

“Won’t We Fix this Issue?” Qualitative Characterization and Automated Identification of Wontfix Issues on GitHub

Sebastiano Panichella^a, Gerardo Canfora^b, Andrea Di Sorbo^b

^aZurich University of Applied Science, Switzerland

^bUniversity of Sannio, Italy

Abstract

Context: Addressing user requests in the form of bug reports and Github issues represents a crucial task of any successful software project. However, user-submitted issue reports tend to widely differ in their quality, and developers spend a considerable amount of time handling them.

Objective: By collecting a dataset of around 6,000 issues of 279 GitHub projects, we observe that developers take significant time (i.e., about five months, on average) before labeling an issue as a wontfix. For this reason, in this paper, we empirically investigate the nature of wontfix issues and methods to facilitate issue management process.

Method: We first manually analyze a sample of 667 wontfix issues, extracted from heterogeneous projects, investigating the common reasons behind a “wontfix decision”, the main characteristics of wontfix issues and the potential factors that could be connected with the time to close them. Furthermore, we experiment with approaches enabling the prediction of wontfix issues by analyzing the titles and descriptions of reported issues when submitted.

Results and conclusion: Our investigation sheds some light on the wontfix issues’ characteristics, as well as the potential factors that may affect the time required to make a “wontfix decision”. Our results also demonstrate that it is possible to perform prediction of wontfix issues with high average values of precision, recall, and F-measure (90%-93%).

Keywords: Issue Tracking, Issue Management, Empirical Study, Machine Learning.

1. Introduction

The complexity of modern software systems is growing fast and software developers need to continuously update their source code [1] to meet users’ expectations and market requirements [2]. In this context, fixing bugs or addressing feature requests and enhancements, reported by users in the form of bug reports [3, 4] and Github issues [5], represents a crucial task of any successful software project [6, 7]. Indeed, during software development and maintenance, issue reports are valuable sources of information for developers interested in improving the quality of the software produced [4, 8].

Nevertheless, software changes that are performed to address user-submitted reports often occur under time pressure [9, 10], with negative effects on the developers’ workloads [11]. Indeed, user-submitted reports tend to widely differ in their quality [12, 13], and software developers have to spend a significant amount of time in handling these reports (e.g., verifying their content or relevance [14, 15] and coordinating the teamwork [16]) for implementing the required changes [17, 18, 19].

In the last decade, research has developed automated solutions to facilitate the issue management and fixing processes, with techniques able to prioritize the requested changes [15,

20], to detect potential issue misclassifications [21, 22] and bug duplications [17]. Hence, most of the proposed tools and prototypes are used to answer critical and relevant questions related to reported issues, e.g., “Who should fix this bug?” [4] or “Is It a Bug or an Enhancement?” [21]. However, to the best of our knowledge, only few works investigated the nature of *wontfix* issues, known as “bugs that will never be fixed” [23]. By analyzing more than 6,000 issues from the history of 279 GitHub projects, we observe that developers require time (i.e., about five months, on average) before closing an issue with the *wontfix* status. This means that, in general, developers take about five months to answer the question: “Won’t We Fix this Issue?”. Starting from this preliminary result, we decided to study the main characteristics of this specific type of issues, thus investigating the main reasons behind a “wontfix decision”. In addition, we further explored potential factors that could be related to the time to close a general wontfix issue and experimented automated approaches to identify with high accuracy the issues that will be labeled as wontfix, by only analyzing issue titles and descriptions. To the best of our knowledge, no prior work proposed approaches to automatically determine whether an issue will be likely marked (or labeled) as a *wontfix*.

Our goal is to support “community members”¹ during the issue management process. As pointed out by Guo *et al.* [24],

Email addresses: panc@zhaw.ch (Sebastiano Panichella), canfora@unisannio.it (Gerardo Canfora), disorbo@unisannio.it (Andrea Di Sorbo)

¹With the term “community members” we refer to developers participating in the discussion of the issues, e.g., by assigning labels, answering to issues

unfixed bugs receive almost the same amount of attention as fixed bugs: *e.g.*, in the Eclipse bug database the average numbers of comments that unfixed bugs and fixed ones receive are 4.5 and 5.8, respectively. Similarly, the average number of comments received by wontfix issues in our dataset is 5.14. Thus, approaches for timely identifying the issues that will likely be not addressed allow to reduce the unproductive effort (and associated costs) required for triaging and resolving such issues [25]. In particular, early identification of issue success can help (i) project managers allocating resources, (ii) developers focusing their attention on the issues that will be actually addressed, and (iii) customers knowing early if their requirements would be satisfied [26]. Indeed, the longer the issue that will be likely not fixed remains open, the more it could catch the attention of developers, making them spending efforts in gathering information for attempting to resolve it [27]. Furthermore, being aware of the reasons why developers decide to not fix specific issues could help understanding the software changes that developers consider less relevant. This information could be very useful for improving issue prioritization and triaging mechanisms, in order to better support developers to focus on the issues that will actually get addressed.

Hence, this paper aims at answering the following research questions:

- **RQ₁: What are the main reasons for closing Github issues with the wontfix status?** In this research question we qualitatively characterize wontfix issues, by manually analyzing a sample comprising 667 wontfix issues extracted from 97 different projects (developed in C#) hosted on GitHub, with the aim of understanding the main reasons behind a “wontfix decision”. As initial outcome, we design two different taxonomies. The first taxonomy encompasses the main reasons that pushed users to open issues that later were marked with the *wontfix* label. The second taxonomy models the main motivations given by developers when they decide to close these issues (as wontfix).
- **RQ₂: What factors relate with the resolution time of wontfix issues?** This research question is a follow-up of the previous one. However, while in RQ₁ we look more at the nature of wontfix issues (*e.g.*, investigating the reasons behind a wontfix decision), here we investigate the factors that could be related to the time to close a wontfix issue, observing also whether different wontfix issue types (*i.e.*, having different resolution motivations) present different characteristics and, consequently, different resolution times. As the time elapsed between the opening and closing of an issue is not necessarily linked to the actual effort spent by developers on the issue itself, RQ₂ is also aimed at better characterizing the effort spent on each wontfix issue type.
- **RQ₃: Can machine learning be used to automatically identify issues that are likely to be closed as wontfix?**

In our work we observed that developers take, on average, about five months to figure out that an issue is not worth to be fixed and therefore be labeled as a *wontfix*. In this research question, we want to explore an automated method to anticipate this decision, thus helping developers recognize wontfix issues earlier in the issue management process. During our investigation (RQ₁ and RQ₂) we observed that most wontfix issues talk about specific aspects (*e.g.*, feature enhancements/request, not critical bugs), and that different wontfix issue types tend to experience different resolution times. We conjecture that the topics discussed in the title and the description of an issue report are discriminant and relevant aspects to consider for the fixing of an issue. Based on this consideration we experimented approaches that leverage textual analysis and machine learning techniques to predict whether an issue will be marked (or labeled) as wontfix, by analyzing only the titles and descriptions of reported issues.

Results of our study provide insights into the nature of wontfix issues, and, in particular, we found that developers mainly tend to close issues (with the *wontfix* label) containing erroneous reports, or requesting features (or changes) that are not relevant or out of the project’s scopes. In addition, the time required to close issues, that developers deliberately decide not to consider, is mainly connected with (i) the issue type, and (ii) the number of participants involved in the related discussions. Finally, our evaluation shows that it is possible to predict whether developers will close an issue as a wontfix by analyzing only the titles and the descriptions of the reported issues and using machine learning and textual analysis techniques. The proposed methodology achieves an average value of precision, recall, and F-measure ranging between 90% and 93%. Pragmatically speaking, the high effectiveness of the experimented ML models in automatically identifying the issues that will likely be closed with the wontfix status could improve the issue management processes allowing developers to focus on more critical issues. Indeed, on GitHub, although the use of labels has been envisaged to mark and manage issue reports [28], no automated mechanisms have been provided to support humans during issue labeling. It is worth pointing out that our work has the only aim of providing developers with a tool to more easily identify issues that developers will deliberately avoid to fix by labeling them as *wontfix*. Other issues may remain open for a long time without any fix [27], or developers could mark them as “invalid”, or close them for any other reason [4]. All these situations are out of the scope of this paper.

We believe that our findings not only shed some more light on the nature of wontfix issues, but have the potential to build and/or improve future recommender systems aimed at prioritizing and supporting the issue fixing and the management processes of modern software projects.

In summary, the main contributions of this paper are:

- Two taxonomies modelling the reasons for opening and closing wontfix issues, along with a manually-labeled dataset (available for replication purposes) of 667 wontfix issues extracted from heterogeneous GitHub projects.

change requests, and identifying also potential issues containing erroneous reports, or requesting features or changes that are out of the project’s purposes.

- Results of our study on the characteristics of the different types of wontfix issues.
- An automated approach (available for research purposes) able to accurately identify the issues that will likely be not fixed.

Paper structure. Section 2 discusses the current issue management cycle (with specific reference to GitHub), while Section 3 illustrates the related literature. Section 4 details the data extraction process and the evaluation methodology adopted to answer our research questions. Section 5 discusses the results achieved in our study, while threats to its validity are reported in Section 6. Finally, Section 7 concludes the paper and outlines directions for future work.

2. Background

An issue tracking system is a repository where users and team members can submit and discuss issues (*e.g.*, bugs and feature requests), ask for advice and share opinions useful for maintenance activities or design decisions [29]. GitHub is a social coding platform hosting more than 57 million of repositories² which provides advanced version control mechanisms and an integrated issue tracker.

Any GitHub user can create an issue in a public repository in order to report bugs, require enhancements, or make other kinds of requests. Thus, issues are the primary mean through which GitHub communities collect user feedback. A typical issue on GitHub is described through a *title* and a *description*. It is worth noticing that, differently from other bug trackers (*e.g.*, JIRA or Bugzilla), the GitHub issue tracker does not provide an explicit description field and the issue description is usually provided in the issue’s first comment. Here and in the remainder of the paper, we refer to the first comment of an issue as the *issue description*. Moreover, one or more predefined *labels* are used to help in categorizing the issue. Each issue is assigned to one *assignee* that is responsible for working on it, but *comments* allow anyone to provide feedback. In order to offer high flexibility, GitHub only provides two issue states (*open* and *closed*), while any other state must be realized via labels. The GitHub issue tracker provides a set of default labels in each repository, including the *wontfix* label which indicates that work will not continue on an issue³. The *wontfix* label is among the most used labels in GitHub projects [30]. Although labeling has a positive impact on the effectiveness of issue processing [31], the labeling mechanism is scarcely used on GitHub [30]. Thus, automated approaches able to predict the correct labels to assign to issues could stimulate the use of such a mechanism.

In this study we are interested in extracting issues from heterogeneous projects hosted on GitHub, having the *closed* status and the *wontfix* label assigned, in order to investigate common characteristics of this kind of documents.

3. Related Work

3.1. Issue Management Process and Practices

Fixing bugs and addressing feature requests or enhancements, reported in the form of bug reports [3, 4] and Github issues [5], are crucial tasks for the success of any software project [6, 7, 32]. For this reason, researchers investigated factors characterizing or affecting the issue management process and practices.

Previous work investigated the aspects that should characterize an informative (or “*good*”) bug report [12, 13, 33]. Specifically, Hooimeijer *et al.* [12] presented a first descriptive model of bug report quality, which is based on a statistical analysis of over 27,000 bug reports of Mozilla Firefox. The evaluation of the model showed its usefulness in reducing the overall cost of software maintenance, suggesting, at the same time, potential features that should be considered when composing bug reports. Bettenburg *et al.* [13] conducted a survey involving 466 developers and reported that there is a huge information mismatch between what developers need and what users provide in the reported issues. Their results suggest that future bug tracking systems should focus on engaging bug reporters, with tools handling bug duplicates. These findings have pushed, in later years, researchers to find solutions to handle the bug duplication problem [17, 34].

Recent research studied socio-technical dynamics [35] concerning the management and fixing of issues [31] or the handling of pull requests [36, 8]. For instance, Aranda *et al.* [16] investigated coordination activities around bug fixing tasks by surveying professional developers and found that, even for simple bugs, an inefficient coordination among developers impacts the efficiency of the issue fixing process. These co-ordination problems [11] are generally caused by the wrong assignment [37] or re-assignment of bug reports [18] to developers. However, in other cases, inefficient bug resolutions are influenced by the length and complexity of issue discussions [38, 39], the actual knowledge/skills of developers [40, 5], and other socio-technical dynamics [41, 42, 43]. In this context, Breu *et al.* [33] pointed out that an active participation of developers represents a crucial aspect for making progress on the reported bugs. For this reason, other work proposed strategies for determining the appropriate person to assign the reported issues [4, 44, 45].

Differently from the aforementioned work, this paper empirically investigates the main characteristics of wontfix bugs and the common reasons behind a wontfix decision.

3.2. Issue Reports Classification and Prioritization

As reported by Cosentino *et al.* [46], even if issues are generally sent to popular projects, the number of pending issues constantly grows [43] and, despite the use of labels has a positive impact on the issue evolution [30], they are scarcely used by GitHub developers. Thus, researchers proposed automated solutions to ease the issue management and fixing processes, with techniques that leverage well-known methods based on textual analysis [17], machine learning [21, 47, 48], natural language parsing (NLP) [49, 50], and summarization approaches [51, 52, 53] to analyze bug reports information.

²<https://goo.gl/JbS6RE>

³<https://help.github.com/articles/about-labels/>

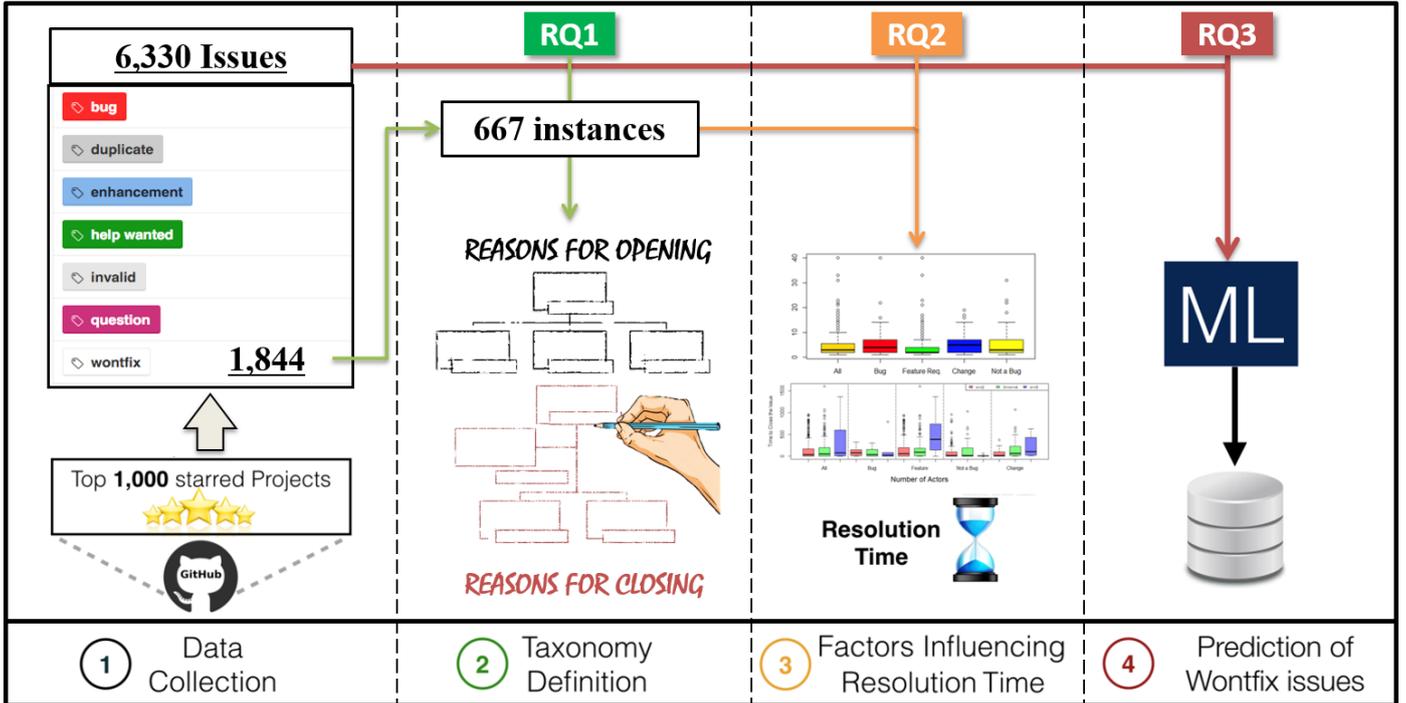


Figure 1: Overview Research Approach

Important results in this direction are related to the definition of approaches that automatically classify or analyze the textual content of reported issues [54] to derive potential misclassifications [21, 22], detecting duplicated bugs [17] or predicting reopened issues [55]. In recent years, tools have been designed to automatically predict the severity of bug reports [14, 56, 57, 58], to support the prioritization of reported issues [15, 20], and to estimate the issue life time [9, 59, 60, 61]. Finally, to facilitate the process of fixing issues, recent strategies have been proposed to translate bug reports into test cases [62], generating auto-fixes [63], or recommending relevant classes [64] for these reports.

In the context of these related studies, this paper empirically investigates the combination of machine learning and textual analysis techniques to automatically predict whether issues will be not fixed, by analyzing (only) the titles and descriptions of reported issues. The closest works to ours are (i) the one by Cabot *et al.* [30] who proposed labels to classify issues in open source projects, and (ii) the one by Guo *et al.* [24] presenting an approach to determine the bugs that will be actually fixed. Finally, recent research [25] started investigating the reasons behind wontfix bugs. By manually analyzing a sample of 600 wontfix bugs (extracted from Bugzilla) pertaining to three open-source projects (*i.e.*, Eclipse, Mozilla, and OpenOffice), the author identified 12 categories of reasons. Similarly to this research, we perform a manual inspection of wontfix issues. However, we identify, for each inspected issue, both the reason behind the issue opening and the motivation for issue closing, to better understand the co-occurrences between the different sorts of reasons. Besides, we (i) investigate wontfix issues pertaining

to 279 heterogeneous projects (hosted on GitHub), identifying further categories of reasons that have not been considered in previous work, (ii) explore the factors that could be related to the time to close a wontfix issue, and (iii) propose an approach able to automatically identify wontfix issues with high effectiveness. To the best of our knowledge, no prior work investigated the nature of wontfix issues on GitHub and proposed approaches to automatically determine whether an issue will be marked as wontfix.

4. Study Design

The *goal* of our study is to shed some light on the nature of wontfix issues, with the *purpose* of building and/or improving recommender systems aimed at prioritizing and supporting the issue fixing and the management processes of modern software projects. Hence, we qualitatively investigate the main reasons behind a “wontfix decision” and explore potential factors that could be correlated with the time to close a general wontfix issue. Finally, we experimented with potential strategies to predict whether an issue will be labeled as a wontfix. Figure 1 depicts the research approach we followed to answer our research questions.

4.1. Data Collection

The *context* of the study consists of 6,330 issues extracted from the history of 279 open source projects hosted on GitHub, whose characteristics are summarized in Table 1. The selection process we applied for this study is based on a “criterion sampling”[65], according to the following steps:

1. **Language selection:** As reported in [66], projects on GitHub developed in C# usually have a higher number of external users, and core developers tend to ignore reports from outsiders [67, 68]. Besides, in a study involving about 100,000 GitHub projects, Bissyandé *et al.* [5] found that GitHub projects developed in C# usually have higher numbers of issues filed (see Figure 7 in [5]). Having a higher number of issues to deal with could increase the likelihood that developers overlook some of these issues due to lack of resources or low priority [69]. Thus, in this type of projects (C# projects), we expect to find higher amounts of wontfix issues. For this reason, we selected projects mostly developed through the C# programming language.
2. **Projects selection:** Recent studies demonstrated that a higher number of issues co-occurs with (i) a higher number of stars received by a GitHub repository[70], and (ii) a faster growth of a GitHub project in terms of stars[71]. Thus, in line with recent empirical studies in software engineering [72, 73, 74, 75], we selected projects relying on stars information. In particular, in order to consider projects with reliable amounts of issues, the 1000 most popular ones (*i.e.*, top starred) have been selected from GitHub.
3. **Metadata extraction:** We collected all closed issues metadata (*e.g.*, URL, title, description, resolution date, etc.) from the projects selected according to the aforementioned criteria.

The aforementioned steps were performed using the R scripts available in our replication package (under the folder “*I_Scripts/I_Data_Collection*”). In particular, a first script was employed for selecting the projects according to the (1) language selection and (2) projects selection criteria. This R script addressed the following technical issues: (i) we selected the projects having the higher number of stars; (ii) we selected projects having issues closed, filtering out projects having no issue labels (not all project had issue labels); (iii) we handled the Github download limits (setting in R a timeout with “*Sys.sleep(40)*”). As result, the script collected the first initial information about the identified projects such as project name, project Github url, project program language, and issue labels. The other three R scripts we implemented are responsible to collect more detailed information about the issues of the selected projects (*e.g.*, issue url, issue title, and issue description), double-checking that no other issues have been closed as wontfix during the extraction analysis (we added few more wontfix issues with this check). It is worth noticing that during our investigation we observed that specific projects may use custom labels for designating issues that will be not addressed (*e.g.*, *status:wontfix*, *Resolution-Won't Fix*, *won't fix*, *resolved:wontfix*, *closed:wontfix*, *wont-fix*, *Won't Fix*, *not-fixing*, *Status-WontFix*, *WontFix*, *status: will not fix* and *Cannot fix*). Thus, our scripts consider issues having the *wontfix* labels described above.

Table 1 reports the number of projects, total number of WONTFIX and NON-WONTFIX closed issues mined from these projects.

Table 1: Characteristics of the analyzed C# projects.

#Projects	#Issues	#Wontfix	#Non-wontfix	Med.#Issues.
279	6,330	1,844	4,486	22.69

We also report, in the last column of Table 1, the median number of issues per project.

Replication package. We make available in our replication package⁴ (i) the scripts developed to extract the data used for this research, (ii) all raw data, used to generate the data and tables reported in the paper. In the replication package, we also include the research prototype we used to answer RQ₃.

4.2. Analysis Method

In the following, we discuss the research methods used to address each specific research question.

4.2.1. RQ₁: What are the main reasons for closing Github issues with the wontfix status?

Answering RQ₁ required, as a first step, to derive a manual labeled *golden set* of wontfix issues, in order to build two taxonomies: (i) a first taxonomy, $M_{opening}$, summarizing the main reasons that pushed users to open issues that later were marked with the *wontfix* label; and (ii) a second taxonomy, $M_{closing}$, modelling the main motivations of developers to close these issues as wontfix. Therefore, with the aim of manually inspecting a representative sample (99% confidence level and a margin of error below 4%), T_{sample} , of the collected wontfix issues (see Section 4.1), we randomly selected 667 issues from the entire set of wontfix issues. These 667 issues belong to 97 projects. Concerning the 667 wontfix issues in our T_{sample} , we observed that 286 of them (42.88%) have been opened by end-users (not contributors to the projects), while the remaining 381 issues have been opened by users with different roles in the analyzed projects. In particular, 61 (16.01%) out of these 381 issues have been opened by repositories’ owners, 225 (59.06%) by organizations’ members, 78 (20.47%) by contributors who had previously committed to the repositories, and 17 (4.46%) by collaborators invited to contribute to the repositories.

To derive the two taxonomies we used *card sort*, a technique to derive taxonomies from input data [76]. We organized card sort in three steps [77]: (i) *preparation*, (ii) *execution*, and (iii) *analysis*.

Preparation: In this step, we created the cards related to each wontfix issue in T_{sample} . Each card represents a wontfix issue and includes: (i) the issue title, (ii) the issue description, (iii) all the messages exchanged in the related discussion, and (iv) all the labels (further to *wontfix*) assigned to the issue by original developers.

Execution: Two authors of the paper analyzed the cards applying *open* (*i.e.*, without predefined groups) card sort. In particular, the two authors performed an iterative content analysis [78],

⁴<https://github.com/wontfixRP/wontfixDetection>

starting with two empty lists (one for $M_{opening}$ and the other for $M_{closing}$) of issue categories. Each time they found a new *wontfix* issue type to add to one of two taxonomies, a new category was added to the connected list. The two authors used *pair-sorting* [77], to discuss discrepancies in their thoughts for each card during the card sorting itself and avoid checking the consistency of the sorting and merging the cards in a later phase.

Analysis: To guarantee the integrity of the emerging categories and remove potential redundancies, the two authors performed a second iteration on all the analyzed cards and redefined some of the categories identified in the previous step. Through the card sorting process 22 reasons for $M_{opening}$ and 26 reasons for $M_{closing}$ emerged. During the sorting process we reflected on how they could be further clustered into higher level groups. At the end of this phase, for each taxonomy we identified five high level groups. The resulting taxonomies are described in Section 5.1, where the set of 667 issues manually validated according to the taxonomies represents our *golden set*.

4.2.2. RQ₂: What factors relate with the resolution time of wontfix issues?

In order to characterize *wontfix* issues and answer RQ₂, we computed the following factors concerning all issues in our *golden set* (i.e., All):

- **descriptionLength:** Issue description length (number of characters);
- **maxAuthorPercentage:** The proportion of messages posted by the author who posted the majority of messages in the issue discussion;
- **majorAuthors:** Number of unique authors who have posted more than one-third of the overall messages present in the issue discussion;
- **meanCommentSize:** Average length of comments (number of characters) in the issue discussion;
- **minorAuthors:** Number of unique authors who have posted less than one-third of the overall messages present in the issue discussion;
- **nActorsT:** Number of distinct authors participating in the issue discussion;
- **nCommentsT:** Number of total comments in the issue discussion;
- **timeToCloseIssue:** Time lapse (in days) between issue opening and closing (with the *wontfix* label);
- **timeToDiscussIssue:** Time lapse (in days) between issue opening and last comment posted in the issue discussion.

These factors allowed us to investigate different issue dimensions, namely (i) the level of participation of the community members to the issue discussion ($nActorsT$, $maxAuthorPercentage$, $minorAuthors$ and $majorAuthors$), (ii) the discussion's size ($descriptionLength$, $nCommentsT$ and $meanCommentSize$),

as well as (iii) timing information about the issue ($timeToCloseIssue$ and $timeToDiscussIssue$).

Moreover, we studied how these factors vary when considering the different $M_{closing}$ categories (i.e., *Bug*, *Feature request/enhancement*, *Not suitable* and *Change*). In particular, to verify whether statistically significant differences could be observed between the different $M_{closing}$ categories, for each metric $m \in \{descriptionLength, maxAuthorPercentage, majorAuthors, meanCommentSize, minorAuthors, nActorsT, nCommentsT, timeToCloseIssue, timeToDiscussIssue\}$, we compared the value distributions obtained for m across the different $M_{closing}$ categories, through the Mann-Whitney U test, a widely used non-parametric test for comparing independent samples [79]. After checking that all the variables of interest (i.e., $descriptionLength$, $maxAuthorPercentage$, $majorAuthors$, $meanCommentSize$, $minorAuthors$, $nActorsT$, $nCommentsT$, $timeToCloseIssue$ and $timeToDiscussIssue$) are not well-modeled by normal distributions (as verified through the Shapiro-Wilk test [80]), we decided to use a non-parametric test (i.e., Mann-Whitney), as the assumptions for this test are satisfied by each considered variable. For coping with multiple comparisons, the Benjamini-Hochberg correction procedure [81] has been adopted to adjust p-values.

In addition, in order to investigate if some of the considered metrics may influence the time required to close the issue, similarly to the work by Linares-Vásquez *et al.* [82], for each metric m we grouped the issues in our *golden set* in different subsets, on the basis of specific values of m (e.g., $nActorsT \leq 2$, $3 \leq nActorsT \leq 4$ and $nActorsT \geq 5$) and verified (through the Mann-Whitney U test) whether statistically significant differences can be observed in the $timeToCloseIssue$ distributions obtained for the different subsets. Again, this investigation has been carried out for (i) all the issues in our *golden set* (i.e., All), as well as (ii) the various $M_{closing}$ categories.

4.2.3. RQ₃: Can machine learning be used to automatically identify issues that are likely to be closed as wontfix?

After investigating the nature of *wontfix* issues, we propose an approach, to automatically predict or classify whether an issue will be labeled as a *wontfix*. For achieving this goal, we considered all non-*wontfix* (4,486) and *wontfix* (1,844) issues in our dataset (see Table 1), collected through the metadata analysis explained in Section 4.1. Specifically, our approach leverages machine learning (ML) techniques and consists of four main steps:

1. **Preprocessing:** All terms contained in the titles and descriptions of all 6,330 (4,486 non-*wontfix* plus 1,844 *wontfix*) issues in our dataset are used as an information base to build a textual corpus that is preprocessed applying stop-word removal (using the English Standard Stop-word list) and stemming (i.e., English Snowball Stemmer) to reduce the number of text features for the machine learning techniques [83]. In addition, we use the R package `textclean`, which contains a function called `replace.html` that allows the automated removal of html tags from the text. The output of this phase corresponds to a Term-by-Document matrix M where each column represents

an issue of our dataset and each row represents a term contained in title and or description of the various issues. Thus, each entry $M_{[i,j]}$ of the matrix represents the weight (or importance) of the i -th term contained in the j -th issue.

2. **Textual Feature Weighting:** Words are weighted using the *tf-idf* score [83], as opposed to simple frequency counts, because it assigns a higher value to rare words (or group of words) appearing in issues, and a lower value to common ones. This allows identifying the most important words in the issue titles and descriptions. The weighted matrix M represents the output of this phase.
3. **Training and Test sets:** For the classification step, we split the matrix related to all issues of our dataset in two parts (50% each), *i.e.*, training and test sets. As training set we considered the sub-matrix, $M_{training}$, obtained by randomly selecting from the original matrix M the columns associated to a half of wontfix issues and a half of non-wontfix issues. Vice versa, for the test set we considered the sub-matrix, M_{test} , obtained by considering from the original matrix M the columns associated to the remaining half of wontfix issues and the remaining half of non-wontfix issues.
4. **Classification:** We automatically classify wontfix issues in the test set by relying on the output data obtained from the previous step consisting of the matrix $M_{training}$ and M_{test} . Specifically, to increase the generalizability of our findings, we experimented (relying on the Weka tool⁵) different machine learning techniques, namely, the standard probabilistic Naive Bayes classifier, the sequential minimal optimization (SMO) algorithm⁶, and the J48 tree model. It is important to note that, the choice of these techniques is not random since they were successfully used for bug reports [21, 84] or vulnerability [85] classification, recent work on user reviews analysis [86, 2, 87], and several works on bug prediction [88, 89].

It is worth highlighting that, as stated in Section 1, for predicting whether an issue will be labeled as *wontfix*, the machine learning models have been trained by exclusively using information that are immediately available at the issue opening (*i.e.*, issue title and description, without considering the other features), this to simulate a more realistic scenario in which the automated classification can really help developers identifying the issues that will likely be not fixed.

To evaluate the performance of the experimented ML techniques, we adopted well-known information retrieval metrics, namely precision, recall, and F-measure [83]. It is important to specify that, as described in the steps 3 and 4 of our approach, we apply a cross-projects setting to train the ML models on data coming from different projects. This choice was made to ensure

Table 2: Motivations for issue opening (only wontfix issues)

Category	Motivation	Freq.
Functional Aspects	Feature enhancement	42.6%
	Reported a bug	29.1%
	Feature Request	24.3%
	Unexpected/unclear functionality behaviour	1.0%
	Unknown crash	0.1%
Problem	Performance Issue	2.1%
	Testing related issue	1.2%
	Security issue	0.6%
	Compilation issue	0.4%
	Browser issue	0.3%
	Persistence issue	0.3%
Configuration	Tool version no longer supported	0.6%
	Configuration Problem	0.1%
	Backup problems	0.1%
Documentation	Question	2.8%
	Request of adding an example in the documentation	0.6%
	Not clear or incorrect code/software examples	0.1%
Other	Development aspects	1.3%
	Comment on the software (Not a real issue post)	0.9%
	Improvement of graphical aspects	0.4%
	Current feature beyond the scope of the library	0.1%
	General comment	0.1%

that a more general classification model is trained. To complement the evaluation process and alleviate concerns related to overfitting and selection bias, we also provide the classification results of the experimented machine learning models, by computing a 10-fold validation strategy.

5. Results

This section discusses the results of our empirical study.

5.1. RQ₁: Reasons for Wontfix Issues

To explore the common motivations for closing issue reports which developers deliberately avoid to consider/fix (*i.e.*, *wontfix*), and answer RQ₁, we performed a manual analysis of a sample of issues extracted from our data collection (as described in Section 4.2.1). More specifically, such *golden set* encompasses 667 closed issues (with the *wontfix* label) extracted from 97 distinct projects hosted on GitHub.

Each issue in the sample has been marked with two labels: (i) the motivation behind the issue opening, as stated by the issue reporter (*i.e.*, the motivation for issue opening, $M_{opening}$), and (ii) the reason for its closing (with the *wontfix* status), as declared by developers within the issue discussion (*i.e.*, the motivation for issue closing, $M_{closing}$).

For $M_{opening}$, we found 22 different motivations (reported in Table 2 along with their frequencies within the analyzed sample), that have been grouped in five distinct categories. It is worth to highlight that 64 (9.6%) issues have been assigned to more than one $M_{opening}$ category, since they have been opened for multiple purposes (this explains why the sum of percentage values in Table 2 is higher than 100%). For $M_{closing}$, we

⁵<https://www.cs.waikato.ac.nz/ml/weka/>

⁶<https://goo.gl/BwMjzS>

Table 3: Motivations for issue closing (only wontfix issues)

Category	Motivation	Freq.
Feature request/ enhancement	Feature request/enhancement already implemented or not needed	47.8%
	Not relevant request	4.5%
	It was fixed in the context of previous bug fixes	0.7%
	Too expensive feature request	0.3%
	Feature request that will be implemented in the near future	0.1%
	Already Implemented feature request by an external contributor of the project	0.1%
Change	Not relevant change	14.5%
	No time to work on this change	1.0%
	Requested change leading to further problems	0.7%
Not suitable	Not a bug	10.8%
	Configuration/backup problem on the user side	3.9%
	Duplicated issue	2.4%
	Unclear wrong/usage of a functionality	1.6%
	Problem already fixed with the new version	1.3%
	Problem fixed updating the new version of a dependent library/tool	0.7%
	Tool version no longer supported	0.4%
	Not replicable bug	0.1%
Bug	Impossible to fix the issue or too expensive change	4.9%
	Not a critical bug	2.5%
	It will be fixed in future	1.8%
	Difficult to fix or to replicate	1.6%
	Unknown crash	0.1%
Other	General comment from a user	0.4%
	Closed by the user	0.1%
	It was a test failure	0.1%
	Updated the documentation on wiki	0.1%

found 26 distinct motivations (reported in Table 3). Such motivations have been clustered in five categories. Only 23 issues (3.4%) have been marked with multiple $M_{closing}$ motivations, this is mainly due to the fact that community members usually tend to provide a precise indication for not fixing an issue.

In most cases, the reasons for opening an issue are related to bugs reporting, feature requests or enhancements, and only in few cases, by other requests (e.g., clarification questions, performance, and testing related aspects). As expected, the majority (i.e., 648, 97.15%) of issues in our sample have been opened in order to signal troubles dealing with functional aspects (see Table 2). As illustrated in Table 2, many of the issues belonging to the *Functional Aspects* category are opened in order to (i) request improvements for specific features (i.e., *Feature enhancement*, 42.6%), (ii) report a bug (i.e., *Reported a bug*, 29.1%), or (iii) require a new feature to be implemented (i.e., *Feature request*, 24.3%). As anticipated, requests for fixing defects not strictly related with the software features (i.e., *Problem in Table 2*), the requests about software documentation (i.e., *Documentation in Table 2*), and configuration problems (i.e., *Configuration in Table 2*) resulted in rare motivations for opening issues.

Community members usually decide to ignore issues (i) containing erroneous reports, or (ii) requesting features or changes that are out of the project’s purposes (e.g., requests of improving the performance or GUI associated to a feature, or adding a functionality that is already present in the system). As a

matter of fact, 319 (47.8%) issues in our sample, have been closed with the motivation that the requested features/enhancement were not needed or had been already implemented, while 142 (21.3%) issues erroneously reported problems, which have been proved to be not suitable (see Table 3). Indeed, only 63 (9.4%) issues revealed actual bugs, which developers decided to not fix, as they have been often evaluated as (i) too expensive to fix (i.e., *Impossible to fix the issue or too expensive change*, 52.4% of issues signaling actual bugs), (ii) not critical (i.e., *Not a critical bug*, 27% of issues signaling actual bugs), or (iii) that will be fixed in the future (i.e., *It will be fixed in future*, 19% of issues signaling actual bugs). In addition, issues reporting change requests (i.e., *Change in Table 3*) are mainly closed (i.e., 14.5% of issues) because the change they propose are judged as not strategically relevant by community members. We argue that results in Table 3 could be useful to implement more accurate analysis (not necessarily based on binary classification) for future work, such as multi-label issue classification [90] and issue prioritization [91].

In Figure 2 we illustrate the frequency with which issues opened with the most recurrent purposes in $M_{opening}$ (i.e., *Feature enhancement*, *Reported a bug* and *Feature Request*) are closed with one of the motivations in $M_{closing}$ (see Table 3). In Figure 2 the thickness of the lines is proportional to the amount of issues opened with a specific reason (on the left) and closed with a specific $M_{closing}$ motivation (on the right). In particular, 146 (89%) issues opened for requiring a new feature (i.e., *Feature Request*) have been closed with the motivation *Feature request/enhancement already implemented or not needed*, while 173 issues (60.1%) requiring feature enhancements (i.e., *Feature Enhancement*) have been closed due to the same motivation. Moreover, 83 (28.8%) issues having the same purpose (i.e., *Feature Enhancement*) have been closed, since they proposed *Not relevant changes*. Finally, issues reporting bugs (i.e., *Reported a bug*) are mainly closed because (i) they do not signal actual bugs (i.e., *Not a bug*, 33.5%), (ii) the bugs reported are too expensive or impossible to fix (i.e., *Impossible to fix the issue or too expensive change*, 12%), or (iii) the signaled defects mainly depend on configuration/backup problems on the user side (i.e., *unset backup or other configurations required on the user side for enabling the main functionalities of the project*, 11.5%).

RQ₁ summary: *Developers mainly tend to close issues (with the wontfix label) containing erroneous reports, or requesting features (or changes) that are not relevant or not needed.*

5.2. RQ₂: Factors Related with the Wontfix Issues Resolution Time

As reported in Table 4, wontfix issues are mainly discussed among limited numbers of major actors and such discussions encompass 4.43 comments, on average. As anticipated, wontfix issues are closed very long time after their opening (i.e., about five months on average) and continue to be discussed even after their closing.

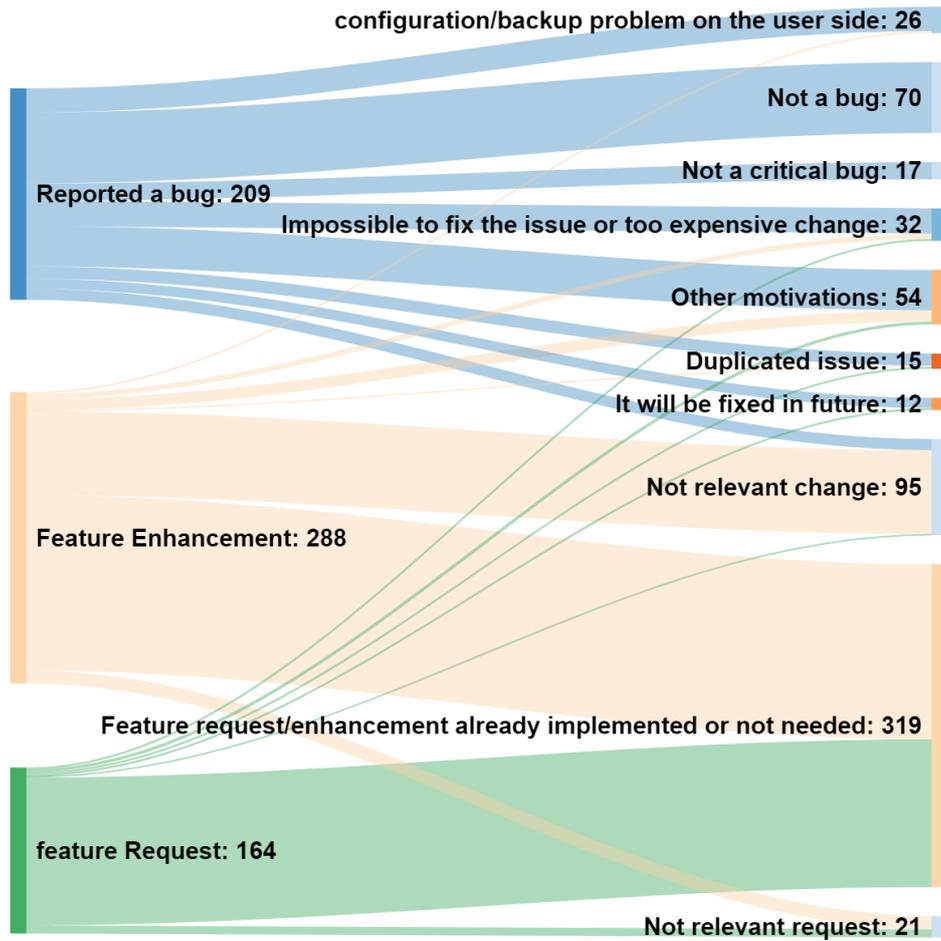


Figure 2: Co-occurrences of motivations for issue opening and closing (only wontfix issues)

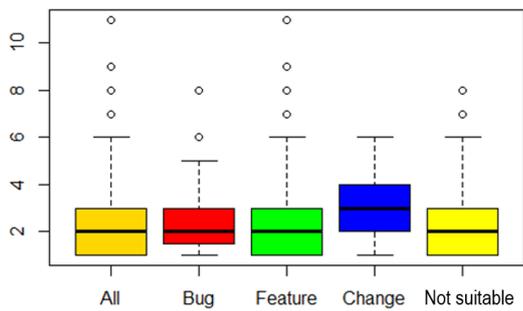


Figure 3: Distributions of $nActorsT$ for the $M_{closing}$ categories, All (median = 2), Bug (median = 2), Feature (median = 2), Change (median = 3), Not suitable (median = 2)

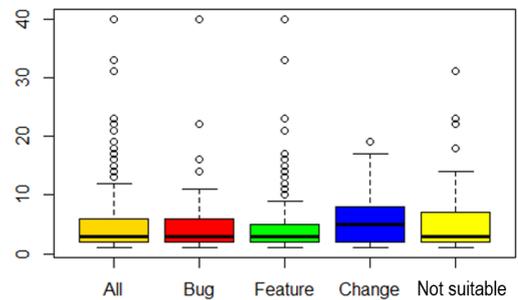


Figure 4: Distributions of $nCommentsT$ for the $M_{closing}$ categories, All (median = 3), Bug (median = 3), Feature (median = 3), Change (median = 5), Not suitable (median = 3)

Table 4: Summary of wontfix issues characteristics

Metric	Median	Average
nCommentsT	3.00	4.43
nActorsT	2.00	2.34
maxAuthorPercentage	0.50	0.63
minorAuthors	0.00	0.82
majorAuthors	2.00	1.54
timeToCloseIssue	42.60	153.59
timeToDiscussIssue	76.05	201.33
DescriptionLength	470.00	847.49
meanCommentSize	380.00	495.68

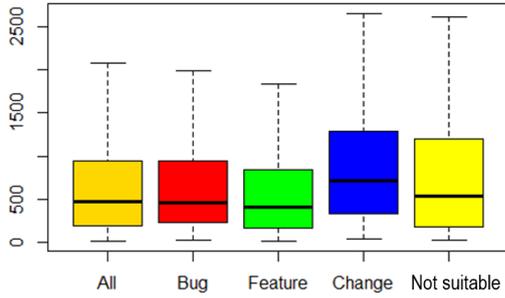


Figure 5: Distributions of *descriptionLength* for the $M_{closing}$ categories, All (median = 470), Bug (median = 459), Feature (median = 408.5), Change (median = 720), Not suitable (median = 535.5)

In order to investigate differences in the different types of wontfix issues, we study the extent to which the collected metrics vary across the specific $M_{closing}$ categories, and discuss the most interesting peculiarities. More specifically, to study the differences occurring between the different kinds of wontfix issues, and verify whether the observed differences are statistically relevant, we tested the following null hypothesis:

H_0 : The distributions of values of the metric m for the populations of wontfix issues of the types t_i and t_j are equal

$\forall m \in \{descriptionLength, maxAuthorPercentage, majorAuthors, meanCommentSize, minorAuthors, nActorsT, nCommentsT, timeToCloseIssue, timeToDiscussIssue\}, \forall t_i, t_j \in \{All, Bug, Not\ suit-$

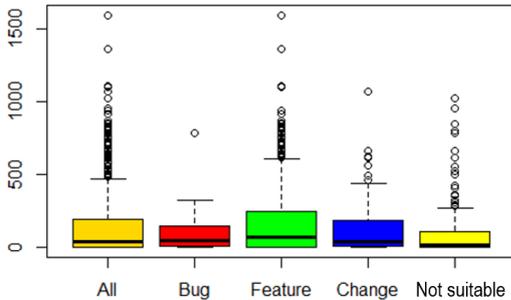


Figure 6: Distributions of *timeToCloseIssue* for the $M_{closing}$ categories, All (median = 42.60), Bug (median = 52.06), Feature (median = 73.59), Change (median = 39.19), Not suitable (median = 15.67)

able, *Feature request/enhancement, Change*) and $i \neq j$.

H_0 has been tested with Mann-Whitney test and the p -value was fixed to .05. Table 5 reports the results of the Mann-Whitney test. This investigation is aimed at understanding whether specific types of wontfix issues exhibit more unproductive (or longer) discussions, before their closing. In particular, for each metric (on the rows) and each pair of wontfix issue types (on the columns) Table 5 reports the p -value obtained when testing the null hypothesis H_0 through the Mann-Whitney U test. The Benjamini-Hochberg correction procedure [81] has been adopted to adjust p -values, since multiple comparisons are performed simultaneously.

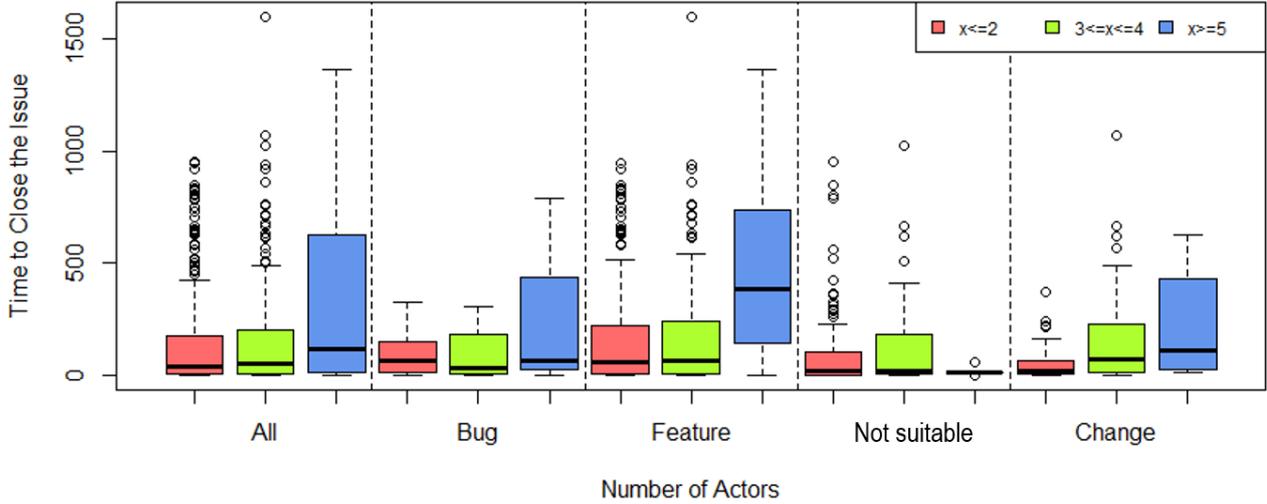
Figures 3, 4, 6 and 5 report the respective distributions of $nActorsT$, $nCommentsT$, *timeToCloseIssue* and *descriptionLength* obtained for the overall issues in our dataset (*i.e.*, *All*), as well as the various $M_{closing}$ categories. *Change* requests (*i.e.*, *Change*) (i) are usually described through longer texts (see Figure 5), (ii) require to be discussed between a greater number of actors (as illustrated in Figure 3) and, consequently, (iii) the related issue discussions comprise a greater amount of comments (see Figure 4), than other kinds of issues. As shown in Figure 6, not suitable reports (*i.e.*, *Not suitable*) are closed much faster than the other types of issues: 50% of issues of this type are closed in less than 16 days, while the 50% of the issues belonging to the other categories require more than 39 days to be closed. Probably, this is due to the fact that developers are more resolute in closing the issue, once verified that the signaled defect is not actually suited to be addressed.

On the contrary, *Feature requests/enhancement* issue types usually require more time to be closed (the median value of *timeToCloseIssue* obtained for issues of this type is 73.59 days), probably because developers have greater uncertainty on deciding if the required improvements could be in line with the project's purposes. In general, the number of participants discussing the issues may influence, with statistical evidence (see Table 6), the time required to close a wontfix issue, while for the other collected factors we do not observe significant relationships. Specifically, as illustrated in Figure 7, when the numbers of actors participating in the issue discussions concerning wontfix issues of the *Feature request/enhancement* and *Change* types increase, we observe a longer *timeToCloseIssue*, while for wontfix issues of the *Bug* and *Not suitable* types, no statistically significant differences between the different subsets are revealed. It is worth noticing that we verified whether significant relationships exist between the other metrics and the *timeToCloseIssue*, by using similar analyses. However, such analyses did not produce noteworthy results.

In a study involving more than 4000 GitHub projects and about 1 million issues, Kikas *et al.* [59] found that the median lifetime of about 70% of the investigated issues is 3.7 days. We observe that the median closing time for wontfix issues is about 11.5 times slower than the median lifetime of most issues investigated in prior work, confirming (i) our intuition that wontfix issues usually remain open for a longer time compared to other types of issues, and (ii) the need for early detection of issues that will probably remain unfixed. Our study also partially confirms some of the findings reported in prior research [59],

Table 5: Hypothesis testing results: All (A), Bug (B), Not suitable (NS), Feature request/enhancement (FR), Change (C)

Metric	A-B	A-NS	A-FR	A-C	B-NS	B-FR	B-C	NS-FR	NS-C	FR-C
nCommentsT	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p < 0.05$	$p < 0.05$
nActorsT	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p < 0.05$
maxAuthorPercentage	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p < 0.05$
minorAuthors	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$0.05 \leq p \leq 0.1$	$p > 0.1$	$p < 0.05$	$p < 0.05$
majorAuthors	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$
timeToCloseIssue	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$
timeToDiscussIssue	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$0.05 \leq p \leq 0.1$	$p > 0.1$	$p > 0.1$
DescriptionLength	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$
meanCommentSize	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$0.05 \leq p \leq 0.1$

Figure 7: Distributions of *timeToCloseIssue* for issues (labeled as wontfix) having different number of actors (x): $x \leq 2$, $3 \leq x \leq 4$, $x \geq 5$ Table 6: Hypothesis testing results for issues having different number of actors (x)

Test	All	Bug	Feature	Not suitable	Change
$(x \leq 2)$ vs $(3 \leq x \leq 4)$	$p < 0.05$	$p > 0.1$	$p > 0.1$	$p > 0.1$	$p < 0.05$
$(x \leq 2)$ vs $(x \geq 5)$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p < 0.05$
$(3 \leq x \leq 4)$ vs $(x \geq 5)$	$p < 0.05$	$p > 0.1$	$p < 0.05$	$p > 0.1$	$p > 0.1$

showing that the number of different actors involved in issue discussion is related to the time to close the issue.

RQ₂ summary: *On average, about five months are required to close issues that developers will label as wontfix. This time is mainly connected with (i) the issue type (issues indicating not suitable reports are closed much faster with respect to other kinds of issues), and (ii) the number of participants involved in the related discussions. Such discussions typically comprise less than 6 messages and involve a limited set of major actors.*

Table 7: Evaluation of Machine Learning Classifiers

ML Technique	Precision	Recall	F-Measure
Naive Bayes	0.731	0.698	0.708
SMO	0.758	0.768	0.760
J48	0.896	0.897	0.894

5.3. RQ₃: Automated Classification of Wontfix Issues

As explained in Section 4.2.3, we experimented with different ML techniques, namely (i) the probabilistic Naive Bayes classifier, (ii) SMO algorithm, and (iii) the J48 tree model. These ML models were trained on the training data (*i.e.*, $\mathbf{M}_{training}$) and evaluated on the test data (*i.e.*, \mathbf{M}_{test}) illustrated in Section 4.2.3. Table 7 provides an overview of the main results obtained through the different ML algorithms. For completeness, in Table 8, Table 9 and Table 10 we also provide the actual corresponding confusion matrices of all the three experimented ML models.

The results in Table 7 highlight that the precision, recall, and F-measure values are very positive for the J48 model, while we observe lower precision and F-measure results for the Naive

Table 12: The 15 text features with the highest information gain scores

Score	Feature
0.1467	X6354_issu
0.1409	X14304_work
0.1396	X9773_provid
0.1150	X13840_version
0.1058	X8222_need
0.1050	X12642_test
0.0986	X13272_type
0.0966	X1772_chang
0.0957	X12341_support
0.0957	X3896_error
0.0940	X2540_creat
0.0910	X7306_make
0.0907	X7615_method
0.0899	X12433_system
0.0859	X12182_string

with the wontfix status (see Table 3) is the presence of not desired features (about 47% of wontfix issues concern requests for features enhancement/addition). To quantitatively corroborate this observation, we computed the information gain [93] for all the text features leveraged by our model and ranked them according to their scores. In Table 12, the top 15 ranked features, along with the related information gain scores are reported. Looking at Table 12, it is easy to observe that many of the features with the highest scores are conceptually linked to requests for enhancement or feature additions (e.g., provide, need, change, support, create, and make). However, this could also represent a limitation of our model. Indeed, as shown in Table 10, while the J48 model is quite effective in recognizing issues of the *non-wontfix* class, a higher likelihood of false negatives is observed for *wontfix* issues (i.e., the recall for this class is about 75%). For this reason, we believe that further efforts and tunings could be aimed at reducing the false-negative rate for this class while keeping low the false-positive rate.

RQ₃ summary:

1. Relying on a tree classifier (i.e., J48) it is possible to automatically detect issues that will be labeled as *wontfix*, with precision, recall, and F-measure values up to 0.93.
2. Consistently with the results of RQ₁ the experimented models select textual features semantically related to requests for features enhancement/addition to classify *wontfix* issues.

6. Threats to Validity

Threats to construct validity. In order to carry out our study, we measure different factors that could be not sufficient to model the whole issue handling process. As pointed out by Kalliamvakou *et al.* [94], many active projects do not conduct all their software development activities in GitHub, and separate infrastructures (e.g., mailing lists, forums, IRC channels,

etc. [95]) could be used to support decision-making processes. This only represents a minor threat in our study, since most of the events of our interest (e.g., opening and closing of the issues, label assignments, comments) are mostly recorded in the issue tracking system, as it is primarily used by development teams to track issue data.

Threats to internal validity. Our results have been obtained by analyzing *wontfix* issues having the *closed* status assigned and such results could be misleading if a significant percentage of such issues will be reopened in the future. However, less than 9% of non-fixed bugs tend to be reopened [96] and one of the root causes for re-opening a bug report resolved as not fixed is due to the difficulty in reproducing the bug [97]. It is worth to notice that in our manual inspection, the *Difficult to fix or to replicate* and *Not replicable bug* $M_{closing}$ motivations have been assigned only to 1.6% and 0.1% of issues in our sample, respectively. In addition, 281 out of 667 (42.13%) *wontfix* issues encompassed in our manually analyzed sample are related to the *aspnet/Mvc* project. This could represent a threat to internal validity, as developers of this project could adopt similar criteria for deciding to avoid addressing a reported issue. In our study we analyzed issues labeled as *wontfix*. However, GitHub developers may indicate that a specific issue will be not resolved directly in the issue title (e.g., by using the [WONTFIX] prefix⁸). To counteract this issue, we estimated the number of GitHub issues in C# projects where the WONTFIX keyword appears in the issue title⁹. Such a search returned a very limited number of issues (i.e., < 10) and most of them (i.e., 75%) had also the *wontfix* label assigned. To avoid any bias in the potential evaluation of the performance of the experimented ML techniques, we adopted well-known information retrieval metrics, namely precision, recall, and F-measure [83] and apply a cross-projects setting to train the ML models on data coming from different projects. However, since the information concerning all the issues in our dataset is used for the ML model construction, we can not exclude that characteristics related to more recent issues are leveraged to predict the resolution of previously submitted issues.

Threats to conclusion validity. In our RQ₂, we analyze different issue clusters having different sizes in terms of issue numbers, and some of the differences we observed could be not significant. To mitigate this threat, we compared the values obtained for each cluster through the Mann-Whitney U test, widely adopted for similar purposes in the software engineering community, and discussed some of the differences which resulted statistically significant (*p-value* < 0.05). Since multiple comparisons are performed simultaneously, the Benjamini-Hochberg correction procedure [81] has been adopted to adjust *p-values* and control the false discovery rate.

Threats to external validity relate to the generalizability of our results. To mitigate this kind of threats our evaluation has been performed on a dataset containing 6,330 issues, extracted from 279 heterogeneous projects. Moreover, we manually ana-

⁸See <https://github.com/waveform80/picamera/issues/657>

⁹We used the following search string: "is:issue is:closed WONTFIX in:title language:C#"

lyzed 667 wontfix issues of 97 different projects. However, such sample could be not adequately representative of all the GitHub projects. All the considered wontfix issues (1,844) are related to projects developed using mainly the C# programming language, and the average number of issues per projects tend to be quite skewed for some projects. This may represent a threat to external validity, since such issues can present common characteristics that ease their identification. For these reasons, in the future, we aim at extending our investigation by evaluating wontfix issues in further projects developed through different programming languages. However, on a positive side, the classifier demonstrated to achieve high performance in identifying wontfix issues, even when not trained on issues related to any specific project, thus this training strategy allows the classifier to be more easily used on projects, different from the ones used in our experimentation, without the need for re-training it. On the other hand, the high classification performance achieved by some ML models could depend on the fact that many wontfix issues encompassed in our dataset concern *requests for features enhancement/addition*. Indeed, the qualitative analysis performed in Section 5.3 highlighted that ML models leverage syntactical features semantically related to this kind of requests to perform the classification. Thus, it is not clear if similar results could be obtained on more balanced datasets, in which the different types of wontfix issues are equally represented.

7. Conclusions

Software maintenance and evolution activities represent crucial tasks of any successful software project, and issues reported by users are a valuable source of information for developers interested in improving their systems. However, developers spend significant time handling issue reports and user requests. To support developers during issue handling processes, researchers conceived effective solutions aimed at prioritizing requested changes, as well as detecting potential issue misclassifications or duplications. However, few prior studies explored the nature of wontfix issues, and none of these studies proposed approaches to automatically determine whether an issue will be marked as *wontfix*. We argue that a timely identification of issues that are likely to be not addressed, could help (i) project managers allocating resources, (ii) developers focusing their attention on the issues that will be actually addressed, and (iii) customers knowing early if their requirements would be satisfied [26]. To this aim, in this paper, by collecting more than 6,000 issues extracted from the history of 279 GitHub projects, we (i) analyzed the common characteristics of wontfix issues, and (ii) proposed an approach leveraging textual analysis and machine learning techniques to predict whether an issue will be resolved as a wontfix. Results of our study show that developers mainly tend to close issues (with the wontfix label) containing erroneous reports, or requesting features (or changes) that are not relevant or out of the purposes of projects (RQ₁). However, developers take a significant amount of time (about five months, on average) to decide whether an issue should be labeled as a wontfix. This time is mainly connected with (i) the issue type (issues containing erroneous bug reports, are closed

much faster), and (ii) the number of participants involved in the related discussions (RQ₂). Last but not least, the proposed approach proved to be accurate (with a F-measure up to 93%) in identifying issues that will be likely labeled as wontfix (RQ₃).

This study helps to better comprehend the issue management dynamics in open source communities. As a future work, we aim at investigating whether different projects tend to have different wontfix characteristics (due to different issue management processes), and the extent to which the automated identification of wontfix issues may impact the results produced by issue prioritization approaches. In addition, in the future, we aim at comparing how issues with/without *wontfix* label perform each other, in order to investigate how the presence of *wontfix* issues may affect the overall issue management process. We intend to also study further wontfix factors useful to automatically identify/predict the actual potential motivations (as it could be useful information for developers) behind an issue that will be closed as a wontfix. Moreover, we plan to compare the results of our approach with machine learning approaches successfully used in the same context [28, 98] and involving other types of labels. Finally, we plan to investigate the usage of historical analysis to provide orthogonal/complementary information that could be combined with the adopted textual features.

Differently from issue driven development, in pull-based development developers use branches to make the desired changes independently, and then create a pull request to ask merging their changes into the main repository [99]. The integrators (usually the members of the project's core team) are asked to reply to such request, evaluating the quality of the contributions, and eventually merging or rejecting the changes [100]. Manually identifying high-quality and desirable pull requests may be challenging [101], especially for popular projects, where tens of pull requests are daily submitted [102, 8]. In the future we plan to verify whether the reasons behind the rejection of specific kinds of pull requests are similar to the ones that have been identified for wontfix issues, with the purpose of better comprehending the team behaviors when managing external requests.

Acknowledgment

Sebastiano Panichella gratefully acknowledges the Horizon 2020 (EU Commission) support for the project *COSMOS* (DevOps for Complex Cyber-physical Systems), Project No. 957254-COSMOS.

References

- [1] M. M. Lehman, Programs, life cycles, and laws of software evolution, Proc. IEEE 68 (9) (1980) 1060–1076.
- [2] A. Di Sorbo, S. Panichella, C. V. Alexandru, J. Shimagaki, C. A. Visaggio, G. Canfora, H. C. Gall, What would users change in my app? summarizing app reviews for recommending software changes, in: 24th International Symposium on Foundations of Software Engineering, FSE 2016, Seattle, WA, USA, November, 2016, 2016, pp. 499–510.
- [3] J. Anvik, L. Hiew, G. C. Murphy, Coping with an open bug repository, in: Proceedings of the 2005 OOPSLA Workshop on Eclipse Technology eXchange, eclipse '05, ACM, New York, NY, USA, 2005, pp. 35–39.

- [4] J. Anvik, L. Hiew, G. C. Murphy, Who should fix this bug?, in: 28th International Conference on Software Engineering (ICSE 2006), Shanghai, China, May 20-28, 2006, 2006, pp. 361–370.
- [5] T. F. Bissyandé, D. Lo, L. Jiang, L. Réveillère, J. Klein, Y. L. Traon, Got issues? who cares about it? A large scale investigation of issue trackers from github, in: IEEE 24th International Symposium on Software Reliability Engineering, ISSRE 2013, Pasadena, CA, USA, November 4-7, 2013, 2013, pp. 188–197.
- [6] S. Just, R. Premraj, T. Zimmermann, Towards the next generation of bug tracking systems, in: Proceedings of the 2008 IEEE Symposium on Visual Languages and Human-Centric Computing, VLHCC '08, IEEE Computer Society, Washington, DC, USA, 2008, pp. 82–85.
- [7] R. Salo, T. Poranen, Z. Zhang, Requirements management in github with a lean approach, in: Proceedings of the 14th Symposium on Programming Languages and Software Tools (SPLST'15), Tampere, Finland, October 9-10, 2015., 2015, pp. 164–178.
- [8] M. I. Azeem, S. Panichella, A. Di Sorbo, A. Serebrenik, Q. Wang, Action-based recommendation in pull-request development, in: ICSSP '20: International Conference on Software and System Processes, Seoul, Republic of Korea, 26-28 June, 2020, 2020, pp. 115–124.
- [9] S. Kim, E. J. Whitehead, Jr., How long did it take to fix bugs?, in: Proceedings of the 2006 International Workshop on Mining Software Repositories, MSR '06, ACM, New York, NY, USA, 2006, pp. 173–174.
- [10] U. Ashraf, C. Mayr-Dorn, A. Egyed, S. Panichella, A mixed graph-relational dataset of socio-technical interactions in open source systems, in: MSR '20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020, 2020, pp. 538–542.
- [11] D. Bertram, A. Voids, S. Greenberg, R. Walker, Communication, collaboration, and bugs: The social nature of issue tracking in small, collocated teams, in: Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10, ACM, New York, NY, USA, 2010, pp. 291–300.
- [12] P. Hooimeijer, W. Weimer, Modeling bug report quality, in: Proceedings of the Twenty-second IEEE/ACM International Conference on Automated Software Engineering, ASE '07, ACM, New York, NY, USA, 2007, pp. 34–43.
- [13] N. Bettenburg, S. Just, A. Schröter, C. Weiss, R. Premraj, T. Zimmermann, What makes a good bug report?, in: Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering, SIGSOFT '08/FSE-16, ACM, New York, NY, USA, 2008, pp. 308–318.
- [14] A. Lamkanfi, S. Demeyer, Q. D. Soetens, T. Verdonck, Comparing mining algorithms for predicting the severity of a reported bug, in: 2011 15th European Conference on Software Maintenance and Reengineering, 2011, pp. 249–258.
- [15] Y. Tian, D. Lo, C. Sun, Drone: Predicting priority of reported bugs by multi-factor analysis, in: Proceedings of the 2013 IEEE International Conference on Software Maintenance, ICSM '13, 2013, pp. 200–209.
- [16] J. Aranda, G. Venolia, The secret life of bugs: Going past the errors and omissions in software repositories, in: Proceedings of the 31st International Conference on Software Engineering, ICSE '09, IEEE Computer Society, Washington, DC, USA, 2009, pp. 298–308.
- [17] X. Wang, L. Zhang, T. Xie, J. Anvik, J. Sun, An approach to detecting duplicate bug reports using natural language and execution information, in: 30th International Conference on Software Engineering (ICSE 2008), Leipzig, Germany, May 10-18, 2008, 2008, pp. 461–470.
- [18] O. Baysal, M. W. Godfrey, R. Cohen, A bug you like: A framework for automated assignment of bugs, in: The 17th IEEE International Conference on Program Comprehension, ICPC 2009, Vancouver, British Columbia, Canada, May 17-19, 2009, 2009, pp. 297–298.
- [19] Y. Zhou, Y. Su, T. Chen, Z. Huang, H. C. Gall, S. Panichella, User review-based change file localization for mobile applications, IEEE Transactions on Software Engineering (2020) 1–1.
- [20] J. Uddin, R. Ghazali, M. M. Deris, R. Naseem, H. Shah, A survey on bug prioritization, Artif. Intell. Rev. 47 (2) (2017) 145–180.
- [21] G. Antoniol, K. Ayari, M. Di Penta, F. Khomh, Y.-G. Guéhéneuc, Is it a bug or an enhancement?: A text-based approach to classify change requests, in: Proceedings of the 2008 Conference of the Center for Advanced Studies on Collaborative Research: Meeting of Minds, CASCON '08, ACM, New York, NY, USA, 2008, pp. 23:304–23:318.
- [22] K. Herzig, S. Just, A. Zeller, It's not a bug, it's a feature: how misclassification impacts bug prediction, in: 35th International Conference on Software Engineering, ICSE '13, San Francisco, CA, USA, May 18-26, 2013, 2013, pp. 392–401.
- [23] Bug status wontfix (2016).
URL https://en.opensuse.org/Bug_Status_WONTFIX
- [24] P. J. Guo, T. Zimmermann, N. Nagappan, B. Murphy, Characterizing and predicting which bugs get fixed: An empirical study of microsoft windows, in: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1, ICSE '10, ACM, New York, NY, USA, 2010, pp. 495–504.
- [25] Q. Wang, Why is my bug wontfix?, in: 2020 IEEE 2nd International Workshop on Intelligent Bug Fixing (IBF), IEEE, 2020, pp. 45–54.
- [26] S. L. Ramirez-Mora, H. Oktaba, H. Gómez-Adorno, Descriptions of issues and comments for predicting issue success in software projects, J. Syst. Softw. 168 (2020) 110663.
- [27] R. K. Saha, S. Khurshid, D. E. Perry, An empirical study of long lived bugs, in: 2014 Software Evolution Week - IEEE Conference on Software Maintenance, Reengineering, and Reverse Engineering, CSMR-WCRE 2014, Antwerp, Belgium, February 3-6, 2014, 2014, pp. 144–153.
- [28] R. Kallis, A. Di Sorbo, G. Canfora, S. Panichella, Ticket tagger: Machine learning driven issue classification, in: 2019 IEEE International Conference on Software Maintenance and Evolution, ICSME 2019, Cleveland, OH, USA, September 29 - October 4, 2019, 2019, pp. 406–409.
- [29] A. Murgia, M. Ortu, P. Tourani, B. Adams, S. Demeyer, An exploratory qualitative and quantitative analysis of emotions in issue report comments of open source systems, Empirical Software Engineering 23 (1) (2018) 521–564.
- [30] J. Cabot, J. L. C. Izquierdo, V. Cosentino, B. Rolandi, Exploring the use of labels to categorize issues in open-source software projects, in: 22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering, SANER 2015, Montreal, QC, Canada, March 2-6, 2015, 2015, pp. 550–554.
- [31] Z. Liao, D. He, Z. Chen, X. Fan, Y. Zhang, S. Liu, Exploring the characteristics of issue-related behaviors in github using visualization techniques, IEEE Access 6 (2018) 24003–24015.
- [32] S. Panichella, Supporting newcomers in software development projects, in: R. Koschke, J. Krinke, M. P. Robillard (Eds.), 2015 IEEE International Conference on Software Maintenance and Evolution, ICSME 2015, Bremen, Germany, September 29 - October 1, 2015, IEEE Computer Society, 2015, pp. 586–589.
- [33] S. Breu, R. Premraj, J. Sillito, T. Zimmermann, Information needs in bug reports: Improving cooperation between developers and users, in: Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, CSCW '10, ACM, New York, NY, USA, 2010, pp. 301–310.
- [34] Y. Yu, Z. Li, G. Yin, T. Wang, H. Wang, A dataset of duplicate pull-requests in github, in: Proceedings of the 15th International Conference on Mining Software Repositories, MSR '18, ACM, New York, NY, USA, 2018, pp. 22–25.
- [35] S. Panichella, G. Canfora, M. D. Penta, R. Oliveto, How the evolution of emerging collaborations relates to code changes: an empirical study, in: C. K. Roy, A. Begel, L. Moonen (Eds.), 22nd International Conference on Program Comprehension, ICPC 2014, Hyderabad, India, June 2-3, 2014, ACM, 2014, pp. 177–188.
- [36] D. Ford, M. Behroozi, A. Serebrenik, C. Parnin, Beyond the code itself: how programmers *really* look at pull requests, in: Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Society, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019, 2019, pp. 51–60.
- [37] P. J. Guo, T. Zimmermann, N. Nagappan, B. Murphy, "not my bug!" and other reasons for software bug report reassignments, in: Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, CSCW '11, 2011, pp. 395–404.
- [38] D. Kavalier, S. Sirovica, V. Hellendoorn, R. Aranovich, V. Filkov, Perceived language complexity in github issue discussions and their effect on issue resolution, in: Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering, ASE 2017, Urbana, IL, USA, October 30 - November 03, 2017, 2017, pp. 72–83.

- [39] G. Destefanis, M. Ortu, D. Bowes, M. Marchesi, R. Tonelli, On measuring affects of github issues' commenters, in: Proceedings of the 3rd International Workshop on Emotion Awareness in Software Engineering, SEMotion '18, ACM, New York, NY, USA, 2018, pp. 14–19.
- [40] P. Rodeghero, D. Huo, T. Ding, C. McMillan, M. Gethers, An empirical study on how expert knowledge affects bug reports, *Journal of Software: Evolution and Process* 28 (7) (2016) 542–564.
- [41] M. Ortu, G. Destefanis, M. Kassab, S. Counsell, M. Marchesi, R. Tonelli, Would you mind fixing this issue? - an empirical analysis of politeness and attractiveness in software developed using agile boards, in: Agile Processes, in Software Engineering, and Extreme Programming - 16th International Conference, XP 2015, Helsinki, Finland, May 25–29, 2015, Proceedings, 2015, pp. 129–140.
- [42] M. Ortu, B. Adams, G. Destefanis, P. Tourani, M. Marchesi, R. Tonelli, Are bullies more productive? empirical study of affectiveness vs. issue fixing time, in: 12th IEEE/ACM Working Conference on Mining Software Repositories, MSR 2015, Florence, Italy, May 16–17, 2015, 2015, pp. 303–313.
- [43] R. Kikas, M. Dumas, D. Pfahl, Issue dynamics in github projects, in: Product-Focused Software Process Improvement - 16th International Conference, PROFES 2015, Bolzano, Italy, December 2–4, 2015, Proceedings, 2015, pp. 295–310.
- [44] T. Zhang, G. Yang, B. Lee, A. T. S. Chan, Guiding bug triage through developer analysis in bug reports, *International Journal of Software Engineering and Knowledge Engineering* 26 (3) (2016) 405–432.
- [45] J. Xuan, H. Jiang, Z. Ren, J. Yan, Z. Luo, Automatic bug triage using semi-supervised text classification, CoRR abs/1704.04769. arXiv: 1704.04769.
- [46] V. Cosentino, J. L. C. Izquierdo, J. Cabot, A systematic mapping study of software development with github, *IEEE Access* 5 (2017) 7173–7192.
- [47] W. W. Cohen, V. R. Carvalho, T. M. Mitchell, Learning to classify email into "speech acts", in: Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing, EMNLP 2004, A meeting of SIGDAT, a Special Interest Group of the ACL, held in conjunction with ACL 2004, 25–26 July 2004, Barcelona, Spain, 2004, pp. 309–316.
- [48] P. Bhattacharya, I. Neamtiu, Bug-fix time prediction models: Can we do better?, in: Proceedings of the 8th Working Conference on Mining Software Repositories, MSR '11, ACM, New York, NY, USA, 2011, pp. 207–210.
- [49] A. Bacchelli, T. D. Sasso, M. D'Ambros, M. Lanza, Content classification of development emails, in: 34th International Conference on Software Engineering, ICSE 2012, June 2–9, 2012, Zurich, Switzerland, 2012, pp. 375–385.
- [50] A. Di Sorbo, S. Panichella, C. A. Visaggio, M. Di Penta, G. Canfora, H. C. Gall, Development emails content analyzer: Intention mining in developer discussions (T), in: 30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA, November 9–13, 2015, 2015, pp. 12–23.
- [51] L. Moreno, A. Marcus, Automatic software summarization: the state of the art, in: Proceedings of the 40th International Conference on Software Engineering: Companion Proceedings, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018, 2018, pp. 530–531.
- [52] S. Panichella, Summarization techniques for code, change, testing, and user feedback (invited paper), in: 2018 IEEE Workshop on Validation, Analysis and Evolution of Software Tests, VST@SANER 2018, Campobasso, Italy, March 20, 2018, 2018, pp. 1–5.
- [53] S. Rastkar, G. C. Murphy, G. Murray, Summarizing software artifacts: a case study of bug reports, in: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1, ICSE 2010, Cape Town, South Africa, 1–8 May 2010, 2010, pp. 505–514.
- [54] Y. Zhou, Y. Tong, R. Gu, H. Gall, Combining text mining and data mining for bug report classification, in: Proceedings of the 2014 IEEE International Conference on Software Maintenance and Evolution, ICSME '14, IEEE Computer Society, Washington, DC, USA, 2014, pp. 311–320.
- [55] X. Xia, D. Lo, X. Wang, X. Yang, S. Li, J. Sun, A comparative study of supervised learning algorithms for re-opened bug prediction, in: Proceedings of the 2013 17th European Conference on Software Maintenance and Reengineering, CSMR '13, 2013, pp. 331–334.
- [56] K. K. Chaturvedi, V. B. Singh, An empirical comparison of machine learning techniques in predicting the bug severity of open and closed source projects, *IJOSSP* 4 (2) (2012) 32–59.
- [57] W. Liu, S. Wang, X. Chen, H. Jiang, Predicting the severity of bug reports based on feature selection, *International Journal of Software Engineering and Knowledge Engineering* 28 (4) (2018) 537.
- [58] Y. Tian, D. Lo, C. Sun, Information retrieval based nearest neighbor classification for fine-grained bug severity prediction, in: Proceedings of the 2012 19th Working Conference on Reverse Engineering, WCRE '12, IEEE Computer Society, Washington, DC, USA, 2012, pp. 215–224.
- [59] R. Kikas, M. Dumas, D. Pfahl, Using dynamic and contextual features to predict issue lifetime in github projects, in: Proceedings of the 13th International Conference on Mining Software Repositories, MSR '16, ACM, New York, NY, USA, 2016, pp. 291–302.
- [60] C. Weiß, R. Premraj, T. Zimmermann, A. Zeller, How long will it take to fix this bug?, in: Fourth International Workshop on Mining Software Repositories, MSR 2007 (ICSE Workshop), Minneapolis, MN, USA, May 19–20, 2007, Proceedings, 2007, p. 1.
- [61] H. Zhang, L. Gong, S. Versteeg, Predicting bug-fixing time: An empirical study of commercial software projects, in: Proceedings of the 2013 International Conference on Software Engineering, ICSE '13, IEEE Press, Piscataway, NJ, USA, 2013, pp. 1042–1051.
- [62] M. Fazzini, M. Prammer, M. d'Amorim, A. Orso, Automatically translating bug reports into test cases for mobile apps, in: Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2018, ACM, New York, NY, USA, 2018, pp. 141–152.
- [63] C. Liu, J. Yang, L. Tan, M. Hafiz, R2fix: Automatically generating bug fixes from bug reports, in: 2013 IEEE Sixth International Conference on Software Testing, Verification and Validation, 2013, pp. 282–291.
- [64] R. Almhana, W. Mkaouer, M. Kessentini, A. Ouni, Recommending relevant classes for bug reports using multi-objective search, in: Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE 2016, 2016, pp. 286–295.
- [65] M. Patton, *Qualitative Evaluation and Research Methods*, Sage, Newbury Park, 2002.
- [66] R. Padhye, S. Mani, V. S. Sinha, A study of external community contribution to open-source projects on github, in: 11th Working Conference on Mining Software Repositories, MSR 2014, Proceedings, May 31 - June 1, 2014, Hyderabad, India, 2014, pp. 332–335.
- [67] L. Heppeler, R. Eckert, M. Stuermer, Who cares about my feature request?, in: Open Source Systems: Integrating Communities - 12th IFIP WG 2.13 International Conference, OSS 2016, Gothenburg, Sweden, May 30 - June 2, 2016, Proceedings, 2016, pp. 85–96.
- [68] J. Dalle, M. den Besten, H. Masmoudi, Channeling firefox developers: Mom and dad aren't happy yet, in: Open Source Development, Communities and Quality, IFIP 20th World Computer Congress, Working Group 2.3 on Open Source Software, OSS 2008, September 7–10, 2008, Milano, Italy, 2008, pp. 265–271.
- [69] Y. Fan, X. Xia, D. Lo, A. E. Hassan, Chaff from the wheat: Characterizing and determining valid bug reports, *IEEE Trans. Software Eng.* 46 (5) (2020) 495–525.
- [70] A. Zerouali, T. Mens, G. Robles, J. M. González-Barahona, On the diversity of software package popularity metrics: An empirical study of npm, in: 26th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2019, Hangzhou, China, February 24–27, 2019, 2019, pp. 589–593.
- [71] H. Borges, M. T. Valente, What's in a github star? understanding repository starring practices in a social coding platform, *Journal of Systems and Software* 146 (2018) 112–129.
- [72] S. Nielebock, R. Heumüller, F. Ortmeier, Programmers do not favor lambda expressions for concurrent object-oriented code, *Empirical Software Engineering* doi: 10.1007/s10664-018-9622-9.
- [73] J. Jiang, D. Lo, X. Ma, F. Feng, L. Zhang, Understanding inactive yet available assignees in github, *Information & Software Technology* 91 (2017) 44–55.
- [74] D. Mazinianian, A. Ketkar, N. Tsantalos, D. Dig, Understanding the use of lambda expressions in java, *PACMPL* 1 (OOPSLA) (2017) 85:1–85:31.

- [75] M. Hilton, T. Tunnell, K. Huang, D. Marinov, D. Dig, Usage, costs, and benefits of continuous integration in open-source projects, in: Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE 2016, ACM, New York, NY, USA, 2016, pp. 426–437.
- [76] D. Spencer, Card sorting: Designing usable categories, Rosenfeld Media, 2009.
- [77] A. Guzzi, A. Bacchelli, M. Lanza, M. Pinzger, A. van Deursen, Communication in open source software development mailing lists, in: Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13, San Francisco, CA, USA, May 18-19, 2013, 2013, pp. 277–286.
- [78] H. Khalid, E. Shihab, M. Nagappan, A. E. Hassan, What do mobile app users complain about?, *IEEE Software* 32 (3) (2015) 70–77.
- [79] W. Conover, Practical nonparametric statistics, Wiley series in probability and statistics: Applied probability and statistics, Wiley, 1998.
- [80] S. S. Shapiro, M. B. Wilk, An analysis of variance test for normality (complete samples), *Biometrika* 52 (3/4) (1965) 591–611.
- [81] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing, *Journal of the Royal statistical society: series B (Methodological)* 57 (1) (1995) 289–300.
- [82] M. Linares-Vásquez, G. Bavota, C. Bernal-Cárdenas, M. Di Penta, R. Oliveto, D. Poshyvanyk, Api change and fault proneness: A threat to the success of android apps, in: Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2013, ACM, New York, NY, USA, 2013, pp. 477–487.
- [83] R. A. Baeza-Yates, B. Ribeiro-Neto, Modern Information Retrieval, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [84] Y. Zhou, Y. Tong, R. Gu, H. C. Gall, Combining text mining and data mining for bug report classification, *Journal of Software: Evolution and Process* 28 (3) (2016) 150–176.
- [85] E. R. Russo, A. Di Sorbo, C. A. Visaggio, G. Canfora, Summarizing vulnerabilities' descriptions to support experts during vulnerability assessment activities, *J. Syst. Softw.* 156 (2019) 84–99.
- [86] S. Panichella, A. Di Sorbo, E. Guzman, C. A. Visaggio, G. Canfora, H. C. Gall, How can i improve my app? classifying user reviews for software maintenance and evolution, in: 2015 IEEE International Conference on Software Maintenance and Evolution, ICSME 2015, Bremen, Germany, September 29 - October 1, 2015, 2015, pp. 281–290.
- [87] A. Di Sorbo, G. Grano, C. A. Visaggio, S. Panichella, Investigating the criticality of user-reported issues through their relations with app rating, *J. Softw. Evol. Process.* 33 (3).
- [88] V. R. Basili, L. C. Briand, W. L. Melo, A validation of object-oriented design metrics as quality indicators, *IEEE Trans. Software Eng.* 22 (10) (1996) 751–761.
- [89] T. Zimmermann, N. Nagappan, Predicting defects with program dependencies, in: Proceedings of the Third International Symposium on Empirical Software Engineering and Measurement, ESEM 2009, October 15-16, 2009, Lake Buena Vista, Florida, USA, 2009, pp. 435–438.
- [90] X. Tan, M. Zhou, Z. Sun, A first look at good first issues on github, in: P. Devanbu, M. B. Cohen, T. Zimmermann (Eds.), ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020, ACM, 2020, pp. 398–409.
- [91] A. B. Dhasade, A. S. M. Venigalla, S. Chimalakonda, Towards prioritizing github issues, in: S. Jain, A. Gupta, D. Lo, D. Saha, R. Sharma (Eds.), ISEC 2020: 13th Innovations in Software Engineering Conference, Jabalpur, India, February 27-29, 2020, ACM, 2020, pp. 18:1–18:5.
- [92] G. C. Cawley, N. L. Talbot, On over-fitting in model selection and subsequent selection bias in performance evaluation, *J. Mach. Learn. Res.* 11 (2010) 2079–2107.
- [93] J. R. Quinlan, Induction of decision trees, *Machine learning* 1 (1) (1986) 81–106.
- [94] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, D. Damian, The promises and perils of mining github, in: Proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, ACM, New York, NY, USA, 2014, pp. 92–101.
- [95] S. Panichella, G. Bavota, M. D. Penta, G. Canfora, G. Antoniol, How developers' collaborations identified from different sources tell us about code changes, in: 30th IEEE International Conference on Software Maintenance and Evolution, Victoria, BC, Canada, September 29 - October 3, 2014, 2014, pp. 251–260.
- [96] Q. Mi, J. Keung, An empirical analysis of reopened bugs based on open source projects, in: Proceedings of the 20th International Conference on Evaluation and Assessment in Software Engineering, EASE '16, ACM, New York, NY, USA, 2016, pp. 37:1–37:10.
- [97] T. Zimmermann, N. Nagappan, P. J. Guo, B. Murphy, Characterizing and predicting which bugs get reopened, in: Proceedings of the 34th International Conference on Software Engineering, ICSE '12, IEEE Press, Piscataway, NJ, USA, 2012, pp. 1074–1083.
- [98] R. Kallis, A. Di Sorbo, G. Canfora, S. Panichella, Predicting issue types on github, *Science of Computer Programming* 205 (2021) 102598.
- [99] G. Gousios, M. D. Storey, A. Bacchelli, Work practices and challenges in pull-based development: the contributor's perspective, in: Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016, 2016, pp. 285–296.
- [100] G. Gousios, A. Zaidman, M. D. Storey, A. van Deursen, Work practices and challenges in pull-based development: The integrator's perspective, in: 37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 1, 2015, pp. 358–368.
- [101] C. Yang, X. Zhang, L. Zeng, Q. Fan, G. Yin, H. Wang, An empirical study of reviewer recommendation in pull-based development model, in: Proceedings of the 9th Asia-Pacific Symposium on Internetware, Internetware 2017, Shanghai, China, September 23 - 23, 2017, 2017, pp. 14:1–14:6.
- [102] Y. Yu, H. Wang, G. Yin, C. X. Ling, Reviewer recommender of pull-requests in github, in: 30th IEEE International Conference on Software Maintenance and Evolution, Victoria, BC, Canada, September 29 - October 3, 2014, 2014, pp. 609–612.