

Beyond Textual Issues: Understanding the Usage and Impact of GitHub Reactions

Hudson Borges

Faculty of Computer Science (FACOM)
Federal University of Mato Grosso do Sul (UFMS)
Campo Grande, Mato Grosso do Sul, Brazil
hsborges@facom.ufms.br

Rodrigo Brito and Marco Tulio Valente

Department of Computer Science (DCC)
Federal University of Minas Gerais (UFMG)
Belo Horizonte, Minas Gerais, Brazil
{britorodrigo,mtov}@dcc.ufmg.br

ABSTRACT

Recently, GitHub introduced a new social feature, named *reactions*, which are “pictorial characters” similar to *emoji* symbols widely used nowadays in text-based communications. Particularly, GitHub users can use a pre-defined set of such symbols to react to issues and pull requests. However, little is known about the real usage and impact of GitHub reactions. In this paper, we analyze the reactions provided by developers to more than 2.5 million issues and 9.7 million issue comments, in order to answer an extensive list of nine research questions about the usage and adoption of reactions. We show that reactions are being increasingly used by open source developers. Moreover, we also found that issues with reactions usually take more time to be handled and have longer discussions.

KEYWORDS

Social Coding, GitHub, Reactions

1 INTRODUCTION

Over the last years, open source software development has become more social and collaborative. In fact, social coding fosters formal and informal collaboration in software development by empowering knowledge exchange between developers [8]. In this context, traditional code management tools (e.g., control version systems, issue tracking, etc.) have been extended with social networks features to connect developers and increase their social interactions. Among the social coding platforms, GitHub is nowadays the most popular one and the world’s largest collection of open source software, with around 36 million users and 104 million repositories.¹

Recently, GitHub introduced a new social feature, named *reactions*, which are “pictorial characters” similar to *emoji* symbols widely used nowadays in text-based communications. Indeed, emojis are increasingly being used in many areas and are a popular mechanism for developers express their emotions. However, little is known about the real usage and impact of GitHub reactions on software development practices.

¹<https://github.com/search>, verified on 01/20/2019.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SBES 2019, September 23–27, 2019, Salvador, Brazil

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-7651-8/19/09...\$15.00

<https://doi.org/10.1145/3350768.3350788>

In this paper, we investigate how developers are using the reactions feature introduced by GitHub. To this purpose, we collect and analyze reactions to more than 2.5 million GitHub issues and 9.7 million comments associated to these issues. In particular, we ask nine research questions, which are grouped into two major themes:

Theme 1: Usage of reactions

RQ #1: *Do developers use reactions?*

RQ #2: *What are the most common reactions?*

RQ #3: *How is the usage of reactions evolving?*

RQ #4: *Do popular projects have more reactions?*

RQ #5: *Do certain types of issues have more reactions?*

Theme 2: Impact of Reactions

RQ #6: *Do issues with more reactions get resolved faster?*

RQ #7: *Do issues with more reactions have more discussion?*

RQ #8: *Do negative reactions inhibit further participation?*

RQ #9: *Do reactions reduce the noise in issue discussions?*

We show that more than 28% of the issues threads (i.e., an issue and its respective comments) receive at least one reaction. Moreover, the usage of reactions is growing fast in the last two years. *Thumbs up* is the most common reaction, with almost 80% of all reactions. By contrast, *Confused* is the least common one, with just 1.7%. Also, we show that popularity does not have a major influence on the reactions whereas programming language and issue type do have.

Regarding the impact on software development, we found that issues with reactions take more time to be handled and have longer discussions. Moreover, negative reactions do not inhibit the participation of developers in further discussions. Finally, we show that reactions are contributing to reduce trivial comments (i.e., those containing only an emoji).

To our knowledge, we are the first to study GitHub reactions. Particularly, our contributions are twofold: (1) a thorough quantitative analysis on the usage and impact of GitHub reactions, focusing on their usage in issues; (2) a list of practical recommendations for using reactions, both for project owners and users. With these recommendations, we hope to demystify the use of reactions. On the one hand, we provide empirical data to convince developers to use reactions in recommended scenarios; on the other hand, we prevent the set up of unrealistic expectations regarding the benefits of reactions. Finally, we are making public the dataset used in this paper, with reactions applied to more than 2.5 million issues and 9.7 million comments.

The remainder of this paper is organized as follows. In Section 2, we present background information about GitHub reactions. Next, in Section 3, we describe the design of the study presented in this paper. Then, in Section 4 we answer the proposed research questions. Our key findings are described in Section 5 and threats to

validity are discussed in Section 6. Finally, we present related work in Section 7 and we conclude the paper in Section 8.

2 BACKGROUND

GitHub implements a vast number of social features. For example, developers can follow other developers to receive notifications about events on the platform. Moreover, developers can fork projects to create a copy they can freely modify without affecting the original code [11]. Finally, GitHub provides a social feature named *stars* [3], which is similar to a *like* in social media sites (e.g., Facebook, Twitter, HackerNews, Reddit, etc.).

In 2016, GitHub introduced a new social feature named *Reactions* to allow developers quickly express their feelings on issues and pull requests (Figure 1). GitHub designers argued there are complex and detailed comments that require textual messages, but sometimes developers can use a simple *emoji* to express their feelings on someone’s comment. Besides that, they argue that reacting to a comment avoids long threads with many comments lacking relevant content.

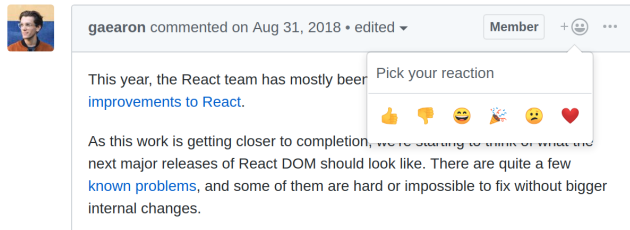


Figure 1: Reacting to a comment on GitHub

GitHub provides six different reactions for use in issues and pull requests, which are showed in Table 1 and described next.²

- Thumbs up*, which is used to express acceptance or approval.
- Thumbs down*, which is used to express rejection or disapproval.
- Laugh*, which indicates a fun situation or happiness. It is also used to express sarcasm or irony in negative situations.
- Hooray*, which is also used to celebrate or express a happy situation. For example, it is commonly used in software releases.
- Confused*, which is proposed to express doubt or some difficulty to understand a specific subject.
- Heart*, which indicates a strong acceptance or approval; it is used when the user really likes something.

Table 1: GitHub reactions

Name	Symbol	Name	Symbol
Thumbs up	👍	Hooray	🎉
Thumbs down	👎	Confused	😞
Laugh	😂	Heart	❤️

²<https://developer.github.com/v3/reactions>, verified on 01/20/2019.

Developers can add only one reaction per textual element. Also, reactions can be added to closed issues and pull requests. However, project owners can lock issues or pull requests to disable both comments and reactions.

3 STUDY DESIGN

In this study, we restrict our analysis to reactions added to issues and their respective comments. Although pull requests can also receive reactions from developers, they typically have different purposes and characteristics, which should be better explored in further work. In the next paragraphs, we describe the criteria to select the projects, issues, and reactions studied in this paper. We also characterize our final dataset.

3.1 Initial projects sampling

To study how developers apply reactions to GitHub issues, we first defined an initial sample of projects. This sample includes the top-5,000 repositories with more stars on GitHub (stars are considered a reliable proxy for the popularity of GitHub projects [3, 4]). To obtain the metadata about these projects, we used GitHub GraphQL API.³ Here, we used the *search* function to navigate through the first 5,000 repositories by stars. For each repository, we retrieved basic metadata, such as programming language, creation date, last update date, and owner.

3.2 Issues collection

After collecting the initial sample of projects, we again used GitHub’s GraphQL API to obtain the issues created in these projects after March 13rd, 2016. We restricted the issues collection to this date because it was when the reaction feature was officially announced.⁴ For each issue, we collected the date when it was created, updated, and closed. In addition, we also collected data about the user who created the issue and his/her role in the project (e.g., collaborator, contributor, owner, or none). Finally, we retrieved all the events in the issue timeline. For example, an event can be a comment, a reference to a commit or another issue, and the action of closing or reopening the issue. Finally, we also collected data about the reactions of each issue and their respective comments. This data includes: (i) the user who reacted; (ii) the reaction (as detailed in Section 2); and (iii) the reaction date.

After the data collection, we found that some projects do not have issues created after March 13rd, 2016. Specifically, we found that 159 out of 5,000 projects (3.18%) have no issues after this date. Probably, they are deprecated or feature-completed projects. Thus, we excluded these projects from this study. In total, we collected 2,544,304 issues and 9,775,928 comments, from 4,841 GitHub projects.

3.3 Final dataset characterization

Our final dataset consists of 4,841 public GitHub projects. Regarding their popularity, the number of stars ranges from 296K stars (FREECODECAMP/FREECODECAMP) to 3,038 stars (SINDRESORHUS/SPEED-TEST). Figure 2 shows the distribution of the number of

³<https://developer.github.com/v4>, verified on 01/20/2019

⁴<https://blog.github.com/2016-03-10-add-reactions-to-pull-requests-issues-and-comments>, verified on 01/20/2019.

stars of these projects. Other popular projects in our dataset include TWBS/BOOTSTRAP (129K stars), VUEJS/VUE (123K stars), and FACEBOOK/REACT (118K stars). In summary, our study is based on relevant and widely used open source projects.




Figure 2: Popularity of the analyzed repositories

Considering the primary programming language of these projects, JavaScript has the highest number of repositories (1,431 repositories), followed by Java (470 repositories) and Python (456 repositories); see details in Figure 3. The prevalence of JavaScript repositories is common on studies based on GitHub projects [3]. Moreover, 366 projects have no programming language, i.e., they are documentation projects, such as books, tutorials, awesome-lists, etc. In total, the projects in our dataset are implemented in 70 different programming languages. Therefore, although JavaScript dominates, we also have a variety of languages, which allows us to study reactions in several of ecosystems.




Figure 3: Top-10 programming language with more repositories

In GitHub, issues have a textual description, which report a bug, feature, enhancement, etc. Project members then discuss the reported issue, by means of comments. Figure 4 shows the distribution of the number of issues and comments in our dataset. The repository with the highest number of issues is MICROSOFT/VSCODE (57K issues), followed by KUBERNETES/KUBERNETES (19.8K issues) and MICROSOFT/TYPESCRIPT (14.9K issues). Regarding the number of comments, MICROSOFT/VSCODE also has the highest result (210K comments), followed by KUBERNETES/KUBERNETES (196K comments) and GOLANG/GO (89.4K comments). On average, each issue has 3.84 comments.




Figure 4: Number of issues and comments per repository

4 RESULTS

Theme 1: Usage of reactions

RQ #1: Do developers use reactions?

In this initial research question, our goal is to quantify the usage of reactions in GitHub issues. To this purpose, we first compute the number of reactions in our dataset. We found that most GitHub issues (87.0%) and comments (89.3%) do not receive reactions. However, when considering whole issues threads, i.e., the initial document that describes an issue and the associated comments, we found that 28.4% have at least one reaction. Figure 5 shows the number of reactions in issues (i.e., on the initial document that describes an issue) and comments (i.e., on further messages that discuss the issue), and in whole issue threads. Furthermore, in the box plot we only consider elements with at least one reaction. The 1st, median, and 3rd quartiles for reactions/issue and reactions/thread are 1, 2, and 4. Considering reactions/comments, the same measures are 1, 1, and 2.




Figure 5: Number of reactions per issue, comment, and thread

Figure 6 shows the most reacted issue in our dataset. This issue, titled “Allow for floating windows”, belongs to MICROSOFT/VSCODE, has 3,654 reactions, and it is opened since August 4, 2016. It is a request to add floating windows in the Visual Studio Code editor. Despite the high number of reactions, the issue has not been fully implemented by the project maintainers. In terms of comments, the most reacted one is issue “#980 – WebAssembly logo voting”, from WEBASSEMBLY/DESIGN. This issue was created with the goal

of selecting a new logo for the project. The candidate logos were presented in the comments and the most reacted one was chosen. This "winner" comment received 1,770 reactions (1,486 *Thumbs up*, 120 *Thumbs down*, 11 *Laugh*, 46 *Hooray*, 34 *Confused*, and 73 *Heart*).

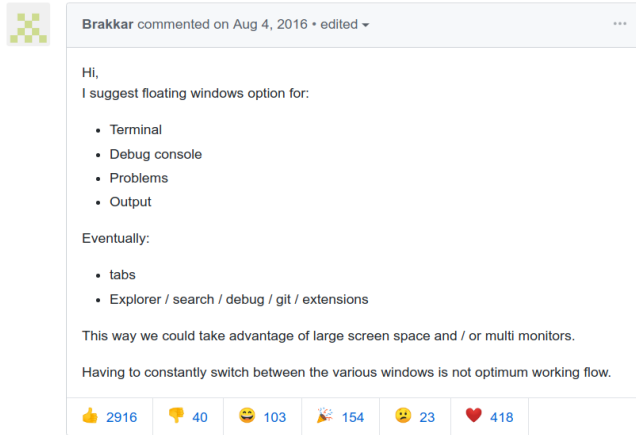


Figure 6: Issue “#10121 – Allow for floating windows” in MICROSOFT/VSCODE. Reactions appear in the bottom.

Finally, we analyzed the usage of reactions by programming language. Here, we computed the total number of reactions applied to a repository (i.e., counting reactions added to issues and comments) and compared the distributions, by programming language. Figure 7 presents the distributions of the top-10 languages with more reactions by repository. We observe that projects implemented in C++ receive more reactions than projects in other languages. By contrast, Objective-C projects are the ones with the lowest number of reactions. By applying the Kruskal-Wallis test to compare multiple samples, we found that these distributions differ in at least one language ($p\text{-value} < 0.001$). Then, we use a non-parametric, pairwise and multiple comparisons test (Dunn’s test) to isolate the languages that differ from the others. Indeed, we found that the distribution of C++ projects is different from all other distributions. Additionally, the distribution of Objective-C projects differ from all others, except HTML. Our hypothesis is that Objective-C projects have less reactions because iOS development is moving to Swift. Therefore, Objective-C projects are gradually becoming deprecated or obsolete.

Summary: Most issues (87%) and comments (89%) do not have reactions. By contrast, there are issues with 3,654 reactions. For whole threads (issues + comments), we found that 28.4% have at least one reaction.

RQ #2: What are the most common reactions?

This second research question aims to reveal the most common reactions and their roles in GitHub issues. Table 2 shows the results. *Thumbs up* is the most common reaction (78.7%), followed by *Hooray* and *Heart*, both with just 5.9%. *Confused* is the least common reaction, with just 1.7%. Table 2 also presents information




Figure 7: Number of reactions per repository

about the issue with the highest number of reactions. For *Thumbs up*, the highest value is found in the issue described in RQ #1 (Figure 6). For *Hooray* and *Heart*, the highest values are 1,031 and 513 reactions, respectively, which are found in the issue “*React 16 RC*”⁵ from FACEBOOK/REACT. The reasons for the high acceptance of this issue includes a complete rework of the framework’s API and the addition of long-standing feature requests. By contrast, the issue “*Changing Redis master-slave replication terms with something else*”⁶, which was created by the repository owner, received the highest number of *Thumbs down* and *Confused* reactions. According to the comments in this issue, the proposal does not add value to the project and causes compatibility issues with previous versions. Finally, issue “*Check continuity on O2 cryogenic tanks before allowing stir*”⁷ from CHRISLGARRY/APOLLO-11 received the highest number of *Laugh* reactions. Clearly, this issue was only created to cause laughter. The repository in question archives the original source code of Apollo-11’s command and lunar modules.

Table 2: Most common GitHub reactions

Reaction	Ratio (%)	Repository	Issue	Max
Thumbs up	78.7	Microsoft/vscode	10121	2,916
Hooray	5.9	facebook/react	10294	1,031
Heart	5.9	facebook/react	10294	513
Thumbs down	5.0	antirez/redis	5335	683
Laugh	2.7	chrislgarry/Apollo-11	3	1,083
Confused	1.7	antirez/redis	5335	155

Summary: Thumbs Up is by far the most common reaction (78.7%), followed by Hooray (5.9%), Heart (5.9%), and Thumbs down (5.0%).

RQ #3: How is the usage of reactions evolving?

In this third research question, we investigate the evolution of the usage of GitHub reactions. Specifically, we compared the number of issues opened with the number of reactions since this feature

⁵<https://github.com/facebook/react/issues/10294>

⁶<https://github.com/antirez/redis/issues/5335>

⁷<https://github.com/chrislgarry/Apollo-11/issues/3>

was introduced in GitHub. Figure 8 shows the evolution of both measures. The curves reveal that reactions usage is growing fast since its introduction, although with a stabilization trend in the last months. By contrast, the number of new issues per month remains almost constant since 2016.




Figure 8: New reactions and issues per month

Finally, Figure 9 shows the growth of each reaction type. As we can see, there is no major difference on the growth pattern of existing reactions. The exception is the *Laugh* reaction, which started as the second most used reaction in 2016 (8.3%) and is now the second least used one (2.3%).




Figure 9: New reactions per month (per reactions type)

Summary: Reactions are gaining momentum; in 2018, developers used 32.5% more reactions than in 2017.

RQ #4: Do popular projects have more reactions?

In this research question, we examine whether issues from popular projects have more reactions. Figure 10 shows a scatterplot correlating the total number of reactions with the number of stars of the 4,841 repositories in our dataset. The result suggests that there is no major correlation between these metrics. We ran Spearman’s rank correlation test and the resulting correlation coefficient ρ is 0.340, with p -value < 0.001 . This ρ value denotes a low correlation.

Summary: The projects popularity (measured by GitHub stars) are weakly correlated to the number of reactions received.




Figure 10: Stars vs reactions correlation

RQ #5: Do certain types of issues have more reactions?

In this research question, we analyze how the use of reactions varies according to the type of issues (e.g., bugs, enhancements, etc.). To this purpose, we rely on the labels provided by developers when reporting issues on GitHub. In total, we found 22,856 distinct labels in 1,252,926 issues (i.e., 49.2% of the issues in our dataset). Table 3 lists the top-20 most common issues labels. As we can see, the top-3 are *bug* (14.4%), *question* (6.8%), and *enhancement* (6.4%).

Table 3: Top-20 most used labels on GitHub issues

Label	Issues	Label	Issues
bug	179,840	feature request	16,307
question	84,589	good first issue	15,524
enhancement	80,270	documentation	13,170
help wanted	42,619	wontfix	12,577
duplicate	28,850	feature-request	12,333
stale	23,788	needs more info	10,894
kind/bug	21,357	support	10,819
type: bug	21,119	verified	10,174
invalid	19,673	discussion	9,663
feature	19,120	*duplicate	9,024

However, since GitHub allows developers to create their own labels, it is common to have labels sharing the same purpose. For example, *bug*, *kind/bug*, and *type: bug* all refer to bugs. Thus, we performed an extra manual step to group labels with the same purpose. After this step, we restrict the analysis to eight main labels: *bug*, *duplicate*, *enhancement*, *invalid*, *question*, *wontfix*, *help*, and *stale*. This selection covers 68.6% of the issues with labels initially considered in this RQ.

Table 4 shows the number of issues that have these labels and the total number of reactions received by them. As mentioned before, *bug* is the most common type of issue (column *Issues*). However, the number of reactions per *bug* (0.35) is almost four time lower

than the number of reactions per *enhancements* (1.39). Interestingly, this finding suggests that GitHub users tend to react more to enhancements than to bugs.

Table 4: Issues labels

Label	Issues (I)	Reactions (R)	R/I
bug	289,386	102,044	0.35
enhancement	213,475	296,653	1.39
question	169,502	63,621	0.38
help	71,821	58,230	0.81
duplicate	48,279	12,465	0.26
stale	30,012	16,143	0.54
invalid	26,081	2,024	0.08
wontfix	13,740	6,516	0.47
Total	862,296	557,696	0.53

Figure 11 shows the percentage of *Thumbs up* (the most common reaction in our dataset) per label, considering only labels with reactions. Excluding *invalid*, almost all labels have at least 80% of *Thumbs up*, reaching the ratio of 94% for *bugs*. Finally, Figure 12 compares the usage of the remaining reactions among the selected labels. *Heart* presents a higher usage on issues broadly related to improvements, such as *question* (4.6%), *help* (5.4%), and *enhancement* (5.3%). *Invalid* issues have a high number of *Thumbs down* (14.3%), *Laugh* (7.7%), and *Confused* (5.9%) reactions, which therefore are used by developers to express their level of disagreement with the issue content. Finally, *wontfix* issues have higher levels of *Thumbs down* (4.3%) and *Heart* (5.8%), which reveals a conflicting sentiment regarding this issue type.




Figure 11: Usage of Thumbs up reaction by issue type




Figure 12: Usage of reactions by issue type

Summary: Developers react more often to *enhancements*. When we consider only *bugs* and *enhancements* with reactions, *enhancements* have 1.39 reactions (on average), against 0.35 reactions in the case of *bugs*. Moreover, there is a correlation between issue and reaction type. For example, *Thumbs down*, *Laugh* and *Confused* are more common in invalid issues.

Theme 2: Impact of Reactions

RQ #6: Do issues with more reactions get resolved faster?

In this research question, we analyze the relation between the time to close a given issue and its number of reactions. As discussed in RQ #5, on GitHub issues can represent different purposes (e.g., bug and enhancements). First, we restrict our analysis to *bugs* and *enhancements* to evaluate how reactions impact the close time of such issues. Also, we restrict this analysis to repositories with at least 102 issues (which is the median number of issues per repository in our dataset). Moreover, since the complexity of issues may vary among repositories, we compute the time as $T = T_I - T_{R(I)}$, i.e., the difference (in days) between the time required to close a given issue (T_I) and the median time to close the other issues in the same repository ($T_{R(I)}$).

Figure 13 presents the distribution of T for *bugs* and *enhancements*, for issues without reactions and for issues with two or more reactions (which is the median number of reactions in our dataset). The results suggest that issues with reactions take more time to be closed than the ones without reactions. For example, *bugs* without reactions take 2.97 more days to be closed, compared with the repository median time, whereas the ones with at least two reactions take 46.8 more days. For *enhancements*, the median values are 20.64 and 128.45 days, respectively. By applying the Mann-Whitney test, we concluded that the distributions are different (p -value < 0.001). By computing Cohen’s d , we found a *medium* effect size for *bugs* and *enhancements* ($d = 0.511$ and 0.572 , respectively).




Figure 13: Time to close bugs and enhancements (with zero or with at least two reactions). Outliers are omitted

In fact, Figure 14 presents the distribution of the time taken to close *bug* and *enhancement* issues in MICROSOFT/VSCODE, which is the repository with the highest number of reactions in our dataset.

In this specific repository, we also observe that *bugs* and *enhancements* with at least two reactions take more time to be closed compared to the ones without reactions. On median, *bugs* without reactions are closed in 3.1 days whereas those ones with at least two reactions are closed in 25.5 days. Regarding *enhancements*, the median times are 43.8 and 230.3 days, respectively.




Figure 14: Time to close bugs and enhancements in MICROSOFT/VSCODE (outliers are omitted)

Summary: Bugs and enhancements with reactions take more time to be closed. We hypothesize that they are more complex and require more time from developers to be resolved.

RQ #7: Do issues with more reactions have more discussion?

In the previous research question, we analyzed the relationship between reactions and the time to close issues. Here, we analyze whether the number of reactions is associated to longer discussions (i.e., more comments). To this purpose, we first compute the number of comments in *bugs* and *enhancements*. Then, we compare the results for issues without reactions and with at least two reactions.

Figure 15 presents these distributions. In all cases, issues with reactions have more comments than those without reactions. For example, *bugs* without reactions have 3 comments against 6 comments (median) in the case of *bugs* with at least two reactions. Besides that, *enhancements* without reactions have 2 comments (median), whereas *enhancements* with reactions have at least 5 comments. By applying a Mann-Whitney test, we detected that the distributions are different ($p\text{-value} < 0.001$). We also found a *large* effect size between the distributions of *bugs* (Cohen’s $d = 0.943$). For *enhancements*, we found a *medium* effect size (Cohen’s $d = 0.626$).

Figure 16 shows the distribution of the number of comments in *bugs* and *enhancements* of a specific repository, MICROSOFT/VSCODE. In this project, we confirm that issues with reactions have more comments. The median number of comments in *bugs* is 2 and 6, respectively. For *enhancements*, the median values are respectively 2 and 5 comments, respectively.




Figure 15: Number of comments in bugs and enhancements (without reactions vs with at least two reactions). Outliers are omitted




Figure 16: Number of comments in bugs and enhancements, in MICROSOFT/VSCODE (outliers are omitted)

Summary: Bugs and enhancements with reactions have longer discussion, which essentially confirms the result of RQ #6, i.e., reactions tend to appear in more complex issues, which demand more discussion.

RQ #8: Do negative reactions inhibit further participation?

In the literature, several authors argue that negative feedback is a factor that inhibits users participation in social platforms [1, 16, 17]. Here, we aim to investigate whether negative reactions (i.e., *Thumbs down*) inhibits further participation through issue reports and comments. Specifically, we compute the number of interactions (i.e., issues and comments) performed by developers *before* and *after* receiving negative reactions. Also, we restrict our analysis to issues and comments that have more *Thumbs down* than the sum of other reactions. When a developer has more than one issue or comment with the required number of *Thumbs down*, we selected the most negative one. Here, we analyze the developers’ participation in a project, i.e., the interactions are computed in the same repository with the negative reactions.

First, Figure 17 presents the distribution of the number of *Thumbs down* reactions received by developers. The distribution shows that most developers receive a low number of negative feedback. However, other developers receive a considerable number of *Thumbs*

down. For example, we have 22,431 developers with at least one issue or comment with more *Thumbs down* than other reactions. Therefore, in this research question, we focus on the top-10% issues and comments, which consists of 2,243 issues and comments with values that range from 11 to 683 *Thumbs down* reactions.




Figure 17: Distribution of negative reactions in issues and comments

In Figure 18, we present the distribution of the number of interactions (i.e., issue reports or issue comments) made *before* and *after* negative reactions. We observe that the values are very close. For example, the median values are 1 and 2 interactions before and after the negative issue reports or comments. By applying the Mann-Whitney test, we detected that these distributions are different (p -value < 0.001), but with a *very small* effect size (Cohen’s $d = 0.136$).




Figure 18: Number of issues and comments made by developers before and after receiving a high number of *Thumbs down*

Summary: Differently from other social platforms, negative reactions do not inhibit GitHub developers participation in further issues of the same repository.

RQ #9: Do reactions reduce the noise in issue conversations?

When the reactions feature was introduced, GitHub argued that one of its purposes was to reduce the noise caused by emojis in issue responses. Indeed, before the introduction of reactions, developers often added comments with a single *Thumbs up*. In this research question, we investigate whether the introduction of reactions reduced the use of emojis (e.g., 👍, 🙌, etc.) in issue conversations. To this purpose, we analyzed the issues created after the introduction of reactions and that have emojis in their body. Specifically,

we searched for comments containing emoji markups in the format `<id>`, where `<id>` refers to the emoji’s code. For example, `:thumbsup:` and `:+1:` refer to the emoji 👍.

Figure 19 presents the monthly usage of emojis in the body of comments, considering all available emojis⁸ and only the emojis representing GitHub reactions (i.e., 👍, 🙌, 🙏, ❤️, 😊, and 😄). The results suggest that the usage of emojis in comments remains the same. In addition, similar results are observed for emojis associated to reactions. The only exception is the emoji `:tada:` (🎉), whose usage increased considerably in the last years. Thus, we can conclude that GitHub reactions did not reduce the usage of emojis in the body of comments.




Figure 19: Usage of emojis in the body of comments after reactions introduction

Finally, we analyzed the comments containing only a single emoji in their body. Specifically, we found that only *Thumbs up* (👍) and *Heart* (❤️) are used in this manner. However, *Heart* is used in only 18 comments, whereas *Thumbs up* is used in 1,650 comments. Figure 20 shows the monthly usage of comments containing a single *Thumbs up*. We can conclude that this practice is in sharp decrease on GitHub.




Figure 20: Usage of only *Thumbs up* in comments after reactions introduction

Summary: GitHub developers are still using emojis in their comments; however they are no longer commenting with a single 👍.

⁸<https://gist.github.com/rxaviers/7360908>, verified on 01/20/2019.

5 KEY FINDINGS AND IMPLICATIONS

We found that most issue reports (87%), comments (89%), and threads (71.6%) do not have reactions (RQ #1). However, it is important to acknowledge that the usage of reactions is gaining momentum; it increased 32.5% in 2018, compared to 2017 (RQ #3). We also reported that *Thumbs up* is by far the most popular reaction (78.7%, RQ #2).

Practical Recommendation #1: We reported interesting use cases of reactions, which might be of the interest of users and owners of GitHub projects. Mostly, we found that reactions can be used (1) to express the importance of fixing a bug or implementing a new feature; (2) to implement surveys with users (e.g., a logo context); (3) to collect user’s feedback about a new release candidate, before it is deployed; (4) to indicate that an issue is invalid. In all these scenarios, reactions are instruments to collect user feedback. But in scenarios (1), (3), and (4), the feedback is spontaneous, i.e., the users are not asked to react. By contrast, in scenario (2), the manifestation is explicitly requested by project managers.

We also found that reactions and stars do not have a strong correlation (RQ #4). Interestingly, we found that users react four times more to enhancements than to bugs (RQ #5). We claim that this finding reveals the pressure suffered by GitHub maintainers to continuously evolve their systems, by providing new features. Furthermore, GitHub users tend to react to more complex issues and enhancements, i.e., the ones that take more time to be handled. For example, enhancements with at least two reactions take 108 more days to be implemented, regarding the median implementation time of other enhancements—without reactions—in a repository.

Practical Recommendation #2: GitHub users should not expect that by reacting to an issue they will get it handled in a few days. Usually, simple issues are handled quickly, even when they do not have reactions. By contrast, it is natural that complex issues require more time to be handled, even when they have many *Thumbs up*. This is exactly the case of the most reacted issue in our dataset, which refers to a new feature which is yet not supported by Microsoft Visual Studio. In other words, reactions should not be viewed by users as an instrument capable of accelerating the implementation of bug fixes or new features. They are not an instrument of pressure; but an instrument to provide feedback.

We also found that issues with reactions have longer discussions (RQ #7) and that negative reactions (*Thumbs down*) do not inhibit developers to comment again (RQ #8).

Practical Recommendation #3: GitHub developers should not view negative reactions as an invitation to abandon projects. Most developers continue to participate in a project, even after receiving a massive number of *Thumbs down*.

Finally, we showed that reactions are not contributing to reduce emojis in comments, including the ones used by the existing GitHub reactions (RQ #9). However, the number of comments with a single *Thumbs Up* is in sharp decline.

Practical Recommendation #4: GitHub should consider supporting other emojis; for example *Clapping hands* (👏), whose usage is facing

a constant increment. Another possible feature is to allow developers to highlight particular sentences of an issue or comment, and then apply a reaction only in the highlighted text. These fine-grained reactions can contribute to reduce the noise caused by emojis in the body of comments (which is not declining, as we showed in RQ #9). Finally, traditional bug tracking systems may also benefit from reactions in their comment systems.

6 THREATS TO VALIDITY

Construct Validity: GitHub has millions of repositories. We built our dataset by collecting the top-5,000 repositories by number of stars, which represents a small fraction in comparison to the GitHub’s universe. However, most GitHub repositories are forks, have very low activity, or have a small community [7, 14, 15]. Despite the repository selection by number of stars, our dataset includes projects from a vast range of programming languages and domains.

Internal Validity: An internal threat relates to the issue types used in RQ #5. In this research question, we relied on the labels provided by the projects maintainers to characterize the issue type. Moreover, we also performed an extra step to group labels with the same purpose. To mitigate this threat, we restricted our analysis to a small group of labels, which is also recommended by default on GitHub. Also, the labels groups were manually verified by the first and second authors of this paper to guarantee that those in the group refer to the same label.

External Validity: The data analyzed in this work is from the GitHub platform. Since the reactions feature is exclusive of this platform, the results may not generalize to other social coding platforms (e.g. GitLab and BitBucket). Finally, during the development of this research, GitHub released two new reactions: Rocket (🚀) and Eyes (👁️). The inclusion of new reactions may influence the usage of reactions in the future. However, we claim these new reactions do not have an immediate impact; and they should only be studied after gaining some adoption [19].

7 RELATED WORK

In the literature, several studies explore emotions on software development. Murgia *et al.* [18] analyzed whether issue reports carry emotional information. Their study confirms that issue reports do express emotions towards design choices, maintenance activity or colleagues. Jongeling *et al.* [13] investigated the agreement between sentiment analysis tools and the sentiment recognized by humans. Their results suggest a sentiment analysis tool specially targeting the software engineering domain is needed. Brants *et al.* [5] conducted a survey in order to determine whether people interpret emojis (specific to emotional states) in a consistent manner. They conclude that there are emojis that are interpreted differently depending on age or gender.

Other studies explore Github projects. Borges *et al.* [3] investigated the characteristics of popular GitHub repositories and found four major popularity growth patterns. Recently, Coelho *et al.* [6] proposed an approach to identify unmaintained projects in GitHub. Their approach apply machine learning techniques and rely on metrics that measure the maintained level to identify and alert users. Machine learning is also used in GitHub to classify user’s

sentiments in the comments of issues [9, 21] and to identify software vulnerabilities [10]. User behavior is also analyzed in other studies. Jiang *et al.* [12] presented an analysis of unfollowing events between GitHub users, identifying potential factors that can motivate this behaviour. Bissyandé *et al.* [2] performed an empirical study about adoption of issues tracker in software projects. The authors investigated the involvement of users and developers when reporting issues and their impact on project success.

Moreover, several studies investigate the impact of reactions—such as *upvotes*, *favorite*, and *like*—in other social networks. In a recent study, Aniche *et al.* [1] investigated two popular news aggregators sites (HackerNews and Reddit). The authors showed that developers view *upvotes* and *comments* as measurement of post’s quality and relevance. Likewise, Sharma *et al.* [20] proposed a recommendation system to support developers learning through URLs posted on Twitter. Their system relies on users actions, as *favorite* and *retweet*, to identify and rank relevant content.

8 CONCLUSION

In this paper, we investigated how developers react to issue reports and their comments on GitHub. For this purpose, we collected and analyzed the data of 4,849 public GitHub projects. We provided answers to nine research questions, referring to two themes: (i) the usage of reactions by GitHub developers, and (ii) the impact of reactions on software development practices. Our key findings are as follows.

- Most issues (87%) and comments (89%) do not have reactions. For whole threads (issues + comments), we found that 28.4% have at least one reaction. Moreover, *Thumbs up* is by far the most common reaction (78.7%).
- Reactions are gaining momentum; in 2018, developers used 32.5% more reactions than in 2017.
- We show that popular projects do not necessarily attract more reactions. We also show that developers react more often to *enhancements* and there is a correlation between issue types and reactions.
- *Bugs* and *enhancements* with reactions take more time to be closed and have longer discussion.
- On the contrary of other social platforms, negative reactions do not inhibit GitHub developers participation in further issues of the same repository.
- Finally, GitHub reactions are contributing to reduce the number of comments containing a single *Thumbs up* emoji. However, developers are still using emojis in their comments.

As future work, we plan to investigate the impact of the profile of the analyzed projects on the usage of reactions. We also plan to conduct studies with GitHub developers to understand their perceptions on the reactions features. First, we intend to conduct a survey to discover what are the factors they consider before reacting to issues and to comments. We also plan to shed light on the importance of reactions for project maintainers through an opened survey with developers from popular projects. Finally, we plan to extend our work to consider reactions on pull requests.

The dataset with the repositories, issues, comments, and reactions used in this paper is publicly available at: <https://doi.org/10.5281/zenodo.2558596>.

ACKNOWLEDGMENTS

This research is supported by grants from FAPEMIG, CAPES, and CNPq.

REFERENCES

- [1] Maurício Aniche, Christoph Treude, Igor Steinmacher, Igor Wiese, Gustavo Pinto, Margaret-Anne Storey, and Marco Aurélio Gerosa. 2018. How Modern News Aggregators Help Development Communities Shape and Share Knowledge. In *40th International Conference on Software Engineering (ICSE) (ICSE '18)*. ACM, 499–510.
- [2] T. F. Bissyandé, D. Lo, L. Jiang, L. Réveillère, J. Klein, and Y. L. Traon. 2013. Got issues? Who cares about it? A large scale investigation of issue trackers from GitHub. In *24th International Symposium on Software Reliability Engineering (ISSRE)*. 188–197.
- [3] Hudson Borges, Andre Hora, and Marco Tulio Valente. 2016. Understanding the factors that impact the popularity of GitHub repositories. In *32nd International Conference on Software Maintenance and Evolution (ICSME)*. 1–11.
- [4] Hudson Borges and Marco Tulio Valente. 2018. What’s in a GitHub Star? Understanding Repository Starring Practices in a Social Coding Platform. *Journal of Systems and Software* 146 (2018), 112–129.
- [5] Wesley Brants, Bonita Sharif, and Alexander Serebrenik. 2019. Assessing the Meaning of Emojis for Emotional Awareness-A Pilot Study. In *2nd International Workshop on Emoji Understanding and Applications in Social Media*.
- [6] Jailton Coelho, Marco Tulio Valente, Luciana L. Silva, and Emad Shihab. 2018. Identifying Unmaintained Projects in GitHub. In *12th International Symposium on Empirical Software Engineering and Measurement (ESEM)*. 15:1–15:10.
- [7] Valerio Cosentino, Javier L. Canovas Izquierdo, and Jordi Cabot. 2017. A Systematic Mapping Study of Software Development with GitHub. *IEEE Access* 5 (2017), 7173–7192.
- [8] Laura A. Dabbish, H. Colleen Stuart, Jason Tsay, and James D. Herbsleb. 2012. Social coding in GitHub: Transparency and collaboration in an open software repository. In *15th Conference on Computer Supported Cooperative Work (CSCW)*. 1277–1286.
- [9] Jin Ding, Hailong Sun, Xu Wang, and Xudong Liu. 2018. Entity-level sentiment analysis of issue comments. In *3rd International Workshop on Emotion Awareness in Software Engineering (SEmotion)*. 7–13.
- [10] Jacob A. Harer, Louis Y. Kim, Rebecca L. Russell, Onur Ozdemir, Leonard R. Kosta, Akshay Rangamani, Lei H. Hamilton, Gabriel I. Centeno, Jonathan R. Key, Paul M. Ellingwood, Marc W. McConley, Jeffrey M. Opper, Sang Peter Chin, and Tomo Lazovich. 2018. Automated software vulnerability detection with machine learning. (2018).
- [11] Jing Jiang, David Lo, Jiahuan He, Xin Xia, Pavneet Singh Kochhar, and Li Zhang. 2017. Why and how developers fork what from whom in GitHub. In *Empirical Software Engineering*. 547–578.
- [12] Jing Jiang, David Lo, Yun Yang, Jianfeng Li, and Li Zhang. 2013. A first look at unfollowing behavior on GitHub. In *Information and Software Technology*. 150–160.
- [13] Robbert Jongeling, Subhajit Datta, and Alexander Serebrenik. 2015. Choosing your weapons: On sentiment analysis tools for software engineering research. In *31st International Conference on Software Maintenance and Evolution (ICSME)*. 531–535.
- [14] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M. German, and Daniela Damian. 2014. The Promises and Perils of Mining GitHub. In *11th Working Conference on Mining Software Repositories (MSR)*. 92–101.
- [15] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M. German, and Daniela E. Damian. 2016. An in-depth study of the promises and perils of mining GitHub. *Empirical Software Engineering* 21, 5 (2016), 2035–2071.
- [16] Aikaterini Katmada, Anna Satsiou, and Ioannis Kompatsiaris. 2016. Incentive Mechanisms for Crowdsourcing Platforms. In *Internet Science*. 3–18.
- [17] Cliff Lampe and Erik Johnston. 2005. Follow the (slash) dot: effects of feedback on new members in an online community. In *5th International Conference on Supporting Group Work (SIGGROUP)*. 11–20.
- [18] Alessandro Murgia, Parastou Tourani, Bram Adams, and Marco Ortu. 2014. Do developers feel emotions? an exploratory analysis of emotions in software artifacts. In *11th Working Conference on Mining Software Repositories (MSR)*. 262–271.
- [19] Everett M Rogers. 2010. *Diffusion of innovations*. Simon and Schuster.
- [20] A. Sharma, Y. Tian, A. Sulistya, D. Lo, and A. F. Yamashita. 2017. Harnessing Twitter to support serendipitous learning of developers. In *24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 387–391.
- [21] Karl Werder and Sjaak Brinkkemper. 2018. MEME: toward a method for emotions extraction from GitHub. In *3rd International Workshop on Emotion Awareness in Software Engineering (SEmotion)*. 20–24.