

# “Condescending, Rude, Assholes”: Framing gender and hostility on Stack Overflow

Siân Brooke

Oxford Internet Institute

University of Oxford

United Kingdom

sian.brooke@oii.ox.ac.uk

## Abstract

The disciplines of Gender Studies and Data Science are incompatible. This is conventional wisdom, supported by how many computational studies simplify gender into an immutable binary categorization that appears crude to the critical social researcher. I argue that the characterization of gender norms is context specific and may prove valuable in constructing useful models. I show how gender can be framed in computational studies as a stylized repetition of acts mediated by a social structure, and not a possessed biological category. By conducting a review of existing work, I show how gender should be explored in multiplicity in computational research through clustering techniques, and layout how this is being achieved in a study in progress on gender hostility on Stack Overflow.

## 1 Introduction

The binarization of gender in computational studies often does not sit well with critical theorists. Treated as the ultimate and most simple categorical variable, 0 = Female and 1 = Male is hardcoded into quantitative approaches from the first introductory text. In contrast, critical scholars see gender as social structure, arguing that it creates opportunities and constraints based on a sex-category. From this standpoint, the so called differences between men and women are entirely social conventions and the male-female binary is a fallacy. From Butler’s (1990) work, scholars have understood gender as performative and existing as a stylized repetition of acts rather than an intense adherence to two distinct classifications. Yet, Butler’s (1990) stylized acts and gendered self are limited by the recursive processes inherent in

gender as a stratification. Risman (2004) argues gender is a social structure, having consequence on the individual level in the development of the self, in interaction, and institutional domains. This paper focuses on the consequences of gender social structures in computational cultures, forming the groundwork of a larger doctoral project into how culture and role-based identities intervene in women’s participation and legitimate interaction in informal coding cultures.

The title for this paper originates from the most common words that women used to describe Stack Overflow, the world’s largest programming forum. In their annual survey in 2019, Stack Overflow asked just under 80,000 users what aspects of the platform they would most like to change – which showed some interesting gender disagreements. The words most likely to differentiate men included “official, complex, algorithm”, whilst the words that differentiated women painted a quite different picture; “condescending, rude, assholes” (Stack Overflow, 2019). This gender difference in participation and perception of the Stack Overflow community is the basis of the project outlined in this paper, showing how hostility in ‘condescension’ and ‘rudeness’ deters women from taking part in programming.

In presenting the findings of this year’s Stack Overflow Developers Survey these results were weighted by gender for the first time. Far from demonstrating an understanding of prejudice and hostility on the platform, the weighting was justified by “characteristics of [the] data” to “correct for demographic skew” (Stack Overflow, 2019). The lack of women in computational cultures is not a simple sampling error or a characteristic of data, but an active gender filter that deters women from taking part. I support a move in data science to infuse computational techniques with the capacity to reflect gendered power relations, moving beyond data based

dismissals. In justifying my stance, I will first outline scholarship on the merits of studying identity and gender within a social context and how research design should acknowledge stereotypes. Next, I will show how women's participation in computational culture effects and is affected by anonymity. Thirdly, I will discuss the difficulties of operationalizing gender and the potential benefits of complicating the binary model. Finally, I will show how clustering has shown to be a promising technique to account for gender structures in online forums and my own proposed study. Overall, this paper argues for complicating the gender binary, forgoing predictive accuracy for representative and messy modelling.

## 2 Identity in Context

Early studies of the Internet heralded its disembodied attributes as liberating and a precursor of equality. It was proposed that anonymity can subjugate gender hierarchies, allowing for free and unhindered expression (Allen, 1995). However, as we make sense of identity online, we often round to the most common attributes, and thus anonymity serves to homogenize participants in online forums as belonging to a singular group. This group is college educated, white, and male (Kendall, 2011; Massanari, 2015). The prevailing voice here amplifies the discourse of an apparently neutral meritocracy – hiding the inequalities of race and gender.

Critical research has a long history of investigating gender, inequality, and interaction. Whilst the principles of social structures are pertinent across contexts, their exact form can change with social locale (Bucholtz & Hall, 2005; Risman, 2004). Wenger & Lave (1991) propose that researchers of such collective identities should focus on communities of practice (CoP) in which members are drawn together by a common interest or that are created deliberately with the goal of gaining knowledge in a certain field. This conception has since been expanded to include virtual communities of practice (VCoP), to show the extension of this anthropological phenomenon online. Stack Overflow can count as one such VCoP, as individuals come together to solve their programming woes. In such communities, Bucholtz (2005) argues that a situational and context-based methodology is fundamental to understanding the gendered social meaning that is

attributed to practices by individuals and cultures. Moreover, the representation of identity in speech should be conceptualized in terms of *communities* of as identity, not collections of individuals (or observations) as the bestowing of agency cannot be segregated from culture. In this manner, one's identity and behavior towards others is shaped by the community in which one participates and interacts – even online.

## 3 Gender in Context

In the initial scholarly discussion of identity formation and interaction, Lakoff (1973) first proposed that men and women differ in how they use words. Whilst gendered meaning necessitates difference, a difference in speech does not directly imply gendered meanings. A man's speech being different to women's means little without context. To ascertain if gendered differences carry meaning one must look at the interrelated layers of the interaction, such as what it means for a woman to be a speaker in this particular scenario (Needle & Pierrehumbert, 2018). What does it mean to be a woman to correct a man in computer science classroom? What does it mean for a man to fail a mathematics class but excel in a gender studies course? Such identity struggles are visible in discourse, or how knowledge creates meaning in interaction as a consequence of social structure (Risman, 2004). It is thus necessary to consider the social context when considering how gender may be presented.

Gender can alter how a community talks about itself and its members. In using the sociolinguistic framing of gender and local context, feminist linguists have pointed to how normative discourse can represent gendered power structures and the male-centric nature of language (Lamerichs & Te Molder, 2003; Tanczer, 2015). This is particularly apparent in the use of 'guys' as a collective. As we move online, the physical markers of gender are invisible in anonymous forums, and male-centricity is amplified to male-by-default (Tanczer, 2015). As a space becomes more masculine and the in-group becomes male, women are framed in terms of stereotypes and identity tropes (Tanczer, 2015). In computational cultures, this communication process cultivates a femininity of technological incompetence and juvenile 'girlness' (Nic Giolla Easpaig & Humphrey, 2016; Shifman & Lemish, 2011). A male dominated masculine

space can therefore lead to understanding women only in terms of the outsider.

There are consequences to stereotypes as they are relational. Gender stereotypes can be internalized and influence the manner in which one conceives of their own abilities and those of others. Risman's (2004) conception of gender as an encompassing social structure permeates online and offline interaction. As gender shapes interactions due to cultural expectations it also shapes one's identity, and there are consequences for institutional domains and technological cultures. As gender power relations are evident in self-presentation and interaction, this in turn affects opportunities in formal settings as stereotypes dictate expectations of others and ourselves (Adams et al, 2006). A popular theory in social psychology, *stereotype threat* refers to being at risk of confirming, as a self-characteristic, a negative stereotype about one's social group (Steele & Aronson, 2000). When one's self is viewed in terms of a salient group membership, performances can be undermined because of concerns about confirming negative stereotypes of one's group. In other words, telling women they can't code because they are women becomes a self-fulfilling prophesy - a false definition of the situation evokes new behavior, which makes the original false conception come true. Ergo, women can't code so there are few women in programming, from here we have 'proof' of the original stance that women can't code.

#### 4 Girls can't code

In negotiating identity in masculine or nerd dominated spaces on, women may purposefully obscure their gender to participate in the social structures of a technical setting. The prominence of stereotypes and the belief that "girls can't code" means that women who show they are women in programming forums often face hostility and harassment (Ford et al., 2016). Nonetheless, Terrell et al (2017) found that women's contributions of code to the repository GitHub were approved at a higher rate than code written by men. In fact, women's contribution acceptance rates were higher than men for every programming language in the top 10 on the GitHub platform (Terrell et al 2017). However, when women's gender was identifiable on their GitHub profile, their acceptance rate dropped to significantly lower than

the average for men (Terrell et al., 2017). This shows not only do women obscure their gender in order to participate, but they are penalised when their gender is known, dropping below the level of men.

Looking to Stack Overflow, Ford et al. (2016) found that impersonal interactions were the main factor that discouraged women from contributing. The women (N = 22) interviewed for the study cited three features of the platform that deterred them from contributing: (1) anonymity was seen to contribute to blunt and argumentative responses on posts, (2) invisibility of women leads to the site feeling like a 'boy's club' full of 'bro humor' (Ford et al., 2016, p. 6), and (3) large communities are intimidating, and not possible in the same way offline. On Stack Overflow we can see a continuation of the theory that anonymous spaces lead to male-by-default interactions. The affordances of anonymity in Computer Mediated Communication (CMC) are evidently more beneficial to an ingroup, and attributes (or language) that might work for a majority group can be barriers for identifying with a community. Building on this, Ford et al. (2017, p. 1) conducted a second study where they developed the concept of *peer parity*: having similar individuals to compare oneself to in a space. The study found that the presence of female-identifying usernames on a thread increased the likelihood that a woman would engage actively with the Stack Overflow community (Ford et. al., 2017). When taken together, Terrell (2017) and Ford (2016; 2017) show that women hide their gender to participate, but this contributes to perceptions of a male-dominated space. This in turn deters women from participating as they do not see anyone like themselves. For women, stereotype threat creates a cyclical self-fulfilling prophesy, as does anonymity in not seeing *someone like me* in technical spaces.

#### 5 Unlikely Allies

The disparity of women's representation in technical culture extends to those capable of computational methods, as only 15% of Data Scientists and computational researchers are women (Miller & Hughes, 2017). Comparatively, and estimated 75% of sociologists who focus on Gender are women (ASA, 2015). There are a number of notable exceptions to the trend, but this does not mean that the overall picture is endangered (See Ford 2016; 2017 as an example).

Whilst Data Science may dismiss inequality and women's lack of representation as a characteristic of the data, those who may provide insight are frequently not invited into the conversation. For Data Scientists, perchance it is not only the stereotype that girls can't code, but maybe also gender theorists.

Research has shown how valuable the social science lens is to computational fields (Kokkos & Tzouramanis, 2014; Nguyen, Doğruöz, Rosé, & de Jong, 2016; Otterbacher, 2013). Researchers at this intersection are aware of the tension between the theoretical framing and empirical methods of their work. Yet, whilst theory must begin with human-orientated ideas, these notions are only valuable if they are confirmed through empirical methods. Far from incompatible, the value placed on creativity and predictive accuracy in computational fields is well matched to the esteemed validity and reliability of the social sciences (Nguyen et al., 2016). This exciting and novel *modus operandi* is beginning to flourish in examining a range of inequalities online.

In computational sociolinguistics text is social data, and the choice of language used signals a performed identity (Nguyen et al., 2016). In a traditional sociological framing, agency occurs in linguistic symbols as social currency. A struggle is evident here, as the parsimonious causality prized by quantitative and computational approaches meets the messiness of the social world. In computational sociolinguistics a balance needs to be sort between language reflecting additional social structures, and language arising from speaker agency (Nguyen et al., 2016). Put simply, not everyone writes in a way that reflects their biology, and thus the agency of speakers should be acknowledged in interpreting findings.

As a case that exemplifies this argument, Otterbacher et al. (2013) examined the anonymous review site Internet Movie Database (IMDb) and found that women's reviews were weight as having less utility than men's. They also found that highly rated woman authors would exhibit "male" characteristics in their writing, such as less pronouns, complexity, and vocabulary richness (Otterbacher, 2013). The agency of the speakers is shown in the increased 'maleness' of language, as well as methodological evidence against biological determinism. Here, reputation voting systems of IMDb meant that female-based writing was downvoted. The reputation system acts as a gender

filter, in which the gender-majority dictates success (Herring et al, 2002). Gender structures clearly mediate online interactions even in contexts that are far less heavily associated with masculine stereotypes that computational cultures.

The proposed study applies this conception that the male-majority dictates the identity performance required to succeed in a given social context and institutional setting (Risman, 2004). Looking to Stack Overflow, we propose that an estimated 89-94% male majority fosters masculine linguistic repertoires where those who don't conform are punished with invisibility – colloquially referred to as being "downvoted into oblivion" (Clark-Gordon et al, 2017). As Hogan (2013) points out, conforming to a male-voice in order to successfully participate in a space is not a characteristic unique to computational culture or online forums. Take for instance the use of male pen names, the Brontë sisters were Currer, Ellis, and Acton Bell and Mary Ann Evans who used the guise of George Eliot (Hogan, 2013). The implication here is that computational methods allow for the mapping of such phenomena. However, before introducing the proposed study, we must first consider how feminine and masculine speech used by both male and female authors complicates the simple binary understanding of gender operationalized in many computational studies.

## 6 The trouble of operationalization

In applying computation, it is crucial that the research design is aptly framed to not recreate inequalities. As noted earlier, gender is often treated as a latent attribute – a implicit assumption that linguistics choices are associated with distinct categories of people (Needle & Pierrehumbert, 2018). The generalization of gender norms in computational research has been shown to contribute to stereotypes – seeing gender as something that people 'have' (not 'do'), neglecting agency to mask ones gender. In defending the binary classification to gender it is important to note that statistical definitions of the accuracy of predicative modelling does not mean that the picture is not oversimplified (Nguyen et al., 2016). Incorporating more critical understandings of gender may decrease predictive accuracy, but as it would include an understanding of social structures the reproducibility of results may benefit. As gender social structures have

consequences in interactions and infrastructure, a critical approach may not overfit a model to gender in a particular context. To build on the aphorism of the statistician George Box, if ‘all models are wrong’, can adding critical gender theory make them more useful?

In discussing the apparently conflict paradigms of social theory and computational methods, Nguyen et al. (2016) point to the value placed on of construct validity in more critical approaches. For the uninitiated, construct validity is “extent to which the experimental design manages extraneous variance effectively” (Nguyen et al., 2016). This can be particular important in how gender is conceived of within a study. As we saw with Otterbacher et al.’s (2013) study into linguistic gender on IMDb, women who exhibited “maleness” in there speech were more highly rated. This shows that whilst a platform appears to be numerical equal, it can still be performatively and legitimately masculine. In not paying due attention to such confounding factors of gender social structures, may leave the results of an investigation to be weak, regardless of the number associated with predictive accuracy. Indeed, the social word is far messier than many predictive models may lead us to believe.

Whilst computational studies into gender differences do valuable work to highlight the dearth of women in technical spaces, they can be guilty of perpetuating the underrepresentation. It can be dangerous to qualify contextual legitimacy or success in terms that are intrinsically gendered. In examining the open source development platform GitHub, Vedres and Vasarhelyi (2018) found that ‘disadvantage is a function of gendered behavior’. In the study the variable of femaleness was qualified by professional ties, level of activity (push/pull requests), and areas of specialization (Vedres & Vasarhelyi, 2018). The study argues that measures of reputation (‘success’ – as starred repositories) and survival (‘time account active’) on the platform were adversely affected by femaleness rather than by categorical discrimination. They found that not only was this true for women, but men and users with unidentifiable gender are also likely to suffer for exhibiting behavior that demonstrated femaleness. The findings of Vedres and Vasarhelyi (2018) are valuable as they show that behavior classified as feminine adversely effects one’s status (in their defined terms), not just listing ‘female’ on a profile.

Nonetheless, as is typical of gender classification studies, the ‘behavioral’ aspect was built from an extrapolation of categorical gender. That is, the features that are defined as ‘femaleness’, are built from behavior associated with a ‘female’ (categorically defined) account. Thus, the causality of gendered performance versus identification is unclear, and not supported by critical studies. The assertion made here that “women are at a disadvantage because of what they do, rather than because of who they are” (Vedres & Vasarhelyi, 2018) oversimplifies acting as a women and being a women into discrete and mutually exclusive categories. Nevertheless, that study shows that the default masculinity is ratified through behavior that generates contextual ‘success’, rather than by the overt presence of men. Vedres & Vasarhelyi's (2018) project reflects one of the significant challenges of critical research with computational methods: the operationalization of gender as a variable in manner that does conflate masculinity with community’s definition of success.

## 7 Beyond the binary

As illustrated above, gender as a binary can miss some vital aspects of community functioning and belonging. As interactions dictate how men and women can act even in identical structural positions, gender can define the capacity for action in a given environment. For instance, Cheryan et al. (2009) found that exposure to stereotypical masculine computer science environments actively deters women from participating, even when the space was populated by women. Extending this, work by Ford and Wajcman (2017) & Schwartz and Neff (2019) shows that the social structures of gender permeate online spaces, as technology’s design and use draws on the cultural and institutional repertoires in a male dominated space.

Whilst incorporated in many traditionally critical gender studies, going beyond a binary understanding can prove challenging to quantitative and computational research. How can gender be operationalized into a variable that accounts for a myriad of gender performances? Much work relies on the idea that the majority of individuals consider gender a binary, so therefore it is binary in studying the social phenomena in which said individuals participate. Whilst there is merit to this rationale, there are simple computational methods that can be used to negate

a constrained and binary understanding of gender. A promising technique is that of Cluster Analysis. In clustering observations are grouped based on similarity, and to show the difference between different groups. Clustering is an unsupervised Machine Learning technique commonly used to gain valuable insights for patterns in data.

In their study into gender, networks, and linguistic style on Twitter, Bamman et al (2014) propose a more nuanced approach to quantitative work on gender. They point to how measures of predictive accuracy do not mean that the model does not distort the social world. Building on Butler's (1990) casting of gender, they take a two step approach to modelling gender.

**Step 1:** Predict gender with a Logit model using lexical features (i.e. Dictionary words, slang, taboo, hashtags)

**Step 2:** Group authors by similarity in word usage and look at the gender breakdown of each cluster.

By looking at which words and lexical features are most associated with users that profile states their gender in Step 1, Bamman et al (2014) take a situated approach to meaning. The stylized reputation of acts that are make up a gender performance vary by context. For instance, if a individual swears, the way that profanity is received by a audience will depends on the characteristics of the speaker (gender, ethnicity, age) and the context in which they are speaking (with friends, family dinner, classroom) and the role they are acting (policeman, mother, priest). As such, a perspective that incorporates situated meaning is the only way to understand the relationship between gender and language. In Step 2 of Bamman et al's (2014) study, Twitter users were grouped by similarities in word usage. Using a clustering algorithm based on the Expectation-Maximization framework, the clusters were built without considering gender yet had strong gender majorities. This approach to clustering allowed for multiple expressions of gender, which the authors speculate may be related to an interaction between age or ethnicity (Bamman et al., 2014). Conducting research in this manner, with gender not treated as the response variable, allowed for findings that were unexpected. For example, whilst taboo terms were generally shown to be preferred by men, several male-associated cultures reversed this trend

(Bamman et al., 2014). Overall, the clustering methodology of this study incorporates the social relation of "male" and "female" categories, going beyond descriptive understandings and acknowledging the normative gender performances that define inclusion and exclusion. Whilst this is not a perfect approach, it does highlight the possibility of clustering to examine how social identity can be evident in data without being determined by demographic markers. However antithetical they may seem, critical gender studies and computational methods can be unlikely, and valuable, allies.

## 8 Stack Overflow: A Research Agenda

Often referred to as the 'programmer's paradise', Stack Overflow is the largest online community of coding knowledge, boasting 9.9 million registered users and 50 million monthly visitors, of whom 21 million are professional developers and university-level students (Ford et al., 2016). Yet, with an estimated population of only 6-11% women, the popular platform is only paradise for some. The approach uses on Butler's conception of gender as enacted, incorporating situational meaning (Lea and Spears' 1991) and considering discourse-in-context, as argued for by Needle and Pierrehumbert (2018), Buckoltz (1999) and Lamerichs & Te Molder (2003). Building on the work of Adam (2003), Edwards (2003), Tanczer (2015), and Sollfrank (1999, 2002), this study examines the visibility of gender in accessible technical spaces. Through a twostep process of Natural Language Process, Machine Learning (*sklearn*), and Cluster Analysis (*Expectation-Maximization framework*), as used by Bamman et al (2014), I will analyse linkages between masculine-linguistic practise and reputation building.

A significant portion of the contribution of the study will be methodological, as I aim to provide a simple road map by which critical research can be conducted with computational methods, accounting for levels of gender visibility. In exploring this, I ask how visible gender is on Stack Overflow, and what situational meaning imbues text with hostility.

## 8.1 Data Collection

In first setting out the data for analysis, I will use the Stack Overflow data dump, hosted on Google BigQuery. Separated into different tables, the information available is posts, users, votes, comments, posts history, and post links. Updated on a quarterly basis, the BigQuery dataset includes an archive of Stack Overflow's user-contributed content, including posts, votes, tags, and badges. This dataset is updated to mirror the Stack Overflow content on the Internet Archive and is also available through the Stack Exchange Data Explorer. Inherent in this data are several challenges of working with big data (~180GB), such as different features of a post stored in separate tables (i.e. 'tags', accepted answers, post content). The Data Dump also contains substantial metadata, meaning data that provides a description of information in the dataset, such as suggested edits and location of users. Datasets such as this provide a wealth of information and contextual CMC that is underutilized in social science research. I will use the location of users to narrow my population to the USA and UK. Whilst this does lead to a Western-focused dataset, it also means that I am not homogenizing gender performances across cultural contexts.

## 8.2 Rudeness and Offence

On the Stack Overflow dataset, I propose to examine how visible gender and what forms hostility can take in context. Informing my analysis with Meta Stack Overflow, I will examine what practices are considered hostile and reduce the visibility (*peer parity* – Ford et al, 2016) of women on the platform. In taking a local and contextual approach to hostile behavior on Stack Overflow candidate features for inclusion were informed by the results of the 2019 Developers Survey and a forum dedicated to studying the research site, Meta Stack Overflow. I will incorporate formal/structural measures of hostility, such as the "Offensive Comment" tag. I will additionally include a subtler element in ascertaining what terms and practices are most associated with this tag. Contextual features that have so far emerged form a reading of Meta Stack Overflow include ratio of code to text, and references to "reading the documentation" in short answers or "not doing your homework for you", and similar sentiments.

The candidate features for inclusion thus reference local and contextual understandings of hostility.

## 8.3 Gender as Tiers

In examining gender as a social structure, I propose to account for both those who clearly identify their gender on their profile as well as those who purposefully obscure it to participate without facing gendered social sanctions. I propose three-tier classifications of gender to map onto the results of the cluster analysis.

- (1) **Self-identified Male or Female:** Identified as a man/woman clearly through their profile (Gender, About Me, Name),
- (2) **Linguistically Masculine or Feminine:** Estimated through a bag-of-word approach using the posts/comments associated with tier 1
- (3) **Neutral:** Unidentified profiles (those users who fall under the conventionally defined 0.8 confidence of tier 2)

Through this distinction, my investigation will not conflate those who identify a gender, with those who perform it. As Otterbacher et al. (2013) show, "maleness" characteristics in speech does not mean that the speaker identifies as a man. This differentiation between claiming and hiding a gender identity in technical cultures will not only be beneficial in terms of building a representative model, but also in not seeing unidentified data as just noise, but rather a potentially purposeful act. These gender classifications will be mapped onto hostility and reputation to see the relationship between gender identification, linguistic-gender and legitimate participation.

Therefore, I will use NLP and clustering techniques to ascertain the gender dimensions of hostile behavior on Stack Overflow, and how this can lead to women's lack of participation. The output of the study will be a categorization of gendered behaviors that mark the space as masculine and create cultural barriers for women's entry into coding forums, even in the anonymous space of programmer's paradise.

## Bibliography

- Adams, G., Garcia, D. M., Purdie-Vaughns, V., & Steele, C. M. (2006). The detrimental effects of a suggestion of sexism in an instruction situation. *Journal of Experimental Social Psychology, 42*(5), 602–615.

- <https://doi.org/10.1016/j.jesp.2005.10.004>
- Allen, B. J. (1995). *Gender and Computer-Mediated Communication. Sex Roles* (Vol. 32).
- ASA. (2015). ASA Sections: Membership by Gender | American Sociological Association. Retrieved June 2, 2019, from <https://www.asanet.org/research-and-publications/research-sociology/trends/asa-sections-membership-gender>
- Bamman, D., Eisenstein, J., & Schnoebelen, T. (2014). *Gender identity and lexical variation in social media 1. Journal of Sociolinguistics* (Vol. 18).
- Bucholtz, M. (1999). “Why be normal?” Language and identity practices in a community of nerd girls Answer the following questions in small groups. *Language in Society*, 28, 203–223.
- Bucholtz, Mary, & Hall, K. (2005, October 1). Identity and interaction: A sociocultural linguistic approach. *Discourse Studies*. SAGE Publications London, Thousand Oaks, CA and New Delhi. <https://doi.org/10.1177/1461445605054407>
- Butler, J. (1990). *Gender Trouble: Feminism and the Subversion of Identity*. New York: Routledge. <https://doi.org/10.1057/fr.1991.33>
- Cheryan, S., Plaut, V. C., Davies, P. G., & Steele, C. M. (2009). Ambient Belonging: How Stereotypical Cues Impact Gender Participation in Computer Science. *Journal of Personality and Social Psychology*, 97(6), 1045–1060. <https://doi.org/10.1037/a0016239>
- Clark-Gordon, C. V., Workman, K. E., & Linvill, D. L. (2017). College students and yik yak: An exploratory mixed-methods study. *Social Media and Society*, 3(2). <https://doi.org/10.1177/2056305117715696>
- Edwards, P. (2003). Nerd Worlds: Computer hackers, unofficial culture and masculin identities, [unpublished paper]. *Program in Science, Technology and Society, Stanford University, Stanford, CA*.
- Ford, D., Smith, J., Guo, P. J., & Parnin, C. (2016). Paradise unplugged: identifying barriers for female participation on stack overflow. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering - FSE 2016* (pp. 846–857). New York, New York, USA: ACM Press. <https://doi.org/10.1145/2950290.2950331>
- Ford, H., & Wajcman, J. (2017). ‘Anyone can edit’, not everyone does: Wikipedia’s infrastructure and the gender gap. *Social Studies of Science*, 47(4), 511–527. <https://doi.org/10.1177/0306312717692172>
- Herring, S., Job-Sluder, K., Scheckler, R., & Barab, S. (2002). Searching for Safety Online: Managing “Trolling” in a Feminist Forum. *The Information Society*, 18(5), 371–384. <https://doi.org/10.1080/01972240290108186>
- Hogan, B. (2013). Pseudonyms and the Rise of the Real-Name Web. In *A Companion to New Media Dynamics* (pp. 290–307). Blackwell Publishing Ltd. <https://doi.org/10.1002/9781118321607.ch18>
- Kendall, L. (2011). “White and Nerdy”: Computers, Race, and the Nerd Stereotype. *Journal of Popular Culture*, 44(3), 505–524. <https://doi.org/10.1111/j.1540-5931.2011.00846.x>
- Kokkos, A., & Tzouramanis, T. (2014). A robust gender inference model for online social networks and its application on LinkedIn and Twitter. *First Monday*, 19(9). <https://doi.org/10.5210/fm.v19i9.5216>
- Lamerichs, J., & Te Molder, H. F. M. (2003). *Computer-mediated communication: from a cognitive to a discursive model* (Vol. 5).
- Lave, J., & Wenger, E. (1991). *Situated learning: legitimate peripheral participation*. Cambridge University Press.
- Lea, M., & Spears, R. (1991). Computer-mediated communication, de-individuation and group decision-making. *International Journal of Man-Machine Studies*, 34(2), 283–301. [https://doi.org/10.1016/0020-7373\(91\)90045-9](https://doi.org/10.1016/0020-7373(91)90045-9)
- Massanari, A. L. (2015). *Participatory Culture, Community, and Play: learning from reddit*. Oxford: Peter Lang.
- Miller, S., & Hughes, D. (2017). *The Quant Crunch: How The Demand for Data Science Skills Is Disrupting the Job Market. BurningGlass, IBM* (Vol. 41). <https://doi.org/10.1177/0032258x6804100305>
- Needle, J. M., & Pierrehumbert, J. B. (2018). Gendered associations of English morphology. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, 9(1), 1–23. <https://doi.org/10.5334/labphon.134>
- Nguyen, D. ., Dođruöz, A. S. ., Rosé, C. P. ., & de Jong, F. . (2016). Computational sociolinguistics: A survey. *Computational Linguistics*, 42(3), 537–593. [https://doi.org/10.1162/COLI\\_a\\_00258](https://doi.org/10.1162/COLI_a_00258)
- Nic Giolla Easpaig, B., & Humphrey, R. (2016). “Pitching a virtual woo”: Analysing discussion of sexism in online gaming. *Feminism & Psychology*, 0(0), 1–9. <https://doi.org/10.1177/0959353516667400>



- Otterbacher, J. (2013). Gender, writing and ranking in review forums: a case study of the IMDb. *Knowledge and Information Systems*, 35(3), 645–664. <https://doi.org/10.1007/s10115-012-0548-z>
- Risman, B. J. (2004). Gender as a social structure: Theory wrestling with activism. *Gender and Society*. <https://doi.org/10.1177/0891243204265349>
- Robin, L. (1973). Language and Woman 's Place. *Cambridge University Press*, 2(1), 45–80. <https://doi.org/10.1017/S0047404500000051>
- Schwartz, B., & Neff, G. (2019). The gendered affordances of Craigslist “new-in-town girls wanted” ads. *New Media & Society*, 146144481984989. <https://doi.org/10.1177/1461444819849897>
- Shifman, L., & Lemish, D. (2011). “Mars and venus” in virtual space: Post-feminist humor and the internet. *Critical Studies in Media Communication*, 28(3), 253–273. <https://doi.org/10.1080/15295036.2010.522589>
- Sollfrank, C. (1999). Women Hackers: A report from the mission to locate subversive women on the net. *Cyberfeminist International*.
- Sollfrank, C. (2002). Not every hacker is a woman. Retrieved April 28, 2017, from [http://www.obn.org/reading\\_room/writings/html/notevery.html](http://www.obn.org/reading_room/writings/html/notevery.html)
- Stack Overflow. (2019). Stack Overflow Developer Survey 2019. Retrieved May 10, 2019, from <https://insights.stackoverflow.com/survey/2019>
- Steele, C. M., & Aronson, J. (2000). Stereotype threat and the intellectual test performance of African Americans.
- Tancer, L. M. (2015). Hacktivism and the male-only stereotype. *New Media & Society*, 1–17. <https://doi.org/10.1177/1461444814567983>
- Terrell, J., Kofink, A., Middleton, J., Rainear, C., Murphy-Hill, E., Parnin, C., & Stallings, J. (2017). Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Computer Science*, 3, e111. <https://doi.org/10.7717/peerj-cs.111>
- Vedres, B., & Vasarhelyi, O. (2018). *Gendered behavior as a disadvantage in open source software development*.