

Adapting Installation Instructions in Rapidly Evolving Software Ecosystems

Haoyu Gao , Christoph Treude , *Member, IEEE*, and Mansooreh Zahedi , *Member, IEEE*

Abstract—README files play an important role in providing installation-related instructions to software users and are widely used in open source software systems on platforms such as GitHub. Software projects evolve rapidly alongside their dependencies in dynamic software ecosystems, requiring frequent updates to installation instructions. These instructions are crucial for users to start with a software project. Despite their significance, there is a lack of systematic understanding regarding the documentation efforts invested in README files and the triggers behind them. To fill the research gap, we conducted a qualitative study, investigating 400 GitHub repositories with 1,163 README commits that focused on updates in installation-related sections. Our research revealed six major categories of changes in the README commits, namely pre-installation instructions, installation instructions, post-installation instructions, help information updates, document presentation, and external resource management. We further provide detailed insights into modification behaviours and offer examples of these updates. We also studied the triggers for the documentation updates, which led to three categories including errors in the previous documentation, changes in the codebase, and need for documentation improvement. Based on our findings, we proposed a README template tailored to cover the installation-related sections for documentation maintainers to reference when updating documents. We further validated this template by conducting an online survey and a pull request study, identifying that documentation readers find the augmented documents based on our template to be generally of better quality, and documentation maintainers find it useful. We further provide recommendations to practitioners for maintaining their README files, as well as motivations for future research directions. These recommendations encompass completeness, correctness and up-to-dateness, and information presentation considerations. The proposed research directions include the development of automated tools, in particular for documentation updates, and conducting empirical studies to enhance comprehension of the needs of documentation users.

Index Terms—Software documentation, README files, qualitative analysis, open source software, installation instructions, software ecosystem.

Received 14 July 2024; revised 7 January 2025; accepted 5 March 2025. Date of publication 18 March 2025; date of current version 18 April 2025. Recommended for acceptance by M. Monperrus. (*Corresponding author: Haoyu Gao.*)

Haoyu Gao and Mansooreh Zahedi are with the School of Computing and Information Systems, The University of Melbourne, Melbourne, VIC 3053, Australia (e-mail: haoyug1@student.unimelb.edu.au).

Christoph Treude is with the School of Computing and Information Systems, Singapore Management University, Singapore 178903.

Digital Object Identifier 10.1109/TSE.2025.3552614

I. INTRODUCTION

GOOD software documentation serves as a vital source of information for different stakeholders involved in software systems. In particular, README files are among the first software documents that users, such as other developers and end-users, and newcomers encounter when accessing an Open Source Software (OSS) project, leaving them with a crucial first impression and providing them with an introductory overview of the project [1], [2]. At the time of this study, GitHub, the most popular open-source code platform, hosts more than 372 million repositories,¹ underscoring the high demand for maintaining high-quality README files.

One of the most common types of content in GitHub README files is related to basic instructions (e.g., project installation guide) [1]. A well-written “How” section could alleviate the effort for developers and users, especially newcomers, to get hands on with the projects. In fact, overcoming the challenge of setting up an environment is one of the hurdles for newcomers to contribute to OSS projects [3].

Software projects evolve rapidly alongside their dependencies in dynamic software ecosystems [4], requiring frequent updates to installation instructions. Therefore, README files need to be carefully maintained; otherwise, various software documentation issues may arise. For example, software tutorials that focus on installing software tools, configuring virtual computing environments, and other related continuous deployment processes have been found to be not executable in more than half of the cases investigated by related work [5]. Furthermore, previous empirical studies identified inappropriate installation instructions, outdated content, and content clarity as frequently appearing documentation issues [6], [7]. These factors collectively constitute the impediments for both developers and project users to use software repositories, while empirical evidence on the current documentation practices and their triggers could pave the way for mitigating these issues from source.

Although previous studies have pinpointed software documentation issues, particularly in installation-related topics [6], [7], [8], their analyses were broad in scope, with the aim of synthesising different aspects of documentation issues from surveys or software documents. In contrast, documentation updates provide a new perspective consisting of the modifications and two states: before and after the modification, while triggers

¹<http://web.archive.org/web/20230424004513/https://en.wikipedia.org/wiki/GitHub>

offer the motivations that drive documentation updates in the rapidly evolving software ecosystem. In fact, increased documentation activities and efforts have demonstrated their efficacy in improving the popularity of and encouraging participation in GitHub projects [9]. Considering the documentation issues and the benefits of investing efforts to maintain high-quality documentation, knowledge of the practices of the README file efforts and the triggers behind them becomes essential. This knowledge reveals documentation activities in the dynamic software ecosystems. It not only provides practitioners with recommendations but also provides researchers with insights for the development of automated documentation tools. However, empirical evidence to support practices related to software documentation efforts remains scarce.

In this work, we conducted an empirical study to qualitatively analyse README file modification activities in installation-related topics to fill the research gap. We mined GitHub repositories from GHArchive that were active between 2015-01-01 and 2023-02-28,² applying stringent filters on different features for the repositories, resulting in 12,908 projects. We examined a statistically significant sample of 1,168 README file commits within 400 randomly selected repositories and adopted qualitative analysis to synthesise the types of README update tasks in the modified sections.

Based on our analysis, we synthesised a comprehensive taxonomy of 189 update behaviours, which are related to the categories of: (1) *pre-installation instructions*, (2) *installation instructions*, (3) *post-installation instructions*, (4) *help information*, (5) *document presentation*, and (6) *external resource management*. We then devised a process to identify the triggers for documentation updates. Our analysis leads to three categories of documentation update triggers: (1) *errors in the previous documentation*, (2) *changes in the codebase*, and (3) *need for documentation improvement*. After that, we proposed a README template tailored to cover the installation-related sections for documentation maintainers to reference when updating documents. The template validation process indicates that readers find README files that apply our template to be of generally increased quality, and the documentation maintainers also find the process of applying the template useful. This serves as a fundamental step towards automated documentation updates. We further discussed their implications for documentation maintainers and software engineering researchers based on our discoveries from perspectives of completeness, correctness and up-to-dateness, and information presentation considerations.

Our key contributions are as follows: (1) A large-scale repository dataset, which can be used for future research. (2) Formalisation of a detailed list of software development tasks related to README file content. (3) A comprehensive taxonomy for the types of installation-related documentation modifications for OSS projects. (4) A README template validated by documentation readers to be of better quality compared to original README files. (5) A categorisation of triggers for documentation updates (6) Actionable recommendations for

OSS documentation maintainers and researchers, in particular the automated documentation updates direction.

The rest of the paper is structured as follows: Section II covers the related work; Section III introduces our study design and methodology adopted to conduct our study; Section IV presents the results of our qualitative analysis for the installation instruction taxonomy; Section V presents the results for triggers for documentation updates; Section VI provides the README template from our study and its validation process with an on-line survey; Section VII discusses the implications of this study to practitioners and researchers; Section VIII covers the threats to the validity of our study; Section IX provides the conclusion. Our replication package is at https://github.com/Haoyu-Gao/README_Empirical_Study, with the dataset in the root folder.

II. RELATED WORK

A. Empirical Studies on Software Documentation

Various empirical studies have been conducted to cover software documentation from different perspectives.

One source of information for researchers is directly from software engineering practitioners. Several studies have investigated the impact of software documentation. Kajiko-Mattson [10] studied the practices of software documentation in Swedish organisations, identifying rudimentary documentation requirements relevant within corrective maintenance. Garousi et al. [11] conducted an industrial case study on the usage and usefulness of different documentation artefacts in the software development life cycle. Their findings indicated variations in documentation usage based on information needs, and the recommendations derived from their study have been integrated into the company's practices where the experiment took place. Software documentation issues are also investigated. Uddin and Robillard [12] conducted a large-scale survey on API documentation failure factors among IBM practitioners, and identified the most pressing issues related to content, as opposed to presentation. Aghajani et al. [7] presented the views of practitioners on various issues of software documentation, along with their proposed solutions to address these issues. Moreover, the quality of software documents is another perspective that has been studied. Treude et al. [13] recruited technical editors and synthesised a framework to assess the quality of the software documentation from the respondents. Their work provides visions for a potential unified quality framework for assessing software documentation.

However, these studies are limited in their context (e.g., practitioners' organisation and knowledge or software artefacts), making it difficult to generalise the results to broader software documentation artefacts. Another source of information for researchers lies in these software artefacts themselves. Different software documentation artefacts have been evaluated, revealing a series of issues. Wen et al. [14] conducted a large-scale empirical study of code-comment inconsistencies, identifying the taxonomy of inconsistencies fixed by developers. Mirhosseini and Parnin [5] evaluated the executability of software tutorials and discussed the categories of errors that can lead to the failure of code block executions. Aghajani et al. [6]

²<https://www.gharchive.org/>

conducted a qualitative study on software documentation issues from Stack Overflow, GitHub issues, GitHub pull requests, and mailing lists, identifying both content and process issues. Gao et al. [15] studied machine learning documentation to highlight the insufficient effort in documenting ethical considerations.

Among the software documentation artefacts, README files play an important role in introducing OSS repositories. They are relatively easy to access, with a rich content editing history. Despite their significance, README files have received relatively less attention in research, and most existing work focuses on more abstract perspectives, such as structure and content topics. Prana et al. [1] conducted a qualitative study involving manual annotation of 4,226 README file sections to categorise README content into different purposes. Liu et al. [16] studied the patterns adopted for README files by GitHub Java repositories and investigated the association between repository popularity and README file structures. The study identified 32 clusters representing common README file structures, and the findings suggested that repositories adhering to GitHub guidelines for README files tend to get more stars. Meanwhile, Wang et al. [17] investigated the factors that influence the popularity of GitHub repositories from the README file angle, and identified a significant difference between popular and non-popular repositories in README file-related features.

The richness of README file edit history provides us with the possibility to investigate software documents from a process perspective. In this work, we provide a more detailed study of README files centred around installation-related instructions, considering both their content and their process, to offer concrete suggestions for documentation maintainers and researchers.

B. Automated Tools for Software Documentation

Inspired by various empirical pieces of evidence regarding the roles and issues of software documentation, there has been much research on automated tool support for software documentation.

Previous research focuses on supporting various software documentation artefacts with non-documentation artefacts, such as source code. To begin with, code summarisation is an actively evolving field, with techniques developed for generating code comments [18], [19], [20], [21], commit messages [22], [23], [24], code reviews [25], [26] and release notes [27], [28]. Technical debt, including code comment inconsistency detection [29], [30], and self-admitted technical debt [31], [32], is also investigated to facilitate software maintenance.

Efforts have also been made to improve and augment software documentation content and navigation across this content. Treude and Robillard [33] investigated a technique for augmenting the content of API documents by incorporating information from Stack Overflow. Treude et al. [34] extracted development tasks from software documents, bridging the gap between software documentation and the information needs of software developers. Gao et al. [35] mined README files from GitHub and performed transfer learning to simplify complicated documentation sentences. In fact, Robillard et al. [36] have

proposed a paradigm shift for automatic documentation generation based on developers' queries, given their contexts.

However, to facilitate more effective automated documentation-writing tools, it is essential to have a deeper understanding not only of the software documentation artefacts but also of the patterns of behaviours associated with their updates. In this work, we conducted a qualitative analysis of installation-related documentation update behaviours along with the triggers behind them, with the aim of gaining a deep understanding and providing motivation for developing automated software tools for documentation writing.

C. OSS Barriers and Project Installation

The term "installation" is widely used in different contexts. A simple installation procedure only requires fetching a remote binary file, while a complicated situation would confirm that minimum system requirements are met, identify and resolve shared library dependencies, verify the integrity of each downloaded component, and other procedures [37]. *Due to the absence of a precise definition of installation-related instructions, we define them to encompass all procedures occurring after the environment setup and the preceding project usage phase in this paper.*

Recent studies [3], [38] have highlighted both documentation problems and the configuration of the development environment as barriers for newcomers when they first contribute to Open Source Software (OSS) repositories. Salerno et al. [39] discovered that student self-efficacy improves as a result of participating in OSS courses, but simultaneously they also perceive an increase in documentation issues. In particular, one of the identified issues revolves around the challenges associated with setting up the development environment.

Moreover, another study [1] pointed out that the major content category of README files is "How", which includes instructions from setting up the environment to installation and running the project. Therefore, README files play a vital role in guiding newcomers and software users with installation-related topics.

In contrast to previous research on the human perspective of OSS barriers for newcomers, this work synthesises empirical evidence on software documentation activities, specifically on parts related to the installation, based on evidence derived from a large-scale dataset that we mined. Such empirical evidence will be useful to mitigate documentation-related barriers for newcomers when engaging with OSS projects.

III. STUDY DESIGN AND METHODOLOGY

The goal of this study is to investigate how installation-related instructions in README files are updated. We formulate three research questions (RQs) to guide our study:

RQ1: How are installation sections in README files modified? This RQ will reveal different modification behaviours for updating the documentation. We studied the documentation modification behaviour rather than the most recent version of README files because the modification behaviours




Fig. 1. Data preparation process.

not only provide the documentation structure and contents but also offer more insights into the documentation maintenance efforts on OSS repositories. By understanding the documentation maintenance practice, we could provide implications to practitioners and researchers for better maintaining the README files. We performed a qualitative analysis of the installation-related sections within our mined data to answer this question. One of the target outcomes of this RQ is a template encompassing all essential components to facilitate better documentation activities for the maintainers.

RQ2: What are the triggers for the README updates in the installation sections? This RQ investigates the reasons that trigger documentation updates in the software system, providing insights into what drives documentation updates in a rapidly evolving software ecosystem. The triggers and update behaviours collectively provide strong empirical evidence for automatic documentation updates as a future direction.

RQ3: Could a template for installation instructions be devised to improve documentation quality? This research question aims to validate our proposed README template, which is the fundamental first step towards automated documentation updates. We conducted a survey asking participants to evaluate the perceived documentation quality before and after applying the README template, based on an established documentation assessment framework [13]. We further conducted a pull request study, similar to Malaquias et al. [40] to evaluate the practitioners' perspective on our template.

A. Data Preparation

In this paper, we only consider GitHub repositories that use Python or Java as their main programming language. Python and Java are two widely used programming languages that span a broad range of applications, from developing easy-to-use tools to large-scale servers. In fact, Python and Java consistently rank among the top three most used programming languages on GitHub,³ following JavaScript.

Although the GitHub API provides access to its data, the rate limit is not enough for collecting all repositories that meet our needs. Therefore, we used GH Archive,⁴ which collects GitHub data based on the Event API. Fig. 1 describes our overall repository collection procedure.

As our focus is on update behaviours, we started with commits rather than repository content. In addition, we want the repositories retrieved to contain a history of updating their installation-related instructions. Therefore, we looked at the “PushEvent” in the data records, and retrieved those with commit messages containing keywords that we pre-defined after tokenisation and lemmatisation. Keywords were derived from four base keywords “instruct”, “guide”, “install” and “deploy”. These terms were chosen to represent the content genres and topics that we aimed to include in our analysis. They were further combined with tags used in Zahedi et al. [41]. We expand the set of keywords by finding both the noun and verb forms of words in WordNet [42]. This procedure resulted in a total of 21,238,948 repositories. After that, we preserved only 5,335,695 repositories that are not forks and have at least one pull request in the history. This filtering step ensures that the selected repositories demonstrate development activity, rather than being solely for archival purposes or mirrors of other open-source platforms such as BitBucket or GitLab [43]. Further refinement narrows down the selection to 886,365 by considering repositories that use Python or Java as their primary programming language. As the aim of this study is to investigate modifications within the README file, we focus only on repositories with at least 500 commits [14] to obtain repositories that went through code changes in different phases of the software development lifecycle. In this way, the README updates ideally capture a broader range of changes. Sampling from the entire population would give higher weight to repositories with fewer commits, potentially reducing diversity in the types of change reflected in README updates. We also used 10 stars as a threshold for eliminating toy projects [44]. This further reduces the number of repositories to 28,896.

The earlier keyword filtering was applied to all commit messages without considering the semantics of changes specifically made to README files. Therefore, we further applied a filter on the README-changing commits, requiring them to have at least one occurrence in the keyword set, which decreased the dataset size to 14,958 records.

There are caveats that need to be taken care of when mining software repositories on GH Archive. First, some repositories are no longer available, either because they have been deleted or made private by their owners. After removing these non-available repositories, we retain 14,485 repositories. Second, due to renaming and other factors, some repositories are redirected by the GitHub API to alternative names, potentially causing duplicate repository issues. We further eliminated these repositories, reducing the number of repositories to 13,245.

In the last step, we removed outliers that fall outside of the scope of this study. We used FastText [45] to filter out 292 repositories with non-English README content. We manually inspected the repositories included after these filtering steps and found that there were some non-software repositories in there. Non-software repositories refer to projects primarily containing content other than traditional software code. These repositories may include, but are not limited to, collections of awesome lists, educational materials, and personal websites. To exclude those systematically, we considered several factors that

³<https://octoverse.github.com/2022/top-programming-language>

⁴<https://www.gharchive.org/>




Fig. 2. Ratio of non-software repositories among top n repositories.

might indicate non-software projects, including (1) README commits to overall commits ratio, (2) number of README commits, (3) number of README commits per day, and (4) repository size. We sorted the repositories by these factors in descending order and manually labelled the top 100 records for each factor, respectively. Fig. 2 illustrates how the percentage of non-software repositories varies as more projects are included. Although the number of README commits can identify the most unrelated repositories, the accuracy is merely 51% for the top 100 entries. Meanwhile, regarding README commits to overall commit ratio, a plateau on the left side of the red vertical line indicates a consistently high accuracy of identifying non-software projects, with the red line representing the percentage of 52.9% of this ratio. Therefore, we rounded the ratio to 50%, used it as a threshold, and discarded these repositories, eventually obtaining a dataset consisting of 12,908 repositories.

B. README Commit Classification

We randomly sampled 400 repositories from the entire population, ensuring a confidence level above 95% with a 5% margin of error. For the 400 repositories sampled, there are 25,209 README commits in total, which need to be classified as “relevant” or “irrelevant” to the installation-related instructions. Given the prohibitively large number for manual annotation, we employed a semi-automated heuristic to significantly reduce the manual workload for identifying the relevant README updates.

Specifically, we extracted the different scales of headers from the sampled README files in their commit histories, with scales ranging from $\langle h1 \rangle$ to $\langle h6 \rangle$. We considered the occurrence of a particular header as one if it appeared at least once in the commit history. All repositories have at least one header, and Table I shows the distribution of header scales. Almost every repository has level 1 and level 2 headers, while level 3 and level 4 headers are also quite common. In contrast, the frequency of level 5 and level 6 headers is scarce. A closer examination identifies that level 5 and 6 do not convey meaningful document structural information; instead, some of them are used merely for emphasis within the text.

TABLE I
COUNTS FOR HEADER SCALES

Header Scale	$\langle h1 \rangle$	$\langle h2 \rangle$	$\langle h3 \rangle$	$\langle h4 \rangle$	$\langle h5 \rangle$	$\langle h6 \rangle$
Count	387	382	312	167	45	8

The basic idea of our heuristics is that README headers work as an overview of the contents beneath them. By examining the section headers, we can discern whether the sections are relevant to our target without delving into the detailed content. Therefore, we extracted the 100 most frequently occurring headers from $\langle h1 \rangle$ to $\langle h4 \rangle$, along with the 100 most frequently occurring word tokens. Three authors independently annotated these headers and tokens into categories of “relevant”, “irrelevant” or “not sure”. Agreement on the annotations was achieved through discussions in regular meetings. In total, we identified 66 relevant headers, 14 relevant keywords, and 177 irrelevant headers. These keywords and headers are also made available in our replication package.

Based on the information obtained from the headers, we developed a semi-automated labelling heuristic. Starting from the $\langle h1 \rangle$ headers, this program recursively examines whether the modified sections in the README commit are in the “relevant” set or not. If any of the modified sections has a “relevant” header, the program will mark the commit as relevant. Conversely, if all modified sections are categorised under the “irrelevant” headers, the commit is marked accordingly as irrelevant. Otherwise, the program cannot determine and proceeds to the next level of headers, repeating the decision process. If it still cannot ascertain the relevance after inspecting the headers at level 4, it assigns a “not sure” label, indicating that human annotation is required.

To verify the accuracy of this program, we manually annotated four repositories, and the program reached an accuracy of 92.4% (133/144). Subsequently, we applied this program to all repositories, resulting in 7,161 relevant, 8,583 irrelevant and 9,465 “not sure” commits generated by the semi-automatic labelling heuristic.

To annotate the data points categorised as “not sure”, the first author randomly selected 50 records and three authors independently performed the annotation. The outcome of the annotation reached a Fleiss’ Kappa score of 0.69, indicating a good level of agreement. The three authors further discussed the disagreed entries and came to an eventual agreement. Subsequently, the first author continued to annotate the remaining commits. Finally, within the dataset of 25,230 README commits, 8,673 commits were identified as relevant and 16,557 as irrelevant. The high accuracy rate (over 92.4%) and high recall (82.6%) indicate that our heuristic is quite capable of categorising the README documentation.

C. Dataset Characteristics

Before the details of the data analysis, we briefly introduce the characteristics of our curated dataset. The 8,673 identified relevant commits from the previous step cover 378 of 400 sampled repositories, that is, 94.5% of the repositories




Fig. 3. Comparison between repositories in our dataset and ones above 2,000 stars.

contain at least one README commit modifying installation-related sections. To assess whether we had missed relevant installation-related documents, for each repository, we further searched “INSTALL.md”, “INSTALLATION.md”, “docs/INSTALL.md”, and “docs/INSTALLATION.md” within its project folder (case insensitive), and consider them as additional installation documents other than the default README file. Only two repositories contain additional installation documents. Based on this observation, the majority of repositories in our analysed dataset use the README file as the primary installation documentation source, and the provided installation steps in the README file thus serve as the primary source for documentation users to access them. This observation further motivates us to focus on the analysis of the README file.

Regarding the representativeness of our analysed repositories, we selected three measures, namely repository size, total number of pull requests, and total number of issues in the repository. The repository size serves as a proxy for the complexity of the project, while the other measures quantify the project’s interaction with its users. We accessed Java and Python repositories on Github with above 2,000 stars using GHS [44]. The 2,000 stars threshold was also used in a previous study [46] when trying to build an Automatic Program Repair dataset, with the aim to collect data from popular repositories. Fig. 3 shows the box plot for these three measures in our data and the repositories having above 2,000 stars. We performed the Mann-Whitney U test for these three measures, resulting in a p-value of 0.38, 0.02, and 0.06 for total size, total issues, and total pull requests, respectively. This indicates that there are no differences between the groups in the population in size and total pull requests. For total issues, we computed the Cohen’s d effect size [47], which reached -0.09, indicating that the difference is negligible. Additionally, the majority of data from both sources ranges between 100 and 1000, suggesting a medium-to-large scale of issues in both sources. This result shows that our analysed projects are comparable to popular GitHub projects in terms of their scale and user interactions.

On average, each repository contains 23 relevant README commits, with a standard deviation of 25.7, indicating a skewed distribution. Specifically, with a minimum of 1 and a maximum of 224, the 25th, 50th, and 75th percentiles are 8, 14 and 31 commits, respectively.




Fig. 4. Average updated sections per commit.

For the qualitative study, we need to obtain a sample from the data. However, due to the skewed distribution of README commits across repositories, a completely random sample would result in a dataset dominated by repositories with a large number of README commits, thereby reducing diversity across repositories.

Therefore, we performed a cluster sampling [48], splitting the dataset into four clusters: repositories with README commits from the minimum to the first quartile, from the first quartile to the second quartile, from the second quartile to the third quartile, and from the third quartile to the maximum. Within each cluster, we performed a random sample with the sample size required to achieve a 95% confidence level (confidence interval: 5), resulting in samples of 206, 280, 324, and 358 README commits for each respective cluster. In total, we conducted a large-scale qualitative study on 1,168 README commits.

We calculated the correlation coefficient between the number of relevant README commits and the number of stars, as well as with the number of years the project has existed. The Pearson correlation coefficient reached 0.19 and 0.01, respectively. This statistic indicates that projects that actively participate in documentation activities, particularly for the installation-related sections, have a positive correlation to the project’s popularity, while the number of relevant README commits does not necessarily grow with the years of the project. However, it is important to note that these correlations are relatively weak, while other factors also contribute to the popularity of repositories.

From the section level, Fig. 4 shows the number of sections that are added, removed, or modified per commit on average. The figure illustrates that the README updates mostly focus on sections in the top two hierarchies of headers, while the modification operation is much more frequent than the addition and removal operations.

Based on the results of Fig. 4, we visualised the top 20 most updated keywords within the section headers at the first two levels to have a better understanding of the updated sections. The heat maps are presented in Fig. 5.




Fig. 5. Update keywords heat-map (A: Added, R: Removed, M: Modified).

From the figures, it is clear that more addition and removal operations are conducted on the level 2 hierarchy, while most sections are related to themes including installation, usage, build, setup and running.

D. Data Analysis

We conducted a qualitative study of the annotated data to answer our research question. Specifically, we synthesised a comprehensive taxonomy of 6 categories, consisting of 189 different codes. The findings provide a comprehensive understanding of the nature of the modifications with rich implications for practitioners and researchers.

We adopted open coding [49] for qualitatively analysing the update events. The first author conducted the data analysis and created spreadsheets containing all raw data and codes. The spreadsheets were shared among all authors, and the second and third authors cross-validated all codes, concepts, and categories to reduce bias and increase the reliability of the findings [50]. Concepts and categories were discussed among authors through regular meetings and finalised after several rounds of revisions.

We started our analysis by searching each commit on GitHub and recording the text summarising the changes from the perspective according to our research question: updates for the installation-related sections. We call this process “key points extraction and getting familiar with the data”. Then we rigorously performed open coding to assign each README commit with one to many succinct descriptions of their update modifications. This process took into account both the key points extracted and the content of the README commits. The additional investigation into the raw README commit data is to address any unclear key points and to double-check whether any part of the modification is missing.

During the coding procedure, we constantly compared the emerging codes within and across different README commits to synthesise a higher level of abstraction called concepts, and later the highest level of abstraction called categories. Fig. 6(a)

illustrates examples of the procedure of generating one of the categories called “Pre-installation Instruction Updates”. Fig. 6(b) and 6(c) further explain the process for how we assign codes for different changes within the commit. We have made our analysed data available. Please refer to the Replication Package section for more details.

IV. RESULTS

Table II presents the hierarchical taxonomy of 189 codes that we identified from the 1,168 data points to answer our RQ1. The updates are categorised into six main groups: (1) *pre-installation instruction*; (2) *installation-related instruction*; (3) *post-installation instruction*; (4) *help information*; (5) *document presentation*; and (6) *external resource management*.

To give readers a better understanding of our analysed result, Fig. 7 synthesises our analysis in a timeline fashion, corresponding to pre-installation, installation, and post-installation, where different software development tasks are included in each phase. Moreover, various information perspectives regarding the README files are also displayed within this figure, including instructions, help information, external resource links, and document presentation. For each of the categories, we describe representative examples and discuss implications for practitioners and researchers.

It is worth noting that, compared to other qualitative studies [1], [6], [7], [12], our work focuses on modifications to Markdown-formatted text, which involves two pieces of documents. Due to the substantial space to illustrate examples of README updates, we will only include figures for a few examples. Additional examples are made accessible through URLs that we have formatted as clickable links in Table II. We have also included the URLs and their corresponding screenshots in our replication package.


A. Pre-Installation Instruction

Pre-installation instructions, also seen as environment setup instructions, contain preparation instructions before installation. In our analysed samples, 628 out of 1,168 performed updates are related to environment setup topics in their README documents. In this category, we synthesised three concepts, which are discussed as follows:

Dependency Setup (527).⁵ Dependencies enable code reuse and require maintenance effort as the project and other dependencies evolve [51]. In our analysed dataset, more than 40% of the documentation updates contain dependency-related instructional modifications, indicating the importance of this perspective.

Among the modifications of the dependency setup, the concept of programming language version compatibility refers to details concerning the target version of the programming language upon which the software is built, ensuring its accurate execution. Fig. 8 includes examples of how the programming language version information in the README files is updated.

⁵Number in the parenthesis represents the frequency of this concept.



(a) Steps for Synthesising Categories from Raw Data

```
34 * Create and activate a virtualenv
35 * sh
36 * python3 -m venv venv
37 * source venv/bin/activate
38 *
39 * Install and upgrade wheel inside the virtualenv
40 * sh
41 * pip install --upgrade pip wheel
42 *
43 * Install orchestra
44 * sh
45 * pip cache remove orchestra
46 * pip install --force-reinstall https://github.com/revng/revng-orchestra/archive/master.zip
47 *
48 * Install dependencies (Ubuntu only)
49 * sh
50 * ./orchestra/ci/install-dependencies.sh
51 *
52 * Install Python-only dependencies (all distributions)
53 * sh
54 * ./orchestra/ci/install-python-dependencies.sh
55 *
```

(b) Example for Commit 245f55bfc7

```
36 ## Dependencies
37
38 * LinuxCNC 2.8^ or master
39
40 * Python 2.7
41 * Python 3.6
42 * PyQt5 or PySide2
43
44 - QtPyVCP is developed and tested using the LinuxCNC Debian 9 and 10 x64 (stretch and buster)
45
46 + QtPyVCP is developed and tested using the Debian 10 and 11 x64 (buster and bullseye)
```

(c) Example for Commit 3830b325fc

Fig. 6. Coding process and examples of assigned codes.

Updates are mostly small, only reflecting the new version of the programming language that their projects are using. However, as seen in Fig. 7(c), when the programming language requirement information is newly added to the documentation, it means that the previous README suffered from incomplete information, and efforts are made to remedy the lack of information in the programming language requirements.

Apart from the programming language, other dependencies setup include external libraries and modules (e.g., API), as well as third-party software (e.g., database) that projects rely on. The inclusion of new features, the resolution of technical debt during project development, and the updates of constantly evolving dependencies collectively contribute to the volatility of the dependency management information within the README files. Indeed, most updates in this aspect involve updating dependency information, as indicated in Example 1, and instructions on installation of dependencies as indicated in Example 2. We observed a common practice of manifesting dependencies by listing them (e.g., in the forms of plain texts, bullet points, or tables), with additions, updates, and deletions to

reflect the corresponding changes. Meanwhile, the patterns for providing the installation instruction of dependencies are quite diverse. Centralised commands (e.g., requirements.txt) or separate installation commands for each dependency are widely used. However, maintaining the separate command style requires more effort, as individual attention to each dependency is needed, making modifications potentially more time-consuming and the process more complicated.

In Example 3, the commit adds configuration instructions for the database on which the project relies. This underscores the importance of incorporating essential configuration details when dependencies are not ready for use following the dependency installation step, which is a common scenario for third-party software dependencies. Interestingly, some examples, such as commit Example 4, add guidelines to verify the success of dependency installation.

Platform Setup (82). Platform setup information includes details for systems on which the software operates, which is the second largest updated concept for the pre-installation instructions category.

TABLE II
EXAMPLES FOR QUALITATIVE ANALYSIS

Category	Concept	# Codes	Frequency	Example ID	Addition	Modification	Removal
Pre-installation instruction	Dependency Setup	23	527 (45.3%)	1, 2, 3, 4	47.8%	34.8%	17.4%
	Platform Setup	11	82 (7.1%)	5, 6, 7	54.5%	18.2%	27.3%
	Virtual Machine and Hardware	6	19 (1.6%)	—	66.7%	16.7%	16.7%
	Total	40	628 (54.0%)		52.5%	27.5%	20.0%
Installation-related instruction	Source code	23	267 (23.0%)	8, 9	43.4%	43.5%	13.0%
	Published package	14	162 (13.9%)	10, 11, 12, 13	57.1%	28.6%	14.3%
	Total	37	429 (36.9%)		48.6%	37.8%	13.5%
Post-installation instruction	Project running instruction	12	185 (15.9%)	14, 15	37.5%	50.0%	12.5%
	Installation check	3	13 (1.1%)	16	40.0%	60.0%	0.0%
	Running Tests	6	32 (2.8%)	—	88.9%	0.0%	11.1%
	Deployment Instruction	3	9 (0.8%)	—	75%	0.0%	25.0%
	Upgrade Instruction	4	6 (0.5%)	—	100.0%	0.0%	0.0%
	Uninstall Instruction	2	3 (0.3%)	—	0.0%	0.0%	100.0%
	Total	30	248 (21.3%)		58.6%	24.1%	17.2%
Help information	Tutorial updates	6	70 (6.0%)	17	50.0%	33.3%	16.7%
	Explanation	14	125 (10.7%)	18, 19	61.5%	38.5%	0.0%
	Notes and troubleshooting	16	70 (6.0%)	20, 21	68.8%	6.3%	25.0%
	Total	36	265 (22.8%)		62.9%	22.9%	14.3%
Document Presentation	Documentation structural modification	7	197 (16.9%)	22, 23	0.0%	100.0%	0.0%
	Content formatting and cleaning	10	119 (10.2%)	24, 25, 26, 27	0.0%	100.0%	0.0%
	Presentation fixing	2	61 (5.2%)	28, 29	0.0%	100.0%	0.0%
	Text rephrasing and editing	5	70 (6.0%)	30, 31	0.0%	100.0%	0.0%
	Documentation debt	4	30 (2.6%)	32	25.0%	0.0%	75.0%
	Total	28	477 (41.0%)		3.6%	85.7%	10.7%
External resources management	Project external documentation	3	120 (10.3%)	33	25.0%	50.0%	25.0%
	Project non-documentation artefact	11	136 (11.7%)	34	44.4%	33.3%	22.2%
	Third-party external resources	4	110 (9.5%)	35	60.0%	20.0%	20.0%
	Total	18	366 (31.5%)		44.4%	33.3%	22.2%




Fig. 7. Development tasks in README files.

In our analysed data, we found that 48 instances involve modification in setting up virtual environments, and the majority of them are related to newly added instructions such as the one indicated in *Example 5*. Furthermore, *Example 6* explicitly mentions the virtual environment as their recommended way of installation of the project.

Operating system information is another platform setup concern. We observed efforts that mostly focused on adding operating system compatibility information (e.g., *BiglyBT comes in several editions for different operating systems. Mac OSX,*

```

1 1 # Abkhazia
2 2
3 3
4 4 --**Abkhazia is no more compatible with python2, please use python3 instead. The
5 5 -- last version compatible with python2 is
6 6 -- [abkhazia-0.3.1](https://github.com/bootphon/abkhazia/releases/tag/v0.3.1)**
7 7
8 8
9 9
10 10
11 11 - ## Online documentation can be found [here]
12 12 (https://coml.lscp.ens.fr/docs/abkhazia)**
13 13 + ## See Online documentation https://coml.lscp.ens.fr/docs/abkhazia
14 14
15 15
16 16
17 17 The Abkhazia project makes it easy to obtain simple baselines for
18 18 supervised ASR (using [Kaldi](http://kaldi-asr.org)) and ABX tasks
19 19 (using [ABXpy](https://github.com/bootphon/ABXpy)) on the large corpora
20 20 of speech recordings typically used in speech engineering, linguistics
21 21 or cognitive science research.
22 22
23 23
24 24
25 25
26 26
27 27
28 28
29 29
30 30
31 31
32 32
33 33
34 34
35 35
36 36
37 37
38 38
39 39
40 40
41 41
42 42
43 43
44 44
45 45
46 46
47 47
48 48
49 49
50 50
51 51
52 52
53 53
54 54
55 55
56 56
57 57
58 58
59 59
60 60
61 61
62 62
63 63
64 64
65 65
66 66
67 67
68 68
69 69
70 70
71 71
72 72
73 73
74 74
75 75
76 76
77 77
78 78
79 79
80 80
81 81
82 82
83 83
84 84
85 85
86 86
87 87
88 88
89 89
90 90
91 91
92 92
93 93
94 94
95 95
96 96
97 97
98 98
99 99
100 100

```

(a)

```

34 34
35 35
36 36
37 37
38 38
39 39
40 40
41 41
42 42
43 43
44 44
45 45
46 46
47 47
48 48
49 49
50 50
51 51
52 52
53 53
54 54
55 55
56 56
57 57
58 58
59 59
60 60
61 61
62 62
63 63
64 64
65 65
66 66
67 67
68 68
69 69
70 70
71 71
72 72
73 73
74 74
75 75
76 76
77 77
78 78
79 79
80 80
81 81
82 82
83 83
84 84
85 85
86 86
87 87
88 88
89 89
90 90
91 91
92 92
93 93
94 94
95 95
96 96
97 97
98 98
99 99
100 100

```

(b)

```

52 52
53 53
54 54
55 55
56 56
57 57
58 58
59 59
60 60
61 61
62 62
63 63
64 64
65 65
66 66
67 67
68 68
69 69
70 70
71 71
72 72
73 73
74 74
75 75
76 76
77 77
78 78
79 79
80 80
81 81
82 82
83 83
84 84
85 85
86 86
87 87
88 88
89 89
90 90
91 91
92 92
93 93
94 94
95 95
96 96
97 97
98 98
99 99
100 100

```

(c)

Fig. 8. Programming language examples.

Linux and Windows use the full BiglyBT-API based on Java 8.) or updating operating system version information as indicated in *Example 7*, which updated the support operating system from Debian 9 and 10 to Debian 10 and 11.

Virtual Machine and Hardware (19). Depending on the nature and usage purpose of different software repositories, hardware requirements may differ for different software projects and sometimes even require deployment on cloud servers. We observed commit updates that include hardware specification information or steps for configuring a virtual machine. It would be important to include this information for certain projects, although this does not generalise to all repositories.

Summary of pre-installation instruction results: The concept of dependency setup includes management of programming language version and other dependencies to reflect software system changes. However, the concept of platform setup majorly includes updates in providing virtual environment instructions and operating system information of the software project. Virtual machine and hardware instructions are less frequently encountered.

B. Installation-Related Instruction

In this category, we synthesised two concepts, which are discussed as follows:

Installation-related Instructions from Source Code Operations (267). Installation from source code requires users to obtain the source code and then locally execute a set of pre-defined operations, encompassing compilation, building, and other procedures, to install the software project. While projects may use various technologies, including Docker, Pip, and Maven, we identified in our analysed data that source code-level installation is the most commonly adopted method that the projects tend to initially provide. In addition, the inclusion of newly added instructions also accounts for various operating systems. Given that commands for Windows systems typically differ from those for Linux, additional information in the README file instructions or even the source code is needed to facilitate a smooth installation experience for different users.

Documentation efforts are also observed to provide better-encapsulated installation alternatives that are easier to install, as well as versions of projects targeting different user groups. For example, in commit *Example 8*, the installation option initially provided includes obtaining the source code and performing compilation. The newly added instruction introduces an alternative approach specifically for building and running a Docker image from the Dockerfile within the source code. This provides an integrated installation experience, alleviating concerns related to compilation and other dependency management issues.

Regarding update behaviours for the installation instructions from source code, a wide range of project-specific patterns are identified. For example, modifications within code instructions covering project names, file paths for executables, and configuration file names are usually trivial, reflecting changes at the

file level and repository level. A simple example in commit *Example 9* migrated the codebase from GitLab to GitHub, while also changing the project name from *atoml-local* to *catlearn*, leading to the corresponding adjustment of the commands in the README file. Updates to the source code installation option also involve adjustments of parameters, with various patterns observed for this type of change. As an illustrative example, commit *Example 9* altered the Docker compose command with an additional “-d” command indicating a detached mode to run the Docker container.

Installation-related Instructions from Published Package Operations (162). In contrast, users can also install software and libraries from package management systems, such as Pip, Conda, and Maven, which enable an easier installation experience. Published package installation instructions are typically not the first option OSS developers incorporate in their system, considering the additional effort required to configure the distribution binary on package management systems. In our analysed data, only 33% of the newly added published package installation instructions do not have other installation alternatives in their previous README files. For example, in commit *Example 10*, a new alternative was added to install the project through *PyPI*, which is a package management system. However, alternative options for source code level installation were already provided in the previous README file, making the newly added published package installation option an additional choice.

Regarding the update operations in the published package installation instructions, we identified a predominant pattern that focuses on updating the project version. Trivial updates are usually found in this type of modification, with merely a version number increment for the repository’s latest stable version, and sometimes combined with the development version. For example, both the stable version and the snapshot version were incremented within the code instruction for the Maven dependency of the repository in this commit *Example 11*.

Apart from version modification, other noticeable update operations for published package installation include changes for package names or package sources and changes for package management technology instructions. For example, the path to the image on DockerHub was updated in commit *Example 12*, reflecting the changes in the source of the project. As an example of the change in technology instruction, in commit *Example 13*, the Gradle command *compile* is updated to *implementation* as the previous keyword was deprecated. The rest of the updates are more specific to reflect the project development status.

Summary of installation-related instruction results: The concept of installation-related instruction from source code operation is the most commonly adopted choice in our analysed data. Their update behaviours vary, covering factors including project meta-data changes and parameter changes. On the other hand, the concept of installation-related instructions from published package operations is less frequent, and they are usually used as additional operations on top of the source code installation alternative.

C. Post-Installation Instruction

After project installation, different components are usually covered within README files. In total, 248 of 1,168 data points are related to post-installation instruction updates. We synthesised six concepts based on their tasks and discuss them as follows:

Project Running Instruction (185). The most important aspect after project installation is to execute the software to accomplish specific tasks, and we observed that the vast majority of modifications in this concept fall into the running instruction changes.

OSS projects often come with different execution options that accept different parameters as input for different use cases. In these scenarios, adding different combinations of parameters with varying but similar functionalities could lead to bloated documents, but we did notice a few instances in practice. Meanwhile, we observe that newly added commands often contain the default setting or offer a template with placeholders for parameters. For example, commit *Example 14* added a template execution command with two brief examples. Moreover, it included a “get help” command, allowing documentation users to access detailed running options.

In addition, OSS projects sometimes contain multiple runnable modules, each serving a specific aspect of the project, such as server/worker modules in distributed system projects and training/inference modules in machine learning projects. Properly documenting each component’s execution instruction is essential for users to make use of the software product. Within our analysed data, we identified in *Example 15* that additional execution commands were added for other blocks of functions.

Installation Check (13). After a complicated software installation process, the installation check commands allow users to verify that the software has been installed and configured successfully. In our analysed data, only a few repositories incorporated the installation check process into their documentation. Represented by commit *Example 16*, the added installation verification process normally contains the execution of a simple predefined code snippet.

Running Tests (43). Tests may not be required from the user’s perspective, but become crucial in the context of installation in development mode. Although **Deployment Instruction (9)**, **Upgrade Instruction (6)**, **Uninstall Instruction (3)** also occurred in our analysed data, their occurrence is relatively infrequent.

Summary of post-installation instruction results: The most dominant concept in this category is project running instruction. Changes in this concept are found to be related to updating parameter options and file paths to the executables. Other concepts are less frequently encountered but they all constitute the installation-related software development tasks.

D. Help Information

Previous sections discuss updates of instructional commands, which focus on providing commands for the reader

of the documentation to follow. In contrast, the help information serves as a resource for documentation readers to get clarification or examples to better understand the project and commands. In total, 265 out of 1,168 data points are related to the Help Information category. In this category, we synthesised three concepts, which are discussed as follows:

Tutorial Updates (70). Limited in space and purpose, README files typically do not offer comprehensive usage instructions, such as API documentation. We observed that most updates within this concept involve providing code demos for the basic usage of a project. Most documentation changes are trivial, merely made to align with the source code modifications. In *Example 17*, updates are made within the code demos, altering all naming paradigms for the methods of using the model.

Explanation Information Updates (125). Instructional commands require users to follow step-by-step commands to complete tasks. In cases where users lack related background knowledge, they may need to blindly follow instructions. In our analysed data, we observed examples where efforts were made to enhance the accessibility of README file instructions by updating explanations.

Explanations can be added as complementary materials to different types of instruction. For example, a new dependency *scikit-learn* was added in commit *Example 18*. However, instead of providing a simple manifest of dependencies, it includes explanations of when and why this dependency is required: *If there is no interest in this level of optimisation/benchmarking, then the codebase is not needed.* This type of information empowers documentation users to make informed decisions about whether to include a particular package when installing the project.

Explanation is not only added where the instructions are optional, when the instructions are not common enough or potentially have side effects, additional explanation is also needed, such as the case in *Example 19*. In this example, an explanation is added for the parameter option “no-deps”, detailing its intended outcome, along with caveats specifying when manual installation of dependencies for other modules may be necessary.

Notes and Troubleshooting (70). In our analysed data, we identified documentation efforts in updating notes and troubleshooting content. A brief definition of “notes” is **additional** information or important considerations that must be taken into account when following the instruction. Meanwhile, troubleshooting is dedicated to providing guidance towards common errors that people may encounter.

For example, the commit *Example 20* added a note, specifying additional steps that need to be done for *linux-armv6* and *linux-armv7* systems. On the other hand, in commit *Example 21*, troubleshooting information was added, including a description of the problem related to an outdated docker-compose version and the solution provided to address this issue. These instructions are crucial for specific user groups, and the inclusion of common notes and troubleshooting information can enhance users’ experiences by reducing the time spent searching for solutions.

```

32 - ## Building
33
34 - For build instructions, refer to the
   [documentation](#documentation).

17 + ## Building and using Eradiate
18
19 + For build and usage instructions, refer to the
   [documentation]
   (https://eradiate.readthedocs.org/).

7 - ## Python modules
8
9 - DepthAI requires [numpy]
   (https://numpy.org/) and [opencv-
   python]

7 + ## Python modules (Dependencies)
8
9 + DepthAI Demo requires [numpy]
   (https://numpy.org/), [opencv-python]
   (https://pypi.org/project/opencv-

```

Fig. 9. Section header name update examples.

Summary of help information results: The most frequent updates in this category are the explanation information updates, followed by notes and troubleshooting updates and tutorial updates. Unlike instructional updates, updates in this category focus more on providing clarifications and examples, but are also less frequent.

E. Document Presentation

Different from the previous sections, document presentation updates involve less content modification. Instead, the emphasis is on organising the documentation to enhance visual presentation and reading comprehension. In total, 477 of 1,168 data points are associated with this category, indicating a relatively heavy documentation effort. In this category, we synthesised five concepts, which are discussed as follows:

Documentation Structural Modification (197). The structure of the documentation contains essential information for the reader to distinguish its source, topic, and other characteristics [52]. In our analysed data, we identified a large amount of effort in document structure updates, reaching around 16% of all data.

Section headers are indicators of the content within the sections. Fig. 9 shows examples for updating the section names to more accurately reflect the content under the section. Fig. 9(a) updates the header name as a result of the inclusion of new distinct contents in the section, while Fig. 9(b) undergoes an update with the aim of improving the clarity of the header of the section.

In addition, when a section is cluttered with excessive contents, it becomes necessary to re-organise the materials in a more logical manner by relocating them into other sections. In commit *Example 22*, the instructions in the previous README file are largely gathered in a single section. The updates, as indicated in the commit, distribute the corresponding parts among different subsections, enhancing the organisation of the documentation for better comprehension. This also happens when instructions are introduced with the same purpose but with different technologies. For example, in commit *Example 23*, the Docker installation instruction is moved under “Installation with Docker” after adding the Pip install instruction.

Other modification operations including adjusting the section header size, section hierarchical changes which organise multiple sections under a single common section, and section order

updates are less commonly found in practice, but all contribute to the organisation of information in README files.

Content Formatting and Cleaning (119). README content formatting involves using Markdown syntax to either highlight or organise the content. Various information is commonly formatted, including large code blocks in *Example 24*, inline code blocks in *Example 25*, and URLs in *Example 26*. All of these formatting operations contribute to enhancing the presentation of the documentation, making the README files easier to follow. Specifically, the formatting on code blocks enables better highlighting of code components, while the formatting on URLs facilitates enhanced accessibility and clearer content display. Plain text is also formatted to highlight important information, as indicated in *Example 27*, with bold and italic texts.

Other formatting includes adding empty lines to separate logically different paragraphs, using bullet points, or an ordered list to organise the instructions within the documentation. These formatting operations organise the content within a section in a more logical manner to improve the reading experience.

Presentation Fixing (61). Presentation fixing in the domain of README documentation presentation involves addressing two perspectives: grammar and typo fixing, as well as Markdown syntax fixing.

Errors in grammar and typography can influence the reading experience, but may not severely impact the documentation’s usability. In commit *Example 28*, the whole commit is dedicated to fixing typos to improve readability. For example, the word *Pyhton* is corrected to *Python*. On the other hand, Markdown syntax errors sometimes prevent the content from correctly rendering, making the documentation difficult to read. For example, in commit *Example 29*, an extra space is inserted between the hash symbol and section header, ensuring the correct display of the header.

Text Rephrasing and Editing (70). Text rephrasing and editing is different from updates in instructions or help information, as the latter two entail different or additional information in the actual guidance, procedures, or technical details, while the first one focuses on modification of the wordings and sentence structures.

Regarding the style of text rephrasing and editing, we observed instances where efforts were made to simplify the sentences or render them more succinct. In commit *Example 30*, the phrase *instead of ‘pip install’* was deleted. Since there already exists the correct code command, providing an incorrect command is unnecessary and could potentially lead to users’ misuse. On the other hand, there are also light updates that perform less apparent modifications. For example, in commit *Example 31*, the phrase *doesn’t necessarily work* is revised to *might not work as expected*. The modified text included the behaviour that the code does not work, highlighting the potential limitations.

Self-admitted Documentation Debt Management (30). Although uncommon, there are projects on which developers explicitly leave unfinished marks in certain sections. In commit *Example 32*, an installation section is added, with only a placeholder “coming soon”. Although in most scenarios, these

self-admitted documentation debts are removed once the contents are added, we identified cases in which the documentation debt is only commented out using Markdown syntax. In this case, the reader can still access the complete information, while the documentation maintainers are aware of the outstanding documentation debt from README files.

Summary of document presentation results: Documentation structural modification includes changes in section headers, and organising content underneath them. The concept of content formatting and cleaning involves using Markdown syntax to either highlight or organise the content. Error fixing includes both Markdown syntax fixing and grammar and typography fixing. In addition, the concept of text rephrasing and editing is primarily related to making text simpler and more succinct. Furthermore, we also identified a few cases where self-admitted documentation debt is managed within README files.

F. External Resources Management

External Resources refer to other documentation and non-documentation artefacts that are not embedded within the project's README files but are accessible only via the provided URL links. In total, 366 of 1,168 data points are related to this category. In this category, we synthesised three concepts, which are discussed as follows:

Project External Documentation Management (120). README files serve as a starting point for an OSS project, which contains essential information for users to learn and start on a GitHub repository [17]. However, including all information within the README files would lead to information overload, hindering documentation users from quickly identifying their desired contents. In our analysed data, we discovered one common activity as managing the project's external documentation, taking up around 10% of the whole data instances.

In commit *Example 33*, the content regarding "how-to-use" became overly lengthy and complicated. This commit reorganised the content into a Wiki document for this project, with a URL provided to easily navigate to this external documentation. Meanwhile, it also provided a URL redirect to another Markdown file with contribution information. Generally, we found that the newly introduced documentation has a wider range of formats, including Markdown files, HTML files, Wiki files, RST files, and so on. There are also update operations for the URL of external documents, which can be rooted to change of documentation location or new version reflection. In fact, we observed that the addition of external documentation resources is often accompanied with the deletion of instructional content within the README, which corresponds to documentation organisation efforts.

Project Non-documentation Artefacts Management (136). Apart from external documents, other artefacts (e.g., source code, release package, data, and configuration file) are closely related to installation-related tasks. Typically facilitated by Git, the most frequent modification for the non-documentation artefact is on source code access. This access

is a crucial step to enable the source code installation method. Release packages, such as the jar package for Java projects, are also actively maintained for users to get access to different versions of the release. Additionally, some projects organise other URLs, such as links to configuration files (e.g., Dockerfile), providing quick access to corresponding files, such as the update in commit *Example 34*. This practice provides quick access to editable files, which helps with local development.

Third-Party External Resources Management (110). Although third-party requirements are important for OSS projects, it is not the responsibility of the README file to provide detailed information for them; doing so may lead to issues of tangled information [12]. Therefore, we identified that the majority of data instances in this concept involve adding or updating URLs for third-party dependencies or other requirements of this project, such as the update in commit *Example 35*. These changes often accompany related instructional changes, and the URLs typically lead to their official websites or documentation sources.

Summary of external resources management results: Project external documentation management is related to providing links for other detailed project documents, which helps avoid lengthy README files. Project non-documentation artefacts refer to URLs for project artefacts including source code, executables and other configuration files. Furthermore, the concept of third-party external resources management corresponds to updates for URLs for external requirements of the software project.

G. Cross-Cutting Concerns

There are several perspectives that we identified during the analysis that are scattered across different categories and important to the README file update operations. Therefore, we dedicate a subsection to discuss them.

(1) **Version Management.** As described in previous sections, version management is a notable area where substantial efforts are invested when updating README files. With the development of the project, new versions are released to incorporate additional features, which leads to the active update for version information. In summary, we identified that version management primarily focuses on project versions, dependency versions, and operating system versions.

In terms of the project version, information is provided for the latest stable version and the developing version. This version information is usually updated within installation-related instructions, project running instructions, URLs for accessing project artefacts, and plain text manifest. Additionally, previous versions are sometimes provided, and outdated versions are included as help information. Regarding dependency version updates, the programming language version will be updated when a newer version is incorporated into the software project, e.g., migrating from Python 2 to Python 3 or adding support for the latest version of JVM. As for other dependencies, although version information is not always displayed in the README

TABLE III
QUANTITATIVE EVIDENCE ON UPDATE CATEGORIES AND INFORMATION PRESENTED

Update	Information				
	Pre-Installation	Installation	Post-Installation	External Resources	Help Information
Pre-installation instruction	27	19	14	17	11
Installation-related instruction	15	28	17	18	7
Post-installation instruction	19	26	29	10	10
External resources management	20	26	15	26	9
Help information	16	24	20	15	20
Document presentation	16	30	17	13	7

file for every repository, their versions are also updated accordingly in the case that they are displayed, to reflect the changes within each project. As indicated in the previous example, operating system version updates reflect project changes that introduce a newer compatible version. In addition, when a new compatible operating system is added to the project, it might be necessary to document specific instructions on different software development tasks.

(2) **Development related Instruction Updates.** Development-related instructions require different settings from normal usage mode, and are documented in README files to facilitate better contribution guidelines.

Starting from pre-installation setup, the development mode requires slightly different or additional dependencies compared to the normal usage mode. For example, *pytest* is only useful for testing and is not required for normal use cases. In terms of installation-related processes, separate instructions are usually added for development mode, with an additional flag in the command parameter to handle extra procedures that are needed. Moreover, some projects also provide a snapshot version of the project that is under development for users to experience. Regarding post-installation instructions, development-specific activities are included to complete the development guidelines, which includes testing and processes to contribute to the project.

(3) **Error Fixing.** Errors are in various forms, as summarised above. The impact on documentation users varies depending on the nature of these errors. Document presentation errors are among the least severe, mainly affecting the readability of README files. In addition, presentation errors are relatively easier to solve by applying additional grammar and Markdown syntax checking tools.

However, we also observed efforts to address instructional errors. Instructional errors are scattered across pre-installation, installation-related, and post-installation sections, and their presence can lead to failure of command execution, thus hindering the usability of README files. Errors also exist in URLs that serve as links to external resources. These types of errors are less apparent and more severe, requiring detailed examination or even execution to identify and a higher level of expertise to rectify.

H. README Update Behaviours

Last, we summarise the overall update behaviours based on the derived codes. The most frequently updated category

is “Pre-installation instruction,” with dependency setup being the most modified concept. Documentation presentation is the second most focused category, where maintainers primarily work on modifying the documentation structure and formatting the documents for better readability. Additionally, updates to instructional commands significantly outnumber updates to help information, with nearly five times as many. These observations indicate that OSS documentation maintainers invest more effort in updating dependencies and documentation structures. They also focus more on providing instructional information rather than help information such as troubleshooting in README files.

Additionally, we distinguish our base codes with the operations of “Addition”, “Modification”, and “Removal”, with the statistics displayed in Table II. Unlike the section-level updates shown in Fig. 4, we categorise them based on more detailed modification behaviours. Specifically, whether the modification introduces new information, updates existing information, or removes old information. Seen from the table, apart from document presentation, the largest proportion of update behaviour involves “addition”. This indicates that more information is being incorporated into README files as the software development process progresses, requiring documentation maintainers to spend time assembling new information to reflect system changes.

We also provide quantitative evidence on information already present in the first version of a README when different categories of updates occur. Specifically, we sampled 30 documentation updates from each category and annotated the first version of each README to identify the occurring information. As approximately 60% of the data involves only one type of update category, and to better illustrate the distribution of sections across independent update behaviours, we limited our sampling to data containing only a single update category. Table III shows the frequency of each section present in the previous README version (totalling 30 samples) when each follow-up modification occurs, highlighting the relationship between section completeness and modification patterns.

From the table, it is evident that all types of information generally exist in the initial version of a README file. Notably, in the first version across various update behaviour categories, installation information is the most frequently included content, except in cases where pre-installation instructions are updated. Moreover, for each update behaviour category (except for document presentation), related content was already present in

the initial version of the README. This finding suggests that practitioners often update documentation incrementally rather than introducing entirely new content.

V. TRIGGERS FOR DOCUMENTATION UPDATES

Additionally, we analysed the factors that triggered documentation updates across the entire set of 1,163 documentation updates. Specifically, we followed a series of steps to identify these triggers, stopping once the trigger was identified or reaching the final step. The steps are as follows: (1) read and understand the document update, (2) review the previous README version for obvious issues (e.g., syntax errors, incomplete instructions, presentation issues), (3) examine the commit message to determine whether it is directly linked to a pull request or issue, or if it includes the motivation for the update, (4) if the commit is associated with a pull request/issue, review the first block of the description and initial discussion of that pull request, (5) examine any additional pull requests linked in one level in the discussion thread (non-recursive), (6) check the most recent release for the repository corresponding to this README version, (7) search for modified keywords such as package names and file paths using the GitHub search function, and (8) if none of the above methods yields an answer, we mark the trigger as “no explicit evidence” as we were not able to identify explicit evidence for what triggered the update.

This procedure for identifying triggers was developed based on an initial set of pilot data, discussed and refined until reaching consensus among all authors. The design of this process aims to identify triggers using the most commonly used channels in open-source projects while avoiding spontaneous and endless searching, ensuring that our method is replicable. After that, the first author analysed the rest of the triggers and flagged uncertain entries for further discussion.

Following this procedure, we identified 614 triggers from the entire 1,168 documentation updates, while the rest of the documentation updates are annotated with “no explicit evidence”. We then performed open card sorting [53] to group similar triggers into the same group and assign meaningful names to it. Specifically, we identified 11 different types of triggers. These 11 triggers are further grouped into 3 bigger themes, namely 1) Errors in the previous documentation, 2) Changes in the codebase, and 3) Need for documentation improvement. In the following, we introduce each type of trigger, indicating in parentheses how many instances we identified for each type.

A. Errors in the Previous Documentation

Presentation issues in the previous version of README (93). Documentation maintainers update their documentation when identifying issues within the previous version of a README that impact its presentation. For example, the commit [54] added a white space between the hashtag and the section title, as the previous version of the README violated the markdown syntax for displaying headers.

Incomplete contents in the previous version of README (78). In this case, documentation maintainers left the

documentation clearly unfinished in the previous version. Documentation debt, such as explicit *TODOs* [55] or documentation with only skeletons of section headers [56], are strong indicators of unfinished jobs, suggesting that new content should be added in accordance with the context provided. Meanwhile, when the README document is added to the project at the beginning (i.e., the first version of the README), a lot of content is needed, and our proposed template might be a good starting point.

Instruction or content issues in the previous version of README (48). Instruction or content errors severely impact downstream users’ experience in installing and using the project, and developers should quickly identify the existence of such issues and fix them. For example, commit [57] removes command `--system-site-packages`. The error in this command is reported in issue 1 in this repository by a user and its fix is later reflected in the document.

Outdated elements in previous README (12). Software documentation gets silently outdated during the software development process. When developers realise the existence of such outdated pieces, they make adjustments in the README documents accordingly. For example, commit [58] updates the Python version to 3, while the migration from Python 2 to Python 3 was done in a pull request three years before this documentation update [59], indicating the existence of such an outdated element in the README file.

B. Changes in the Codebase

Update of software version (131). This is the most frequent trigger. When new releases of a software are made, which can usually be observed in the repository release channel or release log, version information in the documentation gets incremented as well, as showcased in this example [60].

Add or modify software functionality (83). This relates to large changes in the codebase, usually accompanied by adding, updating, or deleting multiple files. As a result of such updates, software functionalities such as installation or running instructions might be updated. This requires the documentation authors to carefully account for the changes to be updated in the README. For example, pull request [61] made large updates with various functionalities, and the new OS support is reflected in the README within the commit [62]. Such information is difficult to summarise from the large amount of code changes, while clear pull request descriptions can help developers easily track what should be documented in the README files.

Command updates or file path updates (76). Simple command and file path updates are short code snippet modifications that can be easily tracked to the corresponding files and directories. When such code elements get updated during the development, README changes are also required to keep the documentation up-to-date. For example, commit [63] moves the directory `configs` from the root folder to the `code2seq` folder, thus triggering updates to the file path present in the documentation.

Update of software dependencies (52). Dependencies evolve as the software evolves. Updates to dependencies also

triggers documentation updates. For example, commit [64] updates the dependency Confluent version to 3.1.2, while such an update can be identified within another commit [65] within the same pull request to update the `build.gradle` file. In fact, the dependency update information can usually be found within the build file or requirement file in the same repository.

C. Need for Documentation Improvement

Self-motivated documentation improvement (50). Some projects inspect their README document periodically to evaluate its quality and propose updates to make it easier to understand. For example, a lot of effort was spent in pull request [66] to make the README file better. Incorporating our proposed template could help mitigate the effort spent on this aspect.

Request of documentation updates by downstream users (13). Sometimes downstream users proactively approach the developers, requesting an explanation or enhancement of the documentation. The GitHub Issue channel is suitable for monitoring such concerns raised by users. For example, issue [67] asked for an installation option on Ubuntu and Debian, and the instruction was added later in commit [68].

VI. README TEMPLATE AND ITS VALIDATION

As we discovered that most categories contain more addition operations and that significant efforts are spent on documentation presentation adjustments by practitioners, we created a README template based on the derived taxonomy. This template encompasses all essential installation-related contents, aiming to provide a comprehensive checklist and reduce practitioners' maintenance efforts for README files. In this section, we describe our proposed template and validate it through a survey conducted on Prolific.

A. README Template Content

Fig. 10 presented the suggested template. It covers basic instructions including prerequisites, installation steps, usage, external documents, version history, and help and support sections. These sections are linked to the taxonomy derived from our RQ1, and we have explicitly labelled them in the figure. In the prerequisites section, we incorporate programming language version and dependencies for the project. URLs for the official websites of the dependencies are also provided. The installation steps include both source code level installation and published package level installation. The usage section covers basic commands for running the project, as well as some advanced options to display the capability of the software. External documents provide links to external comprehensive documents. The version history section provides the version information of this project, while the help and support section provides most common FAQs and channels for users to interact with. In cases where specific functionalities are yet to be implemented, these placeholders can be seen as a form of self-admitted document debt, serving as reminders for document maintainers to complete the necessary information.




Fig. 10. Proposed README template.

B. Template Validation Process

To further verify whether our proposed template is perceived as useful by the documentation users, we conducted a survey on Prolific. We applied our proposed template to the original README documents by performing rewrites. To avoid subjectivity in the rewriting process, we instructed ChatGPT to handle this task. Specifically, we provided the original README file and the template, along with two additional instructions: (1) do not omit any details from the original document, and (2) do not add any information not present in the original document. We used ChatGPT 4o throughout this procedure. Only one entry of instruction is provided to ChatGPT to collect the generation, without further tuning or additional feedback.

For the evaluated samples, we first randomly selected 20 repositories. From each repository, we randomly chose one commit to ensure each document was selected only once per repository. We then instructed ChatGPT to apply our proposed template to these documents. We refer to the original documents as “original” documents and the generated documents as “augmented” documents. In total, we obtained 20 “original” to “augmented” document pairs for evaluation.

Table IV shows the number of projects containing the top two levels of headers in both the “original” and “augmented” documentation. For these two most frequent headers, the

TABLE IV
COUNTS FOR HEADER SCALES FOR SAMPLED
REPOSITORIES

Header Scale	$\langle h1 \rangle$	$\langle h2 \rangle$
Original Version Count	18	16
Augmented Version Count	13	15

TABLE V
DIMENSIONS OF SOFTWARE DOCUMENTATION QUALITY

Dimension	Question
Quality	How well-written is this document (e.g., spelling, grammar)
Appeal	How interesting is it?
Readability	How easy was it to read?
Understandability	How easy was it to understand?
Structure	How well-structured is the document?
Cohesion	How well does the text fit together?
Conciseness	How succinct is the information provided?
Effectiveness	Does the document make effective use of technical vocabulary?
Clarity	Does the document contain ambiguity?

majority of original documents feature these top two header levels, indicating that the original README versions are already structured. The augmented version, however, contains slightly fewer top-level headers. As a result, the evaluation is not biased by having more structure in the augmented version, allowing for a focus on the content and organisation of our proposed template.

We conducted a Prolific survey by dividing the 20 document pairs into 10 different surveys, with each survey taken by three different participants. We performed a sequential survey release as in [35] to prevent the same individual from participating multiple times. Meanwhile, we followed a study design similar to Nadi and Treude [69] by inserting a “quality gate” in the survey. After reading the documents, we asked the participants to answer a simple question, e.g., “Is Gradle used in the above-mentioned project?”. We filtered out participants who did not answer the question correctly.

To evaluate how users perceive the documents, we adopted an evaluation framework proposed by Treude et al. [13], which provides a ten-dimensional framework for assessing software documentation quality. Table V displays these dimensions as well as questions in this paper. At the beginning of the survey, we first provided the definitions of these ten dimensions and asked the participants to read through them. After that, they were requested to read through the original and augmented documents and provide scores on these ten dimensions. To prevent biased judgements, the documents were anonymised regarding their status as original or augmented. Likert scale is used to mark each of the aspects, with a score of 1 for strongly disagreeing and a score of 5 for strongly agreeing. Additionally, we revealed the template at the end of the survey and with an open-ended question aiming to gather participants’ opinions about it. We added a sample as well as all README file pairs in our replication package.

C. Demographics of Annotators

We first applied predefined filters on Prolific, requiring participants to be in the Information Technology (IT) employment

sector and be proficient in English. We asked the same question about their employment again at the beginning of our survey and filtered out 18 participants who did not answer this question consistently with their Prolific registration. Our “quality gate” question further filtered out four participants.

In the survey, we asked participants about their job roles and how many years they have worked in the IT field. For the 30 participants who passed all filters and submitted their responses, they have on average 6.9 years of experience working in IT, with a maximum of twenty years and a minimum of one year. Fig. 12 shows the demographic distribution, with the largest proportion of participants being “developers”, and the legend sorted in descending order based on the number of participants. We computed the inter-rater agreement among the evaluators, which reached 0.34 for Krippendorff’s alpha [70] coefficient overall. This indicates that the evaluators reached moderate agreement, while the evaluation remains somewhat subjective.

D. Template Validation Result

Seen from Fig. 11, the average scores for the “augmented” document are consistently higher than those of the “original” documents. The figure is sorted based on the difference between the average scores of these two types of documents. The most improved dimensions are structure, understandability and conciseness, with a gain of 0.80, 0.53, and 0.53 of the Likert average score, respectively. Meanwhile, the increases in clarity and consistency are less noticeable.

We also performed Wilcoxon signed-rank tests [71] on all dimensions to determine whether the differences are statistically significant. As indicated in the figure, eight out of ten dimensions are with a p-value below 0.05, indicating statistically significant differences. We further computed the Cohen’s d effect size [47] for the ten evaluation dimensions. The results indicate that the adoption of our proposed template significantly enhances the perceived structure with a large effect, while significantly enhances understandability, conciseness, readability, quality, cohesion, appeal, and effectiveness with medium effect. The improvements in clarity and consistency are small.

Therefore, based on the quantitative results of our survey, we find that applying our proposed template increases the quality of the generated documents, particularly in terms of structure, understandability, and conciseness, without compromising clarity and consistency, as perceived by document readers.

E. Practitioners’ Perspectives From Pull Requests

To further assess the proposed README template and methodology, we conducted an additional study to gather developers’ opinions. Specifically, we applied the proposed template to the latest version of README files and submitted pull requests to the respective repositories, requesting the suggested changes. Out of the 400 repositories initially considered, we excluded those without any commits in the past year, narrowing the sample to 219 active repositories. From these, we randomly selected 30 repositories to conduct the experiment.

In the pull request, we included a commit with the updated README file that integrates our proposed template. Along with this, we provided a description of the potential benefits




Fig. 11. Survey result.




Fig. 12. Job roles of annotators.

of adopting the template and included a screenshot of the raw template for reference. At the time of writing, among eight repositories that responded and were comfortable with LLM-generated improvements to their README files, six projects responded positively, while two projects responded negatively. Four repositories expressed reluctance to use LLM-generated content for documentation or questioned the purpose of unsolicited pull requests.

Positive feedback. Fig. 13 shows the merged pull request to project `aerleon/aerleon`.⁶ Similarly, project `spring/uberserver` and `actinia-org/actinia-core` also directly merged our proposed change without further comment.⁷ Project `apache/openmeetings` responded positively with “I really like the new look of the readme” and suggested some customised changes.⁸ The pull request was then merged into the codebase after we made the suggested changes. Additionally, project `actris-cloudnet/cloudnetpy` adopted part of our proposed changes while leaving other parts

⁶<https://github.com/aerleon/aerleon/pull/378>

⁷<https://github.com/spring/uberserver/pull/400>

⁸<https://github.com/apache/openmeetings/pull/196>




Fig. 13. Example of merged pull request.

of their README unchanged.⁹ Specifically, they adjusted section header hierarchies to improve consistency. For project `vespa-engine/pyvespa`,¹⁰ the maintainer commented, “I think we can add some of the suggestions, but some also are not correct/ helpful.”, and further suggested that “do a human review before creating a PR”. This feedback highlights the potential for hallucinations in ChatGPT-generated suggestions and underscores the importance of human validation to ensure correctness when updating documentation.

Negative feedback. Out of the two negative responses, one reason for rejection is the lack of project-specific details in the updated version. For example, project `3dcitydb/importer-exporter` mentioned that “Some

⁹<https://github.com/actris-cloudnet/cloudnetpy/pull/110>

¹⁰<https://github.com/vespa-engine/pyvespa/pull/976>

sections that are important for us, like the partners and supporters and additional information about the technical standards behind the project are missing”.¹¹ Additionally, both repositories viewed the proposed updates as offering insignificant improvement.

Therefore, we see that our proposed method can be adopted by open-source repositories and is endorsed by several merges and positive comments from real-life projects. However, the practitioners also raised critical points regarding the completeness and correctness of the generated documentation. Semi-automated documentation updates with humans in the loop might be a solution for such concerns while reducing manual efforts to a large extent.

It is important to note that developing a directly applicable pipeline for applying the template to README files is not the primary focus of this paper. Instead, the validations obtained through user feedback and this pull request study serve as a proof of concept, demonstrating that the template has the potential to make real-world impact. Future research should focus on further refining the template and developing a comprehensive methodology for its application, ensuring it can be seamlessly integrated into real-world production workflows.

VII. IMPLICATIONS AND DISCUSSION

The ultimate goal of our study is to understand the practices of updating README files in the parts related to the installation, enabling better documentation activities to improve the quality of software documentation and its process. In this section, we discuss the findings of this study and their implications for different stakeholders, including researchers and practitioners.

Completeness considerations. The importance of documentation content completeness issue has been identified in previous work [6], and efforts have been made to use recommendation systems to recommend newly added code elements to be documented [72]. However, more factors constitute the completeness factor in the context of README files. We provide a more detailed discussion of completeness considerations in the context of GitHub README files.

For documentation maintainers, instruction completeness is important. The categories that emerged from our qualitative analysis of pre-installation instructions, installation-related instructions, and post-installation instructions collectively account for a complete process, enabling self-contained documentation for users to follow. In our synthesised codes, the introduction of a new instruction without a pre-existing alternative implies a deficiency in the previous documentation. For example, we found that the programming language and operating system are frequently newly added to the document. This type of update addresses previously absent details. On the other hand, practitioners should also take into account the completeness of help information. Well-managed help information can improve the comprehension of the documentation and reduce the search time for users in case of errors [73]. Based

on our analysed data, it is evident that the frequency of help information is notably lower compared to the combined occurrence of instructions within the three major categories. This observation implies a tendency among practitioners to overlook the inclusion of comprehensive help information. Therefore, practitioners could closely monitor the issues and pull requests to incorporate frequently-asked or confused questions into the README documents. For documentation updates, they could use our proposed template as a checklist before performing the updates.

For researchers, we suggest several directions for investigation to improve the completeness of the README files. First, this work primarily focuses on the update operation with less attention on the state of README file before and after modification. Therefore, identifying whether the updates are complete in reflecting the system change and fixing the previous document's lack of information would help better characterise a “good” or “bad” update. Second, an investigation of the temporal aspect of the absence of various types of instructions within the README file could better quantify the severity of this issue. Furthermore, researchers can also assist in providing essential help information. Automatically supplying help information is a difficult task. However, this can be done by algorithms to detect insufficient information for each instruction. Moreover, previous efforts [33] augmented API documentation with information from Stack Overflow. Similar ideas could be adopted by automatically extracting the most requested help information from various online channels based on certain project characteristics. Lastly, the development of tools to detect completeness issues could help document maintainers identify and address these concerns quickly.

GitHub could also help mitigate this issue by providing the template proposed in this study for installation-related instructions when users create or update the README files. According to the result discussed in the previous section, by simply applying the template to the original document, documentation readers could experience increased documentation quality. In fact, the survey participants left comments on our proposed template: “Good to see that you provide Frequently Asked Questions (FAQ)”, “Covers all the important elements”. Therefore, this template could serve as a good starting point for reminding practitioners of the essential aspects to cover for their README files before performing updates.

Correctness and up-to-dateness considerations. Issues related to correctness and up-to-dateness have been explored in prior studies [12], [14]. Tan et al. [74] introduced a method to identify outdated code elements in README files. In our discussion, we offer a more in-depth exploration, encompassing considerations of correctness within the context of README information and providing a more comprehensive list of up-to-dateness considerations other than code elements.

As indicated in the previous section, errors occur across different categories. Less severe errors encompass grammar issues, typographical errors within the text, and Markdown syntax errors. More severe errors are related to instructions, which require a certain level of expertise to identify. Although not explicitly coded, we identified a few updates fixing errors

¹¹<https://github.com/3dcitydb/importer-exporter/pull/312>

within their instructions. In combination, these two types of errors constitute a noticeable percentage of update behaviours. This type of documentation effort can be avoided by checking the content more carefully.

Concerning considerations of up-to-dateness, delays in updating the README file can result in a temporary misalignment and potentially introduce correctness issues in the document. Codes related to updating various instructions reflect the changes in the software system. As also discovered in [74], code elements are indeed updated frequently in our analysed data. Specifically, the frequent pattern found for code *update code demo for using this project* corresponds to modifying code elements to reflect naming paradigm shifts within the source code. Additionally, change factors regarding metadata (project names, file path names, and technology used) of the project, modification of dependencies for the project, and instructions with specific operating systems, all influence the up-to-dateness of the README file and require close monitoring by document maintainers.

Information presentation considerations. Previous studies investigated how information is presented within the software documentation [6]. Considering the role of README files, which serve as the introductory point for users to participate in OSS projects, we discuss their distinctive characteristics and implications.

As indicated in Table II, document presentation undergoes frequent updates, ranking as the second most frequently updated category in our analysis. This observation underscores the considerable effort practitioners invest in updating this particular aspect of the documentation. However, most of the updates are repetitive and could be avoided by adopting better documentation schemes. Within this category, the code *change section header name* has the highest frequency of 88 occurrences, followed by codes that adjust the hierarchies of the sections and move the contents from one section to another. Therefore, practitioners can adopt more standardised section headers and examine whether there are mismatches between contents and headers before committing the change. For researchers, natural language processing techniques can be applied to examine content-header similarity or summarise the section contents with more appropriate headers [75]. In addition, techniques can be applied to examine and adjust the hierarchical structure for the README sections. Furthermore, this observation also motivates us to propose the README installation-related templates in Fig. 10 to mitigate practitioners' efforts in document structure modifications.

Text rephrasing is also frequently observed within README updates. Two codes dedicated to adjusting the wording within the text or making sentences succinct account for 53 occurrences in total. For document maintainers, we suggest that they consider the audience of README files before writing the document. Specifically, as README files target newcomers and serve as an introductory document, overly complicated sentences and terminologies should be avoided. For researchers, previous research [35], [76] attempted to identify documentation smells or simplify complicated sentences within the documentation. With the advent of ChatGPT [77], text rephrasing

becomes easier, and attention is drawn to how different groups of users perceive useful information.

README files serve as an introductory resource for navigating the OSS projects and should not be regarded as exhaustive documentation. In our analysed data, we observed instances where the concept of project running instructions involved presenting a comprehensive list of all running options with varying parameter combinations. This approach could lead to information overload, making it challenging for documentation users to locate their desired content. In addition, we also identified a pattern of content removal in combination with external document addition, when content within the README files becomes excessively long. For example, there are a few cases where instructions are removed and a link to `INSTALL.md` or `CONTRIBUTE.md` is supplied in the README files.

For practitioners, the above-mentioned external content organisation approach can be adopted to avoid information overload. Furthermore, we identified in code *add code instruction for running this project* that add a comprehensive list of project running instructions, making it difficult for users to identify the desired content. In this scenario, these instructions can be documented in a `RUNNING.md` with README providing a reference to it, or covered by a command line interface with parameter `--help`.

For researchers, detecting information overload helps determine when a documentation refactor of moving content outside the README file is needed. A survey to documentation users on the appropriate amount of information can help better understand the users' needs, while automated tools can also be developed for overload detection.

Automatic documentation updates. By monitoring common open-source channels such as commit messages, pull requests, issues, and previous versions of the documentation, we identified various triggers for documentation updates. An automatic tool for suggesting documentation updates by taking this information into consideration could help developers quickly identify the need for documentation updates during the software development process. Future research could also investigate more from detailed source code level to identify additional triggers for documentation updates.

Additionally, the results in the template validation section indicate that documentation users perceive improvements in four key dimensions – structure, understandability, conciseness, and readability – after applying our template to the original README documents. In short, the application of this template directly mitigates the above-mentioned issues related to documentation structure, targeted audience, and information overload. In particular, one participant commented on the templates *“I like the order of chapters and it is easy to find everything. Just enough information to get started with technology.”*. In addition, our pull request study further shows that applying this template is also perceived as useful by documentation maintainers.

Therefore, to enhance the downstream users' experience, documentation maintainers could manually adopt our template, or use an automatic approach of delegating this task to LLMs. The automatic approach could further alleviate the massive

efforts from practitioners to adjust documentation presentation. However, this process should incorporate a human-in-the-loop approach, involving manual review and adjustments to account for specific details, as a general template may not be suitable for all scenarios.

For researchers, our approach did not extensively adopt prompt engineering or make adjustments to the generated documentation. Refinements in these aspects might further improve the quality of README documents. The usability of the augmented documents could be verified through live installation sessions with documentation users.

Furthermore, integrating documentation update suggestion tools with our proposed template has the potential to automate documentation updates, addressing issues related to up-to-dateness, correctness, and presentation simultaneously. This pipeline could be seamlessly integrated into the GitHub pull request process, where documentation suggestions are generated based on changes in relevant channels and the status of the previous version of the README. Future research could focus on refining these tools or enhancing the template to realise this vision.

VIII. THREATS TO VALIDITY

We consider threats to the validity of our study.

First, we only collect repositories that use Python or Java as their primary programming language, which could affect the generalisability of our findings to other software projects. However, we did not observe any significant differences in the taxonomy between Python and Java projects in our analyzed data. The taxonomy derived in this study includes modifications in the installation process, help information, documentation presentation, and resource management. These factors can serve as a starting point for researchers and practitioners when addressing projects with different programming languages.

Second, as indicated earlier, our semi-automatic labelling heuristic is not completely accurate in distinguishing “relevant” and “irrelevant” commits. In the sampled dataset, we discarded the ones that were misclassified during our manual analysis, and filled each sample bucket with additional randomly sampled commits to ensure the same amount of data is analysed.

Third, the 400 analysed repositories may not be representative of the overall population, potentially introducing sample bias. To address this threat, we ensured that the sampled dataset size exceeded a 95% confidence rate with a 5% margin of error for the overall dataset. Furthermore, our qualitative analysis may inherently carry subjectivity. To mitigate this concern, a rigorous process was implemented, involving all authors in reviewing the codes. Weekly meetings were held to facilitate discussions and address any discrepancies in codes, concepts and categories.

Fourth, this study is conducted on the GitHub platform, and its generalisability is expected to extend to other open-source software (OSS) projects that follow similar processes. However, considering the different procedures and use cases for proprietary and private software projects, it is important to note that we do not assert generalisability beyond the domain of open source.

Fifth, our analysis is limited in README update contents, without examining external referenced links. The results may also only apply to GitHub and are not easily generalisable beyond GitHub projects to other platforms and closed-source projects.

Sixth, our procedure for identifying triggers for documentation updates may miss certain triggers, such as changes embedded in commits outside of pull requests. Additional reasons for updates might emerge from a deeper analysis of the codebase. However, the procedure was designed to monitor the primary channels within open-source repositories. This approach provides a systematic way to identify triggers without requiring exhaustive manual searches. Moreover, this design offers a foundation for automating the linkage between triggers and documentation updates, potentially enabling the identification of more triggers in the future.

Seventh, our template validation was conducted on 20 README documents with 30 participants, which might not represent the entire range of README documentation genres or the broader participant population. To mitigate this limitation, we randomly sampled the evaluated documents and recruited participants who work in IT with an additional “quality gate”. However, the participants were not actual README users; they were only required to evaluate the perceived documentation quality without installing the software according to the guidelines provided in the documents. Additionally, no formal training of participants except the definition is provided before the survey, leading to a relatively lower agreement with some level of subjectivity. Therefore, we only claim an improvement in the documentation reading experience rather than in documentation accuracy. Future work can further investigate this perspective by involving actual README users with live installation sessions.

Lastly, the pull request validation study only covers 30 repositories and may not generalise to all repositories on GitHub. Moreover, this investigation does not provide long-term data about whether the merged changes will remain in the repository. Furthermore, some repositories had not responded to our pull requests at the time of paper submission, preventing us from including those results in the current version of the manuscript.

IX. CONCLUSION

README files serve a critical role of introductory overviews for software projects, among which the installation-related sections are included. Considering various issues these documents may encounter brought by the rapidly evolving software repositories and ecosystem, as well as the benefits of investing efforts to maintain high-quality documentation, knowledge of the practices of the README file efforts becomes essential. To achieve this, we conducted a large-scale qualitative analysis on GitHub README file commit updates and their associated triggers specifically related to installation sections, aiming to gain comprehensive insight into documentation activities.

In this paper, we collected GitHub repositories with several stringent filters to obtain a dataset consisting of 12,908 repositories. From this dataset, we randomly sampled 400 repositories, resulting in 25,209 commits that modified a README file.

We applied semi-automatic labelling heuristics and eventually obtained 8,673 commits that modified sections related to installation topics. Among these, 1,168 instances were cluster sampled for qualitative analysis, resulting in a comprehensive taxonomy of 189 update behaviours.

Our research reveals six categories of updates, namely pre-installation instructions, installation instructions, post-installation instructions, help information updates, document presentation, and external resource management. We further provide detailed insights into modification behaviours and offer examples of these updates. We also investigated the triggers for documentation updates, which led to three categories including errors in the previous documentation, changes in the codebase, and need for documentation improvement.


Based on our findings, we proposed a README template for documentation maintainers when updating documents. The template validation process finds that the augmented documents based on our template are perceived as higher quality by documentation readers. Moreover, our pull request study also shows that documentation maintainers find it useful. This step serves as an initial step towards automated documentation updates to facilitate better documentation activities. Furthermore, we provided suggestions for practitioners and potential directions for researchers from perspectives of completeness, correctness and up-to-dateness, and information presentation considerations. To sum up in an abstract way, documentation maintainers should take care of completeness and correctness for instructions in a succinct manner, while incorporating help information and external resources to make README files more accessible and easier to comprehend. Document presentation is also a critical aspect to consider. Meanwhile, future research directions include identifying triggers for README updates, understanding user information needs in documentation, and developing automated documentation tools. Knowledge of triggering events could reveal the relationship between document and code activities, user information needs help understand the appropriate instructions to include, while all these directions contribute to the eventual automated documentation tools.

REFERENCES

- [1] G. A. A. Prana, C. Treude, F. Thung, T. Atapattu, and D. Lo, "Categorizing the content of GitHub readme files," *Empirical Softw. Eng.*, vol. 24, pp. 1296–1327, 2019.
- [2] M. Koskela, I. Simola, and K. Stefanidis, "Open source software recommendations using GitHub," in *Digit. Libraries Open Knowl.: 22nd Int. Conf. Theory Pract. Digit. Libraries, TPD, Porto, Portugal*. Cham, Switzerland: Springer, 2018, pp. 279–285.
- [3] I. Steinmacher, A. P. Chaves, T. U. Conte, and M. A. Gerosa, "Preliminary empirical identification of barriers faced by newcomers to open source software projects," in *Proc. Brazilian Symp. Softw. Eng.*, Piscataway, NJ, USA: IEEE Press, 2014, pp. 51–60.
- [4] V. Boucharas, S. Jansen, and S. Brinkkemper, "Formalizing software ecosystem modeling," in *Proc. 1st Int. Workshop Open Compon. Ecosyst.*, 2009, pp. 41–50.
- [5] S. Mirhosseini and C. Parnin, "Docable: Evaluating the executability of software tutorials," in *Proc. 28th ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2020, pp. 375–385.
- [6] E. Aghajani et al., "Software documentation issues unveiled," in *Proc. IEEE/ACM 41st Int. Conf. Softw. Eng. (ICSE)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 1199–1210.
- [7] E. Aghajani et al., "Software documentation: The practitioners' perspective," in *Proc. ACM/IEEE 42nd Int. Conf. Softw. Eng.*, 2020, pp. 590–601.
- [8] I. Škorić, B. Pein, and T. Orehovački, "Selecting the most appropriate web ide for learning programming using AHP," in *Proc. 39th Int. Conf. Inf. Commun. Technol., Electron. Microelectron. (MIPRO)*, Piscataway, NJ, USA: IEEE Press, 2016, pp. 877–882.
- [9] K. Aggarwal, A. Hindle, and E. Stroulia, "Co-evolution of project documentation and popularity within GitHub," in *Proc. 11th Work. Conf. Mining Softw. Repositories*, 2014, pp. 360–363.
- [10] M. Kajko-Mattsson, "A survey of documentation practice within corrective maintenance," *Empirical Softw. Eng.*, vol. 10, pp. 31–55, Jan. 2005.
- [11] G. Garousi, V. Garousi-Yusifoglu, G. Ruhe, J. Zhi, M. Moussavi, and B. Smith, "Usage and usefulness of technical software documentation: An industrial case study," *Inf. Softw. Technol.*, vol. 57, no. 0950–5849, pp. 664–682, 2015.
- [12] G. Uddin and M. P. Robillard, "How API documentation fails," *IEEE Softw.*, vol. 32, no. 4, pp. 68–75, Jul./Aug. 2015.
- [13] C. Treude, J. Middleton, and T. Atapattu, "Beyond accuracy: Assessing software documentation quality," in *Proc. 28th ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2020, pp. 1509–1512.
- [14] F. Wen, C. Nagy, G. Bavota, and M. Lanza, "A large-scale empirical study on code-comment inconsistencies," in *Proc. IEEE/ACM 27th Int. Conf. Program Comprehension (ICPC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 53–64.
- [15] H. Gao, M. Zahedi, C. Treude, S. Rosenstock, and M. Cheong, "Documenting ethical considerations in open source AI models," in *Proc. 18th ACM/IEEE Int. Symp. Empirical Softw. Eng. Meas.*, 2024, pp. 177–188.
- [16] Y. Liu, E. Noei, and K. Lyons, "How readme files are structured in open source Java projects," *Inf. Softw. Technol.*, vol. 148, 2022, Art. no. 106924.
- [17] T. Wang, S. Wang, and T.-H. P. Chen, "Study the correlation between the readme file of GitHub projects and their popularity," *J. Syst. Softw.*, vol. 205, 2023, Art. no. 111806.
- [18] J. Fowkes, P. Chanthirasegaran, R. Ranca, M. Allamanis, M. Lapata, and C. Sutton, "Autofolding for source code summarization," *IEEE Trans. Softw. Eng.*, vol. 43, no. 12, pp. 1095–1109, Dec. 2017.
- [19] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation," in *Proc. 26th Conf. Program Comprehension*, 2018, pp. 200–210.
- [20] X. Hu, G. Li, X. Xia, D. Lo, and Z. Jin, "Deep code comment generation with hybrid lexical and syntactical information," *Empirical Softw. Eng.*, vol. 25, no. 1382–3256, pp. 2179–2217, 2020.
- [21] W. Wang et al., "Reinforcement-learning-guided source code summarization using hierarchical attention," *IEEE Trans. Softw. Eng.*, vol. 48, no. 1, pp. 102–119, Jan. 2022.
- [22] S. Xu, Y. Yao, F. Xu, T. Gu, H. Tong, and J. Lu, "Commit message generation for source code changes," in *Proc. IJCAI*, 2019, pp. 3975–3981.
- [23] S. Liu, C. Gao, S. Chen, L. Y. Nie, and Y. Liu, "ATOM: Commit message generation based on abstract syntax tree and hybrid ranking," *IEEE Trans. Softw. Eng.*, vol. 48, no. 5, pp. 1800–1817, May 2022.
- [24] J. Dong et al., "Fira: Fine-grained graph-based code change representation for automated commit message generation," in *Proc. 44th Int. Conf. Softw. Eng.*, 2022, pp. 970–981.
- [25] S.-T. Shi, M. Li, D. Lo, F. Thung, and X. Huo, "Automatic code review by learning the revision of source code," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 1, 2019, pp. 4910–4917.
- [26] H. Y. Lin, P. Thongtanunam, C. Treude, and W. Charoenwet, "Improving automated code reviews: Learning from experience," in *Proc. IEEE/ACM 21st Int. Conf. Mining Softw. Repositories (MSR)*, Piscataway, NJ, USA: IEEE Press, 2024, pp. 278–283.
- [27] L. Moreno, G. Bavota, M. Di Penta, R. Oliveto, A. Marcus, and G. Canfora, "Automatic generation of release notes," in *Proc. 22nd ACM SIGSOFT Int. Symp. Found. Softw. Eng.*, 2014, pp. 484–495.
- [28] L. Moreno, G. Bavota, M. Di Penta, R. Oliveto, A. Marcus, and G. Canfora, "ARENA: An approach for the automated generation of release notes," *IEEE Trans. Softw. Eng.*, vol. 43, no. 2, pp. 106–127, Feb. 2017.
- [29] S. Panthaplackel, J. J. Li, M. Gligoric, and R. J. Mooney, "Deep just-in-time inconsistency detection between comments and source code," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, no. 1, 2021, pp. 427–435.

- [30] Z. Liu, X. Xia, D. Lo, M. Yan, and S. Li, "Just-in-time obsolete comment detection and update," *IEEE Trans. Softw. Eng.*, vol. 49, no. 1, pp. 1–23, Jan. 2023.
- [31] S. Wattanakriengkrai et al., "Automatic classifying self-admitted technical debt using N-GRAM IDF," in *Proc. 26th Asia-Pacific Softw. Eng. Conf. (APSEC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 316–322.
- [32] Y. Li, M. Soliman, P. Avgeriou, and L. Somers, "Self-admitted technical debt in the embedded systems industry: An exploratory case study," *IEEE Trans. Softw. Eng.*, vol. 49, no. 4, pp. 2545–2565, Apr. 2023.
- [33] C. Treude and M. P. Robillard, "Augmenting api documentation with insights from stack overflow," in *Proc. 38th Int. Conf. Softw. Eng.*, 2016, pp. 392–403.
- [34] C. Treude, M. P. Robillard, and B. Dagenais, "Extracting development tasks to navigate software documentation," *IEEE Trans. Softw. Eng.*, vol. 41, no. 6, pp. 565–581, Jun. 2015.
- [35] H. Gao, C. Treude, and M. Zahedi, "Evaluating transfer learning for simplifying GitHub readmes," 2023, *arXiv:2308.09940*.
- [36] M. P. Robillard et al., "On-demand developer documentation," in *Proc. IEEE Int. Conf. Softw. Maintenance Evol. (ICSME)*, Piscataway, NJ, USA: IEEE Press, 2017, pp. 479–483.
- [37] D. J. Barrera, "Securing decentralized software installation and updates," Ph.D. dissertation, Carleton University, 2014.
- [38] I. Steinmacher, M. A. G. Silva, M. A. Gerosa, and D. F. Redmiles, "A systematic literature review on the barriers faced by newcomers to open source software projects," *Inf. Softw. Technol.*, vol. 59, no. 0950–5849, pp. 67–85, 2015.
- [39] L. Salerno, S. de França Tonhão, I. Steinmacher, and C. Treude, "Barriers and self-efficacy: A large-scale study on the impact of oss courses on student perceptions," in *Proc. Conf. Innovation Technol. Comput. Sci. Educ. V. 1*, 2023, pp. 320–326.
- [40] R. Malaquias et al., "The discipline of preprocessor-based annotations-does# ifdef tag n't# endif matter," in *Proc. IEEE/ACM 25th Int. Conf. Program Comprehension (ICPC)*, Piscataway, NJ, USA: IEEE Press, 2017, pp. 297–307.
- [41] M. Zahedi, R. N. Rajapakse, and M. A. Babar, "Mining questions asked about continuous software engineering: A case study of stack overflow," in *Proc. Eval. Assessment Softw. Eng.*, 2020, pp. 41–50.
- [42] G. A. Miller, "WordNet: A lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [43] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian, "The promises and perils of mining GitHub," in *Proc. 11th Work. Conf. Mining Softw. Repositories*, 2014, pp. 92–101.
- [44] O. Dabic, E. Aghajani, and G. Bavota, "Sampling projects in GitHub for msr studies," in *Proc. IEEE/ACM 18th Int. Conf. Mining Softw. Repositories (MSR)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 560–564.
- [45] A. Joulin, E. Grave, P. Bojanowski, M. Douze, H. Jégou, and T. Mikolov, "Fasttext.zip: Compressing text classification models," 2016, *arXiv:1612.03651*.
- [46] Y. Chen et al., "When large language models confront repository-level automatic program repair: How well they done?" in *Proc. IEEE/ACM 46th Int. Conf. Softw. Eng.: Companion Proc.*, 2024, pp. 459–471.
- [47] L. A. Becker, "Effect size (ES)," 2000. Accessed: Dec. 10, 2024. [Online]. Available: <https://www.uv.es/friasnav/EffectSizeBecker.pdf>
- [48] S. Baltes and P. Ralph, "Sampling in software engineering research: A critical review and guidelines," *Empirical Softw. Eng.*, vol. 27, no. 4, 2022, Art. no. 94.
- [49] B. Glaser and A. Strauss, *Discovery of Grounded Theory: Strategies for Qualitative Research*. London, U.K.: Routledge, 2017.
- [50] R. W. Service, "Book review: Corbin, j., & Strauss, a.(2008). basics of qualitative research: Techniques and procedures for developing grounded theory. thousand oaks, ca: Sage," *Organizational Res. Methods*, vol. 12, no. 3, pp. 614–617, 2009.
- [51] C. Bogart, C. Kästner, J. Herbsleb, and F. Thung, "When and how to make breaking changes: Policies and practices in 18 open source software ecosystems," *ACM Trans. Softw. Eng. Methodol.*, vol. 30, no. 4, pp. 1–56, 2021.
- [52] A. Bratko and B. Filipič, "Exploiting structural information for semi-structured document categorization," *Inf. Process. & Manage.*, vol. 42, no. 3, pp. 679–694, 2006.
- [53] T. Zimmermann, "Card-sorting: From text to themes," in *Perspectives Data Sci. Softw. Eng.*, Amsterdam, Netherlands: Elsevier, 2016, pp. 137–141.
- [54] "Missing whitespaces. actris-cloudnet/cloudnetpy@1ecb773 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/actris-cloudnet/cloudnetpy/commit/1ecb7736496115721ecc44eb668b05862433f05b>
- [55] "updating readme with more running and installation instructions. ncasuk/forest@fce94f8 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/ncasuk/forest/commit/fce94f8dd30dda6763dfb80222e1da62b7e62ff8>
- [56] "added documentation for the musify autocomple api HeIlow2/music-downloader@4a042e4 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/HeIlow2/music-downloader/commit/4a042e40f0885df829486db7dcedfc0b1749cea8#diff-b335630551682c19a781afebfc4d07bf978fb1f8ac04c6bf87428ed5106870f5>
- [57] "Update README.md JohnHammond/katana@cc77886 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/JohnHammond/katana/commit/cc778862ad89f4f9665f5bdcacfc82d4c81741f8>
- [58] "CHORE: Update README.md. It was about time. nottinghamtec/PyRIGS@1480ae1 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/nottinghamtec/PyRIGS/commit/1480ae17faea25c3e2356fecfa5b51c9b8f2ab6>
- [59] "Move to Python Issue # 287 nottinghamtec/PyRIGS — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/nottinghamtec/PyRIGS/issues/287>
- [60] "bumping version to 5.1.1 3dcitydb/importer-exporter@0424b98 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/3dcitydb/importer-exporter/commit/0424b981661683bbaf05d5cc73900d16d669c1307>
- [61] "Dev v1.6.1 by kellyjonbrazil Pull Request #17 kellyjonbrazil/jc — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/kellyjonbrazil/jc/pull/17>
- [62] "doc update kellyjonbrazil/jc@5f4c10f — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/kellyjonbrazil/jc/commit/5f4c10ffd5a9cbe4e5e3d2b9adc34f36492eaf46>
- [63] "Update config structure JetBrains-Research/code2seq@b98dac5 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/JetBrains-Research/code2seq/commit/b98dac5e6e4f876a3da0feafe99931af445735ad>
- [64] "Update README.md to Confluent 3.1.2 RADAR-base/RADAR-Backend@43b1497 — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/RADAR-base/RADAR-Backend/commit/43b14978b981f3316ed74223483db6c0601692d6>
- [65] "Bring master upto date with dev by yatharthranjan Pull Request #34 RADAR-base/RADAR-Backend — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/RADAR-base/RADAR-Backend/pull/34/commits/43bd9d56e4aff687b034be52e6804641a3350b86>
- [66] "Additions and Deletions to the README, to make for better documentation by corletim Pull Request #161 GatorEducator/gatorgrader — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/GatorEducator/gatorgrader/pull/161>
- [67] "How to install on ubuntu? Issue #102 stored-project/blivet-gui — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/stored-project/blivet-gui/issues/102>
- [68] "README: Add information about Debian and Ubuntu repository. stored-project/blivet-gui@94a2f8f — github.com," Accessed: Nov. 25, 2024. [Online]. Available: <https://github.com/stored-project/blivet-gui/commit/94a2f8f2fe8709414f0d5d747c1ef66e35dc68ff>
- [69] S. Nadi and C. Treude, "Essential sentences for navigating stack overflow answers," in *Proc. IEEE 27th Int. Conf. Softw. Anal., Evol. Reeng. (SANER)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 229–239.
- [70] K. Krippendorff, "Reliability in content analysis: Some common misconceptions and recommendations," *Human Commun. Res.*, vol. 30, no. 3, pp. 411–433, 2004.
- [71] R. F. Woolson, "Wilcoxon signed-rank test," *Encyclopedia Biostatistics*, vol. 8, pp. 1–2, Jul. 2005.
- [72] B. Dagenais and M. P. Robillard, "Using traceability links to recommend adaptive changes for documentation evolution," *IEEE Trans. Softw. Eng.*, vol. 40, no. 11, pp. 1126–1146, Nov. 2014.
- [73] C. Chen and K. Zhang, "Who asked what: Integrating crowdsourced FAQs into API documentation," in *Proc. Companion Proc. 36th Int. Conf. Softw. Eng.*, 2014, pp. 456–459.
- [74] W. S. Tan, M. Wagner, and C. Treude, "Detecting outdated code element references in software repository documentation," 2022, *arXiv:2212.01479*.

- [75] Y. Liu and M. Lapata, "Text summarization with pretrained encoders," in *Proc. Conf. Empirical Methods Natural Lang. Process. 9th Int. Joint Conf. Natural Lang. Process. (EMNLP-IJCNLP)*, K. Inui, J. Jiang, V. Ng, and X. Wan, Eds. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3730–3740. [Online]. Available: <https://aclanthology.org/D19-1387>
- [76] J. Y. Khan, M. T. I. Khondaker, G. Uddin, and A. Iqbal, "Automatic detection of five API documentation smells: Practitioners' perspectives," in *Proc. IEEE Int. Conf. Softw. Anal., Evol. Reeng. (SANER)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 318–329.
- [77] J. Achiam et al., "OpenAI," Gpt-4 Tech. Rep., 2023.



Haoyu Gao received the bachelor's degree in mathematics and applied mathematics from Fuzhou University, China, and the master's degree in information technology from the University of Melbourne. He is currently working toward the Ph.D. degree with the School of Computing and Information Systems, The University of Melbourne, Australia. His research focuses on supporting developers in documentation maintenance and knowledge acquisition in both traditional and AI-based software systems.



Christoph Treude (Member, IEEE) is an Associate Professor with Singapore Management University (SMU). Before joining SMU, Christoph worked in Australia, Brazil, Canada, and Germany. He serves on the editorial boards of journals including *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING*, *EMSE*, and *JSS*, and will be a PC Co-Chair for FSE 2026. His research goal - getting information to developers when and where they need it - is gaining new significance in the AI era.



Mansooreh Zahedi (Member, IEEE) is an Assistant Professor (aka Lecturer) in software engineering (SE) with the School of Computing and Information Systems (CIS), The University of Melbourne. Her primary research goal is to apply empirical research methods and tools to investigate the role of people, processes, and technologies in different software development paradigms. Her key research interests include human aspects in software engineering, socio-technical aspects of cyber security, and intersection of software engineering and AI. She has conducted extensive field studies with different companies internationally and published empirically grounded findings. Her work has been published in several high-ranking software engineering venues e.g., *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING*, *TOSEM*, *EMSE*, *FSE*, *ASE*, *CSCW*, *MSR*, and *ESEM*. She has served the research community extensively in different capacities e.g., PC Member for ICSE, MSR, and ESEM.