A Deep Dive into the Impact of COVID-19 on Software Development

Paulo Anselmo da Mota Silveira Neto, Umme Ayda Mannan, Eduardo Santana de Almeida, *Senior Member, IEEE*, Nachiappan Nagappan, *Fellow, IEEE*, David Lo, Pavneet Singh Kochhar, Cuiyun Gao, and Iftekhar Ahmed

Abstract—The COVID-19 pandemic is considered as the most crucial global health calamity of the century. It has impacted different business sectors around the world and software development is not an exception. This study investigates the impact of COVID-19 on software projects and software development professionals. We conducted a mining software repository study based on 100 GitHub projects developed in Java using ten different metrics. Next, we surveyed 279 software development professionals for better understanding the impact of COVID-19 on daily activities and wellbeing. We identified 12 observations related to productivity, code quality, and wellbeing. Our findings highlight that the impact of COVID-19 is not binary (reduce productivity versus increase productivity) but rather a spectrum. For many of our observations, substantial proportions of respondents have differing opinions from each other. We believe that more research is needed to uncover specific conditions that cause certain outcomes to be more prevalent.

Index Terms—COVID-19, empirical study, survey and mining software repository (MSR)

1 Introduction

The latest threat to global health is the ongoing outbreak of the respiratory disease that was recently given the name Coronavirus Disease 2019 (COVID-19) [1]. COVID-19 was recognized in December 2019, in Wuhan, a large city in central China [2]. Three months later, on March 11, 2020, the World Health Organization (WHO) characterized COVID-19 as a pandemic.[1] According to WHO:"*We have never before seen a pandemic sparked by a coronavirus. This is the first pandemic caused by a coronavirus. And we have never before seen a pandemic that can be controlled, at the same time.*" In August 2020, COVID-19 was spread out more than 200 countries with more than 500k confirmed deaths.

The COVID-19 pandemic is considered the most crucial global health calamity of the century and the greatest challenge humankind faced since the 2nd World War [3]. It has enormously impacted how we live and interact with each other (social distancing, wearing masks, washing hands frequently, quarantine, etc.) Besides the health problems, the action taken by the countries around the world to manage the COVID-19 pandemic, e.g., restricting travel, shuttering nonessential businesses, and implementing universal social distancing policies, are having drastic economic consequences. In the United States, for example, more than 30 million Americans have filed initial unemployment claims since March 2020.[2]

Business sectors, such as travel and transportation, manufacturing, hotels, restaurants, live entertainment and movie, and sports are strongly impacted and they are having to adapt to this new situation. The software development sector is not an exception. Besides the economic aspects, suddenly, companies had to support and (in some cases) equip office workers who quickly transitioned to a work-from-home set up because of the pandemic.

Working from home is not a new reality for software development. Many companies adopt different approaches for it, such as remote teams that have employees in a single country or even in one city, and some/all of them work from home without having to go to the office every day. Professional websites such as Upwork, Linkedin, and Stack Overflow present several job offers to work remotely. However, working from home during a pandemic is not like regular remote work. Additional difficulties are involved due

1. WHO characterizes COVID-19 as a pandemic. https://www.who.int/emergencies/diseases/novel-coronavirus-2019/events-as-they-happen

2. https://edition.cnn.com/2020/04/30/economy/unemployment-benefits-coronavirus/index.html

- *Paulo Anselmo da Mota Silveira Neto is with the Departament of Computer Science, Federal Rural University of Pernambuco, Recife 52171-900, Brazil. E-mail: paulo.motant@ufrpe.br.*
- *Umme Ayda Mannan is with the Departament of Computer Science, Oregon State University, Corvallis, OR 97331 USA. E-mail: mannanu@oregonstate.edu.*
- *Eduardo Santana de Almeida is with the Department of Computer Science, Federal University of Bahia, Salvador 40170-110, Brazil. E-mail: esa@rise.com.br.*
- *Nachiappan Nagappan and Pavneet Singh Kochhar are with the Microsoft Research, Redmond, WA 98052 USA. E-mail: {nachin, Pavneet.Kochhar}@microsoft.com.*
- *David Lo is with the School of Information Systems, Singapore Management University, Singapore 188065, Singapore. E-mail: davidlo@smu.edu.sg.*
- *Cuiyun Gao is with the Department of Computer Science and Engineering, Chinese University of Hong Kong, Hong Kong. E-mail: gcyydxf@gmail.com.*
- *Iftekhar Ahmed is with the Department of Informatics, University of California, Irvine, Irvine, CA 92697 USA. E-mail: iftekha@uci.edu.*

Manuscript received 3 Nov. 2020; revised 7 June 2021; accepted 8 June 2021. Date of publication 11 June 2021; date of current version 19 Sept. 2022. (Corresponding author: Paulo Anselmo da Mota Silveira Neto.) Recommended for acceptance by A. M. Moreno. Digital Object Identifier no. 10.1109/TSE.2021.3088759

to the lack of proper physical infrastructure, the need to care for children with school and daycare being closed, fear and anxiety of contracting COVID-19, etc.

Previous research studies [4], [5], [6], [7] have investigated remote work and responses for a better understanding of pandemic [8], [9], [10], [11], [12]. However, just a few studies [13], [14], [15] have started to investigate the impact of a pandemic in software development. This aspect can be justified since COVID-19 is the first pandemic after World Wide Web development. Thus, to gain insights into the impact of COVID-19 in software development, our first research question (*RQ1: What is the Impact of COVID-19 on Projects?*) explores how the pandemic impacts open source projects considering different perspectives. In addition to quantitative information about the projects, it is also important to understand the impact on developers' wellbeing; this motivates us to investigate another research question (*RQ2: What is the Impact of COVID-19 on Developers WellBeing?*).

To answer these questions, we conducted a mining software repository study based on 100 GitHub projects developed in Java. We decided to start by mining since it can help us to identify and measure the changes in activity before and during COVID-19.

The Java projects were selected according to different criteria ranging from the last update to the number of commits. We collected ten metrics for analyzing the projects. Since our goal was to gain a deeper understanding about the impact of COVID-19, it is necessary to collect the information from real software professionals along with mining software repositories. So we surveyed 279 software development professionals from 32 countries. The survey asked respondents to provide feedback on the impact of COVID-19 on software projects and their well-being.

Overall, the paper makes the following contributions:

- We perform a large scale quantitative study to investigate the impact of COVID-19 on software development based on ten different metrics.
- We complement this study with a survey of how software development professionals perceive the impact of COVID-19 on daily activities.
- Based on a set of observations from the mining software repository study and survey, we provide some recommendations for practitioners, organizations, and researchers.
- For replication and reproducible research, we make our materials available on our project website. These include our repository mining data (project data, including metrics and time series data) and survey instrument. Our artifacts can be found at the accompanying website.[3]

2 Background and Related Work

In this section, we discuss the main work related to our study.

2.1 COVID-19 Studies

The COVID-19 pandemic originated the development of several multidisciplinary initiatives around the world. The Center for Systems Science and Engineering (CSSE) at Johns Hopkins University, created an interactive web-based dashboard to visualize and track COVID-19 reported cases in real-time [9]. The dashboard illustrates the location and number of confirmed COVID-19 cases, deaths, and recoveries for all affected countries. Zhang *et al.* [12] created, Neural Covidex, a search engine for clinicians, researchers, and other experts trying to better understand COVID-19. The system offers access to the Allen Institute for AI's COVID-19 Open Research Dataset (CORD-19). CORD-19 is a curated public resource of more than 40,000 scholarly articles, medical reports, journal articles, and preprints about COVID-19 and the coronavirus family of viruses. Researchers from MRC Centre for Global Infectious Disease Analysis, from Imperial College London, have developed a set of tools and prediction models based on different scenarios (social distancing, shielding the elderly, and healthcare availability) [10], [11]. Chen *et al.* [8] created a public Coronavirus Twitter dataset with more than 100 million tweets. According to the authors, the dataset can help track scientific coronavirus misinformation and unverified rumors, and contribute towards enabling informed solutions and prescribing targeted policy interventions.

In the software engineering area, Ralph *et al.* [14] conducted a survey with 2225 software developers to understand the effects of the COVID-19 pandemic on developers' wellbeing and productivity. They identified that developers had lower wellbeing and productivity while working from home due to COVID-19. In addition, disaster preparedness, fear related to the pandemic and home office ergonomics all affect wellbeing and productivity; and women, parents, and people with disabilities may be disproportionately affected.

Rahman and Farhana [13] conducted an empirical study with 129 open source COVID-19 projects hosted on GitHub to identify what categories of bugs exist in this kind of project. Initially, they identified seven categories of COVID-19 projects (aggregation, education, medical equipment, mining, user tracking, statistical modeling, and volunteer management). Next, applying open coding on 550 bug reports, they identified eight bug categories for these projects (algorithm, data, dependency, documentation, performance, security, syntax, and user interface). Based on this taxonomy, user interface was the most frequent category using the proportion of bugs across all projects, and documentation was the least frequent category.

Miller *et al.* [16] investigated how team culture and team productivity has been affected due to COVID-19 by conducting two surveys at a large software company. Their exploratory survey during the early months of the pandemic revealed that developers faced challenges reaching milestones and their team productivity had changed. They also identified that the ability to brainstorm with colleagues, difficulty communicating with colleagues, and satisfaction with interactions from social activities are important factors that are associated with how developers report their software development team's productivity. They also provide recommendations on how managers can support sustained team productivity during times of crisis and beyond.

Although many of the prior work investigated the affect of COVID-19 on productivity ([14], [16], all of them analyzed the impact from only one aspect, gathering developers perception through surveys [14], [16]. However, there

3. https://github.com/pamsn/covid-study

are numerous examples of long-held beliefs that proved to be incorrect or outdated when actual evidence was collected through empirical analysis [17], [18]. Given that, to get a better understating about the impact of COVID-19 on productivity, we need to collect not only developer's perception but also check if the perception matches the findings through analyzing the software repositories and activities in them (e.g., number of commits, issues, pull requests, branches, comments, time to fix a pull request and an issue, and the number of active and new contributors). Ours is the first work to investigate the affect of COVID-19 from both aspects.

2.2 Productivity

There are a number of studies [19], [20], [21], [22], [23] that investigate developer productivity. For example, Meyer *et al.* [24] investigated how developer workday looks like and the relationships between their activities and perceived productivity. The study uncovered that productivity is a personal matter, and factors such as emails and meetings are often considered detrimental to productivity. In another work, Murphy-Hill *et al.* [25] surveyed 622 developers across 3 companies regarding self-perceived productivity and its contributing factors. They uncovered that developers perceived productivity is strongly correlated to job enthusiasm, peer support, and reception of useful feedback. Moreover, they reported that compared to other knowledge workers, *ability to work remotely* is more strongly related to perceived productivity. Unlike the aforementioned studies, in this work, we investigate the impact of COVID-19 on software projects.

2.3 Wellbeing

Many prior works investigate wellbeing in the workplace, e.g., [26]. However, not many investigate it in the context of software development. One such work is by Kuutila *et al.* [27] that investigated the relationship between developer wellbeing and software repository metrics. They reported that developers who reported "high hurry" were less productive. Moreover, factors that hamper wellbeing (such as stress, sleeping problems, etc.) are negatively related to the number of chat messages. Another work by Graziotin *et al.* [28] investigated what happens when developers are happy and unhappy. They found that unhappiness impacts the developer's own being (in terms of low cognitive performance, mental unease, and disorder, etc.) and their work products (in terms of low productivity, low code quality, etc.). Happiness leads to the opposite effect (e.g., high cognitive performance, high productivity, high code quality, etc.).

Johnson *et al.* [20] conducted a mixed-method study (two surveys and interviews) with 1159 participants from Microsoft to understand the effect of work environments on productivity and satisfaction of software engineers. They found that personalization, social norms and signals, room composition and atmosphere, work-related environment affordances, work area, and furniture, and productivity strategies were considered important factors for work environments. In addition, the ability to work privately with no interruptions and the ability to communicate with the team and leads were important factors related to satisfaction.

Meyer *et al.* [29] investigated what is a good and typical workday for software developers at Microsoft. They conducted a survey with 5971 developers and identified that on good workdays, developers make progress and value projects they consider meaningful and spend their time efficiently, with little randomization, administrative work, and infrastructure issues.

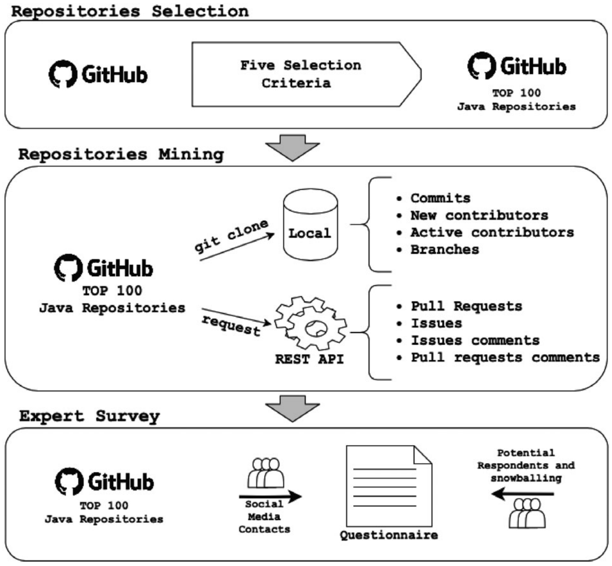

Fig. 1. Research design.

Fucci *et al.* [30] performed a quasi-experiment with 45 undergraduate students to investigate whether, and to what extent, sleep deprivation impacts the performance of novice software developers using the agile practice of test-first development (TFD). The students were divided into two groups, where 23 stayed awake the night before carrying out the tasks, while 22 slept normally. They identified that a single night of sleep deprivation reduced 50 percent in the quality of the implementations. In addition, the study found that sleep-deprived developers make more fixes to syntactic mistakes in the source code.

3 Research Design

This section describes the research design used in our study. First, a *(i) Repository Mining* analysis was performed to understand the impact of COVID-19 pandemic on software projects. It considered different perspectives such as number of commits, issues, pull requests, branches, comments, time to fix a pull request and an issue, and the number of active and new contributors. Next, a *(ii) Survey* with software professionals was performed to get insights on how the pandemic impact on their wellbeing. Fig. 1 shows the overall research design used in our study. Our epistemological stance is empiricism since we do not rely on our own intuition rather on observation through mining and survey [31]. Through data mining, we are taking a positivist [32] stance by assuming that we can observe and infer the impact of COVID-19. There are limitations to this approach; what we are observing is software engineering-specific artifacts such as number of branches, number of active contributors, etc. However, there are external factors that may have an impact on software development. To better understand how COVID-19 has impacted software development and productivity, we took the pragmatic approach of surveying to understand the developers' everyday experience.

TABLE 1
Metrics Used in the Analysis

Metric	Description	Ref.
№ of Active Contributors	The developers that had at least one commit in the study date range.	[19]
№ of New Developers	The developers that committed at first time during the study date range.	[19]
№ of Branches	List of all public branches of each repository.	[33]
№ of Created Pull Requests	Number of pull requests created per month considering the study date range.	[23], [25]
№ of Closed Pull Requests	Number of pull requests updated per month considering the study date range.	[23], [25]
Pull Request created and closed per month	Number of pull requests created and closed in the same month considering the study date range.	[23], [25]
№ of Commits	Number of commits considering the study date range	[19]
№ of Created Pull Requests Comments	Number of comments created related to a previous created pull request.	[34]
№ of Updated Pull Requests Comments	Number of comments updated related to a previous created pull request.	[34]
№ of Bug-fix commits	Number of project commits related to bug fix.	[35]

We acknowledge that the multifaceted and wide-ranging activities of software development would require observing individuals in their work environment. However, this was not possible due to the pandemic, and we did the next best possible thing by surveying the developers.

Since we are collecting developers perception through survey, we made sure the survey does not provide any presupposing agreement which can bias the respondents inputs.

3.1 Repositories Selection and Mining

Metrics. The metrics used in our analysis are presented in Table 1. These metrics were used since they have been used as indicators for contribution by other researchers. For each of the metrics, we collected information for one year and four months (2019-Jan to 2020-May). It was important to show the behavior before and during the pandemic period. We also want to investigate how the metrics evolve month by month.

Selecting and Filtering Repositories. First, we choose GitHub repositories written in the Java language, since it is one of the most popular programming languages.[4] Next, in order to detect pandemic impact on GitHub repositories, we retrieved Java repositories created from 2019 and sorted by their popularity [36]. In addition, as recommended by Munaiah *et al.* [37], we filtered out the noise in such large repositories by applying different inclusion and exclusion criteria, as follows:

- *Inclusion Criteria 1:* The repository has been updated at least once in the last year (2019-Jan to 2020-May);
- *Inclusion Criteria 2:* The repository must have at least 34 commits in the study period (2019-Jan to 2020-May); this corresponds to two commits per month in the 17-month study period. This criteria was used to filter out inactive repositories;
- *Inclusion Criteria 3:* The repository must have at least 10 contributors in the study period (2019-Jan to 2020-May). This criteria was used in order to eliminate irrelevant repositories, c.f., [38], [39], [40];
- *Exclusion Criteria 1:* Repositories that did not have their artifacts and description in English were not considered in the study;
- *Exclusion Criteria 2:* Repositories corresponding to tutorials, books, and classroom materials were also removed from our analysis.

After the filtering step, we selected the Top 100 Java repositories. As described in Fig. 1, the repositories were cloned and the `git log` and `git branch` commands were used to get all projects commits and branches. Using the `.csv` file generated, a Python script was used to extract the number of commits, time to fix a issue and pull request, the number of new contributors (i.e., those who made their first contribution in the study period (2019-Jan to 2020-May)), the number of contributors (i.e., those who collaborated in the study period) and number of remote branches for each repository. When the information could not be collected from local cloned repositories, we used the GitHub API. Specifically, pull requests and their comments were retrieved by GitHub REST API requests.[5]

Next, we present how each metric was identified and measured for the repository mining process.

3.1.1 Number of Active Contributors

Active Contributors Identification and Measurement. The number of active contributors was identified by first cloning each repository and collecting all commits[6] in the study period. Once the commits were collected in a `.csv` file, a Python script was used to identify each developer responsible for each commit. This way, we considered an active contributor to be a developer which has at least one commit in the study period.

3.1.2 Number of New Developers

New Developers Identification and Measuring. The number of new developers was also collected from the commits extracted from each repository. The Python script searched for developers that had their first commit in the repository

4. TIOBE Index. https://www.tiobe.com/tiobe-index/

5. https://developer.github.com/v3/

6. All commits were extracted using the git command: `git log –all –format=H,aE,ci > commit_file.csv`

during the study period. It is important to know if most of the new developers initiated their contribution before or during the pandemic.

3.1.3 Number of Branches

Number of Branches Identification and Measuring. The number of branches was also collected from the cloned repositories throughout the GitHub REST API. It is important to mention that we used an API Endpoint that retrieves all protected and unprotected repository branches. Finally, a Python script was used to collect the branch creation date based on its commits.

3.1.4 Number of Created and Closed Pull Requests

Number of Created, Closed Pull Requests Identification and Measuring. The created and closed pull requests were collected using the GitHub REST API[7] since this information is not available at local cloned repositories. According to GitHub REST API documentation, both issues and pull requests should be retrieved by the same `GET /issues` endpoint. In order to identify pull requests, we need to filter the results searching for the "pull request" string in each issue occurrence. With this information, we were able to identify how many pull requests were created and closed, considering the study date interval. In addition, we also identified the pull requests created and closed in the same months. With this metric, we analyze how the pandemic impacts team productivity.

3.1.5 Number of Commits

Number of Commits Identification and Measurement. The project commits were collected from the cloned repository using `git log` command. Next, the Python script was used to group all commits by month-year in a `.csv` file. This information is important to understand how the pandemic impacts on team productivity over the months.

3.1.6 Number of Created and Updated Pull Requests Comments

Number of Created and Updated Pull Requests Comments Identification and Measuring. The GitHub REST API[8] was also used to collect information regarding to the pull requests comments. This metric collected both new and updated comments related to repository pull requests. The comments show important insights regarding to the project activity over the study interval.

3.1.7 Number of Bug-Fixing commits

Number of Bug-Fixing commits Identification and Measurement. To gather the number of commits related to bug-fixing, we used the GitcProc Tool[9] which uses a keyword search to determine if that commits is related to a bug fixing or not. It searches for related words such as error, bug, defect, and fix within each commit message [35].

7. https://developer.github.com/v3/
8. https://developer.github.com/v3/
9. https://github.com/caseycas/gitcproc

3.2 Survey

Protocol. We created a 20-minute survey designed to understand the impact of COVID-19 on software development from the perspective of projects and developers' wellbeing. It was composed of seventeen closed questions on a Likert scale (Strongly Disagree, Disagree, Neutral, Agree, Strongly Agree, and N/A) and three open questions. We included the option N/A to prevent respondents providing arbitrary ratings to questions that they find to be unclear. The survey also collected demographic information from respondents. Previous surveys related to software productivity [22], well-being [20] and COVID-19 [14] were also consulted. For the design of the survey, we followed the Kitchenham and Pfleeger's guidelines for personal opinion surveys [41].

We piloted our survey with three researchers (2 with Ph.D. degrees and 1 Ph.D. student) with experience in the area to get feedback on the questions and their corresponding answers; difficulties faced to answer the survey and time to finish it. As these pilot respondents were experts in the area, we also would like to know if we were asking the right questions. Next, we conducted another pilot study with two software developers (one from open-source community and one from industry). We conducted several iterations of the survey and rephrased some questions and removed others to make the survey easier to understand and answer. Then, we ended up with 20 questions. Another concern in this stage was also to ensure that the participants could finish it in 20-minutes. The pilot survey responses were used solely to improve the questions, and these responses were not included in the final results. We kept the survey anonymous, but in the end, the respondents could inform their email to receive a summary of the study. The survey instrument can be seen in the accompanying website.[10]

Respondents. We followed a four steps approach to recruit survey respondents: initially, we sent the invitation to 6705 unique contributors from the selected projects (Section 3.1). As we received only 15 answers, we decided to post survey information on personal accounts of social media (Twitter, LinkedIn, Instagram). Next, most authors contacted potential respondents by email (convenience sample) and asked them to share it with other potential respondents (snowballing). Because of this process, we were not able to track the total number of invitations. In total, we received 287 responses. We disqualified six responses without actual data (without responses to any of the survey questions of interest to the study despite responding to basic demographics questions, such as the role of the respondent) and two repeated ones, leading to 279 valid responses that were considered (15 respondents from open-source projects and 264 respondents from industry).

The respondents spread out in 32 countries across four continents. The top three countries where the respondents come from are Brazil, United States, and Germany. The professional experience of these 279 respondents varies from no experience to 50 years, with an average of 12.5 years and a median of 10 years.

Regarding the size of the organizations, 34 percent of the respondents work for companies with 1000 to 9999 employees,

10. https://github.com/pamsn/covid-study

23 percent of the participants work for companies with 100 to 999 employees, 16 percent of the respondents work for companies with 10 to 99 employees, 11 percent of the respondents work for companies with 10000 to 99999 employees, 8 percent work for companies with 100000 or more employees, and 8 percent of the respondents work for companies with 0 to 9 employees. The majority of the respondents have a Bachelor's degree (41 percent) and an advanced degree (41 percent), i.e., Master's or Ph.D.

95 percent of the respondents were paid professional, and 5 percent were volunteers. 90 percent of the respondents work full time, and 94 percent had no disability. On average, the participants live with three people (including himself/herself), and only 32 percent live with children under the age of 12. The home office experience of the 279 respondents varies from no experience to 35 years, with an average of 3.5 years and a median of 1 year. Finally, 82 percent of the participants were working from office before the pandemic and switched to a home office, 17 percent were working the whole time remotely, and 1 percent were working at the office the whole time.

Data Analysis. We collected the ratings that our respondents provided for each question. Next, we converted these ratings to Likert scores from 1 (Strongly Disagree) to 5 (Strongly Agree). We computed the average Likert score of each statement related to productivity and wellbeing perspectives and plot Likert scale graph. A Likert scale graph is a bar chart that shows the number of responses corresponding to strongly disagree, disagree, neutral, agree, strongly agree, and N/A. We applied the Fisher's Test [42] to identify significant difference between the open-source and industry participants answers.

In addition, we applied the coding process [43] to analyze the answers that the survey respondents gave to explain the three open questions related to productivity and wellbeing, and their perceptions about how the COVID-19 pandemic has affected him/her and his/her team. We used a set of first cycle and second cycle coding methods to data analysis [43]. First cycle methods are those processes that happen during the initial coding of data. Second cycle methods, if needed, are advanced ways of reorganizing and reanalyzing data coded through first cycle methods. In our study, the codes created were later on clustered in categories. Based on our data, we considered that the categories defined were appropriated to understand the answers that the surveys participants gave. It is also important to highlight that the development of an original theory, for example, is not always a necessary outcome for a qualitative study [44].

To reduce the subjective bias during the coding process, we assigned the three questions to two authors of this paper. The authors had previous experience and conducted similar studies involving interviews and surveys in different aspects of software engineering. Our process required each author to analyze the answers separately and conduct the coding process. Once all the data were analyzed, the two authors met to resolve differences in coding. After discussing and comparing, we established a "substantial agreement" of 0.65 measured using Cohen's kappa [45].

As the amount of data was not too expressive, the coding process did not use any special CAQDAS tool. All the process was conducted with online tools (Google Docs).

3.3 Statistical Analysis

In the mining part of the study, we calculated the number of each metric mentioned in Table 1 for each project for 17 months (for details, see Section 3.1). To investigate the impact of COVID-19 on these metrics, we partition the time into six time-windows. As all the collected data related to project metrics are time-series [46], we did a time series analysis for each metric in each time frame across the projects to see the changes. We also performed a pairwise two-sample t-test to check if there is any statistically significant difference in the mean of each metric among the time-frames. The time series plots for each metric for each category and for each time-frame is available at the accompanying Website.

4 Results

This section presents our findings from analyzing both mining data and survey responses. The content is organized around our two research questions.

4.1 RQ1: What is the Impact of COVID-19 on Projects?

Our goal is to identify the impact of COVID-19 on software projects with the first research question. For this, we conducted a mining study of 100 Java projects where we investigate the change of 10 different project-related metrics (defined in Table 1) during the COVID-19 period. In order to do this comparison, we divided our dataset of 17 months into six time-windows:

1) *Jan'19-May' 20 (TW1)*: This time-window includes time starting from January 2019 to May 2020.
2) *Jan'19-Jun'19 (TW2)*: This time-window includes the time starting from January 2019 to June 2019.
3) *Jul' 19-Dec'19 (TW3)*: This time-window includes the time starting from July 2019 to December 2019. This time-window includes the initial COVID-19 spread time.
4) *Jan'20-Mar' 20 (TW4)*: This time-window includes time starting from January 2020 to March 2020. During this time-window, WHO (World Health Organization) declared COVID-19 as a global pandemic on March 11,2020 [47]. Moreover, people started to work from home during this time-window.
5) *Apr' 20-May' 20 (TW5)*: This time-window includes the months of April and May 2020.
6) *Jan'20-May' 20 (TW6)*: This time-window includes time staring from January 2020 to May 2020.

After dividing the metrics data into six time-windows (including the window with full-time length TW1), we investigated the changes of each metric (Table 1). We tracked the evolution of each metric for identifying their trends in the projects during the time-windows. For this purpose, we calculated the effect size of the month on each metric using a linear regression model, giving us how much that metric changed for each project per month. Then, based on the effect size, we categorized each project into one of two categories: increasing or decreasing. If the effect was positive, we marked those projects as increasing for that metric. For a negative effect, we marked the project as decreasing. For example, Fig. 2 shows

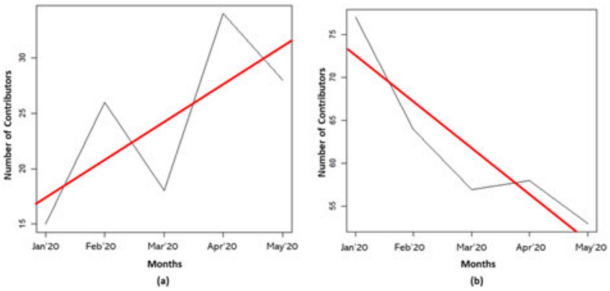

Fig. 2. Month wise trends for Active Contributors across all projects.

that for the metric "number of active contributors", project "Brave" (Fig. 2a) shows an increasing trend for the time-window January 2020 to May 2020 and "ExoPlayer" project (Fig. 2b) shows a decreasing trend according to our time series analysis. All the figures generated from the analysis are provided in the accompanying website. In Table 2, we report the number of projects belonging to each trend for the six time-windows mentioned earlier.

In the following sections, we list the observations pertaining to the impact of COVID-19 on different project activities.

4.1.1 Impact on Bug-Fixing Activity

From our mining analysis, we found that in general, more projects are showing decreasing trends in the number of bug-fix commits per month during the pandemic period. For example, time-window between Jan'20-May'20 (TW6) shows more projects with a decreasing trend in comparison to the time-windows between Jan'19-Dec'19 (See Table 2). Next, we performed a pairwise two-sample t-test to check if there is any statistically significant difference in the mean between these time-windows and found that between the time-windows Jul'19-Dec'19 and Jan'20-May'20, there is a statistically significant difference in their mean (p-value $<$ 0.05). Table 3 shows the summary of the pair-wise t-test for each metric in different time-windows. Table 3 only shows the statistically significant results of our t-test.

In our survey, we asked respondents if they agree/disagree with the statement, "Since you began working from home, the number of bug-fixes is lower than usual". Among the 279 survey respondents, 44.08 percent disagreed, where 15.77 percent agreed that the number of bug-fixes is lower than usual since they began working from home during the pandemic. 42, 81, and 95 respondents strongly disagree, disagree, and neutral with this statement. The average Likert score for this statement is 2.58 (i.e., between "disagree" and "neutral").

We also ran the Fisher's exact test to check the answers provided by OSS and industry participants and no difference between categories was identified (p-value = 0.9245).

> **Observation 1:** Though the majority of projects show decreasing trends in the number of bug-fix commits most of the survey respondents disagree regarding decreased number of bug-fix commits during the pandemic period.

We also asked the respondents about their view on the statement: "The number of bugs in the project has increased since they began working from home". 39, 92, and 107 respondents strongly disagree, disagree, and neutral with this statement. The average Likert score for this statement is 2.46 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.3752).

The following are some comments that refute or confirm the statement:

×*"Overall we've been finishing more features at a higher quality."*

✓*"The overall level of engagement has increased, but I feel that the quality has suffered somewhat."*

> **Observation 2:** 46.60 percent of the respondents conveyed that the number of bugs did not increase during the pandemic period.

Fig. 3 shows the summary of the survey results on the impact of COVID-19 on the bug-fixing activities.

TABLE 2
Summary of Time Series Analysis for Each Category

Metric	TW 1		TW 2		TW 3		TW 4		TW 5		TW 6	
	↑	↓	↑	↓	↑	↓	↑	↓	↑	↓	↑	↓
Active Contributors	42	58	42	58	45	55	31	69	47	50	18	82
New Developers	49	49	41	59	29	69	30	66	49	41	18	70
Branches	93	7	26	28	35	39	91	9	33	29	77	23
Created Pull Requests	81	11	29	25	35	35	73	15	49	21	54	30
Closed Pull Requests	45	54	45	53	49	50	34	64	49	48	25	71
Pull request opened & closed in same month	46	53	42	58	44	55	39	58	52	43	20	75
Commits	45	55	40	60	44	56	32	68	55	44	23	77
Created Pull request Comments	56	41	36	64	38	46	34	55	50	34	23	59
Updated Pull request comments	58	39	38	62	37	45	38	51	49	35	32	51
Bug-fix Commits	40	60	38	62	37	62	32	65	49	47	25	70

TABLE 3
Summary of t-Test Among Projects Between Different Time Range

Metric	TW 1	TW 2	Mean1	Mean2	p-value
Active Contributors	Jul19-Dec19	Jan20-May20	19.13	29.02	0.019
Closed Pull Requests	Jan20-Mar20	Apr20-May20	78.79	46.94	0.047
Created Pull Requests	Jul19-Dec19	Jan20-May20	6.79	12.48	0.04
New Developers	Jan20-Mar20	Apr20-May20	8.88	4.46	0.00012
Pull Request opened & closed in the same month	Jan20-Mar20	Apr20-May20	62.25	35.7	0.048
Bug-fix Commits	Jul19-Dec19	Jan20-May20	69.92	32.54	0.028

4.1.2 *Impact on Discussion*

From our mining part of the study, we found that in case of creating and updating pull request comments (which is a way of discussion among developers), more projects are showing decreasing trends during the pandemic period. For example, time-windows between Jan'20-May'20 (TW6) show more projects with a decreasing trend in comparison to the time-windows between Jan'19-Dec'19 (See Table 2) in both creating and updating comments. Next, we performed a pairwise two-sample t-test to check if there is any statistically significant difference in the mean between these time-windows and did not find any statistically significant difference in mean between any two time-windows (p-value $>$ 00.5, see Table 3 for details).

We asked our survey respondents about the review activity (another form of discussion among developers) in their projects during the pandemic. 54.83 percent of the survey respondents mentioned that, in general, review activities are not lower than usual during the pandemic period. 56, 97, and 60 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.41 (i.e., between "disagree" and "neutral"). In addition, the Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.753).

The following are some comments that refute the statement:

×*"I having more structure timed schedule. Review with the team, coding, review with the team."*

We also asked about the rate of discussion among team members in our survey. 37.63 percent respondents did not think there is an increment in discussion among team members since they began working from home in the pandemic period. 19, 86, and 73 respondents strongly disagree, disagree, and are neutral in this, respectively. The average Likert score

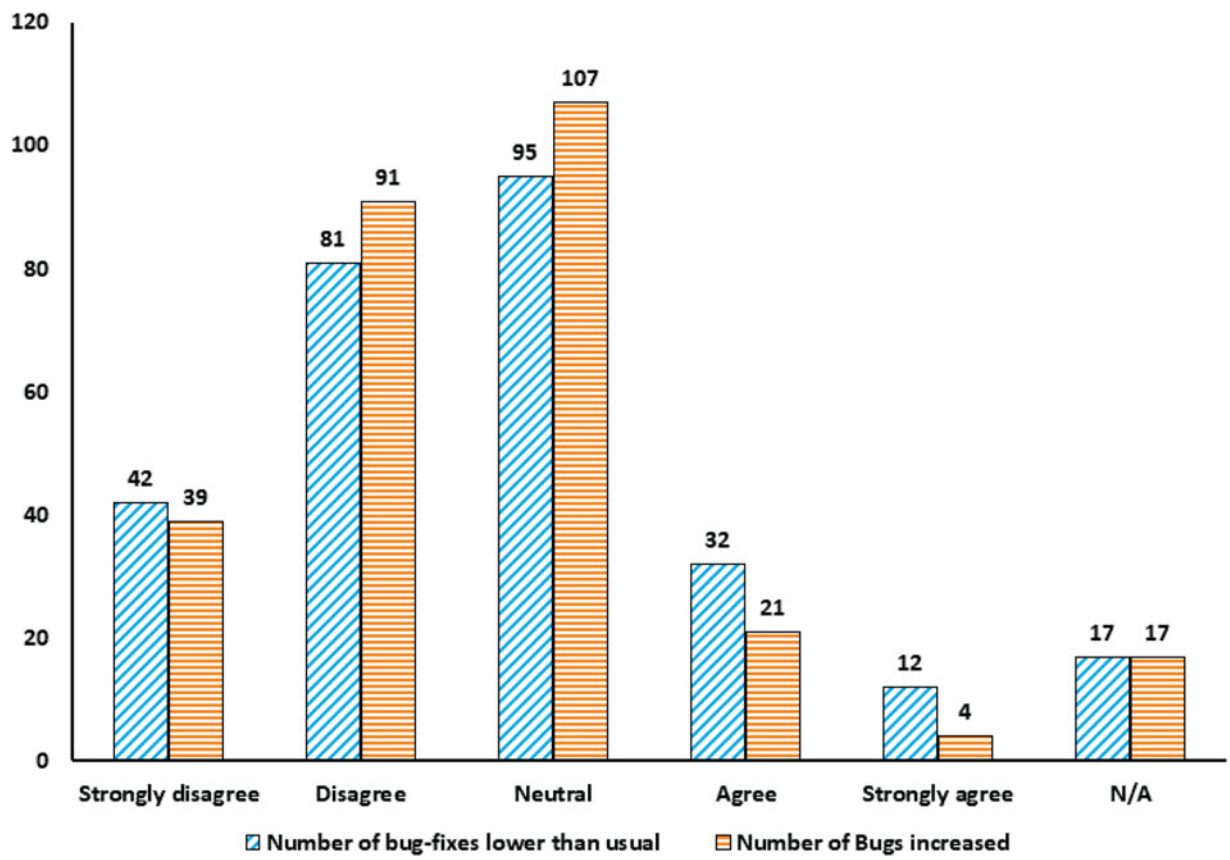

Fig. 3. Summary of the survey results on the impact of COVID-19 on the bug-fixing activities.

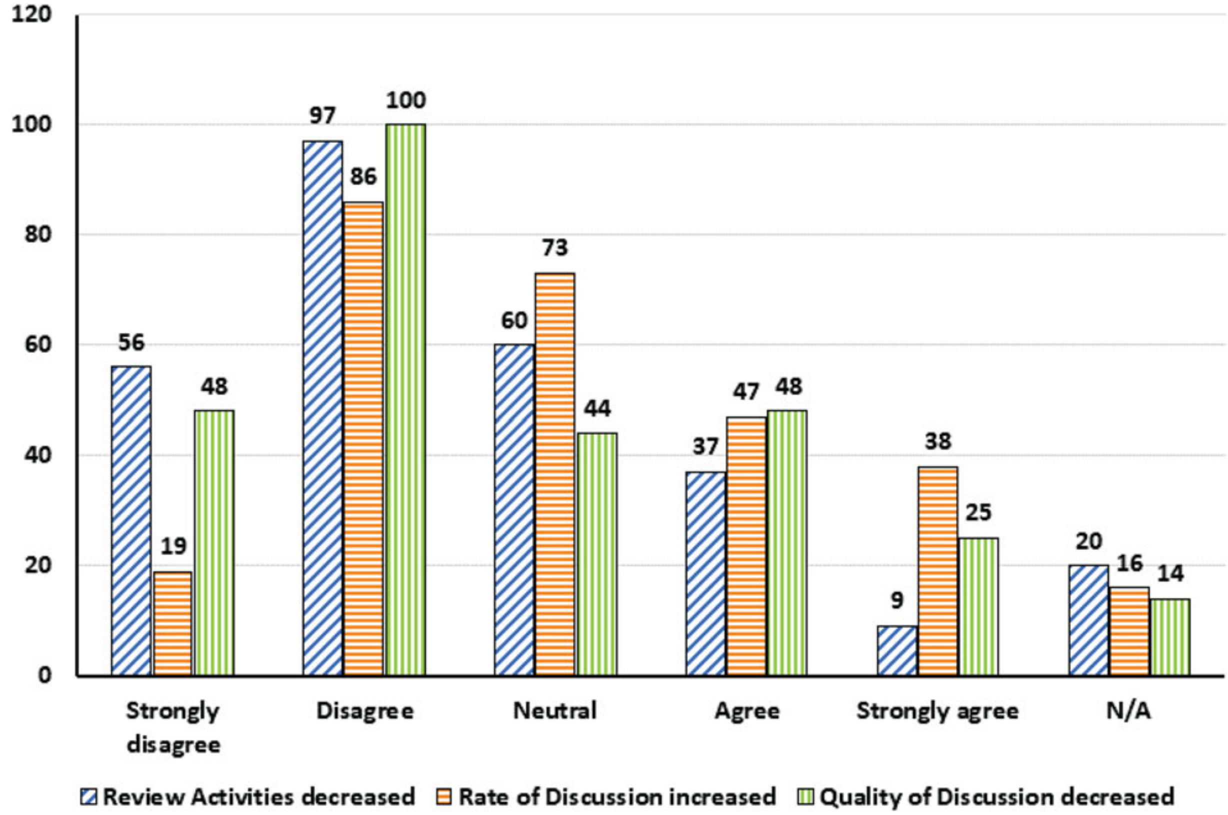

Fig. 4. Summary of the survey results on the impact of COVID-19 on discussions among developers.

for this statement is 3.0 ("neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.4571).

The following are some comments that refute or confirm the statement:

×*"Normally during work days, at least on the team/office I worked with/on we would have some time to discuss about problems that occur to us during the development of some task. Now with everyone so distant, even though we have tools to workaround that problem it still feels like you are significantly distant from that person, both physically and mentally."*

✓*"On our company we have a permanent virtual room opened to attend doubts and have discussions."*

> **Observation 3:** Though the number of pull request comments created and updated decreased over the pandemic period, discussion among developers related to their work and review activity did not decrease during the pandemic.

53.04 percent of the respondents disagree on the quality degradation of discussion among team members since they began working from home, while 26.16 percent agree that there is a degradation in discussion quality. The average Likert score for this statement is 2.63 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.7165).

The following are some comments that refute or confirm the statement:

×*"I make calls with my colleagues the whole time, to discuss and cheap talk, this helps focus and fell like we are at the same room."*

✓*"In general, people are more nervous and stressed, which makes discussions less constructive and more stressful."*

Fig. 4 shows the summary of the survey results on the impact of COVID-19 on discussion among developers.

4.1.3 Impact on Code Contribution

To investigate the impact of COVID-19 in code contribution, we analyzed the number of pull requests created, closed each month during and before the pandemic period.

From our mining analysis, we found that in general, more projects show increasing trends in the number of created and closed pull requests per month during the pandemic period. For example, time-windows between Jan'20-May'20 (TW6) show more projects with an increasing trend in comparison to the time-windows between Jan'19-Dec'19 (See Table 2). Next, we performed a pairwise two-sample t-test to check if there is any statistically significant difference in the mean of the number of created pull requests between these time-windows and found that between the time-windows Jul'19-Dec'19 and Jan'20-May'20, there is a statistically significant difference in their means (p-value < 0.05). For the closed pull requests, we found that between the time-windows Jan'20-Mar'20 and April'20-May'20, there is a statistically significant difference in their means (p-value < 0.05, see Table 3 for details).

We also found that the pull requests that opened and closed in the same month show increasing trends in more projects during the pandemic compared to the time-windows before the pandemic. We found a statistically significant difference in mean for this metric between the time-windows Jan'20-Mar'20 and Apr'20-May'20.

In our survey, we asked our respondents about code contribution. There are different ways of contributing code,

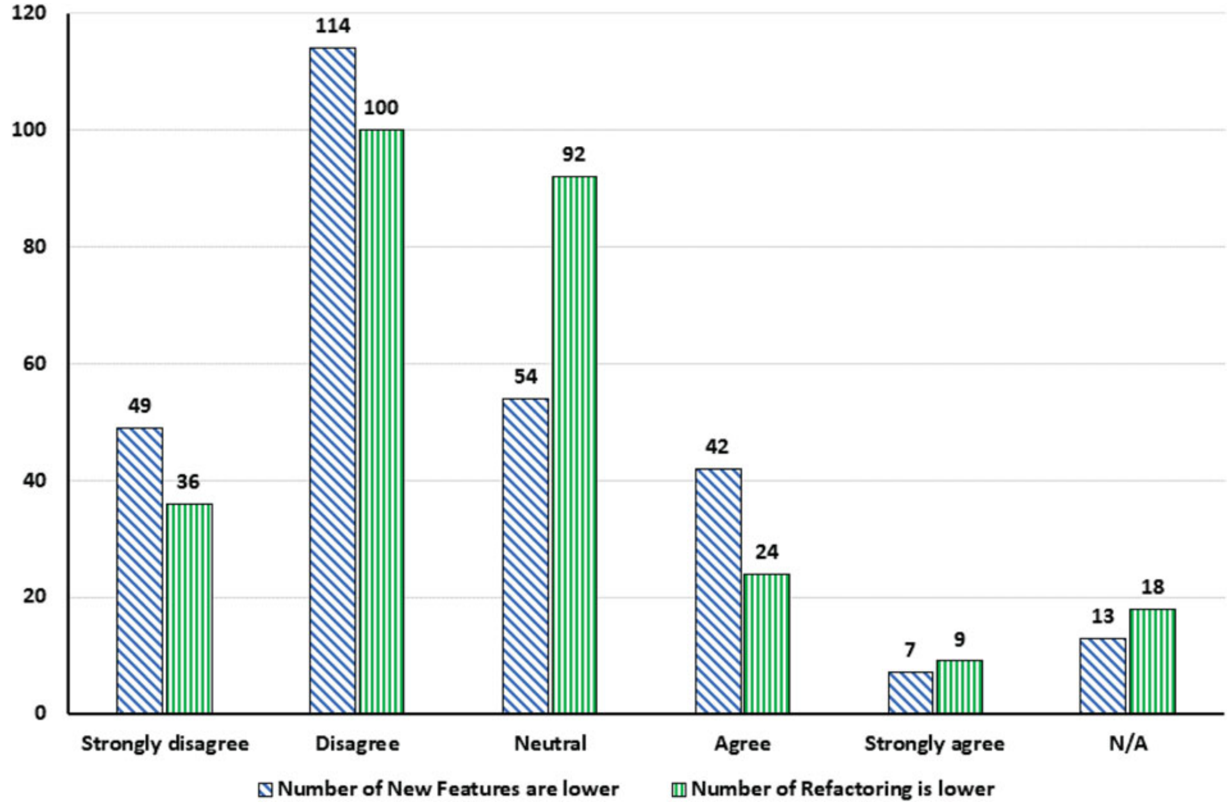

Fig. 5. Summary of the survey results on the impact of COVID-19 on code contribution.

such as adding a new feature, refactoring, etc. 58.42 percent of our survey respondents do not think that working from home during pandemic results in adding fewer new features than usual time. 49, 114, and 54 respondents strongly disagree, disagree, and are neutral with the statement: "The number of new features is lower than usual" statement, respectively. The average Likert score for this statement is 2.41 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.6425).

The following are some comments that refute the statement:

×*"I am doing home office. It's effective to keep the features delivery."*

×*"Overall we've been finishing more features at a higher quality. However, we have limited QA resources (we only have one full-time QA staff member), so the fact that the developers are churning out code is probably making it much more stressful for our QA lead."*

48.75 percent of our survey respondents disagree that the number of refactoring is lower than usual since they began working from home. 36, 100, and 92 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.51 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.833).

The following are some comments that confirm the statement:

✓*"Besides of interpersonal communication be affected, I perceived that long term strategies about the quality and refactoring, for example, are depreciated instead of short term tasks, such as new implementations and tests automation related to the implementations."*

Observation 4: In general, working from home during pandemic period does not impact code contribution.

Fig. 5 shows the summary of the survey results on the impact of COVID-19 on code contribution.

In our survey, we also asked the survey respondents about task completion time, productivity, and quality of their work during the pandemic period.

47 percent survey respondents disagree that it takes longer to complete the tasks than usual since they began working from home. 58, 73, and 52 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.72 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.989).

The following are some comments that refute or confirm the statement:

×*"I've more time to study my tasks, with more attention and patience, I've created step by steps before to start working and this is more effective."*

×*"We scheduled more small conversations during the day and splitted more the tasks. We are keeping the same productivity as before or even higher."*

✓*"Tried using the Pomodoro method of time boxing tasks. Otherwise, nothing. I don't think it helped much."*

✓*" Work more hours to complete necessary work, try to be more organized with tasks and time available (prioritization). Not so effective, I still have to work more hours."*

On the overall productivity, 56 percent of the survey respondents disagree they are less productive since they

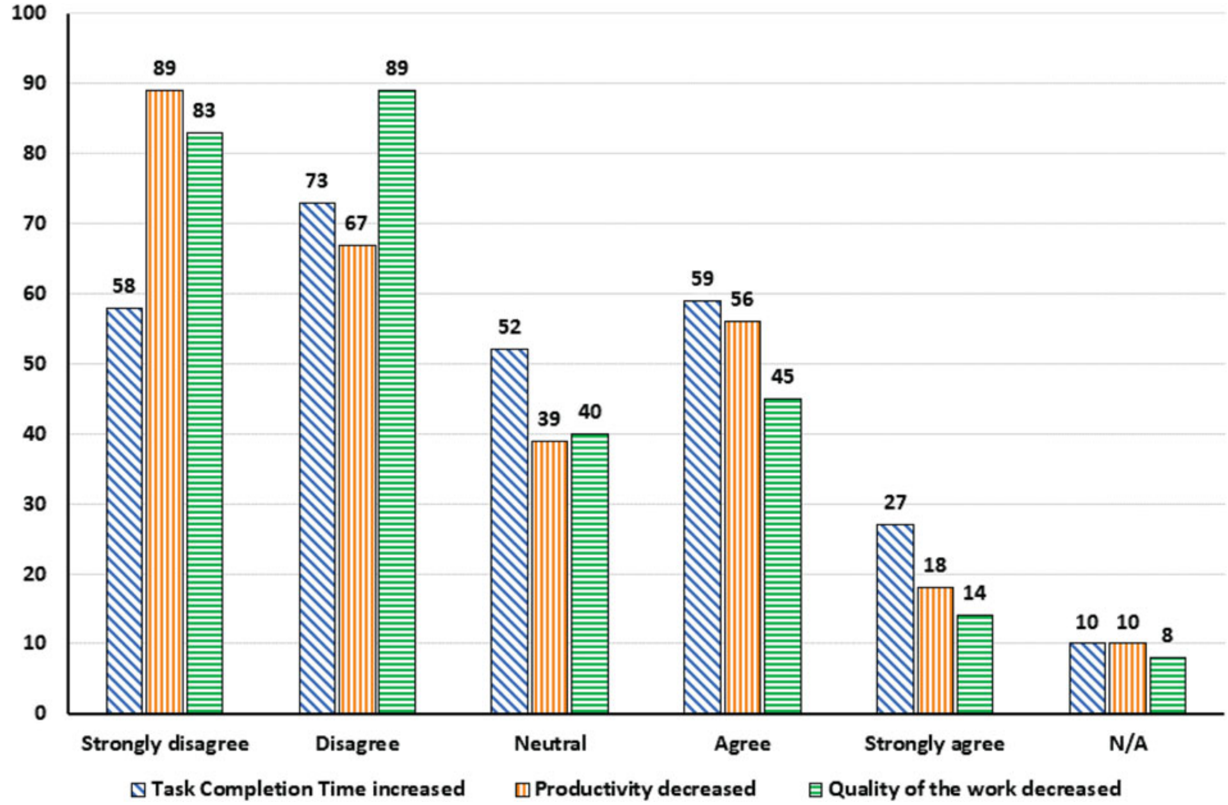

Fig. 6. Summary of the survey results on the impact of COVID-19 on task completion time, productivity, and quality of the work.

began working at home due to the COVID-19 pandemic. 89, 67, and 39 respondents strongly disagree, disagree, and are neutral, respectively. The average Likert score for this statement is 2.43 (i.e., between “disagree” and “neutral”). The Fisher’s exact test did not identify any difference between OSS and industry participants (p-value = 0.7426).

The following are some comments that refute or confirm the statement:

×″*My productivity at work increased a lot in this pandemic and I was even recognized as an QA MVP (Most Valuable Player) for the last 2 months. However, I had to keep track of my working hours because I started working a lot more than I was supposed to.*″

×″*My team became more productive due to the absence of office distractions.*″

✓″*We had been working remotely for years already. However, our productivity still took a big hit - not due to working from home specifically, but due to the cognitive overload we’re all experiencing these days. It’s hard to concentrate and be productive when we’re all worried about our families, our society and our own health.*″

✓″*Decreased one-on-one conversations, low internet bandwidth, irregular power supplies and many other factors have impacted our effectiveness.*″

In general, respondents disagree or are neutral considering that the quality of their work is lower than it should have been since they began working from home. 83, 89, and 40 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.33 (″i.e., between ″disagree″ and ″neutral″). The Fisher’s exact test did not identify any difference between OSS and industry participants (p-value = 0.614).

> **Observation 5:** Overall productivity and task completion time do not decrease during the pandemic period compare to the usual time.

Fig. 6 shows the summary of the survey results on the impact of COVID-19 on task completion time, productivity, and quality of the work.

In general, respondents disagree or are neutral, considering that the amount of testing is lower than usual since they began working from home during the pandemic. 54, 110, and 69 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.28 (i.e., between “disagree” and “neutral”). The Fisher’s exact test did not identify any difference between OSS and industry participants (p-value = 0.1658).

The following is a comment that confirms the statement:

✓″*Besides of interpersonal communication be affected, I perceived that long term strategies about the quality and refactoring, for example, are depreciated instead of short term tasks, such as new implementations and tests automation related to the implementations.*″

Many respondents disagree or are neutral, considering that the code quality has decreased since they began working from home. 62, 112, and 66 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.21 (″i.e., between ″disagree″ and ″neutral″), which is the lowest among all survey questions. The Fisher’s exact test did not identify any difference between OSS and industry participants (p-value = 0.419).

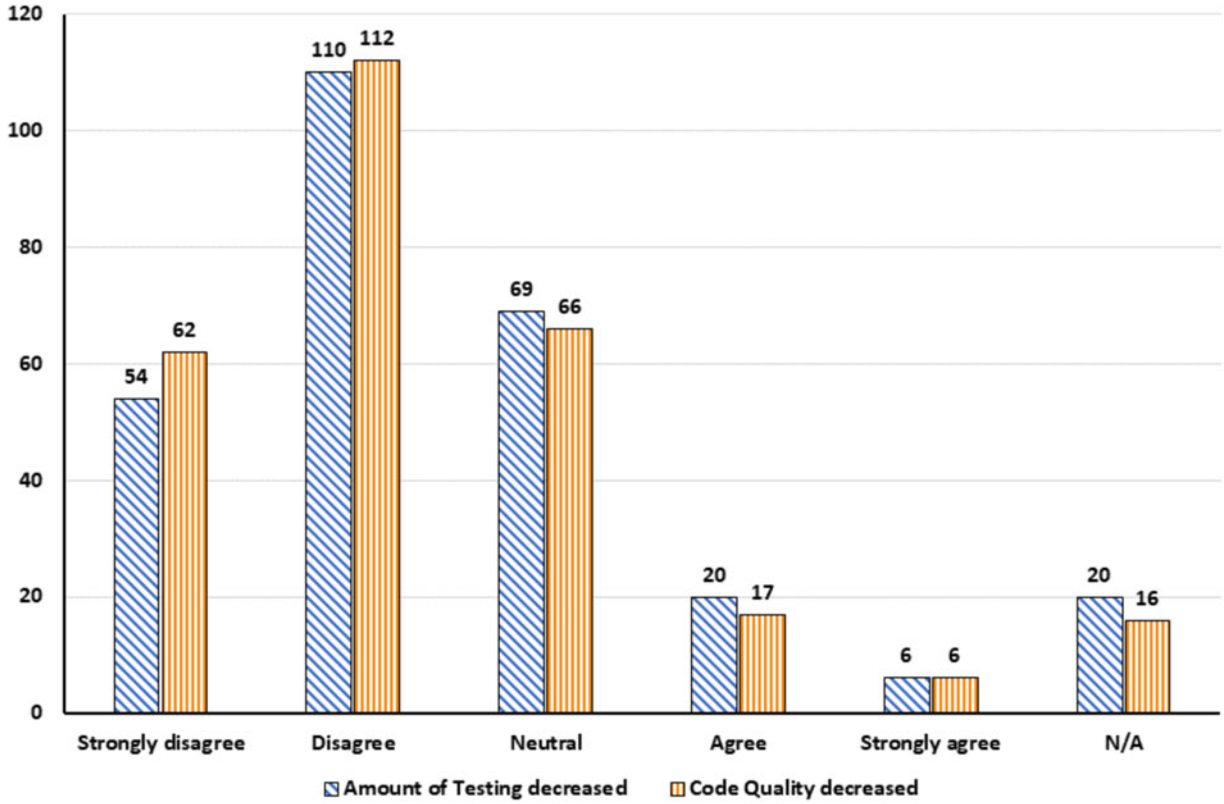

Fig. 7. Summary of the survey results on the impact of COVID-19 on code quality.

Observation 6: According to our survey respondents working from home during pandemic period does not impact code quality.

Fig. 7 shows the summary of the survey results on the impact of COVID-19 on code quality.

4.2 RQ2: What is the Impact of COVID-19 on Developers Wellbeing?

With the second research question our goal is to identify the impact of COVID-19 on developers wellbeing. To achieve this goal, we asked the survey respondents regarding their stress, emotional condition, getting help from others, etc. In the following sections, we present our findings regarding the impact of COVID-19 on developers wellbeing.

4.2.1 Sleep Disorder

Sleep disorder could hamper developer's productivity. In our survey, we asked the respondents if they think their sleeping disorder increased during the pandemic period. Among 279 respondents, 39.40 percent of the respondents agreed that their sleep disorder increased since they began working from home. The average Likert score for this statement is 3.0 (i.e., "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.5156).

The following are some comments that refute or confirm the statement:

×*"This situation is quite complex, especially when we perceive a complete neglect on the part of the federal authorities in Brazil with the pandemic and at the same time see the increasingly devastating advance of the virus in Brazilian territory. The most important action has been to control the amount of information to maintain emotional control. Apparently just for sleep it hasn't been enough."*

✓*"I'm going to sleep at the same time. It was effectively strong."*

✓*"I gave sleep a priority. Discontinues some non-essential activities."*

Observation 7: 39.40 percent of the survey respondents agreed that sleep disorder increased during pandemic period.

4.2.2 Level of Stress

We found that 141 out of 279 survey respondents agreed that the level of stress increased in the pandemic period since they began working from home. A substantial number of respondents choose to be neutral or disagree (44 neutral respondents, and 85 respondents who disagree or strongly disagree). The average Likert score for this statement is 3.33 ("i.e., "neutral") which is the highest among the survey questions. The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.8036).

The following are some comments that refute or confirm the statement:

×*"Constant breaks (10minutes every 50 minutes). No overtime. When my workday is done, I close my laptop and put in a drawer ("Out of sight, out of mind") and get it out only on the next working day. I feel less stressed because I have more time overall (no commuting)."*

×*"The pandemic is serious, but it is far away to be my source of stress. I have low need for social interaction. Sunbathing with the baby is enough to relieve the "quarantine"."*

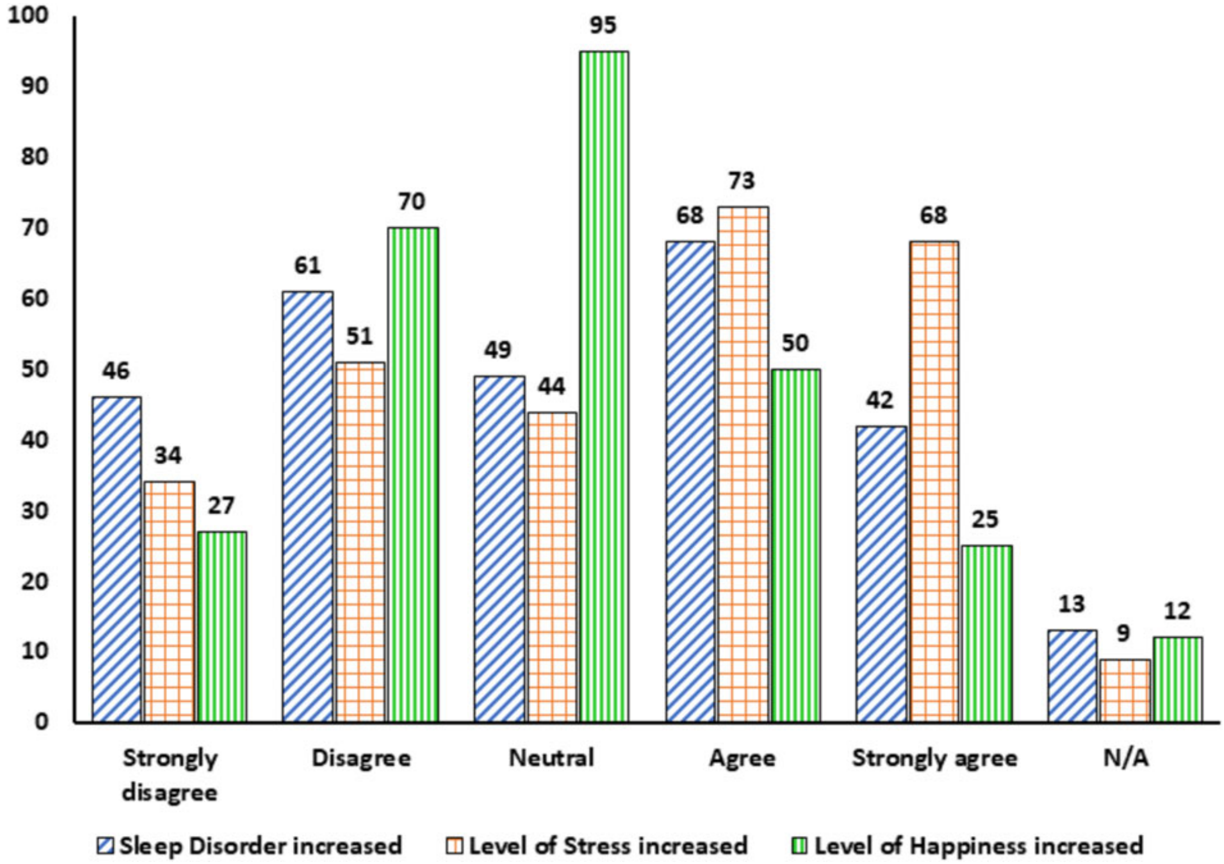

Fig. 8. Summary of the survey results on the impact of COVID-19 on the developer's wellbeing.

✓"*Working more flexible hours to watch family more during day and shift more at night. This has maintained work output and quality but increases stress.*"

✓"*In general, people are more nervous and stressed, which makes discussions less constructive and more stressful.*"

Observation 8: 50.53 percent of survey respondents indicate that their stress level increased since they began working from home during pandemic.

4.2.3 Happiness

Many respondents disagree or are neutral considering that the level of happiness increased since they began working from home. 27, 70, and 95 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.91 ("i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 1).

The following are some comments that refute or confirm the statement:

×"*I am depressed and taking anti-depression pills along with therapy. I check a doctor online.*"

×"*Slightly increased the frequency of my panic attacks.*"

×"*I try to get just enough news to stay informed but other than that I try to avoid news because the more I hear, the more depressed I get. I have a regular schedule of zoom dates with my boyfriend on top of our impromptu calls/facetimes to make sure that we are staying as connected as we can. That helps to some extent because we know we'll get through it, but it still completely sucks that we don't get to see each other in person. To make up for it, we probably text more frequently during work hours than we would have before the pandemic. At work, my team has weekly 30min meetings to just hangout together on zoom. They were good at first to help us stay connected and sane, but now I think we're kind of getting bored with them. We probably don't need them so frequently.*"

✓"*My overall wellbeing was not affected. Just the mood and motivation is down due to the situation. There is not much to do about it. We just have to get through it.*"

Observation 9: 34.76 percent of the survey respondents mentioned that the level of happiness did not increase during the pandemic period.

Fig. 8 shows the summary of the survey results on the impact of COVID-19 on developer's sleep, stress level and happiness.

4.2.4 Mentoring

101 out of 279 survey respondents agreed that the mentoring activities for newcomers has decreased in the project since they began working from home in pandemic. A substantial number of respondents choose to be neutral or disagree (66 neutral respondents, and 81 respondents with disagree or strongly disagree with this statement). The average Likert score for this statement is 3.15 (i.e., "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.8298).

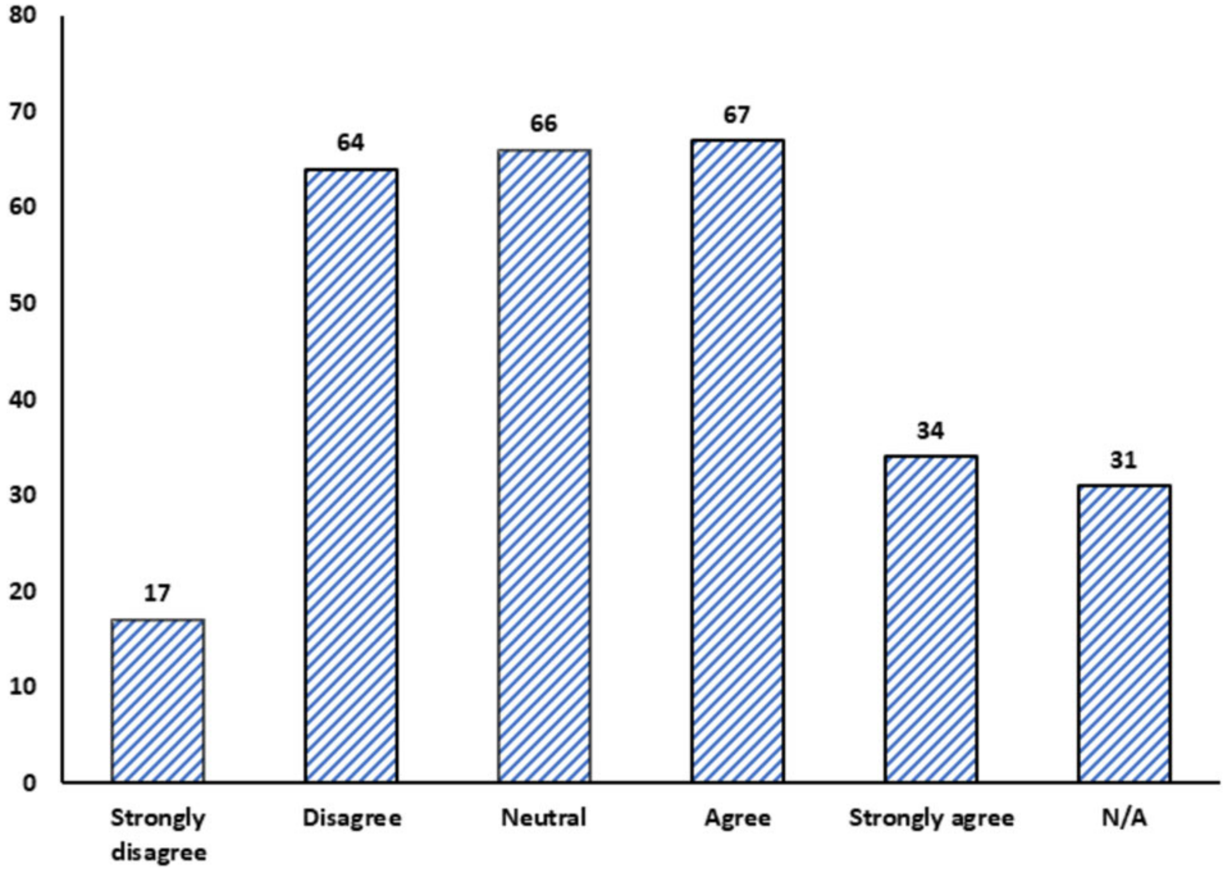

Fig. 9. Summary of the survey results on the impact of COVID-19 on the developer's mentoring activity.

The following are some comments that confirm the statement:

✓"*I think it affected the welcoming of the youngest to the team. Face-to-face activities generate more intimacy between people.*"

✓"*The COVID-19 has influenced the acceptance of new employees. The efficiency of the discussion among team members was much lower than face-to-face discussion.*"

> **Observation 10:** 35.12 percent of the survey respondent mentioned that mentoring activities for newcomers have decreased during the pandemic period.

Fig. 9 shows the summary of the survey results on the impact of COVID-19 on the developer's mentoring activity.

4.2.5 Interruptions

In general, respondents disagree or are neutral considering that newly introduced interruptions along with the old ones are negatively impacting their productivity since they began working from home. 35, 81, 64 respondents strongly disagree, disagree, and are neutral with this statement, respectively. The average Likert score for this statement is 2.87 (i.e., between "disagree" and "neutral"). The Fisher's exact test did not identify any difference between OSS and industry participants (p-value = 0.7762).

The following are some comments that refute or confirm the statement:

×"*Working remotely, others can't interrupt my work so quickly, so this increased my productivity in general.*"

×"*Home office leads to much less interruptions. This helps productivity a lot!*"

✓"*I had to improvise a home office in my daughter's room so I can isolate myself there to avoid too many interruptions or background noises. It helped a little, but my wife and daughter still eventually interrupt me during my working hours.*"

✓"*We tried to schedule a series of video conferences to keep in touch with the team. It helped a lot, but is not as good as talking directly to the persons. The rate of interruptions did increase.*"

> **Observation 11:** 41.60 percent of the survey respondents do not think that newly introduced interruptions along with the old ones are negatively impacting their productivity during the pandemic time.

We posit that interruptions are associated with number of people living in the household. We analyzed the results from the survey and our results show that the average Likert score ranges from 2.9 to 3.4, indicating the perception that the number of people living in the household is not having a major impact.

We also investigated whether interruptions are associated with number of children below 12 years of age living in the household. We analyzed the results from the survey and our results show that the average Likert score ranges from 2.0 to 3.2, indicating the perception that the number of children also does not have a major impact. On the contrary, people mostly disagree with this perception. Fig. 10 shows the summary of the survey.

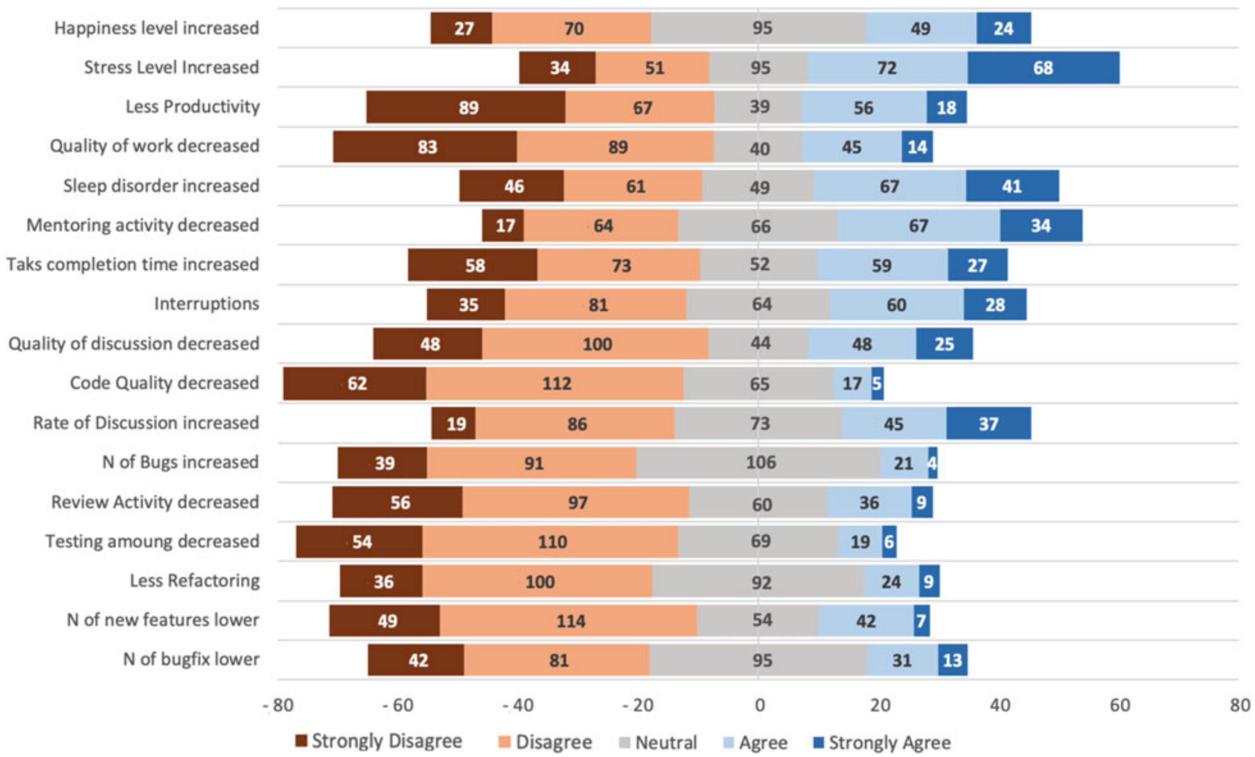

Fig. 10. Survey results summary.

Observation 12: Majority of the respondents do not perceive the number of kids and people living in the house as a source of interruption.

5 Discussion

Software development under COVID-19 introduces some challenges that are valid for any software engineers working from home (even prior to the pandemic). For these challenges, prior strategies that enable people to work effectively from home still apply. However, this pandemic also introduces additional challenges to software development. We highlight the challenges and their implications to practitioners, organizations, and researchers.

Practitioners. From our study, we found several suggestions for the practitioners to maintain a healthy work-life balance during a pandemic. Creating various "forms of entertainment at home" can help. Staying at home provides opportunities for practitioners to expand their knowledge and utilize this time to study and take online courses as several universities and platforms have their offerings. While pandemic has been tough for everyone, it has been positive for many with "better sleep, more time with family, less travel time". Practitioners recommend taking "more small breaks during the day to spend time with family and pets". Others mentioned the benefit of being disciplined in separating work time from rest/family time; for example, one respondent mentions: *"No overtime. When my workday is done, I close my laptop and put in a drawer ("Out of sight, out of mind") and get it out only on the next working day."*

Software development under COVID-19 introduces additional challenges to developers than those who work from home prior to the pandemic. For example, one respondent mentions: *"I had to improvise a home office in my daughter's room so I can isolate myself there to avoid too many interruptions or background noises. It helped a little, but my wife and daughter still eventually interrupt me during my working hours."* Another mentioned: *"I'm used to work from home, but the main problem is that schools are closed since the last week of Feb here in Italy and stay concentrated with a bored 8 years old girl at home is not always easy."* The above highlights that software developers often need to share space with family members, who are unable to go to do their daily activities (school, work, etc.) due to the lock downs, movement restriction policies, closure of schools and day care facilities, etc. The pandemic also causes developers to make a *sudden transition* from work in office to work from home. The sudden transition is often hard; as one's home may be noisy and does not have a suitable space to allow one to work comfortably. To deal with this challenge, many developers have "improvised" their homes in a short period of time, and work with their family members or house mates so that everyone can effectively share their limited space with workable level of distraction. Another respondent mentions: *"I tried to find a space in my house where people wouldn't distract me. Overall my family has been respectful about my work schedule."* This level of adaptability and cooperation is needed for effective work arrangement during this pandemic. Many respondents also mention using the Pomodoro technique to deal with the distractions; it uses a timer to break down work into intervals, and has helped "improve productivity".

The pandemic also causes additional stress beyond the usual workload that a software practitioner needs to bear. One respondent mentions: *"Stopped reading the news, started planning my meals and domestic activities so they wouldn't take a lot of time."* Another mentioned: *"The other problem is that it is unavoidable to keep thinking to the current situation and what's happening around you and then being distracted by it. In a word lack of focus and* *concentration has been my main problem in this period.".* Yet another mentions: *"There are fewer distractions in my home than at the office, so I feel that overall productivity has increased. The exception to this would be an occasional lack of mental clarity due to anxiety from world events, but that is more related to overall well being."* These highlight the stress that comes from worries caused by the bad news, which impact focus and concentration – qualities needed for software development. The first quote also highlights stress that comes from additional chores that one needs to do caused by one's need to stay more at home, and take necessary precautions to stay safe and healthy. To reduce their level of stress, many respondents mention of reducing their exposure to bad news, as well as taking steps to reduce the time needed for the extra chores. One respondent mentions: *"Use grocery delivery services to compensate the extra time spent with objects hygiene. Highly effective."* Moreover, collegiality is also an important factor, and it takes a wider meaning in the time of pandemic. One respondent mentions: *"I organized my schedule better to communicate to my co-workers about what I am doing. We create groups to share what our team is working and talk about amenities, this is a way to be in touch and know if everyone is well or need help."* Another respondent mentions: *"Assured the team that work hours are even more flexible than before. Encouraged team members to take care of themselves, be careful not to work excessive overtime and recognize when they need to take time off. Remind the team to continue using our standard development and quality practices. Talking about it and encouraging feedback have been somewhat effective. Some team members are more productive, a few are less and most are about the same."* Establishing a good communication among colleagues, as well, as talking about things beyond work (e.g., well-being, positive news), encouraging and assuring one another, and offering help go a long way to create a workable environment for everyone to effectively go through this pandemic.

From the responses, we find that the impact of COVID-19 to individual practitioners varies, and depends on one own's specific environment and situation. Some have worked from home for many years and they are less affected ("I've worked remotely since years ago, there is no change on productivity.", "Nothing in particular. We were already working from home a couple of days per week before the pandemic. If anything, productivity increased with permanent working from home."). Others need to make a sudden transition. Some live alone, which introduces another challenge. Some are used with online contact ("Personal relationship-wise I was already used to online contact, so that side didn't affect me much. "), others are not ("people ... are missing personal contact"). Some live in homes that are more conducive to work than others. A good level of understanding and reasonableness go a long way. One respondent commented: *"Be aware that not everybody is going to be prepared for this moment. ... People are alone on their homes, susceptible to mental illness, stress is increasing among familiars. People might be edgy about this situation and you have to understand them. I think these steps proved to be really effective on me during these moments."*

Organizations. Based on our analysis, we found that the COVID-19 situation does not necessarily result in reduced productivity and inferior code quality. Organizations can take active steps to help developers cope with COVID-19 and remain productive and produce high-quality code. We received many inputs from our respondents including the following: *"On our company we have a permanent virtual room opened to attend doubts and have discussions."*, *"At work, my team has weekly 30min meetings to just hangout together on zoom. They were good at first to help us stay connected and sane"*, *"Daily meetings during the beginning and end of the day, it allows visibility for the whole team and visibility about what is being executed. It was too efficient that the team is producing more"*, *"My team opened up a remote call which everybody should be present during all day. This way we 'simulate' the real-world environment which we can talk with the other team members anytime. We noticed that the team collaboration has not decreased"*. From the above, we can note that organizations can positively impact their developers ability to cope well with COVID-19 in various ways to simulate their previous working environment (before the pandemic) and facilitate more interactions between the developers. In such a way, developers can effectively collaborate and feel less isolated in doing their tasks. Some of the arrangements (e.g., "permanent virtual room", "remote call ... during all day") were uncommon prior to the pandemic.

Organizations can also consider making working hour adjustments and allowing for more flexibility. One person commented: *"Adjusted the timings to work because of the kids. Our team has no meeting Fridays. This gives peaceful time."* Another commented: *"We split working hours among my wife and me. So we are working in turns and other one is taking care of kids."* Organizations can also provide professional support even psychotherapy. One respondent mentions: *"Continue psychotherapy through videoconferencing. No success so far due to health insurance provider bureaucracy. I'm now selecting another therapist from ... (subscription-based videoconference therapy service paid by the company). Therapy helped me with many insights about myself that made me improve the way I work."*

Still the organization needs to maintain some reasonable expectations of work quality and quantity; one respondent commented: *"Accountability has decreased in the team. Everyone now has excuses to not deliver or for subpar work "*. Another commented: *"Reduced revenue, reduced workforce, reduced salaries are common in the commercial industry."* If an organization is not careful, it may not be viable anymore and this may result in reduced workforce and salaries. Organizations also need to care especially for people who are new and evolve the existing on-boarding process. One commented: *" I recently started in a new job with much more responsibilities and the pandemic off course is turning the adaptation quite more hard"*. Another commented: *"I think it affected the welcoming of the youngest to the team. Face-to-face activities generate more intimacy between people."*

Organizations also need to care of not introducing too many meetings and micro-managing things, as one respondent lamented: *"The amount of meetings increased a lot and some sort of micromanagement has also increased."* Organizations also want to provide quarantine support especially if they employ migrant workers; one respondent mentions: *"at least one remote worker was infected. At least three live with people that was infected. The company also offered apartments for their quarantine"*. Additional support would also be valued, as one respondent mentioned: *"the company is also lending office equipment and paying a bonus to all workers with lower earnings to help with extra expenses (energy an internet bills, home office equipment, etc)."*

Researchers. Our findings highlight that the impact of COVID-19 is not binary (reduce productivity versus increase productivity) but rather a spectrum. For many of our observations, substantial proportions of respondents have differing opinions from each other. For example, for Observation 3 (*Though the number of pull request comments created and updated decreased over the pandemic period, discussion among developers related to their work and review activity did not decrease during the pandemic.*), 37.63 percent of respondents did not think there is an increment in discussion among team members since they began working from home in the pandemic period, with an average Likert score of 3.0 ("neutral"). This variation may be due to variation in the lifestyles, organizations, countries, cultures and geographies of the participants whom we surveyed. Moreover, the pandemic does not hit everywhere at once and the effects differ, as has become evident when comparing different regions and socioeconomic groups. The responses may also differ depending on various other factors including political ones. Thus, we should not be surprised for the differences in how people perceive the pandemic in their own lives.

More research is needed to uncover specific conditions that cause certain outcomes to be more prevalent. These conditions can correspond to personalities of different developers, their roles, their organization structure, and characteristics, their home conditions (e.g., "low internet bandwidth", "irregular power supplies"), etc. To illustrate the impact of these factors, one of our respondents mentions "... *[I] have a low need for social interaction. Sunbathing with the baby is enough to relieve the "quarantine"* while another mentions "*decreased one-on-one conversations ... impacted our effectiveness*". Yet another mention: "*Slightly increased the frequency of my panic attacks*". This suggests that personality plays a role in the impact of COVID-19 to developers (c.f., [48] that also shows the impact of personality on the developer and team effectiveness prior to COVID-19). Developers with some physical/mental health conditions (e.g., panic attacks) may also be more adversely impacted by COVID-19.

More research is also needed to refine the existing processes to cope with a closer integration between work and family life, which results from this pandemic. For example, research can be done to better answer questions such as: "How to make good arrangements so that there are good overlap among team members' working hours (to allow for synchronous communication), while still accommodating individual team member personal constraints?", "How to design a good onboarding process in this pandemic?", "How effective is Pomodoro? Is there a better method that can help practitioners cope with the high-level of distractions that many need to face?", "How to balance organizational need (completion of projects) with individual need?", "How to deal with communication challenges that may be introduced by not being able to meet face-to-face (as one respondent mentioned: "Also people just don't seem to be as talkative on zoom as in person, so it's a little harder to get people to participate.")?", etc.

The understanding gained from further research can result in: (1) creation of specific guidelines that can help developers or organizations adversely affected by COVID-19 to learn from other developers or organizations that have coped well with COVID-19, (2) Organizations to adopt different strategies to help developers of different personalities and conditions to cope with COVID-19.

We do not view our study as a final definitive study, but rather one of the many that can shed "full" light into COVID-19 (or other pandemics), its impacts, and ways to mitigate those impacts. Some of the future studies can consider performing a smaller scale but more in-depth and focused study on a particular aspect (e.g., the impact of personality on how developers cope with COVID-19).

6 Threats and Limitations

Our research findings may be subject to the concerns that we list below. We have taken all possible steps to reduce the impacts of these possible threats and limitations, but some could not be mitigated and it is possible that our mitigation strategies may not have been effective.

Credibility. The analysis of 100 GitHub projects developed in Java using ten different metrics and survey responses from 279 software development professionals serve as a rich and credible data set for the analysis.

Originality. The paper is the first to investigate the impact of COVID-19 on software development by gathering developers' perceptions through surveys and mining software repositories, and gathering a deeper understanding of the impact.

Internal Validity. The set of analyzed metrics spans across 10 different categories and widely used in literature. However, we cannot guarantee that our set of metric are exhaustive. We plan to expand our metric set in future work. In addition, it is possible that there are defects in the implementation of our mining scripts. To that end, we have extensively tested our implementation, and manually verified sampled counts of different metrics. Regarding the survey, it is possible that the survey participants misunderstood some of the survey questions. To mitigate this threat, we conducted a pilot study with developers with different experience levels from both open-source community and industry. We also conducted a pilot study with survey design experts. The pilot studies were also important to avoid the respondent fatigue bias making sure that the survey could be answered within 20 minutes. We updated the survey based on the findings of these pilot studies.

External Validity. We have not analyzed proprietary repositories, and our samples have been from a single source (GitHub) and a single programming language (Java). This may be a source of bias, and our findings may be limited to open source projects from GitHub. However, we believe that the large number of projects sampled more than adequately addresses this concern. Regarding the survey, we asked respondents from various backgrounds present in any countries. However, our results might not generalize to all the practitioners. Since the respondents of the survey are not solely the developers of the repositories, we found some contradictory results. To make our results widely applicable, we asked a broader set of participants on their perception on the impact of COVID-19 on software development.

Threats to Construct Validity. Threats to construct validity are concerned with the extent to which the setting of the experiment reflects the construct under study, which include the potential problem of evaluation apprehension [49]. It was mitigated by the anonymity of the participants, as well as the guaranty that all information gathered during the survey would be used only by the research team.

7 Conclusion

The COVID-19 pandemic has impacted the whole world in different ways. As software is still "eating the world", it is essential to understand COVID-19's impact on software projects. We conducted a mining software repository study based on 100 GitHub projects developed in Java using ten different metrics. Next, we surveyed 279 software development professionals from 32 countries for gathering more insights about the impact of COVID-19 on daily activities and wellbeing.

Based on our findings, we derived 12 observations that can be used by practitioners, organizations, and researchers. Practitioners can use our recommendations to maintain a healthy work-life balance during a pandemic. Organizations can learn from our survey respondents and take steps to remain productive while creating high-quality code. The research community can explore the social and human aspects to understand the impact of developer personality during a pandemic.

As future work, we intend to expand the analysis of the projects (considering a wider time window) to better understand the impact of a year-long pandemic on software development. We also plan to investigate the impact of COVID-19 using other metrics reported in literature as indicators of productivity.

Acknowledgments

This work was supported in part by INES (www.ines.org.br), CNPq under Grant 465614/2014-0, in part by CAPES under Grant 88887.136410/2017-00, and in part by FACEPE under Grants APQ-0399-1.03/17 and PRONEX APQ/0388-1.03/14. We would like to thank all the participants in our survey.

References

[1] A. Fauci, H. Lane, and R. Redfield, "COVID-19–Navigating the uncharted," *New Engl. J. Med.*, vol. 382, pp. 1268–1269, 2020.

[2] Q. Li *et al.* "Early transmission dynamics in Wuhan, China, of novel Coronavirus–infected pneumonia," *New Eng. J. Med.*, vol. 382, pp. 1199–1207, 2020.

[3] I. Chakraborty and P. Maity, 'COVID-19 outbreak: Migration, effects on society, global environment and prevention," *Sci. Total Environ.*, vol. 728, pp. 138882, 2020.

[4] J. D. Herbsleb and A. Mockus, "An empirical study of speed and communication in globally distributed software development," *IEEE Trans. Softw. Eng.*, vol. 29, no. 6, pp. 481–494, Jun. 2003.

[5] J. D. Herbsleb, "Global software engineering: The future of socio-technical coordination," in *Proc. Future Softw. Eng.*, 2007, pp. 188–198.

[6] M. Cataldo and J. D. Herbsleb, "Communication patterns in geographically distributed software development and engineers' contributions to the development effort," in *Proc. Int. Workshop Cooperative Hum. Aspects Softw. Eng., Leipzig, Germany*, 2008, pp. 25–28.

[7] J. D. Herbsleb, "Technical perspective–Maintaining quality in the face of distributed development," *Commun. ACM*, vol. 52, no. 8, p. 84, 2009.

[8] E. Chen, K. Lerman, and E. Ferrara, "Tracking social media discourse about the COVID-19 pandemic: development of a public Coronavirus twitter data set," *JMIR Public Health Surveill.*, vol. 6, no. 2, p. e19273, 2020.

[9] E. Dong, H. Du, and L. Gardner, "An interactive web-based dashboard to track COVID-19 in real time," *Lancet Infect. Dis.*, vol. 20, pp. 533–534, 2020.

[10] R. Verity *et al.* "Estimates of the severity of Coronavirus disease 2019: A model-based analysis," *Lancet Infect. Dis.*, vol. 20, pp. 669–677, 2020.

[11] P. G. T. Walker *et al.* "The impact of COVID-19 and strategies for mitigation and suppression in low- and middle-income countries," *Science*, vol. 369, pp. 413–422, 2020.

[12] E. Zhang, N. Gupta, R. Nogueira, K. Cho, and J. Lin, "Rapidly deploying a neural search engine for the COVID-19 open research dataset: Preliminary thoughts and lessons learned," 2020, *arXiv: 2004.05125*.

[13] A. Rahman and E. Farhana, "An exploratory characterization of bugs in COVID-19 software projects," *J. Softw. Eng. Res. Develop.*, vol. 9, 2021.

[14] P. Ralph *et al.*, "Pandemic programming," *Empirical Softw. Eng.*, vol. 25, no. 6, pp. 4927–4961, 2020.

[15] D. Ford *et al.* "A tale of two cities: Software developers working from home during the COVID-19 pandemic," 2020, *arXiv: 2008.11147*.

[16] C. Miller, P. Rodeghero, M.-A. Storey, D. Ford, and T. Zimmermann, "'How was your weekend?' Software development teams working from home during COVID-19," in *Proc. Int. Conf. Softw. Eng.*, 2021, pp. 624–636.

[17] P. Devanbu, T. Zimmermann, and C. Bird, "Belief & evidence in empirical software engineering," in *Proc. IEEE/ACM 38th Int. Conf. Softw. Eng.*, 2016, pp. 108–119.

[18] T. Menzies, W. Nichols, F. Shull, and L. Layman, "Are delayed issues harder to resolve? revisiting cost-to-fix of defects throughout the lifecycle," *Empirical Softw. Eng.*, vol. 22, no. 4, pp. 1903–1935, 2017.

[19] D. P. Delorey, C. D. Knutson, and S. Chun, "Do programming languages affect productivity? A case study using data from open source projects," in *1st Int. Workshop Emerging Trends FLOSS Res. Develop.*, 2007, pp. 8–8.

[20] B. Johnson, T. Zimmermann, and C. Bird, "The effect of work environments on productivity and satisfaction of software engineers," *IEEE Trans. Softw. Eng.*, vol. 47, no. 4, pp. 736–757, Apr. 2021.

[21] E. Murphy-Hill *et al.* "What predicts software developers' productivity?," *Trans. Softw. Eng.*, vol. 47, pp. 582–594, 2019.

[22] M. Storey, T. Zimmermann, C. Bird, J. Czerwonka, B. Murphy, and E. Kalliamvakou, "Towards a theory of software developer job satisfaction and perceived productivity," *IEEE Trans. Softw. Eng.*, early access, Sep. 27, 2019, doi: 10.1109/TSE.2019.2944354.

[23] B. Vasilescu, Y. Yu, H. Wang, P. Devanbu, and V. Filkov, "Quality and productivity outcomes relating to continuous integration in GitHub," in *Proc. 10th Joint Meeting Found. Softw. Eng.*, 2015, pp. 805–816.

[24] A. N. Meyer, G. C. Murphy, T. Zimmermann, and T. Fritz, "Retrospecting on work and productivity: A study on self-monitoring software developers' work," in *Proc. ACM Hum.-Comput. Interact.*, 2017, pp. 1–24.

[25] E. Murphy-Hill *et al.* "What predicts software developers' productivity?," *IEEE Trans. Softw. Eng.*, vol. 47, no. 03, pp. 582–594, Mar. 2021.

[26] G. Bhutkar *et al.*, "User experiences and wellbeing at work," in *Proc. Hum.-Comput. Interact. Int. Conf.*, Paphos, Cyprus, 2019, pp. 754–758.

[27] M. Kuutila, M. V. Mäntylä, M. Claes, M. Elovainio, and B. Adams, "Using experience sampling to link software repositories with emotions and work well-being," in Proc. 12th ACM/IEEE Int. Symp. Empirical Softw. Eng. Measure., Oulu, Finland, 2018, pp. 29:1–29:10.

[28] D. Graziotin, F. Fagerholm, X. Wang, and P. Abrahamsson, "What happens when software developers are (un)happy," *J. Syst. Softw.*, vol. 140, pp. 32–47, 2018.

[29] A. Meyer, E. T. Barr, C. Bird, and T. Zimmermann, "Today was a good day: The daily life of software developers," *IEEE Trans. Softw. Eng.*, vol. 47, no. 5, pp. 863–880, May 2021.

[30] D. Fucci, G. Scanniello, S. Romano, and N. Juristo, "Need for sleep: The impact of a night of sleep deprivation on novice developers' performance," *IEEE Trans. Softw. Eng.*, vol. 46, no. 1, pp. 1–19, Jan. 2020.

[31] P. Ralph, "The two paradigms of software development research," *Sci. Comput. Program.*, vol. 156, pp. 68–89, 2018.

[32] S. Easterbrook, J. Singer, M.-A. Storey, and D. Damian, "Selecting empirical methods for software engineering research," in *Guide to Advanced Empirical Software Engineering*. Berlin, Germany: Springer, 2008, pp. 285–311.

[33] E. T. Barr, C. Bird, P. C. Rigby, A. Hindle, D. M. Germán, and P. T. Devanbu, "Cohesive and isolated development with branches," in *Proc. Int. Conf. Fundam. Approaches Softw. Eng.*, 2012, pp. 316–331.

[34] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian, "An in-depth study of the promises and perils of mining GitHub," *Empirical Softw. Eng.*, vol. 21, no. 5, p. 2035–2071, Oct. 2016.

[35] C. Casalnuovo, Y. Suchak, B. Ray, and C. Rubio-Gonz ález, "GitcProc: A tool for processing and classifying GitHub commits," in *Proc. 26th ACM SIGSOFT Int. Symp. Softw. Testing Anal.*, 2017, p. 396–399.

[36] H. Borges and M. Tulio Valente, "What's in a GitHub star? Understanding repository starring practices in a social coding platform," *J. Syst. Softw.*, vol. 146, p. 112–129, Dec. 2018.

[37] N. Munaiah, S. Kroh, C. Cabrey, and M. Nagappan, "Curating GitHub for engineered software projects," *Empirical Softw. Eng*.. vol. 22, pp. 3219–3253, 2016.
[38] A. Agrawal, A. Rahman, R. Krishna, A. Sobran, and T. Menzies, "We don't need another hero?: The impact of 'heroes' on software development," in *Proc. 40th Int. Conf. Softw. Eng. Softw. Eng. Pract.*, 2018, pp. 245–253.
[39] R. Krishna, A. Agrawal, A. Rahman, A. Sobran, and T. Menzies, "What is the connection between issues, bugs, and enhancements?: Lessons learned from 800+ software projects," in *Proc. 40th Int. Conf. Softw. Eng.: Softw. Eng. Practice*, 2018, pp. 306–315.
[40] A. Rahman, A. Agrawal, R. Krishna, and A. Sobran, "Characterizing the influence of continuous integration: empirical results from 250+ open source and proprietary projects," in *Proc. 4th ACM SIGSOFT Int. Workshop Softw. Anal.*, 2018, pp. 8–14.
[41] B. A. Kitchenham and S. L. Pfleeger, "Personal opinion surveys," in *Guide to Advanced Empirical Software Engineering*. London, U.K.: Springer, 2008, pp. 63–92.
[42] R. A. Fisher, "On the interpretation of x^2 from contingency tables, and the calculation of p," *J. Roy. Statist. Soc.*, vol. 85, no. 1, pp. 87–94, 1922.
[43] J. Saldana, *The Coding Manual for Qualitative Researchers*. Newbury Park, CA, USA: Sage, 2015.
[44] J. Mason, *Qualitative Researching*. Newbury Park, CA, USA: Sage, 2018.
[45] J. R. Landis and G. G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, pp. 159–174, 1977.
[46] R. Shumway and D. Stoffer, *Time Series Analysis and Its Applications With R Examples*, New York, NY, USA: Springer, 2011, vol. 9.
[47] D. Cucinotta and M. Vanelli, "Who declares COVID-19 a pandemic," *Acta Bio-medica: Atenei Parmensis*, vol. 91, no. 1, pp. 157–160, 2020.
[48] X. Xia, D. Lo, L. Bao, A. Sharma, and S. Li, "Personality and project success: Insights from a large-scale study with professionals," in *Proc. IEEE Int. Conf. Softw. Maintenance Evol.*, , Shanghai, China, 2017, pp. 318–328.
[49] C. Wohlin, P. Runeson, M. Hst, M. C. Ohlsson, B. Regnell, and A. Wessln, *Experimentation in Software Engineering*. Berlin, Germany: Springer, 2012.

Paulo Anselmo da Mota Silveira Neto received the PhD degree in computer science from the Federal University of Pernambuco, Brazil. He is currently a associate professor with the Departament of Computer Science, Federal Rural University of Pernambuco. He is a senior member of RiSE Labs. His research interests include software product lines (SPL) testing, SPL architecture evaluation and empirical software engineering.

Umme Ayda Mannan received the PhD degree in computer science, focusing on Software engineering from Oregon State University. Her research focuses on identifying factors related to projects, people, and process characteristics to predict design issues in the software using data mining and different machine learning techniques. She finished her BSc in Computer Science and Engineering from Shahjalal University of Science and Technology, Bangladesh.

Eduardo Santana de Almeida (Senior Member, IEEE) received the PhD degree in computer science from the Federal University of Pernambuco, Brazil. He is currently an associate professor with the Department of Computer Science, Federal University of Bahia, where he leads the RiSE Labs. His research interests include software reuse, software product lines, software architecture, and empirical software engineering. He is a senior member of ACM. He is a senior member of ACM and member of Brazilian Academy of Sciences (ABC).

Nachiappan Nagappan (Fellow, IEEE) received the PhD degree from North Carolina State University. He is currently a software engineer of Facebook. When this work was done, he was a Partner researcher with Microsoft Research. He is an ACM fellow.

David Lo received the PhD degree in computer science from the National University of Singapore in 2008. He is currently a professor of computer science with Singapore Management University. He has authored or coauthored in the area of software engineering, AI, and cybersecurity in premier and major conferences and journals. His research interests include the intersection of software engineering and data science, encompassing socio-technical aspects and analysis of different kinds of software artefacts, with the goal of improving software quality and developer productivity. He is an ACM distinguished member.

Pavneet Singh Kochhar received the PhD degree from Singapore Management University. During the PhD, he completed an exchange programme from Carnegie Mellon University. He is currently a software engineer with Microsoft. He has authored or coauthored in several soft engineering top-tier conferences and journals. His research interests include empirical software engineering, software testing, bug localization, and mining software repositories.

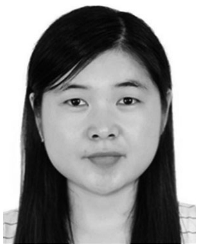

Cuiyun Gao received the PhD degree from the Department of Computer Science and Engineering, The Chinese University of Hong Kong, in 2018. She is currently an assistant professor with the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen, China. She has authored or coauthored more than 20 peer-reviewed publications in conferences and journals in her area of expertise. Her research interests include software repository mining and natural language processing. She was a reviewer for many top-tier conferences and journals.

Iftekhar Ahmed received the BSc degree in computer science and engineering from the Shahjalal University of Science and Technology, Bangladesh, and after working in the industry for four years, received the PhD from Oregon State University. He is currently an assistant professor of informatics with the Donald Bren School of Information and Computer Science, University of California, Irvine. His research interests include software engineering in general and combining software testing, analysis, and data mining to come up with better tools and techniques in particular.