Work Practices and Perceptions from Women Core Developers in OSS Communities

Edna Dias Canedo
University of Brasília (UnB)
ednacanedo@unb.br

Rodrigo Bonifácio
University of Brasília (UnB)
rbonifacio@unb.br

Márcio Vinicius Okimoto
University of Brasília (UnB)
marciobtos@gmail.com

Alexander Serebrenik
Eindhoven University of Technology
a.serebrenik@tue.nl

Gustavo Pinto
Federal University of Pará (UFPA)
gpinto@ufpa.br

Eduardo Monteiro
University of Brasília (UnB)
edumonteiro@unb.br

ABSTRACT

Background. The effect of gender diversity in open source communities has gained increasing attention from practitioners and researchers. For instance, organizations such as the Python Software Foundation and the OpenStack Foundation started actions to increase gender diversity and promote women to top positions in the communities. **Problem.** Although the general underrepresentation of women (a.k.a. horizontal segregation) in open source communities has been explored in a number of research studies, little is known about the *vertical segregation* in open source communities—which occurs when there are fewer women in high level positions. **Aims.** To address this research gap, in this paper we present the results of a mixed-methods study on gender diversity and work practices of core developers contributing to open-source communities. **Method.** In the first study, we used mining-software repositories procedures to identify the core developers of 711 open source projects, in order to understand how common are women core developers in open source communities and characterize their work practices. In the second study, we surveyed the women core developers we identified in the first study to collect their perceptions of gender diversity and gender bias they might have observed while contributing to open source systems. **Results.** Our findings show that open source communities present both horizontal and vertical segregation (only 2.3% of the core developers are women). Nevertheless, differently from previous studies, most of the women core developers (65.7%) report never having experienced gender discrimination when contributing to an open source project. Finally, we did not note substantial differences between the work practices among women and men core developers. **Conclusions.** We reflect on these findings and present some ideas that might increase the participation of women in open source communities.

CCS CONCEPTS

• **Software and its engineering** → **Collaboration in software development**; Open source model; • **Social and professional topics** → **User characteristics**; *Gender*; Women.

KEYWORDS

Gender Diversity, Vertical Segregation, OSS Community, Core Developer

ACM Reference Format:
Edna Dias Canedo, Rodrigo Bonifácio, Márcio Vinicius Okimoto, Alexander Serebrenik, Gustavo Pinto, and Eduardo Monteiro. 2020. Work Practices and Perceptions from Women Core Developers in OSS Communities. In *ESEM '20: ACM / IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM) (ESEM '20), October 8–9, 2020, Bari, Italy*. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3382494.3410682

1 INTRODUCTION

Software development often involves the participation and interaction of many contributors, who do not necessarily share the same physical space, culture, and beliefs [56]. This diversity might positively influence software development practices and achievements. Previous works reported that gender diversity improves not only teams' productivity, but also the quality of software products [16, 29, 56]. Even though gender diversity is valued by many software development organizations [13, 31], the field remains dominated by men, and gender bias has been pinpointed as one of the forces that contribute to the underrepresentation of women in the software industry [31, 57].

Existing studies report the small number of women contributing to OSS communities, especially in leadership positions [20, 27, 32, 32, 37, 42, 48]. Izquierdo et al. [32] analyzed the percentage of women in positions of governance and leadership in the OpenStack Foundation, reporting an increase in the percentage of women in leadership positions—though the number is still low (around 10–12%). Nafus [40] reported that "women were sexualized, hurtful and offensive talk was openly defended, and women were obliged to remind men not to stare and point at them". Wang and Redmiles [57] presented the results of a survey with 142 software engineers in seven OSS organizations and discuss that software engineers regardless their gender implicitly associate software development to a *male activity*. Furthermore, the authors argue that developers express gender biases while taking technical decisions [57]. Finally, Imtiaz et al. [31] have evaluated presence of several gender biases from the sociological literature in OSS projects.

ESEM '20, October 8–9, 2020, Bari, Italy

ACM ISBN 978-1-4503-7580-1/20/10…$15.00
https://doi.org/10.1145/3382494.3410682

So far, in the studies of gender and gender bias in OSS no distinction has been made between more and less experienced contributors. Differences between more senior and more junior women have been observed outside the software engineering realm [22, 26]. We complement the existing literature with the perspective of *vertical gender segregation* [6, 15], which deals with distribution inequalities within organization levels. Note that OSS communities are concerned with both horizontal (i.e., the general underrepresentation of women in OSS) and vertical gender (i.e., the participation of women in high level positions) diversity. The Python community, for instance, has started an effort to increase diversity in its core development team[1] and OpenStack Foundation increased the percentage of women in the Technical Committee from 0% to 15% [20].

The goal of this paper is twofold: first, we explore the issue of vertical segregation in open source communities and, second, we study the work practices and perceptions of gender bias— from the point of view of *women core developers* that contribute to OSS projects. Altogether, we answer the following research questions:

(RQ.1) How common are women core developers in OSS?
(RQ.2) Are there differences in the work practices of women or men core developers?[2]
(RQ.3) How do women core developers perceive gender diversity and gender bias in OSS communities?
(RQ.4) What are the actions women core developers consider important to make OSS communities more inclusive?

In the first two research questions we address the issues of vertical segregation and work practices of core developers. In the third and fourth research questions we address the perceptions of women core developers on gender bias. To answer these questions, we carry out a mixed-method study. We first identify the core developers of open-source systems, by mining the source code history of more than 700 OSS projects. Core developers here are those developers that significantly contribute to the development of a system; and thus the continuity of a project might be compromised in the case they decide not to contribute to the development of a system anymore. We identify core developers using *Truck Factor* [3, 19, 47]. To identify the gender of *core developers*, we leverage two gender classification algorithms: GenderComputer [54] and Namsor [3].

Considering the *intersection* of the results of both gender classification tools, we found 42 women core developers, and we were able to manually confirm the gender of 36. We invited them to answer a survey about gender bias on OSS communities—getting answers from 35 of them (97.22% of response rate). Our study produced a set of findings; we highlight three of them next:

1. While 5.35% of all contributors are women, the percentage of women among core developers drops to 2.30% (characterizing both horizontal and vertical gender segregation.)
2. There is no significant difference between the *work practices* between women and men core developers.
3. Gender bias also occurs among core developers. 34.3% of women core developers surveyed state having observed gender bias at least once while contributing to OSS projects.

[1] http://pyfound.blogspot.com/2019/02/the-north-star-of-pycascades-core.html
[2] We compare women with *men* since OSS is known to be male-dominated [48]. Study of development practices of non-binary software developers should be a topic of a separate study.
[3] https://www.namsor.com/

2 BACKGROUND AND RELATED WORK

Gender diversity in the fields of Science, Technology, Engineering, and Maths (STEM) has been investigated by multiple researchers and gained considerable public attention in recent years. Several educational institutions design programs with the aim of reducing the gender gap among students, which, in a long term, might lead to a positive impact on the gender diversity of teams in the workforce [7, 9, 11, 12, 21, 33, 38, 46]. Although there are some conflicting findings about the effects of diversity on team performance [27, 29, 56], some reports show evidence that promoting gender-balanced teams improves innovation and problem-solving capacity, as well as leads to a healthier work environment [29].

An inclusive work environment should mitigate possible conflicts that might arise due to diversity. Several authors study relation between gender diversity and performance of software development teams [10, 27, 42, 56]. Hui and Farnham [29] seek to understand how interpersonal practices and the use of socio-technical tools can promote gender diversity and help to form more independent innovative teams. Catolino et al. [16, 17] have studied the relation between gender diversity and ineffective communication.

Another line of research has focused on retention of women in OSS. Qiu et al. [45] have shown that involvement in teams using diverse technologies is beneficial for duration of engagement of women in OSS. Balali et al. [4] argued that duration of engagement of women in OSS is negatively affected by differences in the viewpoint of men and women mentors about gender personalities; underestimation of women's capabilities by both open source community and women newcomers themselves; and ignorance of men mentors' about the community being harsh to women.

Yet another group of studies have focused on gender biases in software development [10, 31, 57]. For example, Imtiaz et al. [31] concluded that while the effects of gender bias are virtually invisible on the investigated projects, women restrict their involvement to fewer projects and organizations, in comparison with men developers.

Lee and Carver [36] carried out an investigation of the men and women perspectives on gender relations in Free/Libre OSS projects. The study found, that while some respondents expressed a positive feeling about women's participation, some contributors were strongly opposed to their inclusion. Women reported the difficulty of being accepted in the community and the gender-biased comments of colleagues as major barriers for their participation.

In summary, the existing research on gender diversity in software engineering seeks to investigate the benefits of diversity on software teams and to try to understand the possible causes of the underrepresentation of women. Unlike the previous work, we investigate the vertical segregation problem in a comprehensive number of open source communities. We also study the work practices as well as perceptions of women core developers on gender bias. Finally, we survey the women core developers to identify actions that should be taken to make OSS more inclusive.

3 STUDY SETTINGS

The main goal of this research is to improve our understanding on work practices and gender bias in open source communities, focusing on a particular group of contributors: *women core developers*.

To achieve this goal, we use a mixed-methods approach. First we mine open source repositories to identify women core developers and to understand their work practices when contributing to open source communities. Second, we conduct a survey with women core developers, to understand their perceptions about gender bias in open source communities.

3.1 Settings for the first study: Mining open source repositories

Our approach for mining open source repositories has five steps. **In the first step** we used *purposeful sampling* [5] to build a dataset of open source projects from different domains and written in different programming languages. To this end we use the GitHub API to search for the 100 most popular projects written in the 15 most popular programming languages at GitHub. To operationalize popularity of programming languages we use a recent report[4], of projects—the number of stars [8, 44]. This dataset comprises open source projects of different sizes, targeting different domains (from compilers to mobile apps), and written in a diversity of languages, e.g., scripting languages such as Shell Script, system programming languages such as C and Go, and languages often used for web- and mobile development such as TypeScript and Swift.

As we study core developers, we focus on "sufficiently large" projects with "sufficiently many" committers. To determine the thresholds we compute the first quartiles of the distribution of SLOC and number of committers, and exclude projects having less SLOC or less committers than the thresholds. In this way we preserve 711 projects written in 14 languages[5] with at least 5183 SLOC and 33 committers. Tables 1 and 2 present descriptive statistics.

In the second step, we identify the core developers. To this end we use the notion of the Truck Factor (TF). "TF developers" is the minimal set of developers a project depends on for its maintenance and evolution, i.e., if the "TF developers" abandon the project (e.g., after being hit by a truck) the project maintenance will be heavily affected. We call "TF developers" *core developers*. Indeed, Ricca et al. [47] state that the TF can be used to assess the distribution of project knowledge among developers; and Bosu and Sultana argue that TF is a proxy for identifying "*developers that made significant contributions to guide the development and evolution of the project*" [10]. Several approaches to compute the TF have been proposed in the literature. In our paper, we use the approach of Avelino et al. [3], shown to outperform competing approaches [24]. We have identified 1954 core developers in 711 projects.

In the third step we identified the gender of all core developers using two gender identification tools, GenderComputer [54] and Namsor[6]. In the case of a disagreement between the tools, we assign "Unknown" to a given core developer. From 1954 core developers, the tools disagree in 192 cases (9.82%). After identifying the gender of the core developers, **in the fourth step** we proceeded to collect the contributions from women core developers (WCD). To compare contributions of women with those of men, we randomly select three samples of male core developers (MCD1, MCD2, MCD3)—with the same number of members as WCD. We used the GitHub API to collect all commits and pull requests (PRs) from the contributors in WCD, MCD1, MCD2, and MCD3. We classify the size and the type of contributions using an approach of Hattori and Lanza [28].

Finally, **in the last step (data analysis)** we used Exploratory Data Analysis (EDA) [39] to (a) characterize how common are women core developers in open source projects and (b) understand how women core developers contribute to open source projects. EDA covers different statistics (e.g., median and mean) and graphical methods (e.g., histograms and boxplots) to build a general understanding about the data distribution [39]. In addition, we leveraged a statistical procedure (`nparcomp`) of Konietschke et al. [35] for performing a multiple comparison on the work practices of the experimental groups (MCD1, MCD2, and MCD3) against the control group (WCD). Similarly to a previous work [55], we set the `nparcomp` analysis to use the Dunnett-type contrasts [23] and the *probit* transformation function (as the asymptotic approximation method). All datasets and scripts are available online (https://bit.ly/32ZyBXe).

3.2 Settings for the second study: Survey with women core Developers

To answer the research questions RQ.3 and RQ.4, we conducted a survey with the women core developers we identified in the first study. We follow the recommendations of Kitchenham and Pfleeger [34], and organized our survey in six steps. Regarding **the first step** (planning), our goal is to capture beliefs of women contributors identified as core developers about (a) gender bias in OSS communities and (b) possible ways for OSS communities to become more inclusive with respect to gender.

In **the second step** (defining the target population), we *manually confirm* the gender of the women core developers we identified in the first study, using information from social networks (e.g., Facebook, Google+, and Twitter). This activity was necessary because our goal was to only gather information from women core developers, and thus we followed a conservative approach before inviting our target population to answer our survey. Accordingly, our target population comprises 36 (distinct) women core developers, which we manually confirm the gender. We could not confirm the gender of 6 women core developers (out of 45). Moreover, among the 39 women core developers, three of them are core developers contributing to two projects in our dataset. Figure 1 shows the procedures we followed to define our target population.

In **the third and fourth steps** we designed an online survey and validated it using an iterative approach (the first, third, and fourth authors were responsible for reviewing and validating the questions). The final version of the survey contains 18 questions (14 closed questions using a Likert scale and 4 open-ended questions), organized in three sections: demographics, contribution to open source communities, and perceptions about gender bias in open source communities. In the *demographics section*, the survey covers information such as age and academic degree of the participants. In the *contribution to open source communities* section, the survey presents questions to characterize the engagement of the participants into open source communities, including questions such as *How long have you been contributing to OSS communities?*, *How often do you interact with other team members in OSS communities?*, and *Are you happy with your participation in OSS communities?*. All the

[4]https://octoverse.github.com/2018/projects#languages
[5]No Objective-C projects meet the thresholds.
[6]https://www.namsor.com/

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
Lines of Code	5191	19 523	57 013	259 367.63	195 265	9 442 645
Num. of Contributors	33	80	145	292.77	297	8413
Num. of Forks	54	774	1481	2949.94	3171	64 712
Num. of Watchers	1145	5882	9039	14 284.96	16 418	300 666

Table 1: Descriptive statistics about the projects used in the study

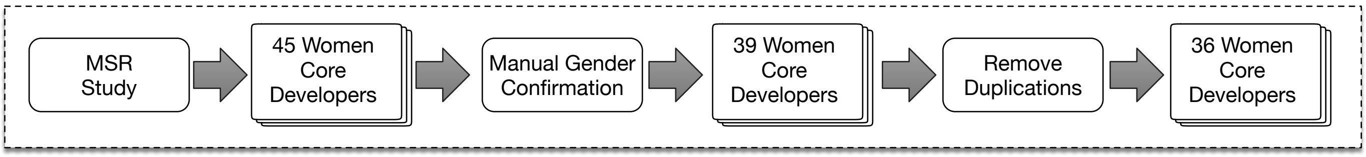

Figure 1: Procedure used to define the target population of the survey

questions are available in the paper website (https://bit.ly/32ZyBXe). In the third section (*perceptions of gender-bias in open source communities*) we look at gender bias in OSS communities, including questions such as *Have you ever felt that one of your contributions was not well received due to your gender?, How often do you feel that your contributions were not well received due to your gender? What would you recommend to increase women's participation in OSS communities?.*

In **the fifth step** (conducting the survey) we contacted (via email) the 36 women core developers that correspond to our target population. Over a period of three weeks, we received answers from 35 women core developers (a response rate of 97.22%). In the email message, we sent the goal of our research and a link to the online survey. Most of the participants agreed to answer the survey without any additional clarification. For instance, one participant answered "*I was really glad to participate, thank you!*" and other perceived the relevance of the research, answering "*Thanks for taking up this issue!*". Nonetheless, other participants requested us to provide additional details before answering the survey, such as our credentials and affiliation. Finally, in the **sixth step** (analysing the results), we leverage exploratory data analysis to consolidate the answers to the Likert scale based questions (in terms of descriptive statistics and plots) while the answers to the survey's open-ended questions were literally quoted (all datasets are available in the paper website https://bit.ly/32ZyBXe).

4 RESULTS OF THE FIST STUDY: MINING SOFTWARE REPOSITORIES

We first report the results of an exploratory data analysis. Figure 2 shows a histogram with the number of core developers in each project: most of the projects have only one core developer. This finding corroborates the work of Avelino et al. [2], which reports that most projects have a small number of TF developers and that the TF algorithm reveals just one core developer in 57% of the projects. This situation might represent a risk, since by definition of the Truck Factor [3], the continuity of a project might be compromised when a single core developer decides to leave it.

Nonetheless, we found projects having more than 5 core developers, including ELASTICSEARCH (17 core developers) and the

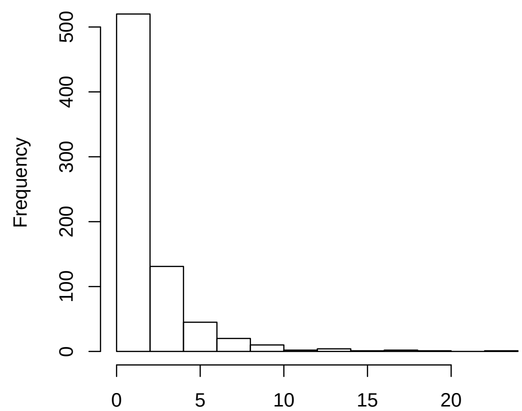

Figure 2: Histogram with the number of core developers

implementation of Python and Go programming languages (15 and 9 core developers, respectively). Considering all 711 projects, 88 projects (12.37%) have at least five core developers; and we found a small correlation between lines of code and the number of core developers of a project (Spearman's $\rho = 0.30$ with p-value = 0.001). Also, there is a moderate correlation between the number of contributors and the number of core developers of a project ($\rho = 0.41$ with p-value < 0.01). Our dataset comprises 1954 core developers, from which 235 developers are core developers in more than one project.

4.1 How common are women core developers in OSS communities?

45 core developers are recognized both by GenderComputer and Namsor as women (2.30%), while 1,717 are recognised as men (87.87%). We could not confirm the gender of 192 core developers (9.82%) due to disagreement between the tools. We found women core developers in 37 (5.24%) out of the 711 GitHub projects considered in our analysis. We also found women core developers in projects written in all programming languages we consider in our study. Interestingly seven projects written in TypeScript (10.93%)

have at least one women as core developer; though only 2.17% of projects using PHP have at least one women as core developer. Table 2 summarizes these findings.

The barplot of Figure 3 presents a different perspective: the percentage of women core developers (over the total of core developers) considering the different programming languages. That is in our dataset less than one percent of core developers in projects using PHP and Shell programming languages are women. Contrasting, more than four percent of the core developers in projects written in Swift and TypeScript are women. Altogether, we answer our first research question ("How common are women key developers in OSS projects?") as follows:

Among 711 GitHub projects, we identified 1954 core developers. 45 core developers (2.30%) are identified as women. Since the percentage of developers identified as women in our dataset is 5.35%, these findings suggest an underrepresentation of women core developers in OSS projects, i.e., vertical gender segregation.

Besides gender bias, other factors (such as the reward model and the possible long term benefits of contributing to open source projects) could contribute to this underrepresentation of women core developers in OSS projects. Considering this quantitative assessment, we can mostly report on the extent of this underrepresentation—and thus we postpone a discussion of possible causes for this underrepresentation (and how to deal with them) to the next section.

4.2 Are there differences in the work practices of women and men core developers?

To better understand the work practices of women core developers, we also explore two additional questions: (a) How do the number, frequency, and size of contributions of women core developers compare to the number, frequency, and size of contributions of men core developers? and (b) How do the *types of contributions* of women core developers differ from the *types of contributions* of men core developers? Accordingly, we mined the commit history from the 36 women core developers (WCD) whose gender we could manually confirm, out of the initial set of 45 we identified (see Figure 1). To counterbalance the effects of randomness in selecting samples, we randomly generated three datasets with men core developers (MCD1, MCD2, and MCD3), and collected their commit history. Each one of these datasets comprise 36 men core developers. We contrast the working practices of the experimental groups (MCD1, MCD2, and MCD3) with our control group (WCD), using the `nparcomp` procedure (see Section 3). Figure 4 shows the total number of contributions (commits) from these sets (MCD1, MCD2, MCD3, and WCD). Figure 4 might suggest a small difference in terms of the distribution of the total number of commits. Nonetheless, the results of the multiple comparison `nparcomp` test do not reveal any significant difference at 5% level.

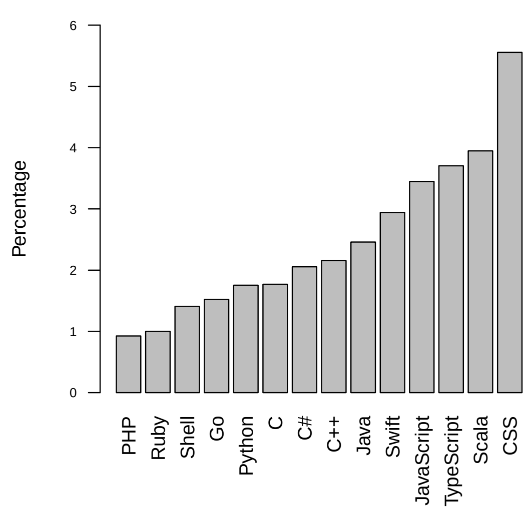

Figure 3: Percentage of women core developers in projects written in different programming languages

To measure the frequency of the commits, we use three auxiliary metrics: *Max Date*, *Min Date*, and *Distinct Dates*. *Max Date* (*Min Date*) corresponds to the date of the last (first) commit of a core developer, in one of the sets MCD1, MCD2, MCD3 or WCD. *Distinct Dates* corresponds to the number of distinct commit dates of a core developer, again, in one of the sets MCD1, MCD2, MCD3 or WCD. Finally, we compute the *Frequency* of commits using Eq. (1).

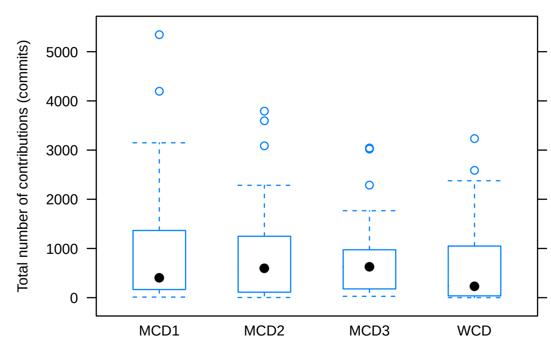

Figure 4: Number of commits from the four sets of core developers. WCD stands for Women Core Developers while MCD* stands for the random sets of Men Core Developers

$$Frequency = \frac{Distinct\ Days}{interval(Min\ Date,\ Max\ Date)} \times 100 \qquad (1)$$

The boxplots of Figure 5 summarizes the frequency of commits per group (WCD, MCD1, MCD2, and MCD3). Descriptive statistics suggest that WCD commit code more frequently than MCD (the median value of the frequency of commits in WCD is 2.34%, while the median value of the frequency of commits in MCD1 is 1.38%, in MCD2 is 1.64%, and in MCD3 is 2.01%). However, these differences are not statistically significant (all p-values reported by the `nparcomp` test exceed 0.3).

Regarding the size of the contributions, we computed the total lines of code (and the number of files) *added*, *changed*, and *deleted*, from the set of contributions of MCD1, MCD2, MCD3, and WCD.

Table 2: Summary of the dataset with core developers. MCD means the number of men core developers, WCD means the number of women core developers, and UCD means the number of core developers we could not identify the gender. PWCD means the number of projects with women core developers, and PPWCD corresponds to the percentage of projects with women core developers.

	Language	Number of Projects	Number of Contributors	MCD	WCD	UCD	PWCD	PPWCD
1	C	50	11 627	99	2	12	2	4.00
2	C#	63	10 150	129	3	14	3	4.76
3	C++	67	16 954	191	5	36	3	4.48
4	CSS	23	2505	30	2	4	2	8.70
5	Go	68	19 232	169	3	25	3	4.41
6	Java	44	9575	109	3	10	2	4.55
7	JavaScript	67	33 899	172	7	24	5	7.46
8	PHP	46	13 315	98	1	9	1	2.17
9	Python	42	20 706	155	3	13	2	4.76
10	Ruby	60	36 064	183	2	15	2	3.33
11	Scala	57	8956	137	6	9	5	8.77
12	Shell	24	5417	63	1	7	1	4.17
13	Swift	36	4129	59	2	7	2	5.56
14	TypeScript	64	20 589	123	5	7	4	6.25
	Total	711	213 118	1717	45	192	37	mean = 5.24, sd = 1.92

Figure 6 summarizes the (log-scale) size of contributions in terms of lines of code. Again, no statistically significant differences could be observed. That is, considering these results, we conclude that there is no difference in terms of the number, frequency, and size of commits with regards the gender of core developers.

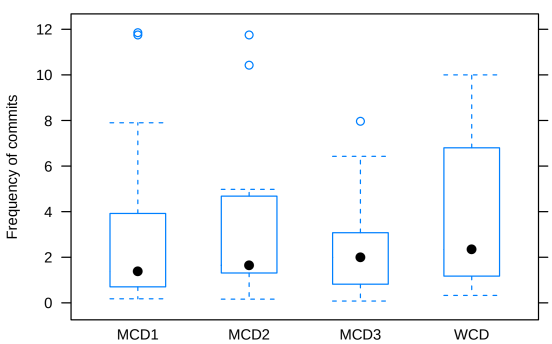

Figure 5: The frequency of commits from core developers. This frequency corresponds to the percentage of days the core developers contribute to a project, considering an interval from the first and last commits to the project.

We could not find statistically significant differences in the practices of women and men w.r.t. the number, frequency, and size of commits.

Finally, we used the approach of Hattori and Lanza [28] to investigate the differences on the type of contributions from MCD1, MCD2, MCD3, and WCD. Their approach classifies the contributions by searching for a set of keywords in the commit message—assigning a commit to a class whenever it finds the first keyword in the commit message. A commit could be classified as *forward engineering*, *reengineering*, *corrective engineering*, and *management*. Regardless of its simplicity, the assessment of this algorithm has shown a good performance (F-measure = 0.70) [28]. Our dataset comprises a total of 115 922 commits (22 326 from WCD and 33 840, 32 818, and 26 938 from MCD1, MCD2, and MCD3, respectively). To avoid unbalanced problems in this analysis, we *undersample* the set of commits in MCD1, MCD2, and MCD3, and thus we only consider 22 326 commits coming from contributors in each set.

We found a statistically significant difference between the size of description of the commits (in terms of number of characters) from women core developers and the size of commits' descriptions from men core developers (with a p-value < 0.0001). That is, based on our dataset of commits we consider here, women core developers tend to present a more detailed message explaining their contribution changes. This might indicate the gender bias symptom named *proving-it-again* [31], which occurs when a group of people that does not align to the default stereotypes has to demonstrate more evidence about their competence. Since we classify the type of a contribution considering the commit message, this result impacts the number of commits that we could not classify using the Hattori and Lanza method, which is lower when we consider the WCD set.

Figure 7 summarizes the class of contributions using a log scale. The boxplots suggest differences in the types of contributions when we consider the different groups. For instance, it seems that the set of WCD contributes more with *Corrective* and *Reengineering* activities; while the sets of men core developers contribute more with *Management* activities and activities that we were not able to classify (*Unknown*—according to the Hattori and Lanza method). There is no much difference in the *Forward Engineering* activities. We actually found that these differences are statistically significant (with a p-value < 0.0001) using the Chi-squared test. This test is useful to investigate if two categorical variables (gender and type of contributions) have a significant correlation.

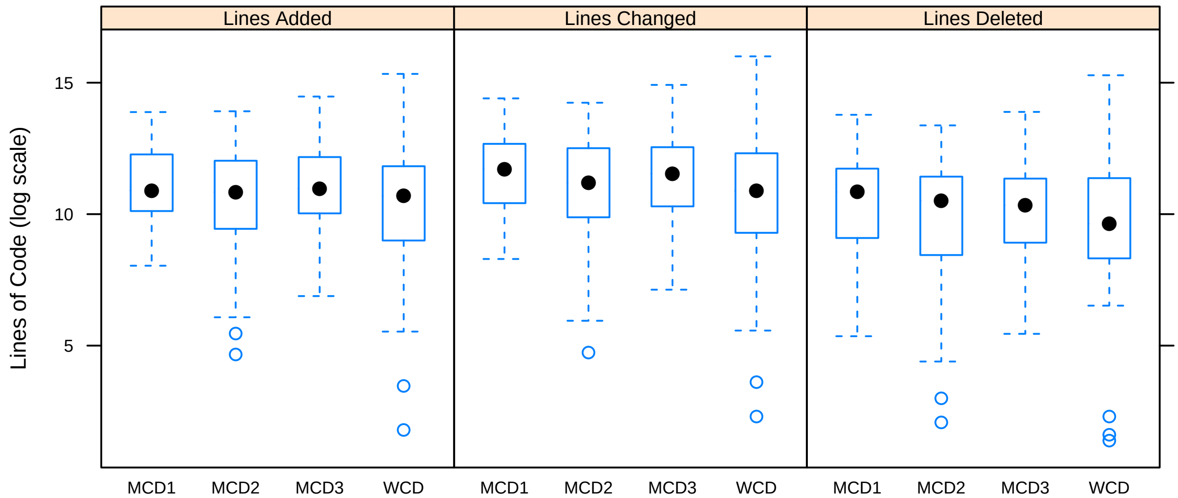

Figure 6: Log-scale of the size of the contributions, in terms of lines of code added, changed, and deleted

We found statistically significant gender-related differences in the kinds of contributions: women core developers tend to contribute more with reengineering tasks.

5 RESULTS OF THE SECOND STUDY: A SURVEY WITH WOMEN CORE DEVELOPERS

When analyzing the answers to our survey, we found that women core developers are in general young: 51.4% of the respondents are between 18 and 25 years old and almost 80% of the respondents are younger than 35 years old. Regarding academic degree, most of them are undergraduate students (31.4%); 34.3% hold a bachelor degree, 11.4%—a master degree, and 22.9%—PhD. 77.2% contribute to open source projects for less than five years: more then twice than percentage of women contribute to open source projects for less than five years reported in the FLOSS 2013 survey [48].

Figure 8-(a) summarizes perceived importance of gender diversity in OSS communities. 68.6% of the women core developers consider gender diversity in OSS communities to be very important. According to the respondents, gender diversity can improve team communication, and attract new contributors: these opinions concur with the findings of Catolino et al. [17]. For instance, one of the participants states that:

> *. . . the interaction between team members, and communication both within the team, as well as with the larger community has to be open and consistent. I believe having more gender diversity can help in this direction, as women may bring a new perspective, and focus more on communication and human aspects. This will help with both building a stronger core team that stays with the project, as well as attracting new contributors. If the team is more welcoming to new members of any gender, then there is a larger pool of potential contributors, and a better chance of them wanting to get involved.*

Other respondent states that diversity can contribute to the design of products based on a broader variety of past experiences, promote empathy and build a safer community.

> *I cannot say what women can bring exactly, but what I know is that everyone has different experiences in life, as women, men, people of color, members of LGBTQ+ communities, people with disabilities, and I am sure that we need people with different experiences to build products, whether in closed source or in open source. Also because in open source you can get anyone in the world to contribute to a project, different people have different sensibility to different ways of working and communicating. Having people from different cultures and walks of life brings more empathy and therefore potentially a safer environment for people to contribute.*

Yet another respondent states that gender diversity might not directly improve the productivity of the teams. However, it might increase the design space when conceiving a product's features, which could also be more generic and inclusive.

> *I have no idea whether diversity improves productivity. However, I have found that having diverse development pools (meant in the widest sense) ensures that the developed software is more fit for purpose and generalizable as the diverse experiences ensure that people are thinking about the design from different angles.*

Figure 8-(b) summarises the answers to the question "*How often do you feel that your contributions were not well received due to your gender?*" Even in a population of core developers, one third of the participants believe that, at least one of their contributions had not been accepted due to gender bias. Moreover, 11.4% (very) often recognize gender bias while someone appraises their contributions. According to one respondent (P19) (see TableSM1 in https://bit.ly/32ZyBXe), gender bias appears whenever a contribution from

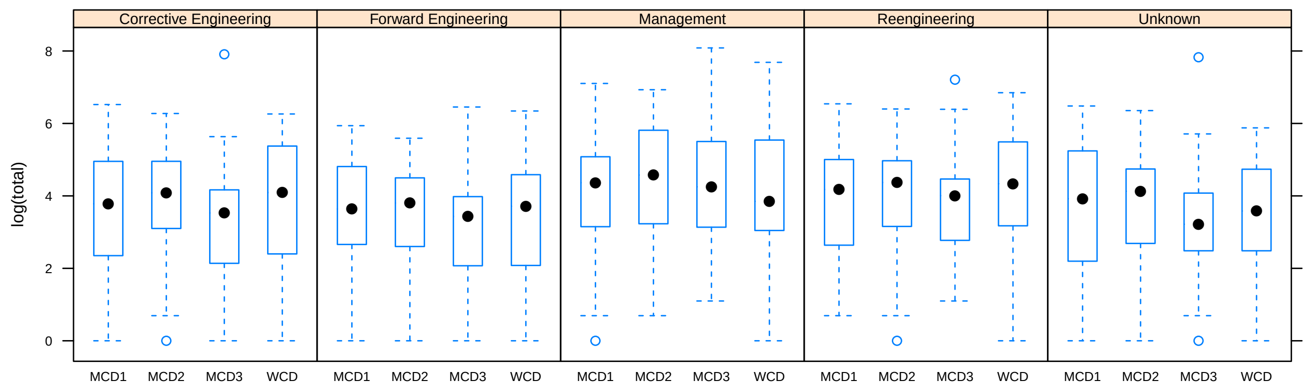

Figure 7: Log-scale of the classes of contributions, according to the approach of Hattori and Lanza [28]

a women developer receive less positive feedback: "*Women are participating, but their performance does not get the same positive reaction as men's*". Gender bias also appears in OSS communities through the language used. P10 (see TableSM2 in https://bit.ly/32ZyBXe) recommends to "*... avoid gender pronouns (e.g.: using 'guys' is very common, and this gives an idea that it is assumed that contributors are mostly men), so moderating language would help*". Nafus [40] also mentioned the use of an inadequate language in OSS communities.

Despite the recognised gender bias, 82.9% of the respondents report being happy in contributing to OSS and 88.5% (strongly) agree that their contributions are well received in open source communities.

We also asked women core developers about the actions that should be taken to create a more inclusive environment for women in OSS. To this end, we use an open-ended question and tabulate the answers verbatim (see TableSM2 in https://bit.ly/32ZyBXe). The respondents suggest promoting *women-specific* mentorship programs and events discussing the relevance of contributing to OSS. Participants also mentioned more specific initiatives (e.g., Outreachy and local Meetups) that might help to attract more women. P9 also emphasizes the relevance of communication, arguing in favor of "*more women promoting open source projects and work, through blogs, forums, public speaking, (and thus) helping to demystify the world of open source*". P23 stressed the importance of increasing confidence as a way of engaging women (cf. recent studies of confidence in context of women in software development [1, 58]):

> *I feel that the solution is to build confidence.... Every approach towards increasing participation of women has a side that increases confidence and another that decreases it. The one that I completely support is building a peer group among girls interested in it. It is not so common to find many girls in technical teams and also, it is not considered cool to be a techie girl. ... So, having people around me with whom I can share everything that I do every day without having a fear of being judged as a freak, has been pretty helpful and encouraging.*

Participant P20 also recommends events by and for women:

> *In the R community there are R-Ladies events—held by and for women. Girl only or (predominantly) events in general, could create an environment that suggests that girls and women are actually wanted to be included in the communities.*

This recommendation agrees with the work of Singh [50] suggesting that women-only spaces in OSS foster discussions, support and empowerment of minorities.

Finally, P2 pointed out that the communities should "*Stop treating women developers as 'women developers' and start treating them as developers*". This quote suggests both presence of gender bias and frustration caused by it. This can also be seen as a call to support code of conducts [53] recommending developers to avoid any behavior that might be understood as non-inclusive. This call has been reiterated by P6: "*OSS Communities should be inclusive not only for women, but for all (men, women, LGBT..., disabled, etc).*". Other recommendations also include promoting more women to senior roles and "*let know the woman not to fear when contributing, she will be treated just like anyone else: good contribution then its accepted regardless of gender, political views or religion, country of origin etc*".

6 DISCUSSION

Below we summarise the insights obtained from our study.

6.1 Observations

Vertical segregation in OSS communities. Although the underrepresentation of women in OSS communities has been studied before, we found that even less women are core developers in OSS projects. This clearly indicates a non-inclusive situation, where women do not appear in top positions, a phenomenon known as *vertical sex segregation* [6]. Some respondents suggest that the underrepresentation of women in OSS communities might be due to the low number of women attending undergraduate courses in Computer

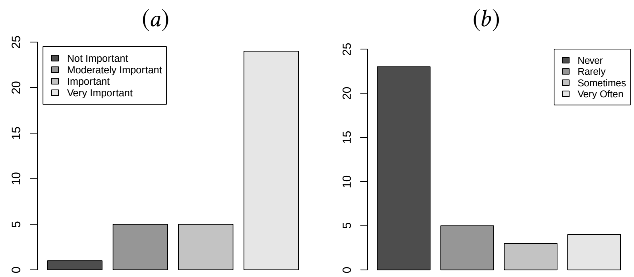

Figure 8: Figure (a) shows the importance of gender diversity in OSS Communities, while (b) shows how often the respondents feel that their contributions were not well received due to their gender.

Science and related subjects. Nonetheless, the percentage of women attending these courses tends to be 10–20% [9, 18, 41, 46], while the participation of women in OSS is ca. 5% [36, 43, 56]. We observed that the participation of women as core developers, according to the Truck Factor, is even lower, ca. 2%. Since participation in OSS communities might help developers to find new job opportunities, it is important to also encourage women to contribute to OSS initiatives.

There are no (statistically significant) gender-related differences in the work practices of core developers. Men and women core developers similarly contribute to the development of OSS. We have also observed women performing managerial activities in some projects: as future work we consider identifying project characteristics that might encourage the participation of women in top positions.

Perception of gender biases by Women core developers. One third of the women core developers faced gender bias at least once. This finding stresses importance of social, as opposed to technical, barriers to the participation of women in OSS [51]. Moreover, we found that 88.50% of core developers stated that their contributions are well received by the open source community regardless of their gender and 65.7% stated that they never had any problems with their contributions and never suffered gender bias. The perception of core developers is different from perception of women developers reported in the previous research [31, 36, 40, 52, 57]. These high percentages might be expected: women not feeling happy contributing to OSS or not feeling appreciated can be expected to be more likely to leave OSS before becoming core developers. At the same time this difference calls for a more careful investigation of the differences between core developers and non-core developers, and differences in perception of their contributions.

6.2 Suggestions to make OSS more inclusive

Survey respondents made several suggestions how to make OSS more inclusive for women. We hope that these suggestions might make OSS more welcoming other minorities as well (cf. [25]):

(1) Promote women-specific mentorship programs, akin to discussed by Hyrynsalmi [30] and Buhnova and Prikrylova [14].
(2) Promote women to senior roles: some communities, e.g., Open Stack, already implement this [32].
(3) Organize women-specific events, such as local meetups, or even tech groups [50]. Such events are organised, e.g., by R ladies group or the Pyladies group.
(4) Avoid gendered language (e.g., using 'guys' when a 'folks' would work).

The four suggestions can be combined and should contribute to women's confidence to contribute to OSS communities [49, 58].

7 THREATS TO VALIDITY

As any empirical work, this work also has many limitations and threats to validity.

Construct Validity. The main construct used in this study is the construct of 'gender". Gender is a complex social construct and no automatic tool can capture its entire complexity. Moreover, the accuracy of gender classifiers is inherently limited by the information developers provide in the software repositories. Many users do not use their real names, so we might not reliably extracted gender information [45]. To minimize this threat we (a) combine the results of two independently developed gender classifiers, and (b) manually validate the gender of women core developers, by checking information publicly available on social networks. We were able to validate the gender from 39 (out of 45) women core developers—without discarding the same women core developers contributing to more than one project. However, manual identification of gender might introduce bias as researchers necessarily can only indicate gender as perceived by the outsiders based on gender expression rather than gender identity. Another construct we use in this study is the construct of a "core developer". To this end we use the notion of a Truck Factor, and specifically the implementation of the Truck Factor detection proposed by Avelino et al. [3]. While the approach of Avelino et al. has outperform the competing techniques in the evaluation study of Ferreira et al. [24], and hence can be seen as state-of-the-art, as any automatic approach it can never be expected to perfectly identify developers whose departure from the project will heavily affect the project maintenance.

Internal Validity. The Truck Factor identification approach we used introduces additional threats. In particular, it only outputs the core developer's name. For this reason, we have to search the name of the contributor in the GitHub API to identify the corresponding user id. We accept the first value returned in the search performed, which can be a threat since we have no way of guaranteeing that the first value returned is, in fact, the login corresponding to that name. Also, we excluded some names that did not return user ids in the search, meaning that the person probably changed the name, or their account was deleted, or their account privacy setting was changed so that the contributor's name was not displayed. We also excluded names where their content was the same as login.

Conclusion Validity. We found that 80% of the women core developers that answered our survey are less than 35 years old and contribute to open source projects for less than five years. This finding brings additional threats to the conclusions of our work. For instance, one might argue that the small number of women core developers contributing to open source projects is due either to lack of experience or lack of interest to continue contributing to the development of open-source systems for long periods.

External Validity. We did not consider all possible open source projects available out there. Although we covered hundreds of projects hosted on GitHub, many other open source projects are hosted on different forges (e.g., GitLab or BitBucket). However, we do not expect major changes in our results, since we consider our projects' population diverse enough (in terms of programming languages used, number of core members, etc.). The survey was answered by 35 developers. This low number is due to the low percentage of women core developers in projects. Therefore, the representativeness of the sample is high, considering that among almost 2000 core developers, women correspond to a really small fraction, only 2.3% of the core developers are women. Thus, we can consider that the response rate was high (35 responses) out of a total of 39 women.

8 CONCLUSION

In this work we studied a different interpretation of gender-bias in open-source communities, i.e., *vertical sex segregation* [6], which occurs when the participation of men and women in top positions is even less balanced. In our study, "top positions" correspond to the core developers of a system, which we identified using the Truck Factor. We found a more significant underrepresentation of women core developers than women developers—only 45 in 711 open-source systems have at least one woman core developer, and only 2.30% of the total number of core developers are women. Nonetheless, this group of women core developers contribute in a similar fashion as other groups of men key developers (considering frequency and size of contributing).

Women core developers believe that gender diversity is important for OSS communities. According to their opinions, gender diversity might contribute to improving the communication among team members and help to generate different ideas while designing a software products. Nonetheless, 34.7% of the women core developers we surveyed in our study report having faced some sort of gender bias (11.4% of them claim often facing gender bias). This group also consider that promoting women specific events can contribute to making open-source communities more inclusive.

ACKNOWLEDGMENTS

This research work has the support of the Research Support Foundation of the Federal District (FAPDF) research grant 05/2018 and CNPq (grant 309032/2019-9).

REFERENCES

[1] Efthimia Aivaloglou and Felienne Hermans. 2019. How is Programming Taught in Code Clubs? Exploring the Experiences and Gender Perceptions of Code Club Teachers. In *Proceedings of the 19th Koli Calling International Conference on Computing Education Research (Koli Calling '19)*. Association for Computing Machinery, New York, NY, USA, Article 22, 10 pages. https://doi.org/10.1145/3364510.3364514

[2] Guilherme Avelino, Eleni Constantinou, Marco Tulio Valente, and Alexander Serebrenik. 2019. On the abandonment and survival of open source projects: An empirical investigation. In *ESEM*. IEEE, 10.1109/ESEM.2019.8870181, 1–12.

[3] Guilherme Avelino, Leonardo Teixeira Passos, André C. Hora, and Marco Tulio Valente. 2016. A novel approach for estimating Truck Factors. In *ICPC*. IEEE Computer Society, IEEE, 1–10.

[4] Sogol Balali, Igor Steinmacher, Umayal Annamalai, Anita Sarma, and Marco Aurélio Gerosa. 2018. Newcomers' Barriers. . . Is That All? An Analysis of Mentors' and Newcomers' Barriers in OSS Projects. *Computer Supported Cooperative Work* 27, 3-6 (2018), 679–714.

[5] Sebastian Baltes and Paul Ralph. 2020. Sampling in Software Engineering Research: A Critical Review and Guidelines. *arXiv preprint arXiv:2002.07764* 1 (2020), 1–21.

[6] Yvonne Benschop. 2006. Of small steps and the longing for giant leaps. *Handbook of workplace diversity* 1 (2006), 274–298.

[7] Shubhasheesh Bhattacharya, Sonali Bhattacharya, and Sweta Mohapatra. 2018. Enablers for Advancement of Women into Leadership Position: A Study Based on IT/ITES Sector in India. *IJHCITP* 9, 4 (2018), 1–22.

[8] Hudson Borges, Andre Hora, and Marco Tulio Valente. 2016. Predicting the Popularity of GitHub Repositories. In *Proceedings of the The 12th International Conference on Predictive Models and Data Analytics in Software Engineering (PROMISE 2016)*. Association for Computing Machinery, New York, NY, USA, Article 9, 10 pages. https://doi.org/10.1145/2972958.2972966

[9] Valeria Borsotti. 2018. Barriers to gender diversity in software development education: actionable insights from a danish case study. In *ICSE (SEET)*. ACM, ACM, 146–152.

[10] Amiangshu Bosu and Kazi Zakia Sultana. 2019. Diversity and Inclusion in Open Source Software (OSS) Projects: Where Do We Stand?. In *ESEM*. ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM), IEEE, 1–11. https://doi.org/10.1109/ESEM.2019.8870179

[11] Carmen Botella, Silvia Rueda, Emilia López-Iñesta, and Paula Marzal. 2019. Gender Diversity in STEM Disciplines: A Multiple Factor Problem. *Entropy* 21, 1 (2019), 30.

[12] Luca Botturi, Chiara Bramani, and Sean McCusker. 2012. Boys are like Girls: Insights in the Gender Digital Divide in Higher Education in Switzerland and Europe. *J. UCS* 18, 3 (2012), 353–376.

[13] Danielle Brown and Melonie Parker. 2019. Google diversity annual report 2019. https://diversity.google/ (Date last accessed 16-April-2019).

[14] Barbora Buhnova and Dita Prikrylova. 2019. Women want to learn tech: lessons from the Czechitas education project. In *Proceedings of the 2nd International Workshop on Gender Equality in Software Engineering, GE@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Ivica Crnkovic, Karina Kohl Silveira, and Sara Sprenkle (Eds.). IEEE / ACM, https://doi.org/10.1109/GE.2019.00013, 25–28. https://doi.org/10.1109/GE.2019.00013

[15] Juan Antonio Campos-Soria, Andrés Marchante-Mera, and Miguel Angel Ropero-García. 2011. Patterns of occupational segregation by gender in the hospitality industry. *International Journal of Hospitality Management* 30, 1 (2011), 91–102.

[16] Gemma Catolino, Fabio Palomba, Damian A. Tamburri, Alexander Serebrenik, and Filomena Ferrucci. 2019. Gender Diversity and Women in Software Teams: How Do They Affect Community Smells?. In *ICSE-SEIS*. IEEE, 10.1109/ICSE-SEIS.2019.00010, 11–20.

[17] Gemma Catolino, Fabio Palomba, Damian A. Tamburri, Alexander Serebrenik, and Filomena Ferrucci. 2020. Gender Diversity and Community Smells: Insights From the Trenches. *IEEE Software* 37, 1 (2020), 10–16. https://doi.org/10.1109/MS.2019.2944594

[18] Michael Christie, Maureen O'Neill, Kerry Rutter, Graham Young, and Angeline Medland. 2017. Understanding why women are under-represented in science, technology, engineering and mathematics (STEM) within higher education: A regional case study. *Production* 27, SPE (2017), 1–9.

[19] Valerio Cosentino, Javier Luis Cánovas Izquierdo, and Jordi Cabot. 2015. Assessing the bus factor of Git repositories. In *SANER*. IEEE Computer Society, 10.1109/SANER.2015.7081864, 499–503.

[20] Nicole Huesman Daniel Izquierdo and Allison Price. 2017. Gender diversity in the OpenStack community: A new report. https://superuser.openstack.org/wp-content/uploads/2017/07/OpenStack-Gender-Diversity-Report_Apr2017.pdf (Date last accessed 16-April-2019).

[21] Hélène de Ribaupierre, Kathryn Jones, Fernando Loizides, and Yulia Cherdantseva. 2018. Towards Gender Equality in Software Engineering: The NSA Approach. In *GE@ICSE*. ACM, ACM, 10–13.

[22] Belle Derks, Colette Van Laar, Naomi Ellemers, and Kim de Groot. 2011. Gender-Bias Primes Elicit Queen-Bee Responses Among Senior Policewomen. *Psychological Science* 22, 10 (2011), 1243–1249. https://doi.org/10.1177/0956797611417258 arXiv:https://doi.org/10.1177/0956797611417258 PMID: 21873568.

[23] Charles W Dunnett. 1955. A multiple comparison procedure for comparing several treatments with a control. *J. Amer. Statist. Assoc.* 50, 272 (1955), 1096–1121.

[24] Mívian M. Ferreira, Marco Tulio Valente, and Kecia Aline M. Ferreira. 2017. A comparison of three algorithms for computing truck factors. In *ICPC*. IEEE Computer Society, IEEE, 207–217.

[25] Denae Ford, Reed Milewicz, and Alexander Serebrenik. 2019. How remote work can foster a more inclusive environment for transgender developers. In *Proceedings of the 2nd International Workshop on Gender Equality in Software Engineering, GE@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Ivica Crnkovic, Karina Kohl Silveira, and Sara Sprenkle (Eds.). IEEE / ACM, https://doi.org/10.1109/GE.2019.00011, 9–12. https://doi.org/10.1109/GE.2019.00011

[26] Judit García-González, Patricia Forcén, and Maria Jimenez-Sanchez. 2019. Men and women differ in their perception of gender bias in research institutions. *PLOS ONE* 14, 12 (12 2019), 1–21. https://doi.org/10.1371/journal.pone.0225763

- [27] Abdul Rehman Gila, Jafreezal Jaafa, Mazni Omar, and Muhammad Zahid Tunio. 2014. Impact of personality and gender diversity on software development teams' performance. In *2014 International Conference on Computer, Communications, and Control Technology (I4CT)*. IEEE, IEEE, 261–265.
- [28] Lile Hattori and Michele Lanza. 2008. On the nature of commits. In *23rd IEEE/ACM International Conference on Automated Software Engineering - Workshop Proceedings (ASE Workshops 2008), 15-16 September 2008, L'Aquila, Italy*. IEEE, dblp computer science bibliography, https://dblp.org, 63–71. https://doi.org/10.1109/ASEW.2008.4686322
- [29] Julie S. Hui and Shelly Diane Farnham. 2016. Designing for Inclusion: Supporting Gender Diversity in Independent Innovation Teams. In *GROUP*. ACM, ACM, 71–85.
- [30] Sonja M. Hyrynsalmi. 2019. The underrepresentation of women in the software industry: thoughts from career-changing women. In *Proceedings of the 2nd International Workshop on Gender Equality in Software Engineering, GE@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Ivica Crnkovic, Karina Kohl Silveira, and Sara Sprenkle (Eds.). IEEE / ACM, https://doi.org/10.1109/GE.2019.00008, 1–4. https://doi.org/10.1109/GE.2019.00008
- [31] Nasif Imtiaz, Justin Middleton, Joymallya Chakraborty, Neill Robson, Gina Bai, and Emerson R. Murphy-Hill. 2019. Investigating the effects of gender bias on GitHub. In *ICSE*. IEEE / ACM, 10.1109/ICSE.2019.00079, 700–711.
- [32] Daniel Izquierdo, Nicole Huesman, Alexander Serebrenik, and Gregorio Robles. 2019. OpenStack Gender Diversity Report. *IEEE Software* 36, 1 (2019), 28–33.
- [33] Antonio M. Lopez Jr., Lisa J. Schulte, and Marguerite S. Giguette. 2005. Climbing onto the shoulders of giants. In *SIGCSE*. ACM, 10.1145/1047344.1047477, 401–405.
- [34] Barbara A. Kitchenham and Shari Lawrence Pfleeger. 2008. Personal Opinion Surveys. In *Guide to Advanced Empirical Software Engineering*. Springer, 10.1007/978-1-84800-044-5_3, 63–92.
- [35] Frank Konietschke, Marius Placzek, Frank Schaarschmidt, and Ludwig Hothorn. 2015. nparcomp: An R Software Package for Nonparametric Multiple Comparisons and Simultaneous Confidence Intervals. *Journal of Statistical Software, Articles* 64, 9 (2015), 1–17. https://doi.org/10.18637/jss.v064.i09
- [36] Amanda Lee and Jeffrey C. Carver. 2019. FLOSS participants' perceptions about gender and inclusiveness: a survey. In *ICSE*. IEEE / ACM, 10.1109/ICSE.2019.00077, 677–687.
- [37] Bin Lin and Alexander Serebrenik. 2016. Recognizing gender of stack overflow users. In *MSR*. ACM, 0.1145/2901739.2901777, 425–429.
- [38] Cristiano Maciel, Silvia Amélia Bim, and Karen da Silva Figueiredo. 2018. Digital Girls Program - Disseminating Computer Science to Girls in Brazil. In *GE@ICSE*. ACM, https://ieeexplore.ieee.org/document/8452748, 29–32.
- [39] John Maindonald and W. John Braun. 2010. *Data Analysis and Graphics Using R: An Example-Based Approach* (3rd ed.). Cambridge University Press, USA.
- [40] Dawn Nafus. 2012. 'Patches don't have gender': What is not open in open source software. *New Media & Society* 14, 4 (2012), 669–683.
- [41] Noelia Olmedo-Torre, Fermín Sánchez Carracedo, M. Nuria Salan Ballesteros, David López, Antoni Perez-Poch, and Mireia Lopez-Beltran. 2018. Do Female Motives for Enrolling Vary According to STEM Profile? *IEEE Trans. Education* 61, 4 (2018), 289–297.
- [42] Marco Ortu, Giuseppe Destefanis, Steve Counsell, Stephen Swift, Roberto Tonelli, and Michele Marchesi. 2017. How diverse is your team? Investigating gender and nationality diversity in GitHub teams. *J. Software Eng. R&D* 5 (2017), 9.
- [43] Stack Overflow. 2019. Developer Survey Results 2019. https://insights.stackoverflow.com/survey/2019 (Date last accessed 18-April-2019).
- [44] Gustavo Pinto, Igor Steinmacher, and Marco Aurélio Gerosa. 2016. More Common Than You Think: An In-depth Study of Casual Contributors. In *IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering, SANER 2016, Suita, Osaka, Japan, March 14-18, 2016 - Volume 1*. IEEE Computer Society, https://doi.org/10.1109/SANER.2016.68, 112–123. https://doi.org/10.1109/SANER.2016.68
- [45] Huilian Sophie Qiu, Alexander Nolte, Anita Brown, Alexander Serebrenik, and Bogdan Vasilescu. 2019. Going farther together: the impact of social capital on sustained participation in open source. In *ICSE*. IEEE / ACM, 10.1109/ICSE.2019.00078, 688–699.
- [46] Penny Rheingans, Erica D'Eramo, Crystal Diaz-Espinoza, and Danyelle Ireland. 2018. A Model for Increasing Gender Diversity in Technology. In *SIGCSE*. ACM, 10.1145/3159450.3159533, 459–464.
- [47] Filippo Ricca, Alessandro Marchetto, and Marco Torchiano. 2011. On the Difficulty of Computing the Truck Factor. In *PROFES (Lecture Notes in Business Information Processing)*, Vol. 6759. Springer, https://doi.org/10.1007/978-3-642-21843-9_26, 337–351.
- [48] Gregorio Robles, Laura Arjona Reina, Alexander Serebrenik, Bogdan Vasilescu, and Jesús M. González-Barahona. 2014. FLOSS 2013: a survey dataset about free software contributors: challenges for curating, sharing, and combining. In *MSR*. ACM, 10.1145/2597073.2597129, 396–399.
- [49] Karina Kohl Silveira, Soraia Raupp Musse, Isabel Harb Manssour, Renata Vieira, and Rafael Prikladnicki. 2019. Confidence in programming skills: gender insights from StackOverflow developers survey. In *ICSE (Companion Volume)*. IEEE / ACM, 10.1109/ICSE-Companion.2019.00091, 234–235.
- [50] Vandana Singh. 2019. Women-only spaces of open source. In *Proceedings of the 2nd International Workshop on Gender Equality in Software Engineering, GE@ICSE 2019, Montreal, QC, Canada, May 27, 2019*, Ivica Crnkovic, Karina Kohl Silveira, and Sara Sprenkle (Eds.). IEEE / ACM, https://doi.org/10.1109/GE.2019.00010, 17–20. https://doi.org/10.1109/GE.2019.00010
- [51] Igor Steinmacher, Marco Aurélio Gerosa, Tayana Uchôa Conte, and David F. Redmiles. 2019. Overcoming Social Barriers When Contributing to Open Source Software Projects. *Computer Supported Cooperative Work* 28, 1-2 (2019), 247–290.
- [52] Josh Terrell, Andrew Kofink, Justin Middleton, Clarissa Rainear, Emerson R. Murphy-Hill, Chris Parnin, and Jon Stallings. 2017. Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Computer Science* 3 (2017), e111.
- [53] Parastou Tourani, Bram Adams, and Alexander Serebrenik. 2017. Code of conduct in open source projects. In *IEEE 24th International Conference on Software Analysis, Evolution and Reengineering, SANER 2017, Klagenfurt, Austria, February 20-24, 2017*, Martin Pinzger, Gabriele Bavota, and Andrian Marcus (Eds.). IEEE Computer Society, https://doi.org/10.1109/SANER.2017.7884606, 24–33. https://doi.org/10.1109/SANER.2017.7884606
- [54] Bogdan Vasilescu, Andrea Capiluppi, and Alexander Serebrenik. 2014. Gender, Representation and Online Participation: A Quantitative Study. *Interacting with Computers* 26, 5 (2014), 488–511.
- [55] Bogdan Vasilescu, Vladimir Filkov, and Alexander Serebrenik. 2013. StackOverflow and GitHub: Associations between Software Development and Crowdsourced Knowledge. In *International Conference on Social Computing, SocialCom 2013, SocialCom/PASSAT/BigData/EconCom/BioMedCom 2013, Washington, DC, USA, 8-14 September, 2013*. IEEE Computer Society, https://doi.org/10.1109/SocialCom.2013.35, 188–195. https://doi.org/10.1109/SocialCom.2013.35
- [56] Bogdan Vasilescu, Daryl Posnett, Baishakhi Ray, Mark G. J. van den Brand, Alexander Serebrenik, Premkumar T. Devanbu, and Vladimir Filkov. 2015. Gender and Tenure Diversity in GitHub Teams. In *CHI*. ACM, 10.1145/2702123.2702549, 3789–3798.
- [57] Yi Wang and David Redmiles. 2019. Implicit Gender Biases in Professional Software Development: An Empirical Study. In *ICSE - Software Engineering in Society*. ACM, https://bit.ly/2Hu42MJ, 1–10.
- [58] Zhendong Wang, Yi Wang, and David F. Redmiles. 2018. Competence-confidence gap: a threat to female developers' contribution on github. In *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Society, ICSE (SEIS) 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, Valérie Issarny and Schahram Dustdar (Eds.). ACM, https://doi.org/10.1145/3183428.3183437, 81–90. https://doi.org/10.1145/3183428.3183437