

# Recognizing Gender of Stack Overflow Users

Bin Lin  
Eindhoven University of Technology  
The Netherlands  
b.lin@student.tue.nl

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

## ABSTRACT

Software development remains a predominantly male activity, despite coordinated efforts from research, industry, and policy makers. This gender imbalance is most visible in social programming, on platforms such as **Stack Overflow**.

To better understand the reasons behind this disparity, and offer support for (corrective) decision making, we and others have been engaged in large-scale empirical studies of activity in these online platforms, in which gender is one of the variables of interest. However, since gender is not explicitly recorded, it is typically inferred by automatic “gender guessers”, based on cues derived from an individual’s online presence, such as their name and profile picture. As opposed to self-reporting, used in earlier studies, gender guessers scale better, but their accuracy depends on the quantity and quality of data available in one’s online profile.

In this paper we evaluate the applicability of different gender guessing approaches on several datasets derived from **Stack Overflow**. Our results suggest that the approaches combining different data sources perform the best.

## Keywords

Stack Overflow, gender identification

## 1. INTRODUCTION

Software development, more than other STEM disciplines, remains predominantly male [6], despite coordinated efforts from research, industry, and policy makers. This gender imbalance, perhaps surprisingly, is most visible in the social programming (open-source) world, on platforms such as **Stack Overflow** and **GitHub** where the fraction of non-male participants is in the single digits [5, 13, 19].

To better understand the reasons behind this disparity, and offer support for decision making, we and others have been engaged in empirical studies of activity in these online platforms, in which gender is one of the variables of interest [10, 12, 15, 16, 19]. However, since gender is not explicitly recorded, it is typically inferred by automatic “gender

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

*MSR’16, May 14-15 2016, Austin, TX, USA*

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4186-8/16/05...\$15.00

DOL: <http://dx.doi.org/10.1145/2901739.2901777>

guessers”, based on cues derived from an individual’s online presence, such as their name. While self-reporting, used in earlier studies, is necessarily limited by the availability of the data or the need to contact the respondents directly, gender guessers can be applied on a larger scale and guess gender of the individuals that cannot be contacted. Accuracy of gender guessers, however, is not perfect and might have affected validity of the previously published results. Hence, in this paper *we evaluate accuracy of different gender guessers*.

While gender guessers can be applied to any social-network-like platform, gender identification on **Stack Overflow** is more challenging. Indeed, as opposed to **Google+**, **Stack Overflow** does not explicitly record gender. As opposed to **Facebook**, **Stack Overflow** does not enforce the “real name policy” and non-real names are not unheard of on **Stack Overflow**. Finally, while a normal activity on **GitHub** (= committing) involves disclosing an email address, this is not the case for **Stack Overflow**. Therefore, in this paper we focus on guessing gender of **Stack Overflow** users.

Most of the automatic gender guessers such as *Gender Guesser*,<sup>1</sup> or *genderChecker*<sup>2</sup> are based on first name lists indicating their frequency; this way one can conclude *e.g.*, that Claire is likely to be female while Lucas is likely to be male. Vasilescu *et al.* [17] and *genderize.io*<sup>3</sup> combine name lists with location information to distinguish between male Andrea from Italy and female Andrea from Germany. Blevins and Mullen introduce the time aspect to gender identification: Morgan born in the 1950s is likely to be a man, while Morgan born in the 1990s is likely to be a woman [3]. The main shortcoming of those approaches is that first-name frequency lists might not be available for certain countries: *e.g.*, Vasilescu *et al.* compiled name lists for thirty countries and only for eleven of them frequency information was available [17]. Therefore, our first research question is:

**RQ1** *How to identify genders of Stack Overflow contributors based solely on their names when no first name frequency lists are available?*

Going beyond the name-based techniques, Argamon *et al.* [1] have developed an approach to gender identification based on gendered differences in writing pertaining to the use of determiners, pronouns, and prepositions. Unfortunately, the technical character of writing on **Stack Overflow** and errors made by non-native speakers are likely to affect the results of gender identification [17].

Rather than working with names one can use image pro-

<sup>1</sup><https://market.mashape.com/montanaflynn/gender-guesser>

<sup>2</sup><http://genderchecker.com/>

<sup>3</sup><https://genderize.io/>

cessing techniques to identify gender based on the user profile pictures [8]. *A priori*, image processing techniques can be beneficial for gender identification only if a face can be recognized in the user profile picture. However, one might wonder how commonly face pictures are used as user profile images, *i.e.*, whether face recognition techniques can be applied at all, and whether **Stack Overflow** contributors that chose to disclose their face are not also likely to disclose their real names as well, *i.e.*, whether it is meaningful to apply image processing *on top* of the name-based heuristics:

**RQ2** *Are image processing techniques beneficial for gender identification of Stack Overflow contributors?*

Finally, research of online personae shows that the combination of information from different social networking sites such as **Facebook** and **LinkedIn** allows one to further the deanonymization attempts, *i.e.*, to discover information about individuals that could not have been discovered by inspecting a single site [21]. We focus on **GitHub**: while this site is complementary to **Stack Overflow** in its purpose (code sharing vs. questions & answers), both are related to programming and their communities are overlapping [18].

**RQ3** *Is information from GitHub beneficial for gender identification of Stack Overflow contributors?*

The remainder of this paper is organized as follows. After introducing the gender identification approaches, datasets, and evaluation metrics in Section 2, we discuss the evaluation results in Section 3 and conclude in Section 4.

## 2. METHODOLOGY

### 2.1 Self-representation on Stack Overflow

**Stack Overflow** is the most popular Q&A site dedicated to computer programming. It has been extensively studied by software engineering researchers [9, 11, 17].

Contributors can present themselves on **Stack Overflow** by means of the user profile page. The user profile page lists representation elements such as the contributor’s name, user profile picture, location, usernames on **Twitter** and **GitHub**, personal website, as well as the number of activity-related elements such as the reputation points and badges.

### 2.2 Gender Identification Approaches

In our experiment, we consider sixteen approaches to gender identification. All gender identification approaches can report “female”, “male” or “unknown”<sup>4</sup>.

The first group of approaches is used as the baseline for comparison with the more advanced techniques. The simplest approach is the *straw man* that always predicts “male”. Next we consider two collections of name-based heuristics: *genderComputer* [17]<sup>5</sup>; and the most popular gender identification tool on Mashape *Gender Guesser*.<sup>6</sup>

To address **RQ1**, *i.e.*, name-based gender identification in absence of first name frequency lists, we propose to exploit a (general-purpose) social network as a source of name frequency information. Indeed, if by consulting such a social network we can establish that most individuals called “Orit” are female, then we can reasonably conclude that “Orit” from

**Stack Overflow** is also female. To support this argument we therefore need to access a popular social network that allows the users to indicate their gender, and this information is publicly available. This turns out not to be the case for **Facebook**, but it is the case for **Google+**. In our implementation we consider the first seven **Google+** users retrieved by the search engine and use the majority vote on female/male/unknown (= undisclosed).

Recently Terrell *et al.* [15] have also used **Google+** to infer gender based on an email address. Since **Stack Overflow** profiles do not include email, we first searched for the **Stack Overflow** username on **GitHub** to obtain the email, and compare it with the MD5 hash of the email recorded in the September 2013 **Stack Overflow** data dump<sup>7</sup>. If the **Stack Overflow** username could not be found on **GitHub** or did not have an email address associated with it, we infer “unknown”. Next, we use **Google+** API to imitate the method of Terrell *et al.* [15]. Since we have discovered that **Google+** API rarely allows one to associate the email address with a **Google+** profile, we also *manually* consult **Google+** to get the most precise answer one could hope for. If the email address could not be associated with a **Google+** profile or the profile did not indicate gender we infer “unknown”.

To address **RQ2** about the applicability of the image processing techniques, we first manually analyze a sample of user pictures of the **Stack Overflow** contributors. We select user profile pictures of 900 **Stack Overflow** contributors such that 450 of them did not indicate their age, 150 are younger than 26, 150 are between 26 and 32, and 150 are 32 or older; moreover, we select the contributors such that a third of them have reputation lower than 200, a third—between 200 and 999, and a third—above 1000. Reputation thresholds have been selected to reflect the **Stack Overflow** “privileges”<sup>8</sup>. 479 out of 900 profile pictures, *i.e.*, slightly more than a half, represent humans. Therefore, we consider image processing techniques promising and evaluate their ability to identify gender of **Stack Overflow** contributors. To this end we use **Face++**, an online image processing tool of Megvii, inc.<sup>9</sup> **Face++** can be used either in isolation or as a post-processing step when the preceding techniques, *genderComputer*, *Gender Guesser* and **Google+**, fail.

Finally, to answer **RQ3** we study whether combining information from **Stack Overflow** user profiles with information from **GitHub** is beneficial for gender identification. Indeed, one can expect that matching information from **Stack Overflow** with information from **GitHub** (cf. [18]) can reveal more about the individuals, and hence be beneficial for gender identification. Therefore, we combine the techniques considered so far with information from **GitHub**: we first search **GitHub** for users whose user name is identical with **Stack Overflow** user name to get the display name (and location), and then feed this information to the subsequent steps of *genderComputer*, *Gender Guesser* and **Google+**, or their combinations with **Face++**. If no matching user can be found on **GitHub**, original information from **Stack Overflow** is used. We do not use profile pictures from **GitHub**.

<sup>4</sup>Some name-based approaches also report “unisex” that we interpret as being equivalent to “unknown”.

<sup>5</sup><https://github.com/tue-mdse/genderComputer>; country-name information by <https://github.com/tue-mdse/countryNameManager>

<sup>6</sup><https://market.mashape.com/montanaflynn/gender-guesser>

<sup>7</sup>We had to use the September 2013 release of the **Stack Overflow** data dump rather than a more recent one, since due to privacy concerns **Stack Overflow** has decided to remove the hashes from the subsequent data dump releases.

<sup>8</sup><http://stackoverflow.com/help/privileges>

<sup>9</sup><http://www.faceplusplus.com/>

## 2.3 Datasets

Since on **Stack Overflow** genders are not included in the profile information, we cannot rely on our own ability to correctly identify gender of the **Stack Overflow** contributors. The situation is further complicated by the “gender swapping”, *i.e.*, men pretending to be women and women pretending to be men [17] (similar “gender swapping” has been observed in an online poker community [22]). Hence, as the ground truth we reuse three datasets obtained via surveys that include gender as reported by the respondents.

The IWC dataset has been conducted to evaluate the approach of Vasilescu *et al.* [17] and records information about 117 respondents (106 male and 11 female). However, since two users cannot be found in the newest **Stack Overflow** data dump<sup>10</sup> or online API, information about 115 respondents is used for the evaluation, 104 male and 11 female.

The FLOSS dataset [13] contains information of 439 individuals. However, 7 did not disclose their gender and we use information about 391 male and 41 female respondents.

Finally, we use the Diversity dataset [20] containing 731 records. However, not all respondents disclose their genders. Moreover, as above, we calculate the MD5 hash of the email address indicated by the survey participant and compare it with the MD5 hash recorded in the September 2013 **Stack Overflow** data dump (cf. [18]). Ultimately, we obtain information about 147 individuals (129 male and 18 female).

The datasets have been collected through different channels and show a very limited overlap: one person (female) was included both in the FLOSS dataset and in the Diversity dataset. Moreover, the  $\chi^2$  test did not observe statistically significant differences between the gender distribution in the three datasets<sup>11</sup>. Thus, we also evaluate the approaches with respect to the combined dataset (624 males and 69 females).

## 2.4 Evaluation Metrics

Since in each dataset more than 87% of the respondents are male, *i.e.*, the dataset is unbalanced, traditional metrics such as accuracy might be misleading [2]: indeed, accuracy of the straw man approach always predicting “male” can be easily higher than of any of the competing approaches.

Therefore, we use three alternative evaluation metrics. The first metrics, the *Adjusted Rand Index (ARI)* [7, 14], measures the correspondence between two partitions of the same data: the partition of the individuals into genders based on the self-reporting and the one induced by the gender guessers. For two partitions  $U$  and  $V$  of a collection of  $n$  individuals let  $a$  be the number of pairs of individuals that  $U$  and  $V$  place in the same group;  $b$ —the number of pairs of individuals that  $U$  places in the same group but  $V$  does not;  $c$ —*vice versa*, and  $d$ —the number of pairs of individuals that neither  $U$  nor  $V$  place in the same group. Then,

$$ARI = \frac{\binom{n}{2}(a+d) - ((a+b)(a+c) + (c+d)(b+d))}{\binom{n}{2}^2 - ((a+b)(a+c) + (c+d)(b+d))}$$

ARI does not exceed 1: the closer the value of ARI is to 1 the better the correspondence between the partitions.

As the second evaluation metrics we consider the *wam*, ratio of the number of women classified as men to the total

number of women. Approach that does not classify women as men ensures that the all individuals claimed by the approach to be men are, in fact, men. Although there might be some men classified as women, if the group with both male and female developers differs in its behavior from the purely male group, we can assure the differences would be more apparent between the group consisting of only women and the group consisting of only men. Ideally, *wam* should be as low as possible. Observe that one could have also considered the *maw*, ratio of the number of men classified as women to the total number of men. An approach optimizing *maw* would favor classifying women as men (rather than men as women) at risk of rendering the women’s contributions invisible.

Finally, we have observed that several **Stack Overflow** users have automatically generated profile images and user names, *e.g.*, “user123456”. For these users, even for humans it is impossible to determine their gender. Therefore, there is an upper bound on what one might reasonably expect from an automatic gender guesser. Hence, the last evaluation metrics we consider is *cut*, the ratio of the number of correctly identified women to the manually established upper bound. Ideally, *cut* should be as close to 1 as possible.

Approaches that prefer to avoid misclassification of women as men by frequently reporting “unknown” indeed reduce *wam* but also reduce *cut* at the same time.

## 3. RESULTS

Results of the evaluation are summarized in Table 1.

We say that approach  $A$  is *preferred* to approach  $B$  on the dataset  $D$  if the ARI of approach  $A$  applied to  $D$  is not lower than of  $B$  applied to  $D$ , *wam* is not higher, *cut* is not lower and at least one of the evaluation metrics of  $A$  is strictly better (*e.g.*, higher for ARI and *cut* and lower for *wam*) than the corresponding metrics of  $B$ . An arrow from approach  $A$  to  $B$  in Figure 1 indicates that  $A$  is preferred to  $B$  for all datasets considered. For the sake of readability Figure 1 omits arrows that are implied by transitivity.

Not surprisingly, almost any approach considered is preferred to the straw man. The basic approaches, genderComputer, Gender Guesser, **Google+** and **Face++**, are incomparable. Among 41 women in the Diversity dataset for 6 all four basic approaches correctly identify their gender while for 12 none of the approaches has successfully done so.

The variants of the approach by Terrell *et al.* [15] are absent from Figure 1 since they are incomparable with the other approaches due to very low agreement ( $ARI < 0$ ). The low ARI and *cut* scores can be explained by few individuals with the verified mail (*e.g.*, 76/432 for FLOSS) and even fewer with the email on **Google+** (*e.g.*, for 21/76 for FLOSS when the API is used and 54/76 for the manual check).

We also observe that integrating **GitHub** is beneficial when name-based heuristics (genderComputer or Gender Guesser) are used: combination of such a heuristics with or without **Face++** extended with the display name and location from **GitHub** is preferred to the corresponding approach without the benefit of **GitHub**. Furthermore, combinations of name-based heuristics and **GitHub** outperform **Face++**. Adding the display name and location from **GitHub** is, of course, beneficial only if they differ from those already contained in **Stack Overflow**; moreover, in several situations **Stack Overflow** records the real name and **GitHub**—a nickname.

We observe that relative preference of the approaches differs for different datasets. Among 16 approaches considered,

<sup>10</sup><https://archive.org/download/stackexchange>, acc. May 19, 2015.

<sup>11</sup> $p \simeq 0.578$  when the person in the overlap has been removed from the FLOSS dataset and  $p \simeq 0.745$  when the person in the overlap has been removed from the Diversity dataset.

Approach	IWC			FLOSS			Diversity			Combined			
	ARI	wam	cub	ARI	wam	cub	ARI	wam	cub	ARI	wam	cub	
<b>Baseline</b>	straw man	0.000	11/11	0/11	0.000	41/41	0/38	0.000	18/18	0/15	0.000	69/69	0/63
	genderComputer	0.112	2/11	5/11	0.109	3/41	16/38	0.024	0/18	5/15	0.089	5/69	28/63
	Gender Guesser	0.039	0/11	3/11	0.044	1/41	10/38	0.065	0/18	8/15	0.046	1/69	19/63
<b>RQ1</b>	Google+	0.264	1/11	6/11	0.192	6/41	22/38	0.217	4/18	8/15	0.205	11/69	35/63
	Terrell <i>et al.</i> [15]—Google+ API	0.115	0/11	1/11	0.003	0/41	2/38	-0.011	0/18	1/15	0.014	0/69	4/63
<b>RQ2</b>	Terrell <i>et al.</i> [15]—Google+ manual	0.010	0/11	1/11	-0.029	0/41	3/38	-0.037	0/18	1/15	-0.025	0/69	5/63
	Face++	0.015	2/11	2/11	0.062	4/41	15/38	0.114	1/18	9/15	0.066	6/63	26/69
	genderComputer + Face++	0.133	4/11	5/11	0.197	5/41	22/38	0.102	1/18	11/15	0.169	9/69	38/63
<b>RQ3</b>	Gender Guesser + Face++	0.110	2/11	4/11	0.154	4/41	19/38	0.155	1/18	10/15	0.149	6/69	33/63
	Google+ + Face++	0.304	2/11	7/11	0.282	6/41	24/38	0.279	4/18	9/15	0.282	12/69	39/63
	genderComputer + GitHub	0.192	1/11	7/11	0.320	3/41	21/38	0.151	0/18	9/15	0.251	4/69	36/63
	Gender Guesser + GitHub	0.167	0/11	5/11	0.249	1/41	19/38	0.128	0/18	9/15	0.204	1/69	32/63
	Google+ + GitHub	0.358	0/11	8/11	0.285	7/41	22/38	0.281	3/18	9/15	0.295	10/69	38/63
	genderComputer + GitHub + Face++	0.229	2/11	7/11	0.416	4/41	26/38	0.271	0/18	13/15	0.344	5/69	40/63
	Gender Guesser + GitHub + Face++	0.221	1/11	5/11	0.362	3/41	25/38	0.249	0/18	13/15	0.308	4/69	42/63
Google+ + GitHub + Face++	0.373	1/11	8/11	0.385	7/41	26/38	0.413	3/18	11/15	0.387	11/69	44/63	

Table 1: Results summary for the four datasets; ARI, *wam* and *cub* are described in Section 2.4.




Figure 1: Comparison of gender identification approaches.

$16 * 15 = 240$  pairs are possible; for 164 the same preference relation is inferred for the four datasets, *i.e.*, for 76 pairs a disagreement is observed. This means that finding the “best approach” for a given dataset might be preposterous; further research into automatic gender identification is required.

Finally, we note that the results obtained by any of the approaches are likely to be improved by manual post-processing [17].

### 3.1 Examples of failures

To obtain more profound insights in the performance of different gender identification approaches, we discuss the cases when those approaches fail to recognize the gender of a **Stack Overflow** contributor of misgender them. To protect the privacy of the individuals involved, we do not use the actual examples from the datasets but construct artificial example showing the limitations of the approach.

A respondent in the FLOSS dataset uses a typically masculine name and profile picture but indicated “female” as her gender. This is why none of the approaches can correctly identify her gender: simpler approaches tend to label her as “unknown” reducing both *wam* and *cub*, more elaborate ones label her as “male” increasing both *wam* and *cub*.

Name-based approaches fail when the user names are not really human names, but can be interpreted as such. For instance, name based approaches would typically recognize “Holly Blue” as a woman, since the first name Holly is more commonly feminine. However, Holly Blue is also a butterfly, *Celastrina argiolus*, and the user name does not necessarily reveal anything about the gender. Another challenge to name-based approaches is presence of names commonly used both by women and men such as Chris. Furthermore, some **Stack Overflow** contributors prefer to use names of TV se-

ries characters whose gender does not necessarily coincide with the gender of the contributor themselves.

Image processing techniques such as **Face++** cannot recognize gender if the user profile picture does not contain a human face, or when the face is not shown frontally. Moreover, in a number of cases the gender has been misclassified due to occlusion, *e.g.*, presence of glasses or sunglasses that hinder localization of the individual’s eyes [23].

## 4. CONCLUSIONS

*Threats to validity* As any empirical study our work is subject to a series of threats to validity, pertaining *e.g.*, the choice of gender identification approaches and datasets. We also made a simplifying assumption of the gender binary.

In this paper we evaluate 16 approaches to automatic gender identification. We apply them to information about **Stack Overflow** contributors obtained from earlier surveys.

We conclude that while individual elements of the gender identification technology such as image recognition and name-based heuristics are readily available, the technology still needs to mature. Existing tools can generate conflicting results and different approaches perform differently on the datasets. A promising direction would involve deeper integration of data from multiple sources including social media sites targeting software developers (*e.g.*, **GitHub**) or general audience (*e.g.*, **Google+**, **Facebook**, and **Twitter**) [4].

## Acknowledgments

We would like to thank Pierre Gaborieau for manual analysis of **Stack Overflow** user pictures, and Bogdan Vasilescu for the feedback on the earlier draft of this paper.

## 5. REFERENCES

- [1] S. Argamon, M. Koppel, J. Fine, and A. Shimoni. Gender, genre, and writing style in formal written texts. *Text*, pages 321–346, 8 2003.
- [2] G. E. A. P. A. Batista, A. C. P. L. F. Carvalho, and M. C. Monard. Applying one-sided selection to unbalanced datasets. In *Mexican International Conference on Artificial Intelligence: Advances in Artificial Intelligence*, pages 315–325, London, UK, UK, 2000. Springer-Verlag.
- [3] C. Blevins and L. Mullen. Jane, John... Leslie? a historical method for algorithmic gender prediction. *Digital Humanities Quarterly*, 9(3), 2015.
- [4] A. Capiluppi, A. Serebrenik, and L. Singer. Assessing technical candidates on the social web. *IEEE Software*, 30(1):45–51, 2013.
- [5] 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*, 20(4):364–398, 2008.
- [6] B. for Labor Statistics. Employed persons by detailed occupation, sex, race, and hispanic or latino ethnicity, Feb. 2015.
- [7] L. Hubert and P. Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218.
- [8] X. Lu, H. Chen, and A. Jain. Multimodal facial gender and ethnicity identification. In *Advances in Biometrics*, volume 3832 of *Lecture Notes in Computer Science*, pages 554–561. Springer, 2005.
- [9] A. Marder. Stack overflow badges and user behavior: an econometric approach. In *Proceedings of the 12th Working Conference on Mining Software Repositories*, pages 450–453. IEEE Press, 2015.
- [10] A. Michailidou and A. Economides. Gender and diversity in collaborative virtual teams. *Computer Supported Collaborative Learning: Best practices and principles for instructors*, pages 199–224, 2007.
- [11] P. Morrison and E. Murphy-Hill. Is programming knowledge related to age? an exploration of stack overflow. In *Mining Software Repositories (MSR), 2013 10th IEEE Working Conference on*, pages 69–72. IEEE, 2013.
- [12] D. Nafus. ‘Patches don’t have gender’: What is not open in open source software. *New Media & Society*, 14(4):669–683, 2012.
- [13] G. Robles, L. Arjona-Reina, B. Vasilescu, A. Serebrenik, and J. M. Gonzalez-Barahona. FLOSS 2013: A survey dataset about free software contributors: challenges for curating, sharing, and combining. In *11th Working Conference on Mining Software Repositories, Data Track*, MSR, pages 396–399. ACM, 2014.
- [14] J. M. Santos and M. Embrechts. On the use of the adjusted rand index as a metric for evaluating supervised classification. In *International Conference on Artificial Neural Networks*, volume 5769 of *LNCS*, pages 175–184. Springer, 2009.
- [15] J. Terrell, A. Kofink, J. Middleton, C. Rainear, E. Murphy-Hill, and C. Parnin. Gender bias in open source: Pull request acceptance of women versus men. *PeerJ*, (4:e1733v1):1–25, 2016. This report has not yet been peer-reviewed, and thus the findings should be considered preliminary.
- [16] S. Turkle. *The Second Self: Computers and the Human Spirit*. MIT Press, 2005.
- [17] B. Vasilescu, A. Capiluppi, and A. Serebrenik. Gender, representation and online participation: A quantitative study. *Interacting with Computers*, 26(5):488–511, 2014.
- [18] B. Vasilescu, V. Filkov, and A. Serebrenik. Stack Overflow and GitHub: Associations between software development and crowdsourced knowledge. In *International Conference on Social Computing, Social Computing*, pages 188–195. IEEE, 2013.
- [19] 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 Conference on Human Factors in Computing Systems, CHI*, pages 3789–3798. ACM, 2015.
- [20] B. Vasilescu, A. Serebrenik, and V. Filkov. A data set for social diversity studies of GitHub teams. In *12th Working Conference on Mining Software Repositories, Data Track*, MSR, pages 514–517. IEEE, 2015.
- [21] M. Veeningen, A. Piepoli, and N. Zannone. Are on-line personae really unlinkable? In J. Garcia-Alfaro, G. Lioudakis, N. Cuppens-Boulahia, S. Foley, and W. M. Fitzgerald, editors, *Data Privacy Management and Autonomous Spontaneous Security*, volume 8247 of *LNCS*, pages 369–379. Springer, 2014.
- [22] R. T. A. Wood and M. D. Griffiths. Why swedish people play online poker and factors that can increase or decrease trust in poker web sites: A qualitative investigation. *Journal of Gambling Issues*, 21:80–97, 2008.
- [23] E. Zhou, Z. Cao, and Q. Yin. Naive-deep face recognition: Touching the limit of LFW benchmark or not? Technical report, Megvii, Inc., 2015.