Analyzing Offline Social Engagements: An Empirical Study of Meetup Events Related to Software Development

Abhishek Sharma*, Gede Artha Azriadi Prana*, Anamika Sawhney*, Nachiappan Nagappan+, and David Lo*
*School of Computing and Information Systems, Singapore Management University, Singapore
abhisheksh.2014,arthaprana.2016,anamikas,davidlo@smu.edu.sg
+Microsoft Research, Redmond, WA USA
nachin@microsoft.com

***Abstract*—Software developers use a variety of social media channels and tools in order to keep themselves up to date, collaborate with other developers, and find projects to contribute to. Meetup is one of such social media used by software developers to organize community gatherings. We in this work investigate the dynamics of Meetup groups and events related to software development. Our work is different from previous work as we focus on the actual event and group data that was collected using Meetup API.**

In this work, we performed an empirical study of events and groups present on Meetup which are related to software development. First, we identified 6,327 Meetup groups related to software development and extracted 250,369 events organized by them. Then we took a sample of 452 events on which we performed open coding, based on which we were able to develop 9 categories of events (8 main categories +"Others"). Next, we did a popularity analysis of the categories of events and found that Talks by Domain Experts, Hands-on Sessions, and Open Discussions are the most popular categories of events organized by Meetup groups related to software development. Our findings show that more popular categories are those where developers can learn and gain knowledge. On doing a diversity analysis of Meetup groups we found 20.46% of the members on average are female, and 20.34% of the actual event participants are female, which is a larger proportion as compared to numbers reported in previous studies on gender representation in software engineering communities. We also found evidence that the gender of Meetup group organizer affects gender distribution of group members and event participants. Finally, we also looked at some data on how COVID-19 has affected the Meetup activity and found that the event activity has dropped, but not stalled. A substantial number of events are now being organized virtually. The results and insights uncovered in our work can guide future studies related to software communities, groups, and diversity-related studies.

I. Introduction

Software development has evolved into an increasingly social activity over past few decades. Social media such as social coding sites, Q&A forums, and microblogs are used extensively by developers for activities such as reusing other projects and tools, keeping up to date, learning new skills, connecting and collaborating with other developers [1]. Storey et al. had done a survey to understand how various social channels shape the participatory culture in software development [2]. One of such channels touched upon in their work is Meetup[1]. Meetup is an online social networking service, which allows people to organize events and gatherings. It allows people to form groups or communities focused on common topics of interest. The organizers of such groups can then organize off-line gatherings or events. The events that are organized range from informal congregations to formal events such as conferences. Liu et al. characterized Meetup as an event-based social network (EBSN) which contains valuable *offline* social interactions in addition to online interactions [3]. It is one of the biggest EBSNs available today with 44 million members spread across 330,000+ groups [4]. Recently, Ingram et al. interviewed the leaders of some technology related Meetup groups and found evidence that software practitioners use Meetup groups and events to stay updated, build local networks, and improve their tacit knowledge by interacting with peers [5].

In this empirical study, taking cue from previous works we analyze what kinds of events are organized in Meetup groups related to software development, and the underlying gender distribution of such events and groups. Our main motivation is to understand if events organized by such Meetup groups can be classified into some specific categories, and how popular such categories are. An understanding of such categories in Meetup groups and events can help in designing tools and techniques which can help software practitioners make better use of knowledge shared in such events and groups. We also looked at the gender distribution in Meetup events and how the onset of COVID-19 has impacted these groups. Such information can be valuable to Meetup organizers to increase participation in their events, irrespective of gender as well as logistical constraints imposed by COVID-19. Our work is different from previous work of Ingram et al. [5] as our focus is to uncover the categories of events organized by Meetup groups and also to see the gender diversity in such groups. Also, the data we have collected is not limited to a particular geography, which was the case with Ingram et al. [5].

First, we found 6,327 Meetup groups whose associated topics (assigned by group organizers) are related to software

[1]https://www.meetup.com/

development. Then, from these groups we extracted a candidate data set of 250,369 events. From this candidate data set, we took a random sample of 100 events, which we analyzed using the open coding methodology [6] in order to develop categories of events related to software development. Then based on these categories we further labeled 400 more randomly sampled events using the methodology used in [7], [8]. In the end we had 452 events labeled into some categories. The final labeled data sample is less than 500 as we dropped some events which labelers found hard to label. The sample of 452 events constitutes a sample with 95% confidence level and 5% error margin. Based on the data collected, we investigate 4 research questions. 1) What are the categories of events organized by Meetup groups related to software development? 2) How popular is each event category? 3) How diverse are Meetup groups with respect to gender? 4) How has COVID-19 impacted Meetup Events? By answering above research questions the major contributions our makes are as follows:

1) We performed an open coding procedure on 100 events and subsequent manual data labeling on 400 more events, to group them into categories, and to find the popular categories of events. We were able to find 9 categories of events (8 main categories +"Others") after performing the open coding procedure. We also found that *Talks by Domain Experts, Hands-on Sessions,* and *Open Discussions* are the most popular categories of events organized by Meetup groups related to software development. These categories indicate the importance software practitioners associate to learning from their peers.
2) We did a diversity analysis of members of Meetup groups and participants of Meetup groups and found that percentage of female members (20.46% for groups and 20.34% for events) is higher as compared to previous studies (such as 3-9% on GitHub [9], [10] and 7% on Stack Overflow [11]) on open source and software communities, especially when the organizers are also female. This serves an insight for group and event organizers on how to attract more female participation.
3) We found that the restrictions due to COVID-19 have reduced the event activity in Meetups, but a lot of events have now moved to virtual setting.

The remainder of this paper is structured as follows. In Section II we give a background of how people join and create groups using Meetup. We list and describe related work in Section III. In Section IV, we describe our research setting. In Section V, we describe our experimental results. We discuss the results in Section VI. Finally, we conclude and mention future work in Section VII.

II. Background

In this section, we give some background of how Meetup groups are formed and how their members organize events. The Meetup website[2], which was launched in 2002, works as an event scheduling and group organization tool in which members can seek, join, or create groups focusing on certain interests or activities, such as art, software development, or travel. It has been characterized as an event-based social network (EBSN) [3], which not only contains online social interactions but also valuable offline engagement and interaction among participants. Few other services similar to Meetup are MEETin[3] and Eventbrite[4]. In this work, we considered Meetup only as there has been previous evidence of it being used by software community to gain knowledge related to software development [2], [5]. In order to create or join a group on Meetup, a person first needs to register as a member on the Meetup website. The members may provide some keywords which represent their topics of interest and also their current location. This helps the website to recommend local groups related to the topics which a member has expressed interested in. Each Meetup member can be a part of one or more groups, and may hold different positions in each group (e.g., organizer, co-organizer, assistant organizer, event organizer, basic member, etc.).

Members in addition to joining existing groups, can also create new groups based on the topics of their interest. During the creation of a group, the creator (also known as "organizer") of the group is prompted to specify the group's location ("hometown") as well as one or more topics associated with the group. The group can be associated with one of the pre-defined Meetup categories, such as "Arts", "Language & Culture", "Tech", etc. In addition to such categories, organizer can also associate fine grained topics related to groups using keywords such as "python", "software development", "machine learning", etc. A full list of topics can be found here[5]. This information also enables Meetup to generate a recommendation of groups that may match the member's interests within a certain distance from his/her location. Meetup groups span a wide variety of interests, and there is also a wide range of group sizes, from few members to tens of thousands of members. Unlike in some other social networks, in which groups are formed based on members' shared interest regardless physical location, in Meetup, there is focus on the creation of groups in which members meet in physical space to participate in group interactions.

Within a Meetup group, the group organizer and the rest of the group's leadership team can plan one-off or recurring events such as a weekly group discussions, workshops, or talks by experts. The events can require participation fees or can be complimentary. The actual content, schedule, and fee of each event is set by its organizer, and each group can have arbitrary combination of one-off and regular events of various types. The event organizers when creating an event provide a one line description as event heading, as well as a detailed description describing details such as who will be speaking at the event, general theme of the event, etc. Group members can

[2]http://www.meetup.com/

[3]https://www.meetin.org/

[4]https://www.eventbrite.com/

[5]https://www.meetup.com/topics/

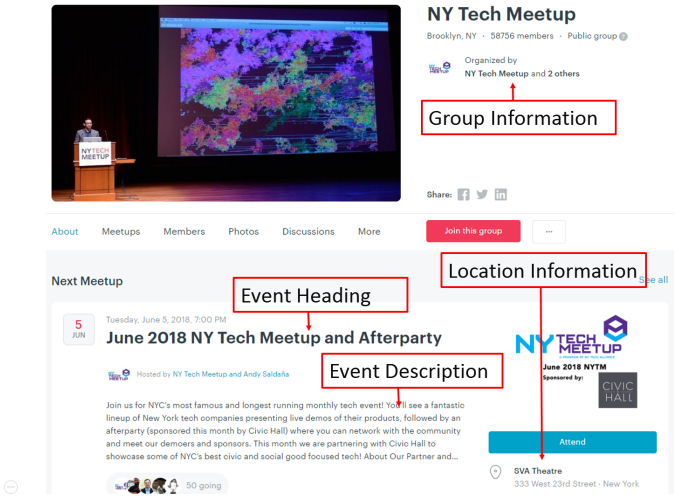

Fig. 1. A Sample Meetup Group Related to Software Development

opt to join any event, provided the event's registration deadline has not passed and the event has not hit its attendee limit.

Figure 1 shows a sample Meetup group (https://www.meetup.com/ny-tech/) and some upcoming events. The top of the page shows *group information* such as group name, group location and the number of members in the group. The *event heading* part gives a very brief overview of the event which is then followed by *event description*. The *event description* describes in detail the agenda of the event. Also, the *location* and *scheduling* information is present on the right side of the page.

III. Related Work

Event Based Social Networks: Sander and Seminar personally attended about 40 social events in one of the first studies conducted on Meetup [12]. Shen et al. demonstrated that Meetup events contribute to the creation of social capital [13]. Liu et al. introduced the term event-based social network (EBSN) to categorize online services such a Meetup, Eventbrite etc. in [14]. In their work they investigated the heterogeneous nature and diffusion patterns inherent across EBSNs. A lot of works have focused on improved event recommendation techniques for recommending events and/or friends in EBSNs. Some location based recommendation approaches have been proposed in [15], [16]. Some other techniques based on context and graph features have been proposed in [17], [18]

The effect of offline gatherings such as Meetup on community participation and social capital contribution has been investigated in [19], [13], [20]. The factors which determine the success of a Meetup group have been analyzed in [21], [22]. Pramanik et al. proposed an algorithm that can predict Meetup group success [23]. In a recent work by Ingram et al., the authors interviewed members of some software engineering Meetup groups in the United Kingdom and found out that the main motivations for people to participate in Meetup events is to stay up to date, learn and develop new skills, and build a local network [5].

Social Media and Software Engineering: As emphasized by Storey et al., social media has revolutionized the way software development is done [1]. In another work, Storey et al. investigated how usages of various social and communication channels affect software development, and found that Meetup is one of the channels used by software developers [2]. The role of social networking in software development was also examined by several other prior works [24], [25]. There have been many works which have analyzed individual sites or channels; we describe some of them below.

Previous works have shown that socially enabled digital channels such Stack Overflow (Q&A forum), GitHub (code hosting) [1], are a rich source of knowledge [26], [27] and used by developers for day to day problem solving and collaboration [28], [29]. For remaining up to date and gaining new knowledge developers primarily use Twitter [30], [31], [32], [33], [34], [35]. Recently developers' usage of channels such as *Youtube* [36], [37], *slack* [38], [39], *Reddit* and *HackerNews* [7] has also been explored. The gender representation in social channels has also been researched in many studies [11], [9], [40].

Our work is different from previous works as we focus on characterizing events in Meetup generated by groups that are related to software development that has not been looked into previous works related to software engineering or EBSNs. It adds to the body of knowledge related to social media channels used by software developers.

IV. Research Setting

In this section, we present our research setting. The overall process that we follow in our empirical study is illustrated in Figure 2. First, we extracted the data of any group related to software development using Meetup API and heuristics leveraging categories and topics in Meetup and tags in Stack Overflow. From the extracted data we took a random sample of the events organized by them, and then analyzed the same using open card sort [41], [42] and subsequent manual labeling [7], [8]. Then we answered a few research questions based on the empirical analysis of the extracted as well as coded data.

A. Research Questions

1) RQ1. What are the categories of events organized by Meetup groups related to software development?: Meetup events are known to help developers in keeping up-to-date [2], [5]. However, to the best of our knowledge no study has yet explored what are the kinds of events held in Meetup groups related to software development. In this research question we have employed manual qualitative analysis to develop some categories of events organized by Meetup groups related to software development. Finding such categories gives some insights into the types of knowledge that is discussed in Meetup groups.

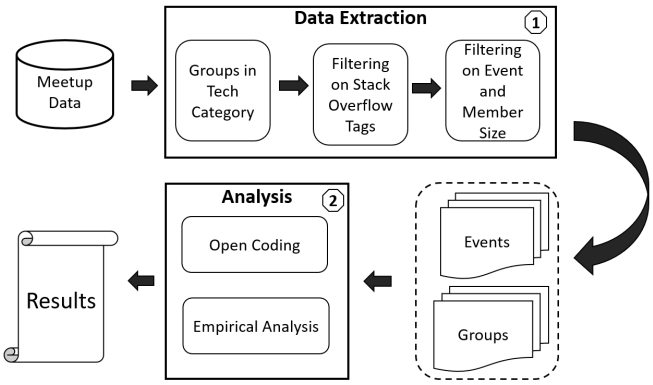

Fig. 2. Overall Process

2) RQ2. How popular is each event category?: We also analyzed the popularity of events based on the categories developed in RQ1. This gives an insight into what categories are more popular among the software development community on Meetup. Building upon the initial insights gathered, in future a detailed study can be done to explore what factors determine the popularity of a Meetup group. This can help Meetup organizers to effectively organize events and make their groups popular.

3) RQ3. How diverse are Meetup groups with respect to gender?: By analyzing the gender diversity of Meetup groups related to software development, we can observe if the results match to diversity ratios in other social media channels. Based on the insights found, future studies can focus on identifying the reasons for difference in diversity if found.

4) RQ4: How has COVID-19 impacted Meetup Events?: In this research question we investigated how COVID-19 situation has affected the event activity in Meetup groups.

B. Data Extraction

For our study we considered the data of groups on Meetup which are categorized under *Tech*[6] category. This category is assigned to a Meetup group by its organizers at the time of group creation. We made use of an open source python client[7] based on Meetup's RESTful API[8] to collect the data in our work, using which we were able to extract 58,460 groups categorized under *Tech* category. The data was extracted during the time period 16.04.2018-19.04.2018 and 58,460 was the absolute total number of groups present in the *Tech* category at that time. The scripts used to download data are present at this link[9].

Next, in order to improve the quality of our data set we applied a heuristic based filtering based on Stack Overflow *tags*[10]. These *tags* are used to describe the topics of questions asked on Stack Overflow, and each tag generally represents a software engineering concept. We were able to get a list of 51,670 *tags* from Stack Overflow archive[11] on 21.04.2018. This list is referred to as *SOTagList* further in the paper unless stated otherwise. Most of the *tags* present in the *SOTagList* are single words, e.g., python, javascript, java. Some tags are composed of multiple words connected by a hyphen (as tags cannot contain space)[12], e.g., visual-studio, apache-spark, ruby-on-rails. In order to filter groups based on *SOTagList* we converted all the *topics* associated with each of the 58,460 groups, as well as *tags* present in *SOTagList* to lower case. Then, we considered only those groups for further processing for which

- *SOTagList* contained any word appearing as a *topic* of a group, e.g., *topic* "python" was present in *SOTagList*
- *SOTagList* contained the hyphen separated form of words associated as a *topic* of a group, e.g., *topic* "big data" when converted to "big-data was present in *SOTagList*
- *SOTagList* contained any word present in the word sequence associated as a *topic* of a group, e.g., "database" in *topic* "database professionals" is present in *SOTagList*

After applying the above heuristic we were left with 56,175 groups. As many of the groups may be very small, we further filtered them by considering only those groups which have at least N members and have organized at least M events. We chose the value of N and M to be 10 for our study in order to focus on groups which have substantial number of events and members. We also excluded the groups whose event and member or data was not publicly visible. After applying this level of filtering, we were left with 17,727 groups. On observing these group descriptions we found out that some of them are not related to software development such as *The Vancouver Blogger Meetup Group*. To address this issue, we further applied one more level of filtering to keep only those groups which also contained a "Software Development" tag among their topics. We were left with 6,327 groups after applying this final level of filtering.

From these 6,327 groups we were able to extract 320,807 events in total. There were some events which repeat after certain intervals of times. As the instances of such periodic events have same description and theme, we remove all but one instance. After removing the repeated occurrences of such events we got a total of 213,477 events. Then we made use of as python package *langdetetct*[13] to keep only those events whose description language contained at least some *English* text. In the end, we were left with 185,758 events. We then manually analyzed and labeled 452[14] events from these 185,758 events into categories which were developed through open coding and subsequent manual labeling. The coding and labeling process will be discussed in detail in Section IV-C. We also extracted member information from 6,308 groups. For

[6]https://www.meetup.com/topics/

[7]http://meetup-api.readthedocs.io/en/latest/index.html

[8]https://www.meetup.com/meetup_api/

[9]https://anonymous.4open.science/r/MeetupDataCollection-9872/readme.md

[10]https://stackoverflow.com/tags

[11]https://archive.org/details/stackexchange

[12]https://stackoverflow.com/help/tagging

[13]https://pypi.python.org/pypi/langdetect/

[14]This corresponds to a 95% confidence level with 5% error margin

19 groups the member extraction failed as there was no group organizer at the time when the query for member information was made. We were able to extract information of 3,123,498 unique members. Among these there were 2,610 members who were not active. We discarded such members and in the end we were left with information of 3,120,888 members. As the original data was extracted in 2018, in order to get some insights into how things might be different for next 2 years for these groups, we extracted the recent data (in May, 2020) for events and member data for the 6,327 groups originally identified. In August, 2019, some restrictions[15] were added on how Meetup's RESTful API[16] could be accessed. Due to these restrictions we could not download the data for all the 58,460 groups. For the data we downloaded in 2020, we just concatenated it with the data collected in 2018, and in the end we had data of 250,369 events and 4,481,670 members which we have analyzed in this work.

C. Content Analysis

In order to find the categories of events organized by Meetup groups related to software development we used the open coding procedure [6]. We performed the card sort procedure on the sample data of 100 events sampled from the larger dataset described earlier in Section IV-B. The technique used in our work is similar to what has been used in various previous studies such as [43], [44], [31], [45]. We first generated a card based on the description of each event. Each card contained the event description, event id, event date, event group, and event's URL link. Then, each card was read and the event description along with other details was discussed and iteratively sorted into categories or groups. In the first iteration, a code was assigned to an event description, and in all the subsequent iterations the codes assigned in previous iterations were analyzed to create higher level concepts or categories. For some event descriptions we were not able to merge them with any other categories, so we merged into a special category *Others*. In the end, we were able to generate 9 categories including *Others*. The first and the third author of the paper together performed the open card sort process.

TABLE I
PROGRESSION OF AGREEMENTS WHILE LABELING

Iteration	Absolute Agreement	Cohen's Kappa	Interpretation
1	0.800	0.700	Substantial
2	0.800	0.705	Substantial
3	0.800	0.738	Substantial
4	0.867	0.832	Almost Perfect

In order to increase our sample size, we further sampled 400 more events. During open coding earlier we had already come up with a coding schema. These 400 events were then coded by the first and second author based on the categories developed earlier. The first author discussed the schema with the second author to arrive at a common understanding and clear any confusions. Then both authors separately labeled 30 events and then met together to discuss and further refine the coding schema if required. The authors continued the iterative process of independently coding and discussing afterwards until they were consistent in labeling. After 4 iterations were completed both the authors were already achieving substantial to perfect agreement, reaching a Cohen's Kappa [46] agreement score greater than 0.7 on all the iterations. After this iteration the rest of the data was split into two sets which were independently coded by first and second authors. The process is similar to what has been followed in [7], [8].

During iterative discussions we came across some event descriptions where very little information was available on what the type of event is. Also, for some events the descriptions were primarily in language other than English (as our filtering criteria was to include any event whose descriptions which had at least some English text) which made it hard to determine the type of events. For some other descriptions it was hard for labelers to assign a single category to the event based on the event description. For all the aforementioned 3 cases the labelers could assign them to an "Unsure" set. After both coders finished the labeling we had in total 52 events which were put in the "Unsure" set. We dropped such events from our final dataset. The final data set is of size 452 events (and not 448 as we had to split 4 events whose descriptions contained two types of events). Note that "Unsure" set is different from the "Others" category where we were able to identify what the event is about but not able to merge it with other categories. We have publicly released our replication package at [47].

D. Data Characteristics

Table II shows some descriptive statistics related to the data used in our work. For *Groups*, the average number of members per group is 1,458, however the median number of members is lower at 834. There are some groups with very high membership count such as *Big Data & AI Introduction* (https://www.meetup.com/CloudxLab/) with 27,217 members. We also looked at the mean and median ratings assigned to the *groups* and there seems to be less deviation in the ratings assigned to groups, with mean and median being close to each other, 4.54 and 4.75 respectively.

For *Events*, looking at their time duration we found that the mean average duration of each event is 4 hours and 48 minutes, whereas the median duration is quite low at 2 hours. There are some events which last over few days, such as a 14 day event related to *Ruby Workshop* (https://www.meetup.com/Girl-Develop-It-Ann-Arbor/events/223889838/) by the group *Girl Develop It Ann Arbor*. For events, the ratings seem to be quite divergent with a mean rating score of 2.02 and a median rating score of 0. Table II also shows a summary of topics associated with groups and members. The mean number of topics associated with groups are 18 per group, which is slightly higher than the 14 topics per member. However, when considering the median scores, most groups only have 24 topics associated with them as compared to 14 topics

[15]https://help.meetup.com/hc/en-us/articles/360028901812-Using-Meetup-s-API

[16]https://www.meetup.com/meetup_api/

associated with members. (The topic data for members is based on data collected in 2018.)

V. Results

In this section, we present the results of our analysis conducted in this study.

A. RQ1. What are the categories of events organized by Meetup groups related to software development?

In this research question we explored if the events organized by Meetup groups related to software development can be grouped into some meaningful categories. We were able to determine 9 categories of events (8 main categories +"Others") using the open coding methodology [6] described earlier in Section IV-C. The categories found are described below along with a relevant event as an example.

Talks by Domain Experts: 200 of the event descriptions analyzed were related to an event where talks or presentations were given by some domain experts. The domain experts consisted of mostly developers, technical managers, entrepreneurs and CEOs. Most of the talks are technical in nature where an experienced developer explains or introduces a core software concept. Sometimes multiple experts came together and participated in panel discussions. In other talks technology management principles such as SCRUM were presented. Other talks especially by entrepreneurs and CEOs focused on a technical product or feature being developed by their companies. An example of an event in this category is shown below:

> *Meetup group* : Boston-Predictive-Analytics
> *Description* : ... Rani Nelken of Outbrain (http://www.outbrain.com/) has graciously offered to present on Bayesian Classification....
> *Event URL:* https://www.meetup.com/Boston-Predictive-Analytics/events/60294452/

Hands-on Sessions: This category contains events where a domain expert does not only give a talk or presentation but is also involved in actively guiding and helping other participants to perform some hands-on tasks or activities related to a topic of presentation. Such kind of events were often marked with a request for the participants to bring their own laptops so as they can practice the exercises that follow a talk. The sessions organized in such events can range from introductory to advanced. Often such events required the participants to pay a fee. An example of such an event is mentioned below:

> *Meetup group* : WaikatoLinuxUsersGroup
> *Description* : ... This is a GNU/Linux-focused workshop where people can bring their PCs, Laptops, Pi's, Android devices etc for trouble-shooting and to learn or try out new skills ...
> *Event URL* : https://www.meetup.com/WaikatoLinuxUsersGroup/events/196352072/

Conferences: The events in this category were created mostly to notify the group members of any upcoming conferences. Different from talks and hands-on sessions, conferences were longer and bigger events that often span multiple days, with participation by many speakers, and covered a more diverse range of topics. Sometimes it also involved call for proposal, participation, or volunteering. An example event is shown below:

> *Meetup group* : jsmeetup
> *Description* : Call for Speakers HTML5DevConf continues to grow as the largest JavaScript and HTML5 conference ...
> *Event URL* : https://www.meetup.com/jsmeetup/events/98573142/

Open Discussions: The events in this category did not have any predefined agenda and no speakers were scheduled to speak in advance. Most of these were open house sessions where any of the participants could speak on any topic loosely related to the topics associated with groups. It included events such as round table discussions, impromptu experience sharing sessions, study groups, code jams, etc. An event categorized into this category is shown below:

> *Meetup group* : london-software-craftsmanship
> *Description* : ... Do you want to discuss an approach pattern or technology and see what others think? Or perhaps discuss a design challenge youre facing? Come along to the Software Craftsmanship round table ...
> *Event URL* : https://www.meetup.com/london-software-craftsmanship/events/16117553/

Social Events: This category includes core networking events where participants were invited to social dining and/or drinking sessions where they could interact with other invitees. Such events include kick-off parties, award ceremonies, etc. An example event is shown below:

> *Meetup group* : Windy-City-Tech-Meetup
> *Description* : Talk with others in Tech, Big Data, Business Intelligence, Open Analytics, etc., and meet new business contacts at River North's Trophy Room on December 28th, while sipping sponsored cocktails and beer.
> *Event URL* : https://www.meetup.com/Windy-City-Tech-Meetup/events/235968811/

Competitions: In the events under this category the par-

TABLE II
DESCRIPTIVE STATISTICS OF OUR DATASET

		Count	Mean	Std Dev	Min	Q1	Median	Q3	Max
Groups	Members	6,327	1,458.82	1,978.15	12	393.50	834	1,765	27,217
	Rating (range 0-5)	6,327	4.54	0.95	0.00	4.61	4.75	4.88	5.00
	Topics	6,327	21.88	7.27	3	15	24	28	34
Events	Duration (HH:MM:SS)	165,414	4:47:56	14:20:20	0:01:00	02:00:00	2:00:00	3:00:00	14 days
	Rating	250,367	2.02	2.35	-1	0	0	4.86	5
Members	Topics	2,370,094	18	15	1	6	14	28	67

ticipants usually formed teams, and then competed with one other on certain technical tasks. The most common type of event under this category was *Hackathon* where teams had to come up with a usable software product in some days or hours. Sometimes the events were held specifically by a technology product company where they offer rewards to the participants in order to find bugs in their products. An example event is shown below:

> *Meetup group* : San-Francisco-Hackathons
> *Description* : ... The Hackathon will consist of 8 teams with up to 6 members. You may register as an individual or bring an entire team. If you register as an individual, we will find a team for you ...
> *Event URL* : https://www.meetup.com/San-Francisco-Hackathons/events/197575472/

Administrative Events: These events were primarily organized to discuss among the group organizers and volunteers the roles and responsibilities each group member would take. Sometimes other organizational aspects such as what kind of events to organize in future were also discussed in such events. The event shown below is one of the events that has been assigned into this category:

> *Meetup group* : Evansville-Technology-Group
> *Description* : ... Come join a round table discussion on technology in Evansville and help us plan our meetups for the year. We need your input to make sure we are providing the topics and events that everyone is interested in ....
> *Event URL* : https://www.meetup.com/Evansville-Technology-Group/events/236824527/

Job Fairs: These are the events which bring together recruiters and job seekers; the focus being on software related jobs. The following event is an example:

> *Meetup group* : Girl-Develop-It-Boulder-Denver
> *Description* : .... The Tech Jobs Tour is coming to Denver and looking to connect with techies and community-focused individuals in the city ...
> *Event URL* : https://www.meetup.com/Girl-Develop-It-Boulder-Denver/events/244491885/

Others: Few events in our sample could not be categorized into any of the 8 above-mentioned categories. Since the remaining events are different from one another, we put all of them in a broad category *Others*. Some of the events moved into *Others* were study tours, marketing events, etc.

B. RQ2: How popular is each event category?

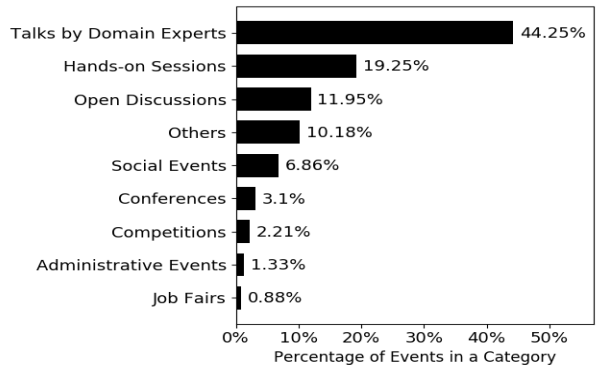

Fig. 3. Popularity of Categories in Sampled Event Data

In this research question we analyzed the popularity of each event category in our labeled data. We define popularity based on how many events are included in a category. To do this we count the number of events categorized into each category and then plot a bar graph of percentage of events occurring in each category. Figure 3 shows a bar graph showing the popularity of each category. The percentage calculation was done on 452 total events. We can observe from the graph that *Talks by Domain Experts* is the most frequent category of events organized by Meetup groups, followed by events related to *Hands-on Sessions* and *Open Discussions*.

We also did a popularity analysis based on the number of people interested in an event. For most events, a field called *yes_rsvp_count* is present, which specifies the total number of people who confirmed participation for an event. We used *yes_rsvp_count* as a proxy for estimating how many people are interested in the event. There is another field known as *rsvp_limit* which specifies the total number of people allowed for the event. For each event, we calculate a metric called $Event\_Attention$ by dividing the value of *yes_rsvp_count* by value of *rsvp_limit* and then averaging it over all the events

for the respective category. We only considered those events for which values of both these fields were present in our dataset. We also ignored events which had the *rsvp_limit* value specified as 1, as it leads to the $Event\_Attention$ value being greater than 1. We were able to find 152 events spread across all 9 categories after removing the events as described above.

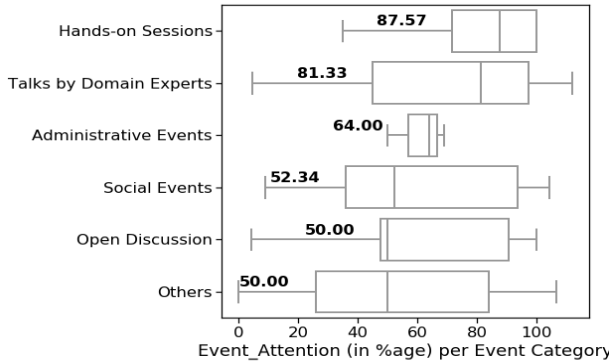

Fig. 4. Popularity of Event Categories based on Interest

Figure 4 shows a graph containing the boxplot of $Event\_Attention$ (in percentage) for 6 categories of events, with median values shown inline. The category *Hands-on Sessions* is the most popular category with the highest median *Event_Attention* of 87.57%. In *Hands-on Sessions*, the participants are generally required to bring their own laptops so that they can work on various tasks and exercises discussed in the event and also ask other people for help. The more interactive nature of such events seems to result in high participation in such events. The percentage participation seems to be less in the categories of *Social Events and Open Discussions*. We have ignored the categories *Competitions, Conference, & Job Fairs* in Figure 4 as only 1 event per category was present in the 152 events described in the last paragraph.

Ingram et al. had conducted surveys with developers and found that the most popular reason why software practitioners join Meetup groups and events is to *learn and develop new skills*. The observations shown in above graph validate the findings of the previous study, as we can clearly see that *Talks by Domain Experts* and *Hands-on Sessions* are the most popular categories of events.

C. RQ3: How diverse are Meetup groups and events with respect to gender?

We also analyzed how different genders are represented among various groups. As discussed in Section IV-B we had extracted information of 4,481,670 members. For these 4,481,670 members, we tried to resolve their gender using the name and country information. For the purpose of resolution, we used the name and country information of members as input to *genderComputer*[17]. The tool was validated in a survey in [11], and the tool has also been used in previous gender related studies [9], [48]. We were able to determine gender of 3,724,973 members which constitutes a fraction of 80.62%. For other cases, either the gender determination failed or it could not be determined if the person is male or female. Based on this gender determination process we were able to find at least 1 female or 1 male member for 6,286 groups. For other groups, we could not successfully resolve gender for any of their members. We then computed the percentage of male and female members based on the total members count present in group which was extracted earlier as described in Section IV-B. The total members count includes members whose gender could not be determined.

[17]https://github.com/tue-mdse/genderComputer

TABLE III
GENDER DISTRIBUTION OVER 6,286 GROUPS

	Mean	Median	Max	Deviation
Male	841.61	483	16,917	1141.80
Male %	57.61	60.12	92	14.58
Female	333.41	157	11,042	569.50
Female %	20.46	18.29	92.08	11.60
Total	1465.07	839.50	27,217	1,982.51

Table III shows some descriptive statistics with respect to gender distribution across the groups. We can see that 20.46% of the members on an average are female (18.29% if median values are considered). This numbers are larger as compared to results reported in previous studies, 3-9% on GitHub [9], [10] and 7% on Stack Overflow [11]. This insight shows that females may be more comfortable to participate in events in Meetup.

For each of the 6,286 groups there was 1 organizer associated per group, so we also did an analysis of the gender diversity among organizers. For 5,027 groups we were able to determine the gender of the organizer of the group and found that for 896 groups (i.e, 17.82%) the organizers were female. Among these 5,027 groups we found that for the groups where the organizer was female, the mean of *female percentage* for such groups is 28.90%, as compared to 18.33% for the groups where the organizer was male. To validate that *female percentage* for groups with female organizers is indeed significantly different from the groups with male organizers, we performed the Mann-Whitney U test [49] on distributions of *female percentage* for both cases. The test gave a p-value which is less than 0.05, and thus we can say that the distributions are significantly different. We also computed the Cliffs Delta [50] statistic for the two distributions and found the delta value to be 0.38 (medium). This suggests that if the organizer of a Meetup group is female, they tend to attract more females to the group.

In addition to looking at gender distribution for groups, we also looked at gender distribution of members who participated in events. We were able find 228,431 events, where after gender determination at least 1 female or 1 male participant was found. We then computed the percentage of female and male participants based on the total members who participated in that event (including members whose gender could not be determined). Table IV shows the descriptive statistics with

TABLE IV
GENDER DISTRIBUTION OVER 228,431 EVENTS

	Mean	Median	Max	Deviation
Male	24.50	13	3,201	35.27
Male %	63.09	68.18	100	23.24
Female	7.31	3	407	12.16
Female %	20.34	14.58	100	20.78
Total Participants	37.36	21	4,382	50.27

respect to gender distribution across these events. We can see that in case of events also 20.34% of participants are female. Here also we checked the effect of gender of the organizer of the group (with which the event was associated) on the percentage of females who participated in events. The Mann-Whitney U test [49] on distributions of *female percentage* participation in events factored by the gender of group organizer, gives a p-value which is less than 0.05. Also the Cliffs Delta [50] statistic for the two distributions in case of event participation gives a value of 0.52 (large). Both these observations suggests that if the organizer of a Meetup group is female, the percentage of females that participated in events of such groups tends to be higher.

D. RQ4: How has COVID'19 impacted Meetup Events?

On 31 January 2020, World Health Organization had declared coronavirus a world health emergency [51], and many countries initiated lockdowns in order to curb the spread of the disease [52]. Such steps have forced a large part of the world population to stay and work from home, including software developers [53]. As discussed earlier Meetup events were primarily organized at a physical location. In this research question we explore how the steps related to curb the coronavirus pandemic have impacted Meetup events.

For evaluating the impact of COVID-19, we considered the events in the months of February, March, and April for a time period of past 5 years. We computed some metrics for the year 2020, and compared it with the average of same metrics averaged over years 2016-2019. We were able to find 45,399 events for these 3 months over the 5 year time period of 2016-2020. The first metric we looked at is *Events Per Day*. From Table V we can see that *Events Per Day* has dropped to 52.48 events per day as compare to 113.94 events per day on average over past 4 years. On performing the Mann-Whitney U test on the *Events Per Day* values for year 2020, and values before 2020, we obtained a p-value less than 0.05 indicating that the drop is statistically significant. The Cliff's delta for the distributions of this metric also gives a value of 0.74 (large).

The next metric we looked at is *Event_Attention* which has been described earlier in Section V-B, which has been computed based on 13,721 events spread across 5 years for the 3 months mentioned earlier. From Table V we can see that the value of this metric has also gone down to 55.49% as compared to 70.36% over years 2016-2019. The drop is statistically significant for *Event_ Attention* metric also, with a p-value less than 0.05 on the the Mann-Whitney U test, and the Cliff's delta value being 0.60 (large).

The above observations suggest that the event activity in Meetup has dropped significantly, but has not completely stalled. In order to see how the events are being organized the first author looked at descriptions of some of the events. There were a total of 4,723 events for the time period 01.02.2020-30.04.2020. By doing a simple text search we found that the description of 2,574 of 4,723 events contains one of the following words: *virtual, online, and video*. We also looked at the top domains present in the URL links that are mentioned in the event description and found domains such as *zoom, altvr, crowdcast, and discord* (tools for online collaboration) constitute 12% of all links mentioned in the description of events in 2020. All these domains were not found in any of the links present in event descriptions of years 2016-2019 for the 3 month period. These results suggest that a lot of events are now being organized virtually.

We also looked if there has been affect of COVID-19 on the percentage of of male and female participants in events being organized. We again only looked at the data of months of February, March, and April for a 5 year period from 2016-2020. From Table V we can see that the percentage of male participants has dropped slightly to 58.07% for events held in year 2020 (with a Cliff's delta value being 0.30 (small)). Interestingly, the number of female participants for events of year 2020 has increased to 24.02% as compared to 21.51% for previous years (Cliff's delta value being -0.35 (medium)). As the effect size is not large more fine grained data may be required to see if the changes in percentages are due to change in way the events are organized during COVID-19 (e.g., online events).

E. Threats to Validity

Threats to *internal validity* relate to errors that may have occurred during experiments and labeling of data. We checked our code multiple times, but there still may be errors that we may have missed out. The labeling process involved 3 persons, 2 PhD students and 1 research engineer having more than 20 years of professional software development experience spread among them. Also filtering the Meetup groups where *English* is the primary language of communication, may limit the results of our study to cultures where *English* is the primary language. Also the result presented in this work (except for RQ3) are valid only for 6,327 groups which were selected in 2018.

We also computed the inter-rater agreement for the labeling task using the measure of Cohen's Kappa [46]. Threats to *external validity* relate to how generalizable our findings are. We have tried to mitigate this threat by randomly sampling events. Also as seen in Table II we can see that our event dataset spans across various topics. Another threat relates to use of names for gender resolution, as users may be using aliases (and not their correct names) on Meetup groups. To address this we use *genderComputer* which has shown to

TABLE V
EFFECT OF COVID-19 ON EVENT METRICS (FOR MONTHS OF FEBRUARY, MARCH, AND APRIL)

	Events Considered	2020	Before 2020	Mann Whitney (p<0.05)	Cliff's Delta
Events Per Day	45,399	52.48 ± 23	113.94 ± 46	True	0.74 (large)
Event_Attention (in %)	13,721	55.49 ± 16	70.36 ± 20	True	0.60 (large)
Female Participation (in %)	42,322	24.02 ± 5	21.51 ± 3	True	-0.35 (medium)
Male Participation (in %)	42,322	58.07 ± 7	60.61 ± 4	True	0.30 (small)

achieve high precision on the gender determination task [11], [9] and has been validated earlier for accuracy [11].

VI. DISCUSSION

From our findings presented in Figure 3, we can clearly see that the most popular category of events is *Talks by Domain Experts*, followed by *Hands-on Sessions* and *Open Discussions*. Also gender analysis on group members shows that female representation in the Meetup groups is higher as compared to other collaborative sites frequently used by developers.

A. Implication for Researchers

We found that female representation in the Meetup groups being 20.46% when membership is considered and 20.34% when event participation is considered. These numbers are larger than what has been reported for other social channels, 3-9% on Github [9], [10] and 7% on Stack Overflow [11]. Females tend to contribute less to open source projects if the technical barrier is too high [54]. So further research is required to validate if the female participation is higher only for Meetup groups which have a low technical barrier or if it is standard across all the Meetup groups. We also found empirical evidence that average female participation in Meetup groups and events is higher if the group organizer is female. If the organizer of a group is female then the social barrier of participation decreases which has been cited as a reason for increased participation [55]. However, further research is required to understand if there are any other factors which may be contributing to higher participation, and if those factors can be replicated elsewhere to improve female participation.

We found that *Talks by Domain Experts* is the most popular category and it highlights the importance software developers give to continual learning. Also from Figure 4, it can be observed that the interest in the events related to category *Hands-on Sessions* is very high. One reason for this phenomena is that in *Hands-on Sessions* participants actually work on exercises and modules, and thus have better understanding of the topic that is being discussed in the event. The popularity of category *Hands-on Sessions* provides additional evidence to previous research on Meetup which found that software practitioners use Meetup groups primarily for staying up to date, learning new things, and improving their tacit knowledge by learning from peers [5]. Further research can be conducted on the data related to organization and participation information of popular event categories, which can then be used to understand what helps developers in their knowledge seeking experience, and if particular types of events have participants from some particular categories only (beginner, experienced, etc.). Such insights can be used by universities and other organizations involved in education (software engineering or otherwise), to understand the knowledge exchange mechanisms in Meetup groups, which can then be used to improve the effectiveness of their programs and courses.

Also a preliminary analysis of events during 2020 shows that that the number of events and *Event_Attention* (i.e. *yes_rsvp_count* of an event divided *rsvp_limit* of the event) have substantially dropped, and moved to online setting, which seem to be result of COVID-19. Further studies can look more deeply into challenges of organizing events in the COVID-19 time, and best practices to mitigate such challenges. These are needed to help event organizers to create successful events in case of restrictions that came in because of COVID-19. Also it would be interesting to study the reasons for the bump in female participation for events held during COVID-19.

B. Implication for Practitioners

The empirical evidence that average female participation in groups is higher if the group organizer is female serves as a cue for various communities such as open source organizations, software development companies etc. to increase the proportion of females in leadership roles in order to encourage more female participation in their organizations. Practitioners can also aim to build tools such as [56] which can perform automated analysis and summarization of discussion and events in relevant Meetup groups.

VII. CONCLUSION AND FUTURE WORK

In this paper, we performed an empirical analysis of the events organized by Meetup groups related so software development. We first randomly sampled 452 events from a candidate set of 250,369 events organized by groups related to software development. We then did a qualitative analysis of these 452 events using open coding procedure and subsequent labeling. After we developed the categories we analyzed the popularity of the event categories, based on how often they are organized. We found that categories *Talks by Domain Experts, Hands-on Sessions, and Open Discussions* are the most popular. This shows that learning from in person interactions with other participants is a popular choice among software practitioners, and also validates the findings of Ingram et al. [5]. We also did a gender based diversity analysis on members of Meetup groups and found that 20.46% of members are female on average, which is a higher proportion as compared to female participation in other social channels related to software development [9], [10]. Finally, we found

that COVID-19 has impacted event activity, but many events are still being conducted in a virtual setting.

One promising direction for future work is to analyze in detail what makes an event or a group popular among software developers. This can be conducted following a the data science based approach followed in [21], [22], which may also be complemented by actually doing a survey with the members querying them if they prefer attending some events while skipping others. The insights gathered from such a work can help organizers better manage their Meetup groups and events. Also we plan to combine data from online social groups such as *Reddit*, *Twitter*, *Hackernews* etc. with the data from offline social groups from resources such as *Meetup*, *EventBrite* etc. for further detailed analysis. The combined data can be used to compare and contrast the strengths and weaknesses of such online and offline interactions and leverage these to aid and assists in software development tasks. We also plan to analyze if the topics of discussion differ in on-line and off-line communities. The identification of categories of events done in this work as well as the analysis of participant's gender diversity is a first step towards the accomplishment of the bigger goal of understanding the mechanisms by which software development communities thrive in offline and online settings and share information with one another.

REFERENCES

[1] M.-A. Storey, L. Singer, B. Cleary, F. Figueira Filho, and A. Zagalsky, "The (r) evolution of social media in software engineering," in *Proceedings of the on Future of Software Engineering*. ACM, 2014.

[2] M.-A. Storey, A. Zagalsky, F. Figueira Filho, L. Singer, and D. M. German, "How social and communication channels shape and challenge a participatory culture in software development," *TSE*, 2017.

[3] X. Liu, Q. He, Y. Tian, W.-C. Lee, J. McPherson, and J. Han, "Event-based social networks: linking the online and offline social worlds," in *KDD*. ACM, 2012.

[4] Meetup, "About meetup." 2020, [Online; accessed 21-May-2020].

[5] C. Ingram and A. Drachen, "How software practitioners use informal local meetups to share software engineering knowledge," in *ICSE*. ACM, 2020.

[6] J. Saldaña, *The coding manual for qualitative researchers*. Sage, 2015.

[7] M. Aniche, C. Treude, I. Steinmacher, I. Wiese, G. Pinto, M.-A. Storey, and M. A. Gerosa, "How modern news aggregators help development communities shape and share knowledge," in *ICSE*, 2018.

[8] G. Uddin and F. Khomh, "Mining api aspects in api reviews," Technical Report. 10 pages. http://swat. polymtl. ca/data/opinionvalue-technical-report. pdf, Tech. Rep., 2017.

[9] B. Vasilescu, D. Posnett, B. Ray, M. G. van den Brand, A. Serebrenik, P. Devanbu, and V. Filkov, "Gender and tenure diversity in github teams," in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. ACM, 2015, pp. 3789–3798.

[10] P. A. David and J. S. Shapiro, "Community-based production of open-source software: What do we know about the developers who participate?" *Information Economics and Policy*, vol. 20, no. 4, pp. 364–398, 2008.

[11] 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, 2013.

[12] T. Sander, "E-associations: using technology to connect citizens: the case of meetup. com." American Political Science Association, 2005.

[13] C. Shen and C. Cage, "Exodus to the real world? assessing the impact of offline meetups on community participation and social capital," *New Media & Society*, 2015.

[14] T. Liu, *Learning to Rank for Information Retrieval*. Springer, 2011.

[15] Y. Lu, Z. Qiao, C. Zhou, Y. Hu, and L. Guo, "Location-aware friend recommendation in event-based social networks: A bayesian latent factor approach," in *CIKM*. ACM, 2016.

[16] Z. Qiao, P. Zhang, Y. Cao, C. Zhou, L. Guo, and B. Fang, "Combining heterogenous social and geographical information for event recommendation." in *AAAI*, 2014.

[17] A. Q. Macedo, L. B. Marinho, and R. L. Santos, "Context-aware event recommendation in event-based social networks," in *RecSys*. ACM, 2015.

[18] T.-A. N. Pham, X. Li, G. Cong, and Z. Zhang, "A general graph-based model for recommendation in event-based social networks," in *ICDE*. IEEE, 2015.

[19] L. F. Sessions, "How offline gatherings affect online communities: when virtual community members meetup," *Information, Communication & Society*, 2010.

[20] D. N. Vaughn, "Meetup and social capital: Building community in the digital age," Ph.D. dissertation, University of Denver, 2015.

[21] C.-H. Lai, "Can our group survive? an investigation of the evolution of mixed-mode groups," *Journal of Computer-Mediated Communication*, 2014.

[22] S. Pramanik, M. Gundapuneni, S. Pathak, and B. Mitra, "Can i foresee the success of my meetup group?" in *Advances in Social Networks Analysis and Mining (ASONAM), 2016 IEEE/ACM International Conference on*. IEEE, 2016.

[23] ——, "Predicting group success in meetup." in *ICWSM*, 2016.

[24] A. Begel, J. Bosch, and M.-A. Storey, "Social networking meets software development: Perspectives from github, msdn, stack exchange, and topcoder," *IEEE Software*, 2013.

[25] M.-A. Storey, C. Treude, A. van Deursen, and L.-T. Cheng, "The impact of social media on software engineering practices and tools," in *FSE/SDP workshop*. ACM, 2010, pp. 359–364.

[26] A. Barua, S. W. Thomas, and A. E. Hassan, "What are developers talking about? an analysis of topics and trends in stack overflow," *ESEM*, 2014.

[27] B. Vasilescu, V. Filkov, and A. Serebrenik, "Stackoverflow and github: Associations between software development and crowdsourced knowledge," in *SocialCom*. IEEE, 2013.

[28] E. Kalliamvakou, D. Damian, K. Blincoe, L. Singer, and D. M. German, "Open source-style collaborative development practices in commercial projects using github," in *ICSE*, ser. ICSE, 2015.

[29] C. Treude, O. Barzilay, and M.-A. Storey, "How do programmers ask and answer questions on the web?(nier track)," in *Proceedings of the 33rd international conference on software engineering*, 2011, pp. 804–807.

[30] A. Sharma, Y. Tian, A. Sulistya, D. Lo, and A. F. Yamashita, "Harnessing twitter to support serendipitous learning of developers," in *SANER*. IEEE, 2017.

[31] A. Sharma, Y. Tian, and D. Lo, "What's hot in software engineering twitter space?" in *Software Maintenance and Evolution (ICSME), 2015 IEEE International Conference on*. IEEE, 2015.

[32] M.-A. Storey, L. Singer, B. Cleary, F. Figueira Filho, and A. Zagalsky, "The (r) evolution of social media in software engineering," in *FOSE 2014 Proceedings of the on Future of Software Engineering*, 2014.

[33] X. Wang, I. Kuzmickaja, K.-J. Stol, P. Abrahamsson, and B. Fitzgerald, "Microblogging in open source software development: The case of drupal and twitter," *Software, IEEE*, 2013.

[34] G. Bougie, J. Starke, M.-A. Storey, and D. M. German, "Towards understanding twitter use in software engineering: preliminary findings, ongoing challenges and future questions," in *Web2SE*, 2011.

[35] L. Singer, F. M. F. Filho, and M. D. Storey, "Software engineering at the speed of light: how developers stay current using twitter," in *ICSE 2014*, 2014.

[36] L. MacLeod, M.-A. Storey, and A. Bergen, "Code, camera, action: How software developers document and share program knowledge using youtube," in *ICPC*. IEEE, 2015.

[37] L. MacLeod, A. Bergen, and M.-A. Storey, "Documenting and sharing software knowledge using screencasts," *Empirical Software Engineering*, 2017.

[38] P. Chatterjee, K. Damevski, L. Pollock, V. Augustine, and N. A. Kraft, "Exploratory study of slack q&a chats as a mining source for software engineering tools," in *ICSE*. IEEE Press, 2019, pp. 490–501.

[39] B. Lin, A. Zagalsky, M.-A. Storey, and A. Serebrenik, "Why developers are slacking off: Understanding how software teams use slack," in *CSCW*. ACM, 2016, pp. 333–336.

[40] M. Burnett, A. Peters, C. Hill, and N. Elarief, "Finding gender-inclusiveness software issues with gendermag: a field investigation," in *CHI*, 2016, pp. 2586–2598.

[41] W. Hudson, "Card sorting," *The Encyclopedia of Human-Computer Interaction, 2nd Ed.*, 2013.
[42] D. Spencer, *Card sorting: Designing usable categories*. Rosenfeld Media, 2009.
[43] T. D. LaToza, G. Venolia, and R. DeLine, "Maintaining mental models: a study of developer work habits," in *ICSE*. ACM, 2006.
[44] D. Lo, N. Nagappan, and T. Zimmermann, "How practitioners perceive the relevance of software engineering research," in *FSE*. ACM, 2015.
[45] T. Zimmermann, "Card-sorting: From text to themes," in *Perspectives on Data Science for Software Engineering*. Elsevier, 2016.
[46] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
[47] Package, "https://github.com/abhishek9sharma/meetupdatacollection."
[48] G. Catolino, F. Palomba, D. A. Tamburri, A. Serebrenik, and F. Ferrucci, "Gender diversity and women in software teams: How do they affect community smells?" in *ICSE-SEIS*. IEEE, 2019, pp. 11–20.
[49] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The annals of mathematical statistics*, pp. 50–60, 1947.
[50] N. Cliff, "Dominance statistics: Ordinal analyses to answer ordinal questions." *Psychological bulletin*, vol. 114, no. 3, p. 494, 1993.
[51] WHO, "Statement on the second meeting of the international health regulations (2005) emergency committee regarding the outbreak of novel coronavirus (2019-ncov)," 2020, [Online; accessed 18-05-2020].
[52] ——, "A third of the global population is on coronavirus lockdown here's our constantly updated list of countries and restrictions," 2020, [Online; accessed 18-05-2020].
[53] P. Ralph, S. Baltes, G. Adisaputri, R. Torkar, V. Kovalenko, M. Kalinowski, N. Novielli, S. Yoo, X. Devroey, X. Tan *et al.*, "Pandemic programming: How covid-19 affects software developers and how their organizations can help," *arXiv preprint arXiv:2005.01127*, 2020.
[54] C. Mendez, H. S. Padala, Z. Steine-Hanson, C. Hildebrand, A. Horvath, C. Hill, L. Simpson, N. Patil, A. Sarma, and M. Burnett, "Open source barriers to entry, revisited: A sociotechnical perspective," in *ICSE*. IEEE, 2018, pp. 1004–1015.
[55] A. Lee and J. C. Carver, "Floss participants' perceptions about gender and inclusiveness: a survey," in *ICSE*. IEEE Press, 2019, pp. 677–687.
[56] G. Uddin and F. Khomh, "Automatic summarization of api reviews," in *ASE*. IEEE Press, 2017.