

Perceptions of Diversity on GitHub: A User Survey

Bogdan Vasilescu
University of California, Davis
vasilescu@ucdavis.edu

Vladimir Filkov
University of California, Davis
filkov@cs.ucdavis.edu

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

Abstract—Understanding one’s work environment is important for one’s success, especially when working in teams. In virtual collaborative environments this amounts to being aware of the technical and social attributes of one’s team members. Focusing on Open Source Software teams, naturally very diverse both socially and technically, we report the results of a user survey that tries to resolve how teamwork and individual attributes are perceived by developers collaborating on GITHUB, and how those perceptions influence their work. Our findings can be used as complementary data to quantitative studies of developers’ behavior on GITHUB.

I. INTRODUCTION

Software development is technical and knowledge-intensive, but also human-centric and collaborative, benefiting from the social attributes of the people involved. Open Source Software (OSS) communities, in particular, tend to be quite diverse, with contributors ranging from professional developers to volunteers, all with varied personalities, educational and cultural backgrounds, age, gender, and expertise. Yet, despite participating in a very decentralized process, and despite this diversity, OSS teams often succeed to work together effectively and productively [1], [2].

Understanding one’s environment, be it work, social or natural, is essential for success and survival, and hinges on the quick and effective perception of it [3]. In the modern world, and in particular in virtual environments, this typically simplifies to being aware of the variance in the social attributes of people in the community, *i.e.*, being aware of the social diversity [4]. Diversity arises from attributes that differentiate people, demographic (*e.g.*, age, gender, ethnicity) or otherwise (*e.g.*, role, expertise, personality). In OSS teams diversity can be desirable, resulting in varied backgrounds and ideas, which provide the team with access to broader information and enhanced problem solving skills [5]. On the other hand, due to greater perceived differences in values, norms, and communication styles, members in more diverse teams become more likely to engage in stereotyping, cliquishness, and conflict [6].

Recently we studied social diversity in GITHUB teams [7], the largest and most popular online collaborative coding platform, focusing on gender and tenure (experience). Using regression modeling on data from more than 23,000 GITHUB projects, we showed that after controlling for team size and other technical confounds, both gender and tenure diversity are significant and positive predictors of productivity, together explaining a small but significant fraction of the data variability. Although numerous studies of GITHUB and developers there have sprouted over the past few years (*e.g.*, [8]–[13], few have addressed the importance of individual programmer

attributes (*e.g.*, gender, tenure, political views) on the overall work environment. Our previous study [7] was, to the best of our knowledge, the first to consider effects of gender diversity on productivity and turnover in OSS communities, and one of the very few studies of diversity in general in OSS or other online peer production systems (*e.g.*, [14]–[16]).

In this paper we offer a qualitative perspective of diversity in software teams: we report the results of a user survey that tries to resolve how teamwork and individual attributes are perceived by developers collaborating on GITHUB, and how those perceptions influence their work. We address a number of research questions, as discussed next.

OSS teams are typically more fluid and less tangible than their offline counterparts. They tend to form and dissolve organically around the task at hand, facing high turnover [17], while interactions between members are often limited to online channels [18]. In addition, GITHUB’s implementation of the pull-based development model [19] enables anyone to submit changes to any repository with minimal effort, through pull requests (the so-called “drive-by” commits [13]). We wish to understand whether this unprecedented low barrier to entry for potential contributors is changing perceptions of teams (**RQ1**) and team dynamics (**RQ2**) in GITHUB teams.

RQ1. What do people perceive constitutes a team?

RQ2. How does team composition change with time?

The extent to which individual characteristics are salient impacts how team members react to diversity [20]. Demographic features such as ethnicity or gender, often a source of social categorization and stereotyping in offline settings, are expected to become less salient in OSS [21]. Instead, OSS communities should function as meritocracies [22], with sustained, high-quality contributions as the main drivers of impression formation, reputation building, and trust [8], [10]. Still, OSS is often criticized for sexism [23], [24], suggesting a prominent role of demographics (gender in particular) in impression formation. We sought to understand whether diversity attributes are recognized by GITHUB contributors in their team members (**RQ3**), and what mechanisms contribute to increased awareness of these attributes (**RQ4**).

RQ3. Do individuals recognize differences among others on their team? Which differences are more prominent?

RQ4. What mechanisms contribute to increased awareness of diversity attributes among team members?

While numerous studies (mostly from offline groups) report on the relationship between diversity and team outcomes, the effects are not always positive [6]. Although there is evidence

that a team’s social diversity may improve its technical performance [25], even from GITHUB where a team’s gender and tenure diversity may improve its productivity [7], the “hacker” OSS culture still has episodes of discrimination, unfriendliness to newcomers, and conflict. We wish to understand how GITHUB developers perceive collaboration in inherently diverse teams, and whether/how they benefit from it (RQ5).

RQ5. How is diversity perceived to influence collaboration?

II. SURVEY DESIGN

We designed an online survey¹ about perceptions of teams, perceptions of diversity attributes among team members, and experiences with working in diverse teams. Our survey consisted of twenty questions, three of which were open-ended. The survey started with a series of questions about demographics (age, gender, nationality, country of residence, years of IT/programming experience, and occupation) and experience with GITHUB (when they started using GITHUB and how they use it, e.g., by pushing commits, reporting issues, submitting pull requests, participating in discussions, and following other developers). Then we asked participants to select one repository among the ones they contributed to, and answer the remaining questions with respect to that repository. We also asked whether they owned this repository or just contributed to it, and what their contributions consisted of.

Following are the questions related to our main goals. In the first group of questions we asked (i) whether respondents considered themselves part of a team when contributing to that repository (*Yes; No; I prefer not to answer*); (ii) whom they consider part of their team (with multiple choices associated with different levels of contribution, from *People who work on my particular feature/branch*—arguably the most exclusive option, through *People who contribute code frequently*, or *People who report issues*, or *People who submit pull requests*, or *People who participate in discussions*, to *Everyone who does something in this repository* (e.g., *pushes code, submits pull requests, reports issues*)—the most inclusive option); (iii) whether and how the composition of their team changed with time (open-ended); and (iv) what mechanisms they used for communicating with team members (*Comments on commits, pull requests, or issues; Emails sent directly; A mailing list; IRC or instant messaging; In person; Other mechanisms*).

In the second group of questions we asked (i) which characteristics of their team members (diversity attributes) respondents are aware of (*Programming skills; Social skills; Gender; Ethnicity; Overall GITHUB experience; Reputation as programmer; Country of residence; Personality; Age; Educational level; Real name; Hobbies; Employment; Political views*) and how often (i.e., for how many team members) this happens; and (ii) when and how they became aware of these attributes, and whether this changed with time (open-ended).

Finally, in the third group of questions we asked (i) how they would rate the overall experience of working in diverse teams (*Positive; Negative; Sometimes positive, sometimes negative; Not applicable*); and (ii) why they selected this rating (open-

ended; respondents were asked to describe circumstances that led to their positive or negative experiences).

To distribute the survey, we compiled a stratified random sample of 4,500 GITHUB contributors. For each combination of gender (as inferred by our name-based tool `genderComputer` [26], i.e., either male, female, or unknown) and number of projects contributed to (one project—the majority of contributors; many projects—7 or more, distributional outliers; and few projects—between 2 and 6), we randomly selected 500 individuals with known email addresses in our data set [7] (in total $3 \times 3 \times 500 = 4,500$). The data [7] is based on the GHTorrent [27], [28] dump dated 1/2/2014.

We contacted all of them individually by email inviting them to participate in the survey. Participation was voluntary and confidential, and was estimated to take about ten minutes. We received 816 responses (71.2% of them within 24 hours; 91.3% within a week; see Figure 1) and 236 automatic replies (failure to deliver) until we closed the survey after one month. Responses were anonymous. The response rate computed after ignoring the automatic replies was slightly higher than 19%.




Fig. 1. Survey response times (logarithmic y-axis).

III. RESPONDENT DEMOGRAPHICS

199 respondents (24.4%) indicated their gender as *female*, 611 (74.9%) as *male*, 4 (0.5%) as *other*, and 2 did not answer this question. The numbers suggest an under-representation of women (at least one third of the invitations had been sent to individuals whose gender has been inferred as female by our tool). The respondents’ age ranged from 14 to 66 (median 29; mean 30), and their IT experience ranged from zero to 44 years (median 8; mean 10.5). The largest group of respondents reside in the USA (264), followed by Germany (53), France (43), UK (40), and others (Table I, left). Nationality rankings follow closely: USA (219), followed by Germany (56), France (54), Russia (41), and others (Table I, right).

To evaluate the representativeness of the respondents, we aggregate country information to macro-regional level (Africa, Asia, Australia and New Zealand, Eastern and Southern Europe, Latin America, North America, Western and Northern Europe). A χ^2 test comparing the macro-regional distribution of our respondents with that reported in a previous study of GITHUB users [29] could reveal no difference ($p > 0.9$).

The most popular occupation reported was web developer (59.7%), followed by manager / team leader (21.5%), student (20.6%), desktop developer (21.3%), and others (Table II). 46.3% of respondents reported a single occupation, 23.5% reported two occupations, and the rest reported three or more.

¹http://bvasiles.github.io/papers/diversity_survey.pdf

TABLE I
TOP TEN COUNTRIES OF RESIDENCE AND NATIONALITIES.

Residence	#	%	Nationality	#	%
USA	258	32.35	USA	213	26.84
Germany	53	6.50	Germany	56	6.86
France	43	5.27	France	54	6.62
UK	40	4.90	Russia	41	5.02
Canada	34	4.17	UK	34	4.17
Russia	32	3.92	Canada	30	3.68
Brazil	22	2.70	India	27	3.31
India	20	2.45	Brazil	25	3.06
Sweden	20	2.45	Italy	22	2.70
Netherlands	18	2.21	Poland	21	2.57

TABLE II
RESPONDENT OCCUPATIONS (MULTIPLE POSSIBLE PER PERSON).

Occupation	%
Web developer	59.70
Manager / Team leader	21.50
Student	20.64
Desktop software developer	21.25
Mobile application developer	19.16
IT staff / System administrator	15.48
Academic	13.51
Other	13.14
Database administrator	9.95
Embedded application developer	9.46
I don't work in tech	2.58

We wish to understand whether there are significant differences between *female* and *male* respondents (the only groups with sufficient data for statistical tests) with respect to different demographic attributes: response time, age, occupation, nationality, IT/programming experience, and ownership of GITHUB repositories. Since the groups are unbalanced (female 24%; male 75%), we randomly subsample the male group 1000 times, and report median statistical significance results and median effect sizes after the 1000 runs.

We found significant differences between the **occupations** reported by female and male respondents (χ^2 tests of independence male vs. female, per occupation), with sizable albeit small effects, as described next. In decreasing order of Cramér’s V measure of effect size, male respondents are over-represented as IT staff/system administrators ($V=0.24$; odds ratio $OR=0.18$; $p^{***}=5.45\times 10^{-6}$), mobile app developers ($V=0.19$; $OR=0.32$; $p^{***}=1.90\times 10^{-4}$), desktop developers ($V=0.16$; $OR=0.43$; $p^{**}=2.74\times 10^{-3}$), database administrators ($V=0.11$; $OR=0.43$; $p^*=0.04$), managers/team leaders ($V=0.11$; $OR=0.57$; $p^*=0.04$), or web developers ($V=0.11$; $OR=0.65$; $p^*=0.04$). Female respondents are over-represented as academics ($V=0.12$; $OR=1.92$; $p^*=0.03$). There are no significant gender differences among students ($p=0.39$), embedded software developers ($p=0.15$), and other occupations ($p=0.15$). In terms of **number of occupations** reported per person, we found significant differences between female and male respondents (WMW $p^{***}=4.43\times 10^{-6}$), but negligible effect sizes (Hodges-Lehmann point estimate for median difference $\hat{\Delta}=3.28\times 10^{-5}$).

We did not find significant **age** differences between genders (median female=male=29; $p=0.72$); a slight difference in **response time** (female: 7.66h; male: 6.64h; $\hat{\Delta}=-1.07$; $p^\circ=0.08$); and a significant difference in number of **years of IT/programming experience** (male: 9; female: 6; $\hat{\Delta}=2.00$; $p^{***}=2.06\times 10^{-4}$). In previous work we found that females

account for only 9% on GITHUB [7]. These results suggest that women are not only underrepresented on GITHUB, but also under-experienced compared to men. That is, the $\hat{\Delta}=2.00$ value means that the difference in experience between male and female respondents, as seen by the WMW rank-sum test, would only disappear if the experience of *each* female would increase by 2 years (or, equivalently, the experience of each male would decrease by 2 years).

Genders are also represented differentially across **nationalities**. Since there is insufficient data for statistical tests, we present the distribution of the fraction of female respondents among the countries with 20 or more respondents visually in Figure 2. Four countries stand out: in USA (37.1% female; 213 respondents) and Brazil (44%; 25) the fraction of female respondents is more than one median absolute deviation (MAD) higher than the median; in contrast, in Germany (14.3%; 56) and Russia (14.6%; 41) the fraction of female respondents is more than one MAD lower than the median.




Fig. 2. Fraction of female respondents, per nationality. Only countries with at least 20 respondents shown. Median=0.238 (solid). MAD=0.083 (dotted).

Finally, we investigated gender differences in usage of GITHUB. We found that female respondents own significantly fewer **public repositories** than males do ($V=0.13$; $OR=0.45$; $p^*=0.01$), although there is no difference in ownership of **private repositories** (WMW $p=0.36$). Furthermore, we investigated gender differences in usage of different GITHUB features (pushing commits; submitting pull requests; reporting issues; reviewing and closing pull requests; participating in discussions on commits, pull requests, or issues; watching; forking; and following) among the top users of each feature (*i.e.*, respondents who reported using that feature *frequently*). We could only find a significant difference with respect to watching repositories: significantly fewer female respondents watch repositories frequently compared to males ($V=0.16$; $OR=0.46$; $p^{**}=0.002$). We conjecture that the gender disparity in having public repositories on GITHUB, despite the otherwise indistinguishable feature usage of active female and male users (except for watching, as mentioned), may be related to the so-called “impostor syndrome” that women programmers reportedly suffer from [24]: despite being knowledgeable and professionally well-settled, women may be more reluctant to publicly display their work. This is especially plausible on GITHUB, where heavy users are known to have a “clear awareness of the audience for their actions”, which influences how they behave and how they construct their actions [10].

IV. TEAMS AND TEAM DYNAMICS

Two-thirds of respondents (541) indicated they consider themselves part of a team when working on a repository. Women score significantly higher than men ($V=0.14$; $OR=1.87$; $p^{**}=0.006$), consistent with their more socio-emotional behavior in teams [30], [31]. The discussion in the remainder of this paper is restricted to these 541 respondents.

A. Perceptions of Teams (RQ1)

When asked whom they consider part of their team, the most popular answer was also the most inclusive: “everyone who does something in this repository (e.g., pushes code, submits pull requests, reports issues)” (72%). Differences in choosing this answer based on gender (male vs. female χ^2 $p \simeq 0.80$), age (below vs. above median $p \simeq 0.53$), programming experience (below vs. above median $p \simeq 0.97$), ownership of said repository (yes vs. no $p \simeq 0.72$), and role (pushing commits vs. others $p \simeq 0.56$; reviewing and closing pull requests vs. others $p \simeq 0.15$) were not statistically significant. Less inclusive answer choices, e.g., “people who contribute code frequently” (53%, second-most popular), “the repository owner and others who can push commits directly” (50%), and “people who work on my particular feature/branch” (38%), were less popular.

OSS communities are known to have an “onion” structure [32]. Developers that have been around the longest and that drive the work typically make up the *core*, the group associated with the highest reputation. All others, who support the core group by reporting issues, submitting patches, and contributing documentation, are viewed as *peripheral*. It is through sustained participation, socialization, and high-quality contributions that they can build their reputation [33], [34] and advance through the ranks, eventually becoming part of the core. We stress that even though similar distinctions can be made between contributor roles on GITHUB, differences between the preference for the most inclusive answer “everyone who does something in this repository” based on the respondents’ role in the repositories were not statistically significant, i.e., core developers pushing commits and integrating pull requests, who may be expected to only consider other core developers as part of the team, chose this answer as frequently as peripheral developers did. As one respondent puts it, “the team is whoever makes good suggestions” [R99].

B. Communication Mechanisms

GITHUB contributors have two main communication needs: getting help from other team members when encountering problems, and discussions around specific coding items (e.g., code review) [35]. To handle the latter, GITHUB offers integrated comments on either issues, commits, or pull requests, which is also the most popular communication mechanism among our respondents: 73.2%; slightly underused by women ($V=0.11$; $OR=0.62$; $p^\circ=0.075$). Respondents also use email communication (directly—63.2%; mailing lists—32.7%), IRC or instant messaging (54.2%), in-person communication (52.7%), and other channels (7%). Significantly more women prefer in-person communication than men ($V=0.25$; $OR=2.76$; $p^{***}=2.95 \times 10^{-5}$). There were no other gender

differences. Other communication mechanisms mentioned by respondents include Skype, Google Hangouts, Mumble (all for videoconferencing), the JIRA and Bugzilla issue tracking systems, the project’s web forums and message boards, Basecamp, the Pivotal Tracker project management software, the Freenode IRC network, Slack chats, Twitter, and SMS.

C. Team Dynamics (RQ2)

Approximately half of the respondents (267 out of 541 who indicated being part of a team; 49.4%) commented on changes in the composition of their teams with time. Coding the open-ended answers resulted in the following themes (one response could not be coded).

a) *Fluid Teams*: Most respondents (142; 53.2%) describe a fluid team formation process, characteristic of OSS, with voluntary developers that “*come and go as their interest waxes and wanes*” [R384]. Becoming part of the team requires contributing to the project (“*people are free to contribute, by contributing they can become members of the team*” [R476]), often by submitting pull requests (“*if the project is being used by many people, a growing number of users will submit pull requests or issues, thus growing the team*” [R219]). Commitment to the project is viewed as important, and sustained contributions lead to higher status (“*a frequent contributor buys a long length of time before we stop considering them part of the team*” [R384]) or becoming part of the core (“*core developers are added based on a history of positive development contributions*” [R507]), and insufficient participation—to friction (“*active contributors tend to drop off, while more people create issues without creating corresponding pull requests. The repo becomes noisier while decreasing in substance*” [R423]).

Respondents also recognize high turnover: “*open source developers come and go much [more] frequently than in a company*” [R81]. Yet, despite this, the general mix of roles tends to remain constant, creating a “*rotating door effect where everyone gets a time to be a main team member, to lead the group, as well as a time where they contribute minimally and work for the group’s goals*” [R156]. Most turnover occurs, as expected, at the periphery: “*there is a core team of people (the owners of the repo and a few large contributors that comment and submit pull requests frequently). Other than this core, the composition of the team is very dynamic, most members are active for a month at most*” [R529].

Multiple factors are associated with changes in GITHUB project teams: varying interest in the project or different personal work focus (“*it is natural for people to be excited and then become complacent [...] People move on when their interest lulls.*” [R69]); changes in project focus (teams “*will expand and contract as demanded and as time becomes available to work on features*” [R51]; the team has changed as the project “*has been handed around via forks on GITHUB*” [R752]); being users of the code/software (“*people get brought in as they have reason to use the code*” [R299]); desire to learn (“*it is common for team members to learn, build skills, and move on to other open source projects when they are comfortable*” [R613]); availability (developers “*drop out as life becomes busy*” [R380]); changes in employment (“*some people change*

job” [R327]); and reliance on fixed-term external contributors, such as student interns (“*our project participates in Google Summer of Code, which attracts students*” [R640]).

b) *Commercial Teams*: 50 respondents (18.7%) described experiences from commercial projects hosted on GITHUB, with different dynamics. Changes in team composition here are mostly due to the company’s human resourcing practices (“*the composition of the team changes as we hire and fire people*” [R733]); outsourcing (“*we frequently engage third-party developers to specialize in particular areas*” [R191]); and internal task reassignment (“*team members rotate frequently*” [R742], “*often on a weekly basis*” [R67]).

c) *Academic Teams*: 17 respondents (6.4%) referred to academic projects hosted on GITHUB, where teams are not as fluid as generally in OSS. Team dynamics here are governed either by “*student placements starting and ending*” [R643], graduate students joining a particular lab, and team assignments made by the course lecturers.

d) *Stable Teams*: A significant fraction of the respondents (21.3%) mentioned static or relatively stable teams. 32 respondents considered that their team’s composition did not change with time at all, given a particular project. Only “*with new repos may come new teams*” [R364]. 10 respondents chose to talk about personal projects, where they work alone or in very small teams (1-2 members). Their teams do not change with time, although they are open to outside contributions: “*If someone became interested in the project and started submitting a bunch of pull requests we thought were good we’d probably add them as a committer and discuss the project with them outside of GITHUB*” [R305]. Finally, 15 respondents recognize small changes in their teams with time, but they qualify these as rare, e.g., “*people can join at any time, but this is an infrequent occurrence*” [R11], or “*as long as I have been part of the team there hasn’t been much change*” [R667].

V. DIVERSITY AWARENESS

A. Perceptions of Diversity Attributes (RQ3)

Almost three quarters of the respondents (Table III) are aware of the *programming skills* of “most other” team members (recall the answer choices: “most other”; “few other”; and “none other” team members), making this the most visible diversity aspect we considered. This is hardly surprising as GITHUB contributors can be expected to interact mostly on technical issues, and demonstrated programming skill is one of the main drivers behind impression formation [8].

More surprisingly, 48.2% of the respondents indicated they are aware of the *gender* of most of their teammates, making it the second-most visible attribute; the awareness of the other social attributes is given in Table III, on the left. This salience of gender, real name, and country of residence among team members on GITHUB, especially given the all-inclusive notion of a “team” (recall the discussion in Sec. IV-A), contradicts earlier claims of obscurity of demographics in OSS [21] and, we believe, is related to the “evolution of the social programmer” [36]: social media (of which GITHUB as a “social coding” platform is an example) has changed how

TABLE III
INDIVIDUAL ATTRIBUTES THAT RESPONDENTS ARE AWARE OF FOR MOST TEAM MEMBERS. χ^2 MEDIAN DIFFERENCES MALE VS. FEMALE AFTER 1000 REPETITIONS WITH BALANCED RANDOM SAMPLES.

Attribute	#	%	Direction	V	OR	p
Progr. skills	401	74.12				0.51
Gender	261	48.24	Female*	0.14	1.74	0.02
Real name	245	45.29	Female*	0.15	1.82	0.02
Social skills	227	41.96	Female*	0.13	1.69	0.04
Country of residence	217	40.11				0.24
Personality	212	39.19	Female*	0.13	1.71	0.04
Reputation as progr.	168	31.05				0.22
Ethnicity	162	29.94	Female**	0.16	1.97	0.007
Employment	162	29.94	Female***	0.20	2.34	0.0007
GitHub experience	150	27.73				0.44
Educational level	142	26.25	Female*	0.14	1.83	0.02
Age	124	22.92	Female*	0.12	1.74	0.04
Hobbies	62	11.46				0.38
Political views	23	4.25				0.77

developers communicate, collaborate, forage for information, form impressions, and coordinate their work.

It is interesting to note significant gender specific differences in the perception of social vs. technical diversity attributes. Female respondents are more frequently aware of the *gender*, *real name*, *social skills*, *personality*, *ethnicity*, *employment*, *educational level*, and *age* of most of their team members, with sizable effects (Table III, right). *Residence country* is a notable exception, being a social diversity attribute for which we could not detect gender differences in awareness.

B. Awareness Mechanisms (RQ4)

In open-ended answers, 320 respondents chose to describe how they became aware of the above attributes of team members. Coding their answers revealed the following themes:

a) *In-person Interactions*: Respondents frequently know their team members in person, allowing the unmasking of their demographic features. In-person interactions, as mentioned in 167 answers (52.2%), are due to either *affiliation with the same organization*, e.g., a commercial company (52 answers) or academic institution (19), some even sharing office space (19); meeting each other at *conferences*, *meetups*, *project sprints*, etc. (25); or *being offline friends or acquaintances* (73).

b) *GitHub-enabled Interaction*: Team members learn about their programming skills from each other’s *code contributions*, *issues*, and *pull requests* (24); they use *comments on commits*, *issues*, and *pull requests* to make inferences about each other’s social skills (24); finally, demographic features (name, photo, gender, location, and employment) become apparent on *GITHUB profile pages* (35).

c) *Email*: 25 respondents use email exchange features (*address*, *signature*, and *contents*) for information. As much as a person’s name, “*country of residence*, *educational level*, and *employment can be inferred via email address and/or email signature*” [R490] alone, while other personal characteristics are discovered through repeated email communication, when discussions “*diverge from talking strictly about the repository to more informal conversation*” [R390]. A special case of email communication is communication on *mailing lists*, *message boards*, and *forums* (10), which may store historical traces of participants’ personal information.

d) *Instant Messaging*: Text messaging, via *IRC* (21) or other *instant messengers* (16), and combined audio, video, and text messaging, via *Skype* (10) or other channels, are also used to uncover diversity attributes among team members.

e) *Social Media*: 26 respondents mentioned using social media, e.g., *Twitter* (11) and *blogs/personal websites* (8), and other online resources to make inferences about collaborators.

C. Perceptions of Diversity Effects (RQ5)

When asked to give an overall rating of the collaborative experience in a diverse GITHUB team, respondents most frequently chose *positive* (62.5% of 541), followed by *sometimes positive, sometimes negative* (30%). The rest chose *not applicable* or gave no rating. None marked it as *negative*.

Coding the open-ended answers of those who also described their experiences resulted in the following themes.

a) *Positive Effects*: Many acknowledge the positive effects team diversity has had on themselves and their teams. At team level, diversity can provide new ideas, perspectives, skills, and approaches to solve problems, e.g., “*having more people with more backgrounds and varying attributes simply helps strengthen the community and helps it grow in a positive light*” [R709]; access to different networks, e.g., “*it is positive to work with people from different countries and cultures [...] [and] benefit from the network of the contributors to spread the word and also to learn how the project could adapt*” [R690]; lively discussions around issues and pull requests, e.g., “*diverse viewpoints often lead to lively discussions and new ideas*” [R284]; and ultimately better code. Positive effects of diversity on software quality are recognized along many axes, ranging from design, e.g., “*the diversity of our community [...] [is] part of our process for ensuring all needs are considered and if possible catered for*” [R674]; through usability, e.g., “*diversity in the body of folks willing to interact and contribute works to strengthen the usability of the library [...] if folks who are not terrifically experienced can make sense of it and use it*” [R597]; to localization, e.g., “*geographic diversity [...] is helpful for finding all sorts of important but otherwise invisible issues, such as time zone bugs, cultural issues*” [R436].

Individually, diversity leads to increased understanding, e.g., “*diversity brings understanding, whether in code or not*” [R310]; learning and mentoring opportunities, e.g., “*teaching or guiding naive users is an excellent way to clarify my own thinking. I also get to learn an enormous amount from programmers who are more experienced than me*” [R719]; and an otherwise enriching experience, e.g., “*in general it is always enriching to communicate with someone different*” [R621].

b) *No Effects*: Some respondents adhere to the meritocratic model of OSS, maintaining that it is “*more about the contributions to the code than the ‘characteristics’ of the person*” [R472], and that “*any demographic identity is irrelevant*” [R490], since “*code sees no color or gender*” [R701]. Still, besides high-quality contributions that adhere to the project’s development practices (i.e., “*clean code*” [R245]), respondents identify the following individual attributes as important for the success of collaborations in OSS: passion, motivation, and

interest towards a common goal; professionalism; respectfulness; politeness; openness to new ideas, agreeableness, and welcomeness; helpfulness; friendliness; communication skills; fostering a positive, constructive, democratic atmosphere and mutual trust; patience and a cool head; being well organized.

c) *Negative Effects*: Related to **programming experience**, examples include: improper contributions by newcomers (“*sometimes newer team members have trouble with code health and best practices*” [R530]), that require more effort to integrate (“*sometimes, pull requests [from less experienced developers] cannot be merged automatically, making my life harder*” [R22]); lack of familiarity with git ([there] “*seems to be a steep learning curve for new developers as far as editing comments, deleting commits, and working on branches*” [R405]); and expectations mismatch (“*it can be frustrating when people have different expectations of what will be done on the project, and in what timespan, usually due to their level of programming experience or personality*” [R656]).

Role diversity within a team is another source of friction, e.g., developers that collaborate with experts from the client’s side, viewing the latter as “*not always a positive asset to the project*” [R625]; pull request submitters having their contributions ignored (“*it’s frustrating if other developers don’t react*” [R99]), mistreated (“*I’m willing to invest my own time to create a pull request and improve or even fix some parts a project and all I get is snifty response from the one who’s in charge of accepting pull requests.*” [R174]), or rejected (“*working on an idea in a branch and having it rejected from inclusion after a discussion*” [R592]). Similarly, differences in **ideology** between team members, e.g., “*disagreements over how and what to implement*” [R478], whether to “*focus on detail/correctness versus efficiency, [or] innovation versus using existing patterns*” [R240] “*can be time consuming with little return value*” [R478], potentially resulting in conflict.

Respondents also mentioned differences in **personality** and **social skills** between team members as a common source of frustration. Personalities may clash “*in the always delicate world of commenting someone else’s code in a code review*” [R527]. As a result, the experience becomes negative: working with “*belligerent developers can be disheartening*” [R483].

Negative experiences may also arise from **gender** differences. Women report that being the only female developer “*can sometimes be frustrating*” [R653], that “*despite having proved [their] competency in certain areas of the code/project, [their] opinion is rarely or never asked for*” [R519], that they “*have used a fake GitHub handle [...] so that people would assume [they] were male*” [R305], or that gender-related incidents have “*caused [them] to leave a project*” [R556].

Differences in **national origin** and **language** may lead to difficulties in communication or “*occasional confusion over the use of idioms and misinterpreted emotion*” [R359], causing “*a lot of stress for the rest of the team*” [R605].

VI. THREATS TO VALIDITY

The generalizability and validity of our findings may be limited by our purposive sampling of participants (stratified according to gender and number of projects contributed to

but not, *e.g.*, according to team size) and self-selection bias. Although we found our sample of respondents to be representative of the overall GITHUB population in terms of geographic distribution, the general population may have other different characteristics and different opinions. Still, we found our results to be quite robust, *e.g.*, perceptions of teams are independent of demographics and project roles. Furthermore, the relatively delicate subject matter and a potential tendency of respondents to appear in a positive light may have influenced the accuracy of the answers.

VII. CONCLUSIONS

We presented a qualitative perspective of diversity in software teams, based on results from a user survey with 816 responses from GITHUB project contributors. As a central hub for collaborative coding, GITHUB offers several features that have shaped software development, especially in OSS, *e.g.*, very low barrier to entry for newcomers via pull requests, and public displays of information about one's history of contributions. In this very dynamic, public, and social environment, we tried to resolve how team composition and individual social and technical attributes (such as gender, nationality, and experience) are perceived by collaborators, and how team diversity along these attributes may influence their work.

In addition to offering a broad perspective on collaboration in GITHUB teams, we uncovered some surprising perceptions, *e.g.*, of team composition (the all-inclusive *everyone who does something in the repository*, regardless of how small the contribution is), and salience of demographics (especially *gender*, the second-most visible attribute, much more frequently salient than, say, one's level of GITHUB experience). We also found that developers have embraced the inherent diversity from GITHUB teams and, by and large, they benefit from it.

This study is complementary to our quantitative study of diversity of gender and tenure in GITHUB [7]. The findings from the two together portray a multifaceted and convincing picture of diversity's importance in OSS projects.

REFERENCES

- [1] K. Crowston, K. Wei, J. Howison, and A. Wiggins, "Free/Libre open-source software development: What we know and what we do not know," *ACM Computing Surveys (CSUR)*, vol. 44, no. 2, p. 7, 2012.
- [2] D. M. German, "The GNOME project: a case study of open source, global software development," *Software Process: Improvement and Practice*, vol. 8, no. 4, pp. 201–215, 2003.
- [3] D. Schacter, D. Gilbert, and D. Wegner, *Psychology*. Worth Publ., 2010.
- [4] B. Wellman and M. Gulia, "Net surfers don't ride alone: Virtual communities as communities," *Networks in the global village*, pp. 331–366, 1999.
- [5] S. E. Jackson and A. Joshi, "Diversity in social context: a multi-attribute, multilevel analysis of team diversity and sales performance," *J. Organ. Behav.*, vol. 25, no. 6, pp. 675–702, 2004.
- [6] S. K. Horwitz and I. B. Horwitz, "The effects of team diversity on team outcomes: A meta-analytic review of team demography," *J. Manag.*, vol. 33, no. 6, pp. 987–1015, 2007.
- [7] B. Vasilescu, D. Posnett, B. Ray, M. G. J. van den Brand, A. Serebrenik, P. Devanbu, and V. Filkov, "Gender and tenure diversity in GitHub teams," in *CHI*. ACM, 2015, to appear.
- [8] J. Marlow, L. Dabbish, and J. Herbsleb, "Impression formation in online peer production: activity traces and personal profiles in GitHub," in *CSCW*. ACM, 2013, pp. 117–128.
- [9] B. Vasilescu, A. Serebrenik, and V. Filkov, "Stack Overflow and GitHub: Associations between software development and crowdsourced knowledge," in *SocialCom*. IEEE, 2013, pp. 188–195.
- [10] L. Dabbish, C. Stuart, J. Tsay, and J. Herbsleb, "Social coding in GitHub: transparency and collaboration in an open software repository," in *CSCW*. ACM, 2012, pp. 1277–1286.
- [11] N. McDonald and S. Goggins, "Performance and participation in open source software on GitHub," in *CHI*. ACM, 2013, pp. 139–144.
- [12] G. Gousios, M. Pinzger, and A. van Deursen, "An exploratory study of the pull-based software development model," in *ICSE*. ACM, 2014, pp. 345–355.
- [13] R. Pham, L. Singer, O. Liskin, and K. Schneider, "Creating a shared understanding of testing culture on a social coding site," in *ICSE*. IEEE, 2013, pp. 112–121.
- [14] S. Daniel, R. Agarwal, and K. J. Stewart, "The effects of diversity in global, distributed collectives: A study of open source project success," *Inform. Syst. Res.*, vol. 24, no. 2, pp. 312–333, 2013.
- [15] L. S. Wang, J. Chen, Y. Ren, and J. Riedl, "Searching for the goldilocks zone: trade-offs in managing online volunteer groups," in *CSCW*. ACM, 2012, pp. 989–998.
- [16] J. Chen, Y. Ren, and J. Riedl, "The effects of diversity on group productivity and member withdrawal in online volunteer groups," in *CHI*. ACM, 2010, pp. 821–830.
- [17] G. Robles and J. M. Gonzalez-Barahona, "Contributor turnover in libre software projects," in *Open Source Syst*. Springer, 2006, pp. 273–286.
- [18] A. Guzzi, A. Bacchelli, M. Lanza, M. Pinzger, and A. v. Deursen, "Communication in open source software development mailing lists," in *MSR*. IEEE, 2013, pp. 277–286.
- [19] E. T. Barr, C. Bird, P. C. Rigby, A. Hindle, D. M. German, and P. Devanbu, "Cohesive and isolated development with branches," in *FASE*. Springer, 2012, pp. 316–331.
- [20] M. E. Zellmer-Bruhn *et al.*, "When and how do differences matter? an exploration of perceived similarity in teams," *Organ. Behav. Hum. Dec.*, vol. 107, no. 1, pp. 41–59, 2008.
- [21] C. M. Riordan and L. M. Shore, "Demographic diversity and employee attitudes: An empirical examination of relational demography within work units," *J. Appl. Psychol.*, vol. 82, no. 3, p. 342, 1997.
- [22] J. Feller and B. Fitzgerald, "A framework analysis of the open source software development paradigm," in *ICIS*. Association for Information Systems, 2000, pp. 58–69.
- [23] D. Nafus, J. Leach, and B. Krieger, "FLOSSPOLs Deliverable D 16 Gender: Integrated Report of Findings," http://www.flosspols.org/deliverables/D16HTML/FLOSSPOLs-D16-Gender_Integrated_Report_of_Findings.htm, 2006.
- [24] D. Nafus, "'Patches don't have gender': What is not open in open source software," *New Media & Society*, vol. 14, no. 4, pp. 669–683, 2012.
- [25] C. P. Earley and E. Mosakowski, "Creating hybrid team cultures: An empirical test of transnational team functioning," *Acad. Manag. J.*, vol. 43, no. 1, pp. 26–49, 2000.
- [26] B. Vasilescu, A. Capiluppi, and A. Serebrenik, "Gender, representation and online participation: A quantitative study," *Interacting with Computers*, vol. 26, no. 5, pp. 488–511, 2014.
- [27] G. Gousios, "The GHTorrent dataset and tool suite," in *MSR*. IEEE, 2013, pp. 233–236.
- [28] G. Gousios, B. Vasilescu, A. Serebrenik, and A. Zaidman, "Lean GHTorrent: GitHub data on demand," in *MSR*. ACM, 2014, pp. 384–387.
- [29] Y. Takhteyev and A. Hiltz, "Investigating the geography of open source software through GitHub," 2010.
- [30] S. G. Rogelberg and S. M. Rumery, "Gender diversity, team decision quality, time on task, and interpersonal cohesion," *Small Group Research*, vol. 27, no. 1, pp. 79–90, 1996.
- [31] A. Michailidou and A. Economides, "Gender and diversity in collaborative virtual teams," *Computer Supported Collaborative Learning: Best practices and principles for instructors*, pp. 199–224, 2007.
- [32] N. Ducheneaut, "Socialization in an open source software community: A socio-technical analysis," *CSCW*, vol. 14, no. 4, pp. 323–368, 2005.
- [33] M. Gharehyazie, D. Posnett, B. Vasilescu, and V. Filkov, "Developer initiation and social interactions in OSS: A case study of the Apache Software Foundation," *Empirical Software Engineering*, pp. 1–36, 2014.
- [34] K. R. Lakhani and E. Von Hippel, "How open source software works: 'free' user-to-user assistance," *Research Policy*, vol. 32, no. 6, pp. 923–943, 2003.
- [35] E. Kalliamvakou, D. Damian, L. Singer, and D. M. German, "The code-centric collaboration perspective: Evidence from GitHub," DCS-352-IR, University of Victoria, Tech. Rep., 2014.
- [36] C. Treude, F. Figueira Filho, B. Cleary, and M.-A. Storey, "Programming in a socially networked world: the evolution of the social programmer," in *The Future of Collaborative Software Development*, 2012, pp. 1–3.