT from GitHub officially defined PRT files paths.

Data selection To conduct our study experiments, we selected repositories to study based on the following criteria. First, we only selected repositories that define a PRT. Second, we only selected repositories that use GitHub PR feature. We looked at the number of opened PRs within the time period of our study. We only selected repositories with more than 600 PRs opened during the Jun-Dec of 2021 inclusive. Third, we only selected repositories with more than five PRT headers in their prescribed templates. This filtration resulted in 128 repositories with 87K pull requests.

Data preprocessing: To shape the data into a form necessary for correlation study and statistical models, we parsed, labeled, and encoded some dataset fields. We used a markdown parser to parse prescribed repositories PRT and individual PR descriptions utilizing the templates. The parsing of PRT and PR description (PRT usage) allowed us to identify the template parts and their content, splitting the markdown PRT/PR into header-content pairs. After the parsing step was completed, we ran a script to compare parsed PR description

⁴<https://docs.github.com/en/rest>

⁵<https://docs.github.com/en/communities/using-templates-to-encourage-useful-issues-and-pull-requests/creating-a-pull-request-template-for-your-repository>

against parsed repositories prescribed PRT, allowing us to label if individual PR description has filled individual template headers.

In addition to the fill or not PRT header labeling, we manually coded prescribed PRTs headers using Li et al. [17] taxonomy of elements included in issue/PR templates. This coding allowed us to account for different wording of headers in different repositories.

The last step of our data preprocessing was encoding categorical and boolean features (Table 4.1) in our dataset. We relied on the sklearn preprocessing module ⁶ to run the encoding tasks. We used “MultiLabelBinarizer” to encode multi-valued categorical features (*i.e.*, Patch_file_types). We used “OneHotEncoder” to encode single-valued categorical features (*e.g.*, Dev_assoc). Lastly, we used “LabelEncoder” to encode boolean features (*e.g.*, PRT_H_is_optional).

4.3.2 Correlation analysis

In this section, we introduce our correlation analysis process. We only used one month of data (June 2021) for the correlation study.

Features list Our data-centric analysis is driven by the data associated with pull requests. We selected a list of features (Table 4.1) based on available information in GHArchive and GitHub API. Our set of features is similar to the recent prediction tool in the pull request context. Zhang et al. [45] used a similar list of factors to predict the influence on pull request decisions (accept or reject). Similarly, Azeem et al. [46] used similar features to predict pull request action (accept, respond, and reject.).

⁶<https://scikit-learn.org/stable/modules/classes.html#module-sklearn.preprocessing>

Table 4.1: Features extracted for each pull request

Category	Feature	Feature Description	Values
Developer Characteristics	Dev_assoc	PR creator affiliation	Categorical
	Dev_per_patch	Number of unique developers who contributed to the PR changes	Numbers
	Dev_per_patch (is_multiple)	More than one developer contributed to the PR changes	True or False
Patch Characteristics	Patch_ch_files_per_comt	Median number of files per commit in the PR patch	Numbers
	Patch_commits	Number of commits in the PR patch	Numbers
	Patch_comts_additions	Number of additions in all PR patch commits	Numbers
	Patch_comts_deletions	Number of deletions in all PR patch commits	Numbers
	Patch_comts_duration	Difference in days between the PR patch's newest and oldest commit	Numbers
	Patch_comts_messages_len	Total length of patch commit messages	Numbers
	Patch_comts_modification	Number of modifications in all PR patch commits	Numbers
	Patch_contains_bug_fix	“bug” and/or “fix” keywords exist in the PR patch files names and/or commit messages	True or False
	Patch_contains_config	“config” keyword exists in the PR patch files names and/or commits messages	True or False
	Patch_contains_docs	“doc” keyword exists in the PR patch files names and/or commits messages	True or False
	Patch_contains_test	“test” keyword exists in the PR patch files names and/or commits messages	True or False
	Patch_file_types	PR patch files types	Categorical
	Patch_file_types	Number of unique PR patch files types	Numbers
	Patch_files_added	Number of files added in the PR patch	Numbers
	Patch_files_changed	Number of files changed in the PR patch	Numbers
	Patch_files_deleted	Number of files deleted in the PR patch	Numbers
	Patch_files_paths	Number of files paths in the PR patch	Numbers
Patch_files_renamed	Number of files renamed in the PR patch	Numbers	
Patch_has_code_changes	PR patch changes modify code	True or False	
Patch_has_code_comments	PR patch changes modify code comments	True or False	
Patch_has_text_changes	PR patch changes modify English text	True or False	

Table 4.1: Features extracted for each pull request (cont.)

Category	Feature	Feature Description	Values
Repository Character- istics	Repo_age (head)	Number of months from PR changes' source repository creation to PR creation	Numbers
	Repo_forks (head)	Number of forks of the PR changes' source repository	Numbers
	Repo_names_match	PR changes' source repository matches the PR repository	True or False
	Repo_open_issues (head)	Number of open issues of the PR changes' source repository	Numbers
	Repo_stars (head)	Number of stars of the PR changes' source repository	Numbers
Template Character- istics	PRT_H_is_optional	The prescribed PRT header contains the words "optional" or "if applicable"	True or False
	PRT_H_order	The prescribed PRT header order (1st, 2nd, ..., Nth)	Numbers
	PRT_H_type	The prescribed PRT header's content type	Categorical

Correlation measures To measure the correlation between a PRT header of rationale being filled and individual features, we used different correlation measures depending on the feature type (number vs. binary). We used phi correlation coefficient ϕ for binary features and point bi-serial correlation coefficient r_{pb} for numerical continuous features. We also used Chi-square test to measure the statistical significance of correlations.

The phi correlation coefficient, denoted as ϕ , is a measure of association between two binary variables. It is calculated using a 2x2 contingency table, where the rows represent the levels of one variable, and the columns represent the levels of the other variable. The phi coefficient ranges from -1 to 1, with 0 indicating no association, positive values indicating positive association, and negative values indicating negative association.

The point biserial correlation coefficient, denoted as r_{pb} , is a measure of association between a continuous variable and a binary variable. It is calculated when one variable is continuous, and the other variable is dichotomous (binary). The point biserial correlation coefficient is an extension of the Pearson correlation coefficient (r). It measures the degree and direction

of the linear relationship between the continuous and dichotomous variables. It ranges from -1 to 1, where 0 indicates no association, positive values indicate a positive association, and negative values indicate a negative association.

The Chi-square test assesses whether the observed frequencies in a contingency table differ significantly from those expected if the variables were independent. The test statistic, denoted as χ^2 (chi-square), is calculated by comparing the observed frequencies to the expected frequencies under the independence assumption.

4.3.3 Statistical Modeling

We use supervised learning algorithms to predict if a PRT header of rationale will not be filled in individual PRs. Those supervised learning models are created from labeled historical training data. The classification task here is for discrete class labels: filled (1) or not filled (0). The goal is to predict the class labels of new instances based on past observations. We train a handful of different algorithms utilizing SKLearn supervised learning library ⁷. In order to train and select the best-performing model, we use the default parameters of supervised learning library algorithms.

To determine whether our models not only perform well on the training set but also generalize well to new data, we also want to randomly divide the dataset into a separate training and test set. We use the training set to train and optimize our machine learning model, while we keep the test set until the end to evaluate the final model. After a model has been fitted on the training dataset, we use the test dataset to estimate how well the model performs on unseen test data to estimate the generalization error.

⁷https://scikit-learn.org/stable/supervised_learning.html

4.3.4 Models Evaluation & Analysis

To measure the performance of the different statistical models, we use a 10-fold cross-validation method and the computation of the standard classification evaluation measures, including precision, recall, and F-value.

Precision measures the proportion of positive predictions (PRT header of rationale is not filled) correctly identified by the classifier. It is calculated by dividing the number of true positive predictions by the sum of true positives and false positives. Precision is valuable when the impact of false positive (PRT header of rationale is filled) predictions is significant.

Recall, also known as sensitivity or true positive rate, measures the proportion of actual positive instances (PRT header of rationale is not filled) correctly identified by the classifier. It is calculated by dividing the number of true positive predictions by the sum of true positives and false negatives. Recall is especially useful when detecting positive instances (PRT header of rationale is not filled) is critical.

F-value, also known as F1-score, is a single metric that combines precision and recall. It is calculated as the harmonic mean of precision and recall, providing a balanced assessment of the classifier's performance. The F-value ranges from 0 to 1, with 1 indicating perfect precision and recall.

Baseline In addition to precision, recall, and F-value evaluation measures, we train a dummy classifier that randomly guesses whether a PRT header is filled or not is used as the baseline. We compare against the dummy classifier evaluation measures to evaluate the success of the predictions.

Feature group performance: We measure the classification performance of different sets of feature combinations (*e.g.*, the patch characteristics) to compare their performance.

4.4 RQ: How are different pull request characteristics correlated with filling a pull request template rationale header?

Table 4.2 shows statistically significant correlation values (phi ϕ or point bi-serial r_{pb} correlation coefficient) between each pull request characteristic (feature), and PRT rationale headers fill observations. Although correlation coefficient values are low overall, the correlations are statistically significant in our one-month data sample. Statistical significance suggests that the observed correlation is unlikely to have occurred by chance and may still have practical implications. The correlation values provide us with initial insights into the nature of the relationship between a pull request feature and the PRT rationale headers fill. Furthermore, the correlation table (Table 4.2) shows complex analysis dimensions (*i.e.*, various features and multiple headers). Low correlation coefficients between individual features and headers may still contribute to a larger predictive model. Multiple weak correlations combined can provide valuable information for understanding complex situations.

Table 4.2: Correlation between features and PRT rationale header fill

Feature	PRT Rationale Header				Feature	PRT Rationale Header			
	Testing	Motivation	Side effect	List of main changes		Testing	Motivation	Side effect	List of main changes
PRT_H_is_optional	0.05	0.20	0.05		Patch_has_.asciidoc	0.04	0.04		
PRT_H_order	0.07	0.26	0.18	0.24	Patch_has_.cpp	0.08	0.07	0.08	
PRT_H_type_checklist	0.05			0.24	Patch_has_.cs	0.06		0.07	
PRT_H_type_question	0.08	0.03	0.08	0.08	Patch_has_.go	0.05	0.06	0.07	0.20
PRT_H_type_yes_no_question			0.26		Patch_has_.gradle	0.07		0.05	0.06
Repo_age (head)	0.08	0.19	0.02	0.05	Patch_has_.h	0.02	0.06	0.09	0.06
Repo_forks (head)	0.03	0.20	0.02	0.05	Patch_has_.html		0.03	0.03	
Repo_names_match	0.10	0.29	0.04	0.19	Patch_has_.java		0.11	0.18	0.05
Repo_open_issues (head)		0.18			Patch_has_.jpg		0.03		
Repo_stars (head)	0.01	0.23	0.04	0.04	Patch_has_.js	0.04	0.20	0.06	
Dev_assoc_COLLABORATOR	0.07	0.18	0.07	0.14	Patch_has_.json		0.06		0.04
Dev_assoc_CONTRIBUTOR		0.18	0.04		Patch_has_.jsx	0.06	0.03	0.04	0.03
Dev_assoc_MEMBER		0.08	0.03	0.08	Patch_has_.kt	0.10			0.07
Dev_assoc_NONE	0.04		0.06		Patch_has_.lock		0.08		0.05
Dev_per_patch (ct)	0.02	0.03	0.01	0.03	Patch_has_.md	0.05	0.07	0.08	0.07
Dev_per_patch (is_multiple)	0.03	0.05	0.04		Patch_has_.mod	0.03	0.03	0.04	0.05
Patch_ch_files_per_comt (mdn)	0.01			0.04	Patch_has_.php	0.13	0.08	0.17	
Patch_commits (ct)	0.03	0.04	0.02	0.05	Patch_has_.png	0.02	0.04	0.08	0.05
Patch_comts_additions (sum)	0.01		0.02	0.01	Patch_has_.py			0.04	0.33
Patch_comts_deletions (sum)	0.01		0.01	0.02	Patch_has_.rb	0.06		0.04	0.04
Patch_comts_duration (days)			0.01		Patch_has_.rs			0.03	
Patch_comts_messages_len	0.02		0.02	0.04	Patch_has_.scss	0.03	0.05	0.08	0.06
Patch_comts_modification (sum)	0.01		0.01		Patch_has_.sh	0.03	0.05	0.06	0.07
Patch_contains_bug_fix			0.05	0.10	Patch_has_.sql	0.11			0.11
Patch_contains_config	0.05		0.05		Patch_has_.sum	0.02			0.05
Patch_contains_docs	0.08	0.08	0.07	0.07	Patch_has_.toml	0.03			0.04
Patch_contains_test		0.05	0.12	0.11	Patch_has_.ts	0.11	0.03		0.06
Patch_file_types (ct)	0.02	0.04	0.01	0.01	Patch_has_.tsx	0.04	0.11		0.13
Patch_files_added (ct)	0.02	0.04		0.02	Patch_has_.txt	0.02	0.03	0.05	0.10
Patch_files_changed (ct)	0.04			0.02	Patch_has_.vue	0.06	0.10	0.04	0.08
Patch_files_deleted (ct)	0.01	0.02		0.02	Patch_has_.xml	0.03		0.10	0.06
Patch_files_paths (ct)	0.04	0.04	0.01	0.03	Patch_has_.yaml	0.09		0.08	
Patch_files_renamed (ct)			0.02	0.02	Patch_has_.yml	0.08	0.07	0.07	0.08
Patch_has_code_changes	0.05	0.04		0.07					
Patch_has_code_comments		0.04	0.08	0.06					
Patch_has_text_changes		0.03		0.09					
					Correlation	.01	.16	.25	.33
						-.01	-.16	-.25	-.33

4.5 RQ: To what extent can classification algorithms predict if a pull request’s template rationale header will be filled?

Figure 4.1 shows six algorithms’ performance in predicting if a PRT header of rationale will not be filled. We used five months of PRs data (87K PRs) for our selected repositories (128 repos). With individual pull requests broken down to PRT header-content-fill observation, the total number of observations for all PRT headers is around 500K observations. We used 10-fold cross-validation to build and test predictors of PRT headers of rationale, a predictor for each repository’s PRT headers of rationale. Figure 4.1 shows each algorithm’s precision, recall, and f1-score in predicting if PRT headers of rationale will not be filled in repositories PRs.

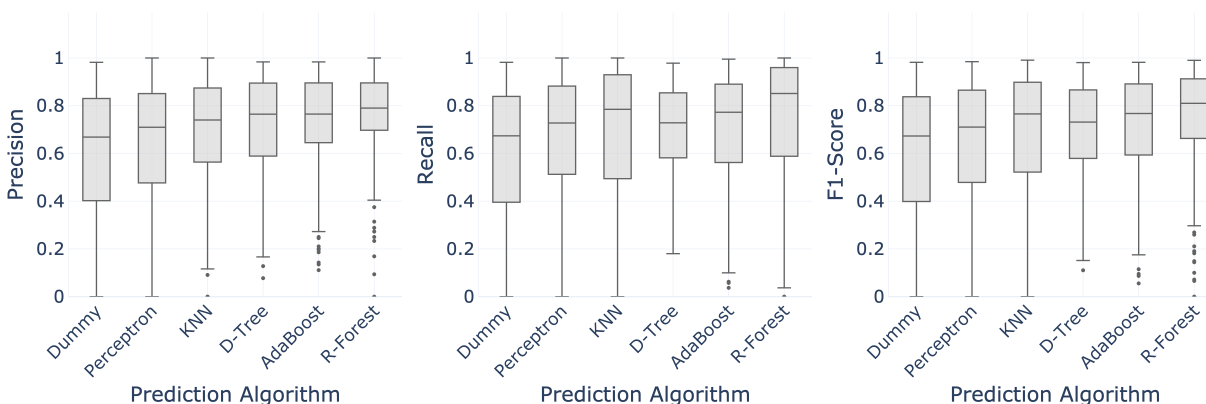


Figure 4.1: Classification results of classifiers trained using 10-fold cross-validation

In the median case, all prediction algorithms perform better than the baseline (dummy prior predictor). The dummy predictor chose predictions based on the empirical class (fill) distribution. Random forest is the algorithm that provides the best performance in both precision and recall. This experiment indicated the ability to predict if a PRT header of rationale will be filled. Using such a predictor in practice could help software developers be

more efficient (save time) by skipping the documentation of the PRT header of rationale if the predictor finds the header likely not to be filled.

4.6 RQ: How effective are different feature groups in predicting if a pull request’s template rationale header will be filled?

To understand which features are better in predicting if a PRT header of rationale will be filled, we compare the prediction performance using only a subset of features at a time. We select the best predictor (Random Forest) from the previous research question and apply a feature permutation experiment. We used 10-fold cross-validation to build and test five Random Forest predictors of PRT headers of rationale, a predictor for each repository’s PRT headers of rationale. Each predictor uses a different set of features: all features, template features, patch features, developer features, and repository features, respectively. Figure 4.2 shows each feature group’s precision, recall, and f1-score in predicting if PRT headers of rationale will not be filled in repositories PRs.




Figure 4.2: Classification results of *Random Forest* classifier trained with different set/groups of features

This experiment shows that different groups of features can improve or decrease prediction performance. In the median case, template features perform better than the prediction with all features. Other features groups predictors perform worse than predicting with all features. The repository features notably perform worse than other feature groups. The negative correlation between the PRT motivation header and repository features may explain this lower performance of repository features. The dummy predictor chose predictions based on the empirical class (fill) distribution.

4.7 Threats to validity

To design our statistical model, we used a set of features that may only partially describe the aspects of pull requests. There are many other aspects of code changes (*e.g.*, continuous integration information) that we did not consider. To mitigate this issue, the features considered in our work describe diverse aspects of the pull requests. Each feature falls into one of four categories: patch, repository, developers, or template characteristics. We have selected our features considering features of recent studies in pull request context [45, 46].

To construct our dataset, we took careful measures to not introduce duplicate observations that might lead to over-fitting or bias in our statistical model building and analysis. We only studied PR events at “open” time, avoiding duplicate PRs information with “close” or “reopen” time events. Another careful measure we employed is dividing the data used for correlation and prediction parts of our study. We used one month of data for the correlation data and the five remaining months for the prediction task data.

Our study is specific to open-source GitHub projects that prescribe and use pull request templates. Those repositories might not be representative of all repositories that use pull request templates. For instance, close-sourced repositories and repositories that use different

code review workflows are not considered in our study. Our study results may not generalize to those repositories.

4.8 Summary

To assist developers in crafting PR descriptions, GitHub repositories prescribe Pull Request Templates (PRT) that ask for rationale for code changes components [17]. Developers will likely not fill all PRT headers for every pull request, including rationale components headers. To assist the developers' practice of filling PRT when opening a pull request, we experimented with building different predictors. The predictor's objective is to suggest if a PRT header of rationale will not be filled. Such predictions could help software developers work more efficiently, skipping the documentation of PRT headers of rationale that will likely not be filled.

Our experimentation showed that all built predictors can predict if a PRT header of rationale will not be filled with better precision and recall than our baseline. In the median case, random forest outperforms other prediction algorithms. We also find that different groups of features can improve or decrease prediction performance. Interestingly, template features perform better than other feature groups. Our observations motivate the implementation of such predictors in practice. Such predictor could assist in (1) filtering likely to be empty fields and/or (2) ranking/reordering template fields.

Chapter 5

Discussion and Implications

This section discusses our research findings and their implications.

5.1 Is the need for rationale of code commits different than the need for rationale in other contexts?

Some of the components of the rationale of code commits that we discovered are also relevant for rationale in software requirements, design, and architecture. These are: *constraints* [19, 47, 48, 49, 50], *alternatives* [47, 50, 51, 52], and *validation* [47, 48, 50]. However, we also discovered components that are specific to the scope of code commits: *committer*, *time*, *location*, and *modifications*. They generally refer to performing the code change. Our participants indicated that these components are more needed and important than those that also show in rationale in other contexts (see figures 3.4A and 3.6A). Similarly, some components were not mentioned in our study and were reported by previous work as part of the design rationale. Examples are: design *assumptions* [19, 48, 49] and *weaknesses* [19, 48].

5.2 Providing a better understanding of the need for rationale of code commits

The observations in this study illustrate the importance of supporting developers in documenting and finding the rationale of code commits. Software developers regularly need the rationale of code commits (see figure 3.1B) and spend a significant amount of time searching for it (see figures 3.2C and 3.2D). Our observations also allow us to speculate about the root cause of this problem. We observed that most components of the rationale of code commits are frequently not recorded (see figure 3.4C and see figure 3.6B), not found (see figures 3.4B-i and see figure 3.6A), or difficult to find (see figure 3.4B-ii). This may be because finding rationale for code commits could be very easy when the rationale is well documented, or the change owner is easily available to provide the rationale.

We also observed that finding the rationale of code commits may cause productivity loss and sometimes without gain. Furthermore, even after investing high effort searching for it, a variety of factors may lead them to abandon the search (see Table 3.1), making their productivity loss more serious, since they did not get any value from the time spent. Furthermore, after giving up their search for rationale, developers may resort to speculating their own understanding, which may lead to introducing code errors.

“At one point, I gave up understanding why they did what they did. I was confused as to why [input check was done] redundantly and gave up trying to figure it out. I removed the [redundant check] later on as it made more sense [to me].”

5.3 Our vision: how should practitioners document the rationale of code commits?

We do not believe that developers should necessarily document every component of rationale all of the time. In fact, our participants mentioned concern about such an approach.

“I know it might not be doable or possible because no one will ever answer all these in a commit. However, it is a good model.”

Instead, the goal of our model is to provide a superset of the possible components that could answer a question about the rationale of code changes. Developers would then choose which components are relevant for which code change. We believe that developers may seek different components of rationale at different times, but not necessarily all of them every time. In fact, we observed that they seek different components with different frequencies (see figure 3.4A-i).

We see our decomposition of the rationale of code commits as an artifact to support developers in documenting it (by reminding them of all the components they may want to document), not as a template that they would always be forced to fill completely. We believe that different components will be relevant for different types of code changes, and developers should judge which ones are worth documenting on each occasion. Our experimentation with predicting the documentation of rationale components in pull request (Chapter 4) aligns with this belief. Our predictor and model work as a *check* to assess documenting the rationale of code changes; they could be very useful for developers to ensure that all the relevant components are documented. A similar approach was successfully applied in the area of bug reports [53, 54], not only to assist in the documentation of various components but also to measure their quality. Checklists are a known powerful mechanism to ensure processes are performed correctly [55].

Another measure that would benefit developers is the *lazy* documentation of the rationale of code changes — *i.e.*, to document it opportunistically after they have needed it. Our observations show that developers often give up seeking the rationale of code changes due to not finding it in the code or documentation. Other times, they do not find it in code or documentation, but end up finding it in different ways, *e.g.*, asking a colleague:

“If I cannot figure it out, I ask someone to help me out to understand the code because I need to work.”

We believe that this signals an opportunity: if developers document the rationale of code changes in centralized documentation (maybe in the code comments themselves) after identifying it by other means, the next developer will easily find it when searching. This incremental approach to documenting may be preferred by many developers since the workload may feel more manageable that way, *e.g.*, as is the case with other practices like incremental testing [56, 57].

Developer teams are diverse [58], and some developers may want to improve the documentation of rationale of their past code commits in a more *exhaustive* fashion, as opposed to opportunistically. Some developers may also want to improve the current documentation of rationale of their past code commits, in a more *exhaustive* fashion —as opposed to opportunistically. Our observations also provide feedback on how to prioritize such efforts. A starting point for improving the rationale documentation of existing commits would be tackling the areas of improvement that we identified. We observed which components of the rationale of code commits are frequently not recorded (see figure 3.4C and see figure 3.6B), not found (see figures 3.4B-i and see figure 3.6A), or difficult to find (see figure 3.4B-ii).

5.4 Our vision: how could tools support developers in documenting the rationale of code commits?

Our observations provide some advice on designing future support for developers to document and find the rationale of code commits. We observed that the need for the rationale of code commits has a wide reach: it affects most software development tasks and subtasks (figure 2.2A). Therefore, any future support should be accessible and convenient, independent of the developer's task.

We also observed that the rationale of code changes is needed for code changes in software projects both within the developer's company and external to it (figure 2.2). This teaches us that the rationale of code changes should be documented by also keeping in mind developers that do not necessarily belong to the software development team. Some examples of additional considerations that developers could have to be inclusive of external developers in their documentation efforts are: including clarifications about vocabulary that is specific to the software project (or adding links to a centralized legend); including assumptions that may be clear to the development team, but not to outsiders (*e.g.*, assumptions about the pre-set configuration of the OS or underlying libraries); documenting best practices that are common within the development team, but that may not be intuitive to outsiders (*e.g.*, unconventional code writing habits, such as special casing, indentation, exception handling, logging, etc.).

Our observations of why developers give up their search for rationale also inform how future techniques can be designed. We observed that the major factors why developers gave up were: (1) the poor state of the code and documentation, or (2) to avoid impacting the productivity of their team, or (3) their own. Therefore, it would be beneficial if future efforts to support the documentation of the rationale of code commits address these issues.

In particular, we believe that future efforts should support (1) recording the rationale of code commits when it is needed. For example, documenting rationale in pull request descriptions as preparation for code review , so that if rationale was missing from individual commits, it could be added. We also believe that future efforts should support (2 & 3) asynchronous communication mechanisms, so that the productivity of the seeker and the people eventually providing the rationale can be impacted as little as possible. Our participants also echoed the benefit of connecting software development artifacts that are related, so that their later comprehension is made easier:

“At my company, typically we link each commit, branch, and pull request (PR) to a ticket number, which should include details and discussion about the change and why it is being made.”

They also commented on the opposite: the difficulties they face when these centralized connections break.

“What usually makes the search more difficult is when code has been moved around, and the version control system lost track of its origin.”

Such support of rationale documentation could provide additional benefits. This practice of documenting needed components can be especially useful before somebody leaves the company — they already documented all the code changes that they performed. This could contribute to reducing the problem of not being able to find rationale that can only be provided by somebody who has already left the team.

Another way to support rationale documentation is having a tool to support conversations about the rationale of code commits. Such a tool would make it easier to document rationale after the discussions have happened — because the conversations themselves could help and simplify the process of producing structured documentation (as in bug-tracking

systems). Such a system could also make it easier to save the documentation in the right place, connected to the code change discussed.

Finally, tools could be developed to make it easier to find the rationale of code changes that have already been documented. If we had a system that kept track of the conversations about and documentation of the rationale of code changes, it would be easy to offer that information in easy-to-access ways in software development environments. The retrieved rationale of code changes could also be represented in different ways to fit the different tasks in which developers need to find it — similar to how the WhyLine system provides already-available program-slicing information in a much more user-friendly format [59]. Such varied representations may also require varied analyses, *e.g.*, rationale finding triggered by a debugging task might require in-depth navigation of the history of relevant code changes (*e.g.*, [60, 61, 62, 63, 64]), whereas rationale finding triggered by learning about an external tool might require a broad search for similar codes across the external codebase. Finally, these efforts should provide the right amount of information to developers, since giving too many recommendations can have hidden costs [65].

An example of such an enhancement of an IDE would be to provide the rationale of code changes while the user is actively debugging a feature. The IDE plugin could gather rationale components from previous commits, where the specific feature under debugging was changed. One of our participants commented on the importance of understanding the evolution of code (*i.e.*, “why the code is this way?”) while debugging:

“I believe that understanding the evolution of the code is just as important as understanding the current code. If you know where the code has been, you can get a sense of where it needs to evolve for the next release and to be able to avoid the pitfalls of past bugs.”

5.5 Our vision: how could the documentation of the rationale of code commits be automated?

The previous section describes how future software tools could support developers in the process of documenting and finding the rationale of code commits. Next, we discuss how some parts of that process could be more strongly automated.

First, we see opportunities to automate the detection of insufficiently documented rationale of code changes. This could be either code changes completely lacking the documentation, or lacking relevant components. Relevant components can be predicted by relying on code change characteristics. Automating this problem would be akin to predicting which code changes developers are likely to face the need to understand rationale components that are not currently documented. Such a tool could base its predictions on observing which characteristics are common among the code changes for which certain components are requested. We believe that this tool would be really valuable for developers to prioritize their efforts of documenting the rationale of code commits *i.e.*, starting with those for which their rationale is likely to be requested.

Next, we also see opportunities for automating the documentation of the rationale of code changes. Techniques could be developed to capture the common characteristics among the code changes with a specific answer to a given component of the rationale of code changes. If such common characteristics could be captured, future code changes with similar characteristics could be automatically labeled with the same answer to the rationale component. Similar mechanisms have been successfully applied to similar goals, *e.g.*, applying machine learning to detect design patterns [66]. Other approaches may also be promising, such as applying techniques for summarizing and documenting code changes like [67, 68], tools answering “what” and “why” questions about code changes like [38, 39, 59, 69, 70], and studies

of impact and risk of changes like [71, 72, 73].

Any efforts to automate or support the process of documenting or finding the rationale of code commits will benefit from the rich understanding provided by our model of the specific pieces of information (components) that developers may seek when they need it. Our model now allows future research efforts to generate *targeted* pieces of information to generate to improve the documentation of rationale.

5.6 Our vision: what other benefits could arise from good documentation of the rationale of code commits?

In a longer timeline, we also anticipate further benefits from having a codebase in which the rationale of code commits is carefully documented. In projects with well-crafted documentation of the rationale of code commits, many future useful analyses may be possible. Next are some of the kinds of analyses that we anticipate.

First, we expect that the documentation of the rationale for code commits may be useful data to improve the traceability of software requirements to areas of the source code. This is a well-known problem and research area, and having richer data may help improve the accuracy of existing techniques to infer traceability or inspire future ones.

Second, the documentation of rationale for code commits may reveal hidden properties of the system. For example, the documentation of code changes *side effects* component may reveal hidden properties of the system changes conflict [74, 75], energy consumption [76], and upgrade impact [77]. Those properties may not otherwise be visible in software documen-

tation. The documentation of rationale for code commits may also reveal assumptions that are not documented elsewhere, or development team habits that are otherwise not visible. Furthermore, the documentation of the rationale of code changes in highly successful software projects, *i.e.*, those with high quality, may reveal best practices that are not necessarily documented (are implicit in their decision-making), and from which other practitioners could learn.

The first step to studying the promise of these (or other) ideas for useful software analytics would start by identifying a current software project that already assigns high importance to documenting their code changes.

Chapter 6

Related Work

Existing work highlights the importance of rationale management throughout the software development life-cycle [1, 2]. Thus, multiple approaches and systems have been proposed to integrate rationale management in the process of software requirements engineering, software design, and software architecture [48, 78].

We focus on the rationale for code commits in the context of software evolution and maintenance. In this context, existing work studied part of the experience of developers seeking the rationale of code changes, finding that it was considered important and sometimes hard to find [5]. In contrast, this dissertation presents an in-depth study of the rationale of code changes. We identified an extensive list of the **tasks** in which developers need to learn the rationale of code changes. We decomposed the rationale of code changes into multiple **components**, and surveyed the developers **experiences** seeking the **rationale and its individual components**. Finally, we experimented with **predicting rationale documentation** in the pull request context.

In the following sections, we discuss the related work in several areas:

6.1 Rationale Management in Software Requirements, Design, and Architecture

Multiple extensions of requirements models were proposed to encourage the capture of rationale within them [79, 80]. In addition to these models, tools have been proposed to manage rationale of software requirements [78, 81, 82].

Many schemes have also been proposed to capture design and architecture rationale. The schemes can be divided into two categories: decision-centric *e.g.*, Lee and Lai [52] [52] and usage-centric approaches *e.g.*, Burge et al. [1] [1]. The decision-centric approaches [19, 83] focus on capturing the rationale as a decision-making process utilizing Toulmin’s model of argumentation [84] and Rittel’s Issue-Based Information System (IBIS) [51]. The usage-centric approaches focus on capturing rationale without representing the decision-making process [19, 48, 49, 50].

The usage-centric approaches “recognize that organizing rationale around decisions is not the best way to elicit and characterize some of the rationale needed for making appropriate design decisions” [1]. Jarczyk et al. [85] provided a survey of the systems developed to support design rationale, all of which were based on Toulmin’s model or IBIS [85]. Design rationale has also been supported by multiple tools [47, 86], which can help detect inconsistencies, omissions, and conflicts [87]. Mehrpour and LaToza [88] surveyed tools support for working with design decisions in code. Some recent tools focus on capturing design rationale from software artifacts like IRC discussions [89, 90, 91], or user reviews [92]. Our work, in turn, focuses on the rationale in software maintenance.

6.2 Rationale in Software Evolution and Maintenance

This section discusses the related work in which *rationale of code changes* were prescriptively decomposed and discovered to be a major information need.

6.2.1 Components of the Rationale of Code Changes

Burge et al. [1] *prescriptively* enumerate a few questions that may answer rationale in software maintenance [1]. We, in turn, provide a *descriptive* model of rationale in the context of software maintenance, from the perspective of what developers need to find when they seek it. Our descriptive approach extends Burge et al. [1]’s prescriptive approach as we discovered more questions that may answer rationale in software maintenance. For example, our change objective questions (*What did you want to achieve?*, *Why did you need to achieve that?*, and *What is the benefit of what you want to achieve?*) were not mentioned in Burge et al. [1] study. Our extended set of questions covers all of Burge et al. [1] prescriptively enumerated questions. Because our set of questions is descriptive of the developers’ needs while seeking and finding rationale for code changes, we believe these questions will help the finding and recording activities.

6.2.2 Tasks that need the Rationale of Code Changes

Past research focused on discovering various developer information needs within a given developer task (*e.g.*, collaborating [6], understanding code [11], understanding bug reports [93], understanding the life of bugs [94], or reviewing code [9]), and uncovered a breadth of aspects about them, such as how early [93] or frequently [11] each piece of information was needed. Our goal is instead to discover the tasks in which this particularly important information

need (the rationale for code commits) exists (*i.e.*, our study takes the opposite direction), and we study various aspects of it. Our results in this theme validate the findings of some previous studies, since we found that the rationale of code changes was needed in *learning* [1], *code review* [9, 13, 95], and *mentoring* [15].

Tao et al. [5] *prescriptively* proposed a list of seven “development scenarios” (*e.g.*, refactoring, developing new features, and fixing bugs) in which they expected developers to need to understand code changes. They asked developers to choose which ones they encountered most often, and found that the most often encountered one was “reviewing others’ changes” [5]. Our work had a different goal: to *discover* the tasks for which developers need to understand the rationale of code changes, as stated by developers, *i.e.*, in a *descriptive* fashion. As a result, we discovered a much more exhaustive list: we observed that developers needed to understand the rationale of code changes in eight tasks and 25 sub-tasks (as opposed to Tao et al. [5]’s seven scenarios).

6.2.3 Experience with the Rationale of Code Changes

Studies involving software history and developers’ information needs in the last decade establish a strong demand for rationale [1, 3, 4, 5, 6, 7, 9, 11, 12, 13, 15, 95, 96, 97]. Our work is motivated by these empirical studies highlighting the importance of rationale for code commits.

The most closely related study to ours is Tao et al. [5]. Tao et al. [5] found that the most important information needs for understanding code changes is rationale, which is sometimes easy and sometimes difficult. We replicated these two questions (importance and difficulty) of their study. We also studied additional questions in three additional contexts: needing, finding, and recording rationale. Our study validates their results since our partic-

ipants reported similar ratings for the importance and difficulty of finding the rationale for code changes. This similarity of results also shows that we studied a similar population of developers.

Our study extends Tao et al. [5] by finding the individual pieces of information that compose rationale; the experiences of developers needing, finding, and recording those individual pieces; and recommendations to improve their documentation. In particular, this finer level of granularity (*i.e.*, rationale of code commits) enables us to provide a possible explanation for one of the main phenomena observed by Tao et al. [5]: that rationale is easy to find when it is *well documented*. We posit that rationale is deemed well documented when it contains the specific components that the developer is seeking at that moment. For instance, a recent qualitative study Liang et al. [98] highlights multiple types of implementation decisions made by developers, rationale may be deemed well documented when it contains the *alternatives* decision sought by the developer.

6.3 Rationale Documentation Support

The software engineering community is strongly pushing for improving software documentation [99, 100, 101, 102, 103, 104]. Studies have been carried out to explore documentation issues [100, 101], lost knowledge in informal documentation [102], management of decision knowledge [103], and new documentation schemes [104]. These studies advocate the development of automatic techniques to assist and improve software documentation. Our research effort is motivated by the software documentation literature. The following sections discuss the related work to supporting the documentation of rationale for code changes.

6.3.1 Rationale documentation in code review context

Code review is one of many tasks for which rationale of code changes is needed. Inspired by recent research studies of code review, we chose code review context to build our experimental rationale documentation prediction approach. Recent code review research notices the need for rationale for code changes. Rationale is noticed by researchers who performed studies to understand code review as a process and identify its challenges and best practices [9, 13, 22, 105, 106, 107, 108]. Another group of researchers empirically studied code review data to discover communicative intention in code review questions [10], frequency and nature of design discussions [109], identifying large-review-effort code changes [110], and confusion in code reviews [14, 111, 112]. More specifically, recent research work empirically studied pull request templates to explore their contents, impacts, and developers' perceptions [17]. Some *rationale components* were present in Li et al. [17] exploration of PRT contents. Our results add to this body of knowledge and provide strategies that will contribute to the improvement of overall changes in documentation and code review practice.

6.3.2 Tools to support software changes documentation

Researchers developed techniques to summarize and document code changes, [68], clarify the rationale behind code changes [113, 114], answer “what” and “why” questions about code changes [59, 69, 70, 115, 116, 117], generate pull request descriptions [118], or generate release notes [119]. Other researchers mined software repositories to characterize commits [120, 121, 122, 123], and others studied the impact and risk of changes [72, 73, 124]. Repository maintainers prescribe *Pull Request Templates (PRT)* [16] to guide software developers in describing their code changes for code review. These technique target components of rationale, among other change information, in their efforts. These research papers target ra-

tionale components like *motivation*, *validation*, and *side effects* of code changes [69, 117, 118]. Our technique is motivated by these techniques and their effort in supporting software developers. Our work on prediction rationale pull request templates headers documentation, complement existing tools.

Chapter 7

Conclusions

The dissertation offers an in-depth examination of the rationale behind code changes in software systems. The aim is to better understand and document the rationale of code changes. To achieve our goal, We performed three studies to understand software developers' need for rationale of code changes, surveying software developers' experiences with code changes, and predicting the documentation of rationale components in pull request templates.

We found that software developers need to investigate code commits to understand their rationale when working on diverse tasks (and subtasks): *programming, working on bugs, communication, tools, documentation, project management, testing, and specifications..* We also found that software developers decompose the rationale of code commits into 15 separate components that they could seek when searching for rationale: *goal, need, benefits, constraints, alternatives, selected alternative, dependencies, committer, time, location, modifications, explanation of modifications, validation, maturity stage, and side effects.*

Our survey of developers' experiences with rationale of code changes uncovered issues and challenges software developers encounter while searching and recording rationale for code changes. We discovered that software developers have different experiences with different components of rationale. We present *human, team, and project* factors leading software developers to give up their search for rationale. For example, *changes in project personnel* make it hard for software developers to find the rationale behind code changes.

Our experimentation with building a predictor for rationale components documentation in pull request templates was effective in achieving high accuracy and recall. We found that the random forest prediction algorithm can perform better than other prediction algorithms. We also found that different feature groups (*i.e.*, template features) lead to better performance in the prediction task.

Our findings provide a better understanding of the need for rationale of code commits. In light of our findings, we emphasize the potential for improving software maintenance and quality through a better understanding and documentation of rationale. We presented our vision for rationale of code commits practitioners' documentation, tools support, and documentation automation. In addition, we presented the benefits of analyses that could arise from good documentation of rationale for code commits. Overall, this document contributes to the field of software development by shedding light on the rationale behind code changes and offering insights to enhance the development process.

Bibliography

- [1] J. E. Burge, J. M. Carroll, R. McCall, and I. Mistrik, *Rationale-Based Software Engineering*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008. [Online]. Available: <https://doi.org/10.1007/978-3-540-77583-6>
- [2] A. H. Dutoit, R. McCall, I. Mistrik, and B. Paech, *Rationale Management in Software Engineering*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006. [Online]. Available: <https://doi.org/10.1007/978-3-540-30998-7>
- [3] M. Codoban, S. S. Ragavan, D. Dig, and B. Bailey, “Software history under the lens: A study on why and how developers examine it,” in *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, Sep. 2015, pp. 1–10.
- [4] T. Fritz and G. C. Murphy, “Using information fragments to answer the questions developers ask,” in *Proceedings of the 32Nd ACM/IEEE International Conference on Software Engineering - Volume 1*, ser. ICSE ’10. New York, NY, USA: ACM, 2010, pp. 175–184. [Online]. Available: <http://doi.acm.org/10.1145/1806799.1806828>
- [5] Y. Tao, Y. Dang, T. Xie, D. Zhang, and S. Kim, “How do software engineers understand code changes?: An exploratory study in industry,” in *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, ser. FSE ’12. New York, NY, USA: ACM, 2012, pp. 51:1–51:11. [Online]. Available: <http://doi.acm.org/10.1145/2393596.2393656>
- [6] A. J. Ko, R. DeLine, and G. Venolia, “Information needs in collocated software development teams,” in *29th International Conference on Software Engineering (ICSE’07)*, May 2007, pp. 344–353.

- [7] T. D. LaToza and B. A. Myers, “Hard-to-answer questions about code,” in *Evaluation and Usability of Programming Languages and Tools*, ser. PLATEAU '10. New York, NY, USA: ACM, 2010, pp. 8:1–8:6. [Online]. Available: <http://doi.acm.org/10.1145/1937117.1937125>
- [8] T. Roehm, R. Tiarks, R. Koschke, and W. Maalej, “How do professional developers comprehend software?” in *Proceedings of the 34th International Conference on Software Engineering*, ser. ICSE '12. Piscataway, NJ, USA: IEEE Press, 2012, pp. 255–265. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2337223.2337254>
- [9] L. Pascarella, D. Spadini, F. Palomba, M. Bruntink, and A. Bacchelli, “Information needs in contemporary code review,” *Proc. ACM Hum.-Comput. Interact.*, vol. 2, no. CSCW, pp. 135:1–135:27, Nov. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3274404>
- [10] F. Ebert, F. Castor, N. Novielli, and A. Serebrenik, “Communicative intention in code review questions,” in *Proceedings - 2018 IEEE International Conference on Software Maintenance and Evolution, ICSME 2018*, 2018, pp. 519–523. [Online]. Available: <https://goo.gl/Bpoqj6>
- [11] W. Maalej, R. Tiarks, T. Roehm, and R. Koschke, “On the comprehension of program comprehension,” *ACM Trans. Softw. Eng. Methodol.*, vol. 23, no. 4, pp. 31:1–31:37, Sep. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2622669>
- [12] S. Srinivasa Ragavan, M. Codoban, D. Piorkowski, D. Dig, and M. Burnett, “Version control systems: An information foraging perspective,” *IEEE Transactions on Software Engineering*, pp. 1–1, 2019.
- [13] A. Ram, A. A. Sawant, M. Castelluccio, and A. Bacchelli, “What makes a code change easier to review: An empirical investigation on code change reviewability,” in

- Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ser. ESEC/FSE 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 201–212. [Online]. Available: <https://doi.org/10.1145/3236024.3236080>
- [14] F. Ebert, F. Castor, N. Novielli, and A. Serebrenik, “Confusion in Code Reviews: Reasons, Impacts, and Coping Strategies,” *SANER 2019 - Proceedings of the 2019 IEEE 26th International Conference on Software Analysis, Evolution, and Reengineering*, pp. 49–60, 2019.
- [15] B. e. l. e. m. Dagenais, H. Ossher, R. K. E. Bellamy, M. P. Robillard, and J. P. de Vries, “Moving into a new software project landscape,” in *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1*, ser. ICSE ’10. New York, NY, USA: Association for Computing Machinery, 2010, p. 275–284. [Online]. Available: <https://doi.org/10.1145/1806799.1806842>
- [16] “Creating a pull request template for your repository - GitHub Docs — docs.github.com,” <https://docs.github.com/en/communities/using-templates-to-encourage-useful-issues-and-pull-requests/creating-a-pull-request-template-for-your-repository>, [Accessed 30-May-2023].
- [17] Z. Li, Y. Yu, T. Wang, Y. Lei, Y. Wang, and H. Wang, “To follow or not to follow: Understanding issue/pull-request templates on github,” *IEEE Transactions on Software Engineering*, pp. 1–16, 2022.
- [18] M. Zhang, H. Liu, C. Chen, Y. Liu, and S. Bai, “Consistent or not? An investigation of using Pull Request Template in GitHub,” *Information and Software Technology*, vol. 144, p. 106797, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0950584921002354>

- [19] A. Tang, M. A. Babar, I. Gorton, and J. Han, “A survey of architecture design rationale,” *Journal of Systems and Software*, vol. 79, no. 12, pp. 1792 – 1804, 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0164121206001415>
- [20] M. Hilton, N. Nelson, T. Tunnell, D. Marinov, and D. Dig, “Trade-offs in continuous integration: Assurance, security, and flexibility,” in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2017. New York, NY, USA: ACM, 2017, pp. 197–207. [Online]. Available: <http://doi.acm.org/10.1145/3106237.3106270>
- [21] J. Singer, S. E. Sim, and T. C. Lethbridge, *Guide to Advanced Empirical Software Engineering*, F. Shull, J. Singer, and D. I. K. Sjøberg, Eds. London: Springer London, 2008. [Online]. Available: <https://doi.org/10.1007/978-1-84800-044-5>
- [22] D. Spadini, M. Aniche, M.-A. Storey, M. Bruntink, and A. Bacchelli, “When testing meets code review: Why and how developers review tests,” in *Proceedings of the 40th International Conference on Software Engineering*, ser. ICSE '18. New York, NY, USA: ACM, 2018, pp. 677–687. [Online]. Available: <http://doi.acm.org/10.1145/3180155.3180192>
- [23] R. Pohl and R. Pohl, “Confirmation bias,” in *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. Psychology Press, 2004, ch. 4, pp. 79–96. [Online]. Available: <https://books.google.com/books?id=k5gTes7yyWEC>
- [24] P. Biernacki and D. Waldorf, “Snowball sampling: Problems and techniques of chain referral sampling,” *Sociological Methods & Research*, vol. 10, no. 2, pp. 141–163, 1981. [Online]. Available: <https://doi.org/10.1177/004912418101000205>

- [25] B. L. BERG, “Qualitative research methods for the social sciences,” 2001.
- [26] S. Adolph, W. Hall, and P. Kruchten, “Using grounded theory to study the experience of software development,” *Empirical Software Engineering*, vol. 16, no. 4, pp. 487–513, 2011. [Online]. Available: <https://doi.org/10.1007/s10664-010-9152-6>
- [27] A. Begel and B. Simon, “Novice software developers, all over again,” in *Proceedings of the Fourth International Workshop on Computing Education Research*, ser. ICER '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 3–14. [Online]. Available: <https://doi.org/10.1145/1404520.1404522>
- [28] B. Saunders, J. Sim, T. Kingstone, S. Baker, J. Waterfield, B. Bartlam, H. Burroughs, and C. Jinks, “Saturation in qualitative research: exploring its conceptualization and operationalization,” *Quality & Quantity*, vol. 52, no. 4, pp. 1893–1907, 2018. [Online]. Available: <https://doi.org/10.1007/s11135-017-0574-8>
- [29] M. M. Hennink and B. N. Kaiser, “Saturation in qualitative research,” Sep 2019. [Online]. Available: <http://dx.doi.org/10.4135/9781526421036822322>
- [30] Y. Liu, S. Li, and E. Tilevich, “Toward a better alignment between the research and practice of code search engines,” in *2022 29th Asia-Pacific Software Engineering Conference (APSEC)*, Dec 2022, pp. 219–228.
- [31] D. Spencer and J. Garrett, *Card Sorting: Designing Usable Categories*. Rosenfeld Media, 2009. [Online]. Available: https://books.google.com/books?id=_h4D9gqi5tsC
- [32] A. Koriat, M. Goldsmith, and A. Pansky, “Toward a psychology of memory accuracy,” *Annual Review of Psychology*, vol. 51, no. 1, pp. 481–537, 2000, pMID: 10751979. [Online]. Available: <https://doi.org/10.1146/annurev.psych.51.1.481>

- [33] E. G. Jelihovschi and J. Faria, “Scottknott: A package for performing the scott-knott clustering algorithm in r,” *TEMA (São Carlos)*, vol. 15, 03 2014.
- [34] K. Herzig and A. Zeller, “Untangling changes,” 2011.
- [35] K. Herzig and A. Zeller, “The impact of tangled code changes,” in *2013 10th Working Conference on Mining Software Repositories (MSR)*, May 2013, pp. 121–130.
- [36] M. Dias, A. Bacchelli, G. Gousios, D. Cassou, and S. Ducasse, “Untangling fine-grained code changes,” in *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, March 2015, pp. 341–350.
- [37] E. Aghajani, C. Nagy, O. L. Vega-Márquez, M. Linares-Vásquez, L. Moreno, G. Bavota, and M. Lanza, “Software documentation issues unveiled,” in *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, May 2019, pp. 1199–1210.
- [38] S. Jiang, A. Armaly, and C. McMillan, “Automatically generating commit messages from diffs using neural machine translation,” in *Proceedings of the 32Nd IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE 2017. Piscataway, NJ, USA: IEEE Press, 2017, pp. 135–146. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3155562.3155583>
- [39] S. Jiang and C. McMillan, “Towards automatic generation of short summaries of commits,” in *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*, May 2017, pp. 320–323.
- [40] D. Graziotin, F. Fagerholm, X. Wang, and P. Abrahamsson, “On the unhappiness of software developers,” in *Proceedings of the 21st International Conference on Evaluation and Assessment in Software Engineering*, ser. EASE’17. New York, NY,

- USA: Association for Computing Machinery, 2017, p. 324–333. [Online]. Available: <https://doi.org/10.1145/3084226.3084242>
- [41] A. Murgia, P. Tourani, B. Adams, and M. Ortu, “Do developers feel emotions? An exploratory analysis of emotions in software artifacts,” *11th Working Conference on Mining Software Repositories, MSR 2014 - Proceedings*, pp. 262–271, 2014.
- [42] D. Graziotin, F. Fagerholm, X. Wang, and P. Abrahamsson, “On the unhappiness of software developers,” *ACM International Conference Proceeding Series*, vol. Part F128635, pp. 324–333, 2017.
- [43] G. Gousios, M. Pinzger, and A. v. Deursen, “An exploratory study of the pull-based software development model,” in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 345–355. [Online]. Available: <https://doi-org.ezproxy.lib.vt.edu/10.1145/2568225.2568260>
- [44] G. Gousios, M.-A. Storey, and A. Bacchelli, “Work practices and challenges in pull-based development: The contributor’s perspective,” in *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*, May 2016, pp. 285–296.
- [45] X. Zhang, Y. Yu, G. Georgios, and A. Rastogi, “Pull Request Decisions Explained: An Empirical Overview,” *IEEE Transactions on Software Engineering*, pp. 1–1, 2022, conference Name: IEEE Transactions on Software Engineering.
- [46] M. I. Azeem, S. Panichella, A. D. Sorbo, A. Serebrenik, and Q. Wang, “Action-based recommendation in pull-request development,” *Proceedings of the International Conference on Software and System Processes*, 2020.
- [47] T. R. Gruber and D. M. Russell, “Design rationale,” T. P. Moran and J. M.

- Carroll, Eds. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 1996, ch. Generative Design Rationale: Beyond the Record and Replay Paradigm, pp. 323–349. [Online]. Available: <http://dl.acm.org/citation.cfm?id=261685.261725>
- [48] F. Gilson and V. Englebert, “Rationale, decisions and alternatives traceability for architecture design,” in *Proceedings of the 5th European Conference on Software Architecture: Companion Volume*, ser. ECSA ’11. New York, NY, USA: ACM, 2011, pp. 4:1–4:9. [Online]. Available: <http://doi.acm.org/10.1145/2031759.2031764>
- [49] J. Tyree and A. Akerman, “Architecture decisions: demystifying architecture,” *IEEE Software*, vol. 22, no. 2, pp. 19–27, March 2005.
- [50] T. R. Gruber and D. M. Russell, “Design knowledge and design rationale : A framework for representation , capture , and use,” 1991.
- [51] W. Kunz and H. Rittel, *Issues as Elements of Information Systems*, ser. California. University. Center for Planning and Development Research. Working paper, no. 131. Institute of Urban and Regional Development, University of California, 1970, no. 131. [Online]. Available: <https://books.google.com/books?id=B-MaAQAAMAAJ>
- [52] J. Lee and K.-Y. Lai, “What’s in design rationale?” *Hum.-Comput. Interact.*, vol. 6, no. 3, pp. 251–280, Sep. 1991. [Online]. Available: http://dx.doi.org/10.1207/s15327051hci0603&4_3
- [53] S. Davies and M. Roper, “What’s in a bug report?” in *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, ser. ESEM ’14. New York, NY, USA: ACM, 2014, pp. 26:1–26:10. [Online]. Available: <http://doi.acm.org/10.1145/2652524.2652541>
- [54] N. Bettenburg, S. Just, A. Schröter, C. Weiss, R. Premraj, and T. Zimmermann,

- “What makes a good bug report?” in *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. SIGSOFT '08/FSE-16. New York, NY, USA: ACM, 2008, pp. 308–318. [Online]. Available: <http://doi.acm.org/10.1145/1453101.1453146>
- [55] G. Dziekan, “Checklists save lives,” *Bulletin of the World Health Organization*, vol. 86, no. 7, 2008.
- [56] A. M. Kazerouni, C. A. Shaffer, S. H. Edwards, and F. Servant, “Assessing incremental testing practices and their impact on project outcomes,” in *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 407–413. [Online]. Available: <https://doi.org/10.1145/3287324.3287366>
- [57] A. M. Kazerouni, J. C. Davis, A. Basak, C. A. Shaffer, F. Servant, and S. H. Edwards, “Fast and accurate incremental feedback for students’ software tests using selective mutation analysis,” *J. Syst. Softw.*, vol. 175, p. 110905, 2021. [Online]. Available: <https://doi.org/10.1016/j.jss.2021.110905>
- [58] A. Gautam, S. Vishwasrao, and F. Servant, “An empirical study of activity, popularity, size, testing, and stability in continuous integration,” in *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, May 2017, pp. 495–498.
- [59] A. Ko and B. Myers, “Debugging reinvented,” in *2008 ACM/IEEE 30th International Conference on Software Engineering*, May 2008, pp. 301–310.
- [60] F. Servant and J. A. Jones, “History slicing,” in *2011 26th IEEE/ACM International Conference on Automated Software Engineering (ASE 2011)*, Nov 2011, pp. 452–455.
- [61] —, “History slicing: Assisting code-evolution tasks,” in *Proceedings of the ACM*

- SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, ser. FSE '12. New York, NY, USA: Association for Computing Machinery, 2012. [Online]. Available: <https://doi.org/10.1145/2393596.2393646>
- [62] —, “Chronos: Visualizing slices of source-code history,” in *2013 First IEEE Working Conference on Software Visualization (VISSOFT)*, Sep. 2013, pp. 1–4.
- [63] —, “Fuzzy fine-grained code-history analysis,” in *Proceedings of the 39th International Conference on Software Engineering*, ser. ICSE '17. IEEE Press, 2017, p. 746–757. [Online]. Available: <https://doi.org/10.1109/ICSE.2017.74>
- [64] F. Servant, “Supporting bug investigation using history analysis,” in *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, Nov 2013, pp. 754–757.
- [65] X. Jin and F. Servant, “The hidden cost of code completion: Understanding the impact of the recommendation-list length on its efficiency,” in *2018 IEEE/ACM 15th International Conference on Mining Software Repositories (MSR)*, May 2018, pp. 70–73.
- [66] R. Barbudo, A. Ramírez, F. Servant, and J. R. Romero, “GEML: A grammar-based evolutionary machine learning approach for design-pattern detection,” *J. Syst. Softw.*, vol. 175, p. 110919, 2021. [Online]. Available: <https://doi.org/10.1016/j.jss.2021.110919>
- [67] L. F. Cortés-Coy, M. Linares-Vásquez, J. Aponte, and D. Poshyvanyk, “On automatically generating commit messages via summarization of source code changes,” in *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*, Sep. 2014, pp. 275–284.
- [68] R. P. Buse and W. R. Weimer, “Automatically documenting program changes,”

- in *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE '10. New York, NY, USA: ACM, 2010, pp. 33–42. [Online]. Available: <http://doi.acm.org/10.1145/1858996.1859005>
- [69] S. Rastkar and G. C. Murphy, “Why did this code change?” in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 1193–1196. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2486788.2486959>
- [70] A. W. Bradley and G. C. Murphy, “Supporting software history exploration,” in *Proceedings of the 8th Working Conference on Mining Software Repositories*, ser. MSR '11. New York, NY, USA: ACM, 2011, pp. 193–202. [Online]. Available: <http://doi.acm.org/10.1145/1985441.1985469>
- [71] S. Zhang and M. D. Ernst, “Which configuration option should i change?” in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: ACM, 2014, pp. 152–163. [Online]. Available: <http://doi.acm.org/10.1145/2568225.2568251>
- [72] S. Jiang, C. McMillan, and R. Santelices, “Do programmers do change impact analysis in debugging?” *Empirical Software Engineering*, vol. 22, no. 2, pp. 631–669, Apr 2017. [Online]. Available: <https://doi.org/10.1007/s10664-016-9441-9>
- [73] C. Rosen, B. Grawi, and E. Shihab, “Commit guru: Analytics and risk prediction of software commits,” in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2015. New York, NY, USA: ACM, 2015, pp. 966–969. [Online]. Available: <http://doi.acm.org/10.1145/2786805.2803183>
- [74] B. Shen, M. A. Gulzar, F. He, and N. Meng, “A characterization study of merge

- conflicts in java projects,” *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 2, mar 2023. [Online]. Available: <https://doi-org.ezproxy.lib.vt.edu/10.1145/3546944>
- [75] S. S. Towqir, B. Shen, M. A. Gulzar, and N. Meng, “Detecting build conflicts in software merge for java programs via static analysis,” in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE ’22. New York, NY, USA: Association for Computing Machinery, 2023. [Online]. Available: <https://doi-org.ezproxy.lib.vt.edu/10.1145/3551349.3556950>
- [76] Y.-W. Kwon and E. Tilevich, “Reducing the energy consumption of mobile applications behind the scenes,” in *2013 IEEE International Conference on Software Maintenance*, 2013, pp. 170–179.
- [77] T. Dumitraş, P. Narasimhan, and E. Tilevich, “To upgrade or not to upgrade: Impact of online upgrades across multiple administrative domains,” *SIGPLAN Not.*, vol. 45, no. 10, p. 865–876, oct 2010. [Online]. Available: <https://doi-org.ezproxy.lib.vt.edu/10.1145/1932682.1869530>
- [78] A. V. Lamsweerde, “Goal-oriented requirements engineering: A guided tour,” in *Proceedings Fifth IEEE International Symposium on Requirements Engineering*, Aug 2001, pp. 249–262.
- [79] A. I. Anton, “Goal-based requirements analysis,” in *Proceedings of the Second International Conference on Requirements Engineering*, April 1996, pp. 136–144.
- [80] H. Kaiya, H. Horai, and M. Saeki, “Agora: Attributed goal-oriented requirements analysis method,” in *Proceedings IEEE Joint International Conference on Requirements Engineering*, Sep. 2002, pp. 13–22.

- [81] ITU-T. (2018, nov) User requirements notation (urn) – language definition. [Online]. Available: <http://handle.itu.int/11.1002/1000/13711>
- [82] D. Amyot, “Introduction to the user requirements notation: learning by example,” *Computer Networks*, vol. 42, no. 3, pp. 285 – 301, 2003, iITU-T System Design Languages (SDL). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128603002445>
- [83] C. Potts and G. Bruns, “Recording the reasons for design decisions,” in *Proceedings of the 10th International Conference on Software Engineering*, ser. ICSE '88. Los Alamitos, CA, USA: IEEE Computer Society Press, 1988, pp. 418–427. [Online]. Available: <http://dl.acm.org/citation.cfm?id=55823.55863>
- [84] S. E. Toulmin, *The Uses of Argument*, 2nd ed. Cambridge University Press, 2003.
- [85] A. P. J. Jarczyk, P. Loffler, and F. M. Shipmann, “Design rationale for software engineering: A survey,” in *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*, vol. ii, Jan 1992, pp. 577–586 vol.2.
- [86] J. E. Burge and D. C. Brown, “Software engineering using rationale,” *Journal of Systems and Software*, vol. 81, no. 3, pp. 395 – 413, 2008, selected Papers from the 2006 Brazilian Symposia on Databases and on Software Engineering. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0164121207001203>
- [87] A. Tang, Y. Jin, and J. Han, “A rationale-based architecture model for design traceability and reasoning,” *Journal of Systems and Software*, vol. 80, no. 6, pp. 918 – 934, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0164121206002287>
- [88] S. Mehrpour and T. D. LaToza, “A survey of tool support for working with design

- decisions in code,” *ACM Comput. Surv.*, jul 2023, just Accepted. [Online]. Available: <https://doi-org.ezproxy.lib.vt.edu/10.1145/3607868>
- [89] R. Alkadhi, T. Lata, E. Guzman, and B. Bruegge, “Rationale in development chat messages: An exploratory study,” in *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, May 2017, pp. 436–446.
- [90] —, “Rationale in development chat messages: An exploratory study,” in *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, May 2017, pp. 436–446.
- [91] R. Alkadhi, M. Nonnenmacher, E. Guzman, and B. Bruegge, “How do developers discuss rationale?” in *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, March 2018, pp. 357–369.
- [92] Z. Kurtanović and W. Maalej, “Mining user rationale from software reviews,” in *2017 IEEE 25th International Requirements Engineering Conference (RE)*, Sep. 2017, pp. 61–70.
- [93] S. Breu, R. Premraj, J. Sillito, and T. Zimmermann, “Information needs in bug reports: Improving cooperation between developers and users,” in *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*, ser. CSCW ’10. New York, NY, USA: Association for Computing Machinery, 2010, p. 301–310. [Online]. Available: <https://doi.org/10.1145/1718918.1718973>
- [94] J. Aranda and G. Venolia, “The secret life of bugs: Going past the errors and omissions in software repositories,” in *2009 IEEE 31st International Conference on Software Engineering*, 2009, pp. 298–308.

- [95] F. Ebert, F. Castor, N. Novielli, and A. Serebrenik, “Confusion in code reviews: Reasons, impacts, and coping strategies,” in *2019 IEEE 26th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, Feb 2019, pp. 49–60.
- [96] T. Roehm, R. Tiarks, R. Koschke, and W. Maalej, “How do professional developers comprehend software?” in *Proceedings of the 34th International Conference on Software Engineering*, ser. ICSE '12. Piscataway, NJ, USA: IEEE Press, 2012, pp. 255–265. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2337223.2337254>
- [97] F. Ebert, F. Castor, N. Novielli, and A. Serebrenik, “Communicative intention in code review questions,” in *2018 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, Sep. 2018, pp. 519–523.
- [98] J. T. Liang, M. Arab, M. Ko, A. J. Ko, and T. D. LaToza, “A qualitative study on the implementation design decisions of developers,” 2023.
- [99] J. Zhi, V. Garousi-Yusifoglu, B. Sun, G. Garousi, S. Shahnewaz, and G. Ruhe, “Cost, benefits and quality of software development documentation: A systematic mapping,” *Journal of Systems and Software*, vol. 99, pp. 175 – 198, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0164121214002131>
- [100] E. Aghajani, C. Nagy, O. L. Vega-Marquez, M. Linares-Vasquez, L. Moreno, G. Bavota, and M. Lanza, “Software Documentation Issues Unveiled,” *Proceedings - International Conference on Software Engineering*, vol. 2019-May, pp. 1199–1210, 2019.
- [101] E. Aghajani, C. Nagy, M. Linares-Vasquez, L. Moreno, G. Bavota, M. Lanza, and D. C. Shepherd, “Software documentation: The practitioners’ perspective,” in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, ser. ICSE '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 590–601. [Online]. Available: <https://doi.org/10.1145/3377811.3380405>

- [102] M. á. r. Lima, I. Ahmed, T. Conte, E. Nascimento, E. Oliveira, and B. Gadelha, “Land of Lost Knowledge: An Initial Investigation into Projects Lost Knowledge,” *International Symposium on Empirical Software Engineering and Measurement*, vol. 2019-Septemer, pp. 1–6, 2019.
- [103] A. Kleebaum, J. O. Johanssen, B. Paech, and B. Bruegge, “How do practitioners manage decision knowledge during continuous software engineering?” *Proceedings of the International Conference on Software Engineering and Knowledge Engineering, SEKE*, vol. 2019-July, pp. 735–740, 2019.
- [104] M. P. Robillard, A. Marcus, C. Treude, G. Bavota, O. Chaparro, N. Ernst, M. A. Gerosa, M. Godfrey, M. Lanza, M. Linares-Vásquez, G. C. Murphy, L. Moreno, D. Shepherd, and E. Wong, “On-demand developer documentation,” in *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, Sep. 2017, pp. 479–483.
- [105] A. Bosu, J. C. Carver, C. Bird, J. Orbeck, and C. Chockley, “Process aspects and social dynamics of contemporary code review: Insights from open source development and industrial practice at microsoft,” *IEEE Transactions on Software Engineering*, vol. 43, no. 1, pp. 56–75, Jan 2017.
- [106] L. MacLeod, M. Greiler, M. Storey, C. Bird, and J. Czerwonka, “Code reviewing in the trenches: Challenges and best practices,” *IEEE Software*, vol. 35, no. 4, pp. 34–42, July 2018.
- [107] A. Bacchelli and C. Bird, “Expectations, outcomes, and challenges of modern code review,” *Proceedings - International Conference on Software Engineering*, pp. 712–721, 2013.

- [108] N. Fatima, S. Nazir, and S. Chuprat, “Understanding the impact of feedback on knowledge sharing in modern code review,” in *2019 IEEE 6th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*, Dec 2019, pp. 1–5.
- [109] F. E. Zanaty, T. Hirao, S. McIntosh, A. Ihara, and K. Matsumoto, “An empirical study of design discussions in code review,” *International Symposium on Empirical Software Engineering and Measurement*, 2018.
- [110] S. Wang, C. Bansal, and N. Nagappan, “Large-scale intent analysis for identifying large-review-effort code changes,” *Information and Software Technology*, p. 106408, 2020. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0950584920300033>
- [111] F. Ebert, F. Castor, N. Novielli, and A. Serebrenik, “An exploratory study on confusion in code reviews,” *Empirical Software Engineering*, 2020.
- [112] —, “Confusion detection in code reviews,” *Proceedings - 2017 IEEE International Conference on Software Maintenance and Evolution, ICSME 2017*, pp. 549–553, 2017.
- [113] L. F. Cortés-Coy, M. Linares-Vásquez, J. Aponte, and D. Poshyvanyk, “On automatically generating commit messages via summarization of source code changes,” in *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*, Sep. 2014, pp. 275–284.
- [114] M. Linares-Vásquez, L. F. Cortés-Coy, J. Aponte, and D. Poshyvanyk, “Changscribe: A tool for automatically generating commit messages,” in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 2, May 2015, pp. 709–712.
- [115] S. Jiang, A. Armaly, and C. McMillan, “Automatically generating commit messages from diffs using neural machine translation,” in *Proceedings of the 32Nd*

- IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE 2017. Piscataway, NJ, USA: IEEE Press, 2017, pp. 135–146. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3155562.3155583>
- [116] S. Jiang and C. McMillan, “Towards automatic generation of short summaries of commits,” in *2017 IEEE/ACM 25th International Conference on Program Comprehension (ICPC)*, May 2017, pp. 320–323.
- [117] J. Shen, X. Sun, B. Li, H. Yang, and J. Hu, “On automatic summarization of what and why information in source code changes,” in *2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1, June 2016, pp. 103–112.
- [118] Z. Liu, X. Xia, C. Treude, D. Lo, and S. Li, “Automatic generation of pull request descriptions,” *Proceedings - 2019 34th IEEE/ACM International Conference on Automated Software Engineering, ASE 2019*, pp. 176–188, 2019.
- [119] L. Moreno, G. Bavota, M. Di Penta, R. Oliveto, A. Marcus, and G. Canfora, “Automatic generation of release notes,” in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, ser. FSE 2014. New York, NY, USA: Association for Computing Machinery, 2014, p. 484–495. [Online]. Available: <https://doi.org/10.1145/2635868.2635870>
- [120] A. Alali, H. Kagdi, and J. I. Maletic, “What’s a typical commit? a characterization of open source software repositories,” in *2008 16th IEEE International Conference on Program Comprehension*, June 2008, pp. 182–191.
- [121] L. P. Hattori and M. Lanza, “On the nature of commits,” in *Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE’08. Piscataway, NJ, USA: IEEE Press, 2008, pp. III–63–III–71. [Online]. Available: <https://doi.org/10.1109/ASEW.2008.4686322>

- [122] A. Hindle, D. M. German, M. W. Godfrey, and R. C. Holt, “Automatic classification of large changes into maintenance categories,” in *2009 IEEE 17th International Conference on Program Comprehension*, May 2009, pp. 30–39.
- [123] R. Goyal, G. Ferreira, C. Kästner, and J. Herbsleb, “Identifying unusual commits on github,” *Journal of Software: Evolution and Process*, vol. 30, no. 1, p. e1893, 2018, e1893 smr.1893. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/smr.1893>
- [124] S. Zhang and M. D. Ernst, “Which configuration option should i change?” in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014. New York, NY, USA: ACM, 2014, pp. 152–163. [Online]. Available: <http://doi.acm.org/10.1145/2568225.2568251>