

NRC Publications Archive Archives des publications du CNRC

Gender gap in natural language processing research: disparities in authorship and citations

Mohammad, Saif M.

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

For the publisher's version, please access the DOI link below. / Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

Publisher's version / Version de l'éditeur:

<https://doi.org/10.18653/v1/2020.acl-main.702>

Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pp. 7860-7870, 2020-07

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=0bb14d73-96f3-4527-9a5d-62a19bab347d>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=0bb14d73-96f3-4527-9a5d-62a19bab347d>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

Gender Gap in Natural Language Processing Research: Disparities in Authorship and Citations

Saif M. Mohammad

National Research Council Canada

Ottawa, Canada

saif.mohammad@nrc-cnrc.gc.ca

Abstract

Disparities in authorship and citations across gender can have substantial adverse consequences not just on the disadvantaged genders, but also on the field of study as a whole. Measuring gender gaps is a crucial step towards addressing them. In this work, we examine female first author percentages and the citations to their papers in Natural Language Processing (1965 to 2019). We determine aggregate-level statistics using an existing manually curated author–gender list as well as first names strongly associated with a gender. We find that only about 29% of first authors are female and only about 25% of last authors are female. Notably, this percentage has not improved since the mid 2000s. We also show that, on average, female first authors are cited less than male first authors, even when controlling for experience and area of research. Finally, we discuss the ethical considerations involved in automatic demographic analysis.

1 Introduction

Gender gaps are quantitative measures of the disparities in social, political, intellectual, cultural, or economic success due to one’s gender. Gender gaps can also refer to disparities in access to resources (such as healthcare and education), which in turn lead to disparities in success. We need to pay attention to gender gaps not only because they are inherently unfair but also because better gender balance leads to higher productivity, better health and well-being, greater economic benefits, better decision making, as well as political and economic stability (Skjelsboek and Smith, 2001; Woetzel et al., 2015; Hakura et al., 2016; Mehta et al., 2017; Gallego and Gutiérrez, 2018).

Historically, gender has often been considered binary (male and female), immutable (cannot change), and physiological (mapped to biological

sex). However, those views have been discredited (Hyde et al., 2019; Richards et al., 2017; Darwin, 2017; Lindsey, 2015; Kessler and McKenna, 1978). Gender is complex, and does not necessarily fall into binary male or female categories (e.g. nonbinary people), and also does not necessarily correspond to one’s assigned gender at birth.

Society has often viewed different gender groups differently, imposing unequal social and power structures (Lindsey, 2015). The World Economic Forum’s 2018 Global Gender Gap Report (which examined data from more than 144 countries) highlighted the gender gap between men and women in Artificial Intelligence as particularly alarming (WEF, 2018). It indicated that only 22% of the professionals in AI are women and that this low representation in a transformative field requires urgent action—otherwise, the AI gap has the potential to widen other gender gaps. Other studies have identified substantial gender gaps in science (Håkanson, 2005; Larivière et al., 2013; King et al., 2017; Andersen and Nielsen, 2018).

Perez (2019) discusses, through numerous examples, how there is a considerable lack of disaggregated data for women and how that is directly leading to negative outcomes in all spheres of their lives, including health, income, safety, and the degree to which they succeed in their endeavors. This holds true even more for transgender people. Our work obtains disaggregated data for female Natural Language Processing (NLP) researchers and determines the degree of gender gap between female and male NLP researchers. (NLP is an interdisciplinary field that includes scholarly work on language and computation with influences from Artificial Intelligence, Computer Science, Linguistics, Psychology, and Social Sciences to name a few.) We hope future work will explore other gender gaps (e.g., between trans and cis people). Measuring gender gaps is a crucial step towards addressing them.

We examine tens of thousands of articles in the ACL Anthology (AA) (a digital repository of public domain NLP articles) for disparities in female authorship.¹ We also conduct experiments to determine whether female first authors are cited more or less than male first authors, using citation counts extracted from Google Scholar (GS).

We extracted and aligned information from the ACL Anthology and Google Scholar to create a dataset of tens of thousands of NLP papers and their citations as part of a broader project on analyzing NLP Literature.² We refer to this dataset as the *NLP Scholar Dataset*. We determined aggregate-level statistics for female and male researchers in the NLP Scholar dataset using an existing manually curated author–gender list as well as first names that are strongly associated with a gender.

Note that attempts to automatically infer gender of individuals can lead to harm (Hamidi et al., 2018). Our work does not aim to infer gender of individual authors. We use name–gender association information to determine aggregate-level statistics for male and female researchers. Further, one may not know most researchers they cite, other than from reading their work. Thus perceived gender (from the name) can lead to unconscious effects, e.g., Dion et al. (2018) show that all male and mixed author teams cite fewer papers by female authors than all female teams. Further, seeing only a small number of female authors cited can demoralize young researchers entering the field.

We do not explore the reasons behind gender gaps. However, we will note that the reasons are often complex, intersectional, and difficult to disentangle. We hope that this work will increase awareness of gender gaps and inspire concrete steps to improve inclusiveness and fairness in research.

It should also be noted that even though this paper focuses on female–male disparities, there are many aspects to demographic diversity including: representation from transgender people; representation from various nationalities and race; representation by people who speak a diverse set of languages;

¹<https://www.aclweb.org/anthology/>

²Mohammad (2019) presents an overview of the many research directions pursued, using this data. Notably, Mohammad (2020a) explores questions such as: how well cited are papers of different types (journal articles, conference papers, demo papers, etc.)? how well cited are papers published in different time spans? how well cited are papers from different areas of research within NLP? etc. Mohammad (2020c) presents an interactive visualization tool that allows users to search for relevant related work in the ACL Anthology.

diversity by income, age, physical abilities, etc. All of these factors impact the breadth of technologies we create, how useful they are, and whether they reach those that need it most.

Resources for the NLP Scholar project can be accessed through the project homepage.³

2 Related Work

Pilcher (2017) shows that names function not only to identify individuals, but also to manage gender throughout one’s life. There is a strong cultural norm in various parts of the world to assign a first name to newborns as per their category of sex (Pilcher, 2017; Barry III and Harper, 2014; Lieber-son et al., 2000; Alford, 1987). Pilcher (2017) argues that, throughout their life, the first name plays a role in repeatedly categorizing a person as being male or female. People may change their name and appearance to manage their gender (Connell, 2010; Pilcher, 2017). People may choose a name that is not associated with male or female categorizations (Connell, 2010).

The strong normative tendency to use names to signal gender has led to a large body of work on automatically determining gender by one’s first name, not just for scientometric analysis discussed below, but also for language studies, social sciences, public health, and commerce. However, this can also lead to *misgendering*, which can cause significant pain and harm (Hamidi et al., 2018). (Misgendering is when a person—or in this case, a machine—associates someone with a gender with which they do not identify.) Further, work that does not explicitly consider gender to be inclusive of trans people can reinforce stereotypes such as the dichotomy of gender.

We expect gender disparities to be different depending on the groups being compared: female–male, trans–cis, and so on. Our work does not aim to infer gender of individual authors. We obtain disaggregated statistics for women, specifically, to study the disparities between female and male NLP researchers. We discuss ethical considerations further in Section 6. See also Mihaljević et al. (2019) for a discussion on ethical considerations in using author name to estimate gender statistics in the Gender Gap in Science Project—a large ongoing project tracking gender gaps in Mathematical and Natural Sciences.⁴

³<http://saifmohammad.com/WebPages/nlp scholar.html>

⁴<https://gender-gap-in-science.org>

Most studies on gender and authorship have found substantial gender disparities in favor of male researchers. They include work on ~1700 articles from journals of library and information science (Håkanson, 2005), on ~12 million articles from the Web of Science (for Sociology, Political Science, Economics, Cardiology and Chemistry) (Ghiasi et al., 2016; Andersen and Nielsen, 2018), on ~2 million mathematics articles (Mihaljević-Brandt et al., 2016), on ~1.6 million articles from PubMed life science and biomedical research (Mishra et al., 2018), on ~1.5 million articles from fifty disciplines published in JSTOR (King et al., 2017), and on ~0.5 million publications from US research universities (Duch et al., 2012). There also exists some work that shows that in fields such as linguistics (LSA, 2017) and psychology (Willyard, 2011), female and male participation is either close to parity or tilted in favor of women.

In NLP research, Schluter (2018) showed that there are barriers in the paths of women researchers, delaying their attainment of mentorship status (as estimated through last author position in papers). Anderson et al. (2012) examine papers from 1980 to 2008 to track the ebb and flow of topics within NLP, and the influence of researchers from outside NLP on NLP. Vogel and Jurafsky (2012) examined about 13,000 papers from 1980 to 2008 to determine basic authorship statistics by women and men. Gender statistics were determined by a combination of automatic and manual means. The automatic method relied on lists of baby names from various languages. They found that female authorship has been steadily increasing from 1980 to 2008. Our work examines a much larger set of NLP papers (1965–2019), re-examines some of the questions raised in Vogel and Jurafsky (2012), and explores several new questions, especially on first author gender and disparities in citation.

3 Data

We extracted and aligned information from the ACL Anthology (AA) and Google Scholar (GS) to create a dataset of tens of thousands of NLP papers and their citations. We aligned the information across AA and GS using the paper title, year of publication, and first author last name. Details about the dataset, as well as an analysis of the volume of research in NLP over the years, are available in Mohammad (2020b). We summarize key information below.

3.1 ACL Anthology Data

The ACL Anthology is available through its website and a `github` repository.⁵ We extracted paper title, names of authors, year of publication, and venue of publication from the repository.⁶

As of June 2019, AA had ~50K entries; however, this includes some entries that are not truly research publications (for example, forewords, prefaces, programs, schedules, indexes, invited talks, appendices, session information, newsletters, lists of proceedings, etc.). After discarding them, we are left with 44,894 papers.⁷

Inferring Aggregate Gender Statistics: The ACL Anthology does not record author demographic information. To infer aggregate statistics for male and female authors, we create two bins of authors: A-Mname (authors that have self-reported as males or with names commonly associated with males) and A-Fnames (authors that have self-reported as females or with names commonly associated with females).⁸ We made use of three resources to populate A-Mname and A-Fname:

1. A manually curated list of 11,932 AA authors and their genders provided by Vogel and Jurafsky (2012) (VJ-AA list) (3,359 female and 8,573 male).⁹
2. A list of 55,924 first names that are strongly associated with females and 30,982 first names that are strongly associated with males, that we generated from the US Social Security Administration’s (USSA) published database of names and genders of newborns.¹⁰
3. A list of 26,847 first names that are strongly associated with females and 23,614 first names that are strongly associated with males, that we generated from a list of 9,300,182

⁵<https://www.aclweb.org/anthology/>
<https://github.com/acl-org/acl-anthology>

⁶Multiple authors can have the same name and the same authors may use multiple variants of their names in papers. The AA volunteer team handles such ambiguities using both semi-automatic and manual approaches (fixing some instances on a case-by-case basis). Additionally, AA keeps a file that includes canonical forms of author names.

⁷We used simple keyword searches for terms such as *foreword*, *invited talk*, *program*, *appendix* and *session* in the title to pull out entries that were likely to not be research publications. These were then manually examined to verify that they did not contain any false positives.

⁸Note that a person may have a name commonly associated with one gender but belong to a different gender.

⁹<https://nlp.stanford.edu/projects/gender.shtml>

¹⁰<https://www.ssa.gov/oact/babynames/limits.html>

firstname–gender association list	P	R	F
from USSA data	98.4	69.8	81.7
from PUBMED data	98.3	81.4	89.1

Table 1: Precision (P), Recall (R), and F-score (F) of how well the first name and gender association matches information in the VJ-AA list.

PUBMED authors and their genders (Torvik and Smalheiser, 2009; Smith et al., 2013).¹¹

We acknowledge that despite a large expatriate population, the US census information is not representative of the names from around the world. Further, Chinese origin names tend not to be as strongly associated with gender as names from other parts of the world. However, it should be noted that Vogel and Jurafsky (2012) made special effort to include information from a large number of Asian AA authors in their list. The PUBMED list is also noted for having a substantial coverage of Asian names (Torvik and Smalheiser, 2009).

We determined first name–gender association, by calculating the percentages of first names corresponding to male and female genders as per each of the PUBMED and USSA fullname–gender lists. We consider a first name to be strongly associated with a gender if the percentage is $\geq 95\%$.¹² Table 1 shows how well the first name and gender association matches with the VJ-AA list.

Given the high precision (over 98%) of the USSA and PUBMED lists of gender-associated first names, we use them (in addition to the VJ-AA list) to populate the M-names and F-names bins. Eventually, the A-Mname and A-Fname bins together had 28,682 (76%) of the 37,733 AA authors. Similarly, we created bins for Papers whose First Author is from A-Mname (P-FA-Mname), Papers whose First Author is from A-Fname (P-FA-Fname), Papers whose Last Author is from A-Mname (P-LA-Mname), and Papers whose Last Author is from A-Fname (P-LA-Fname) to estimate aggregate-level statistics for papers with male and female first and last authors. P-FA-Mname and P-FA-Fname together have 37,297 (83%) AA papers (we will refer to this subset as AA*), P-LA-Mname and P-LA-Fname together have 39,368 (88%) AA papers (we will refer to this subset as AA**).

¹¹<https://experts.illinois.edu/en/datasets/genni-ethnea-for-the-author-ity-2009-dataset-2>

¹²A choice of other percentages such as 90% or 99% would also have been reasonable.

NLP Academic Age as a Proxy for Experience in NLP: First author percentage may vary due to experience, area of research within NLP, venue of publication, etc. To gauge experience, we use the number of years one has been publishing in AA; we will refer to this as the *NLP Academic Age*. So if this is the first year one has published in AA, then their NLP academic age is 1. If one published their first AA paper in 2001 and their latest AA paper in 2018, then their academic age is 18. Note that NLP academic age is not always an accurate reflection of one’s research experience. Also, one can publish NLP papers outside of AA.

3.2 Google Scholar Data

Google Scholar allows researchers to create and edit public author profiles called *Google Scholar Profiles*. Authors can include their papers (along with their citation information) on this page.

We extracted citation information from Google Scholar profiles of authors who published at least three papers in the ACL Anthology.¹³ This yielded citation information for 1.1 million papers in total. We will refer to this dataset as the *NLP Subset of the Google Scholar Dataset*, or *GScholar-NLP* for short. Note that GScholar-NLP includes citation counts not just for NLP papers, but also for non-NLP papers published by authors who have at least three papers in AA.

GScholar-NLP includes 32,985 of the 44,894 papers in AA (about 75%). We will refer to this subset of the ACL Anthology papers as AA’. The citation analyses presented in this paper are on AA’.

4 Gender Gap in Authorship

We use the datasets described in §3.1 (especially AA* and AA**) to answer a series of questions on female authorship. Since we do not have full self-reported information for all authors, these should be treated as estimates.

First author is a privileged position in the author list that is usually reserved for the researcher that has done the most work and writing.¹⁴ In NLP, first authors are also often students. Thus we are especially interested in gender gaps that effect them. The last author position is often reserved for the most senior or mentoring researcher. We explore last author disparities only briefly (in Q1).

¹³This is allowed by GS’s robots exclusion standard.

¹⁴A small number of papers have more than one first author. This work did not track that.

Q1. What percentage of the authors in AA are female? What percentage of the AA papers have female first authors (FFA)? What percentage of the AA papers have female last authors (FLA)? How have these percentages changed since 1965?

A. Overall, we estimate that about 29.7% of the 28,682 authors are female; about 29.2% of the first authors in 37,297 AA* papers are female; and about 25.5% of the last authors in 39,368 AA** papers are female. Figure 1 shows how these percentages have changed over the years.

Discussion: Across the years, the percentage of female authors overall is close to the percentage of papers with female first authors. (These percentages are around 28% and 29%, respectively, in 2018.) However, the percentage of female last authors is markedly lower (hovering at about 25% in 2018). These numbers indicate that, as a community, we are far from obtaining male–female parity. A further striking (and concerning) observation is that the female author percentages have not improved since the year 2006.

To put these numbers in context, the percentage of female scientists worldwide (considering all areas of research) has been estimated to be around 30%. The reported percentages for many computer science sub-fields are much lower.¹⁵ The percentages are much higher for certain other fields such as psychology (Willyard, 2011) and linguistics (LSA, 2017).

Q2. How does FFA vary by paper type and venue?

A. Figure 2 shows FFA percentages by paper type and venue.

Discussion: Observe that FFA percentages are lowest for CoNLL, EMNLP, IJCNLP, and system demonstration papers (21% to 24%). FFA percentages for journals, other top-tier conferences, SemEval, shared task papers, and tutorials are the next lowest (24% to 28%). The percentages are markedly higher for LREC, *Sem, and RANLP (33% to 36%), as well as for workshops (31.7%).

Q3. How does female first author percentage change with NLP academic age?

A. In order to determine these numbers, every paper in AA* was placed in a bin corresponding to NLP academic age: if the paper’s first author had an academic age of 1 in the year when the

paper was published, then the paper is placed in bin 1; if the paper’s first author had an academic age of 2 in the year when the paper was published, then the paper is placed in bin 2; and so on. The bins for later years contained fewer papers. This is expected as senior authors in NLP often work with students, and students are encouraged to be first authors. Thus, we combine some of the bins in later years: one bin for academic ages between 10 and 14; one for 15 to 19; one for 20 to 34; and one for 35 to 50. Once the papers are assigned to the bins, we calculate the percentage of papers in each bin that have a female first author. Figure 3 shows the results.

Discussion: Observe that, with the exception of the 35 to 50 academic age bin, FFA% is highest (30%) at age 1 (first year of publication). There is a period of decline in FFA% until year 6 (27.4%)—this difference is statistically significant (t-test, $p < 0.01$). This might be a potential indicator that graduate school has a progressively greater negative impact on the productivity of women than of men. (Academic age 1 to 6 often correspond to the period when the first author is in graduate school or in a temporary post-doctoral position.) After year 6, we see a recovery back to 29.4% by year 8, followed by a period of decline once again.

Q4. How does female first author percentage vary by area of research (within NLP)? Which areas have higher-than-average FFA%? Which areas have lower-than-average FFA%? How does FFA% correlate with popularity of an area—that is, does FFA% tend to be higher- or lower-than-average in areas where lots of authors are publishing?

A. We use word bigrams in the titles of papers to sample papers from various areas.¹⁶ The title has a privileged position in a paper. Primarily, it conveys what the paper is about. For example, a paper with *machine translation* in the title is likely about machine translation. Figure 4 shows the list of top 66 bigrams that occur in the titles of more than 100 AA* papers (in decreasing order of the bigram frequency). For each bigram, the figure also shows the percentage of papers with a female first author. In order to determine whether there is a correlation between the number of papers corresponding to a bigram and FFA%, we calculated the Spearman’s rank correlation between the rank of a bigram by

¹⁵<https://unesdoc.unesco.org/ark:/48223/pf0000235155>

¹⁶Other approaches such as clustering are also reasonable; however, results with those might not be easily reproducible.

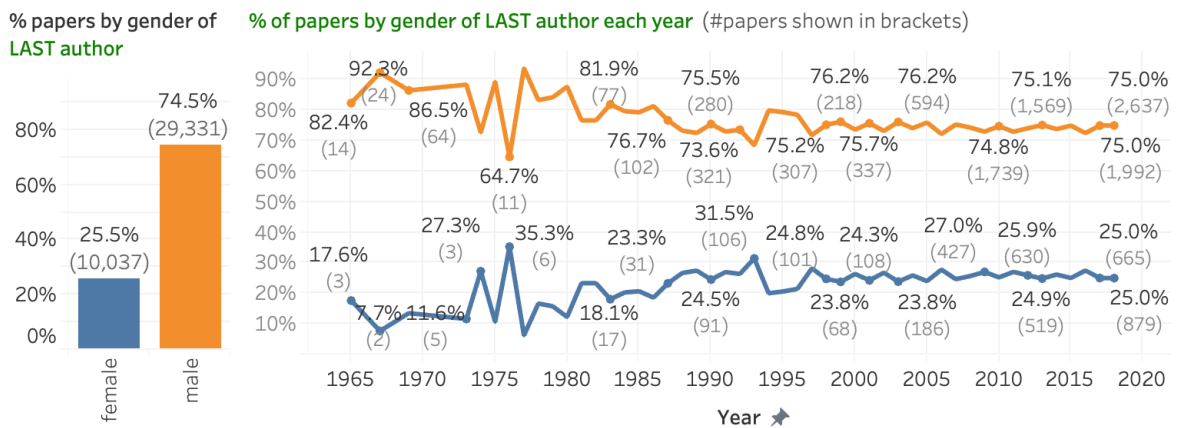
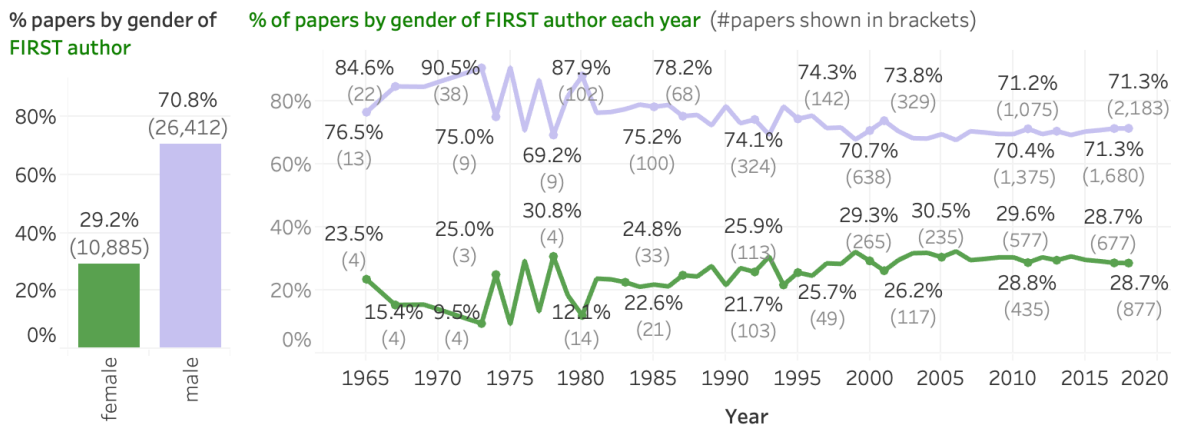
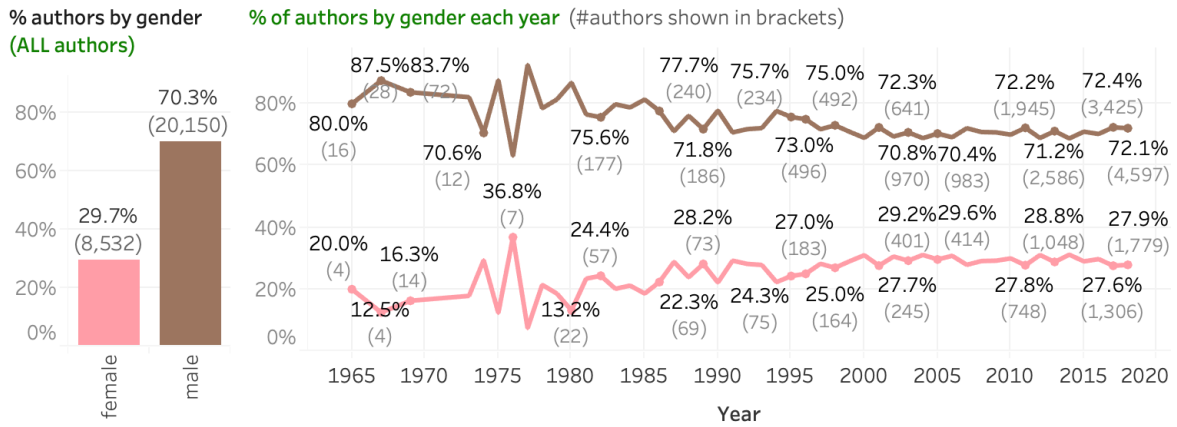


Figure 1: Female authorship percentages in AA over the years: overall, as first author, and as last author.

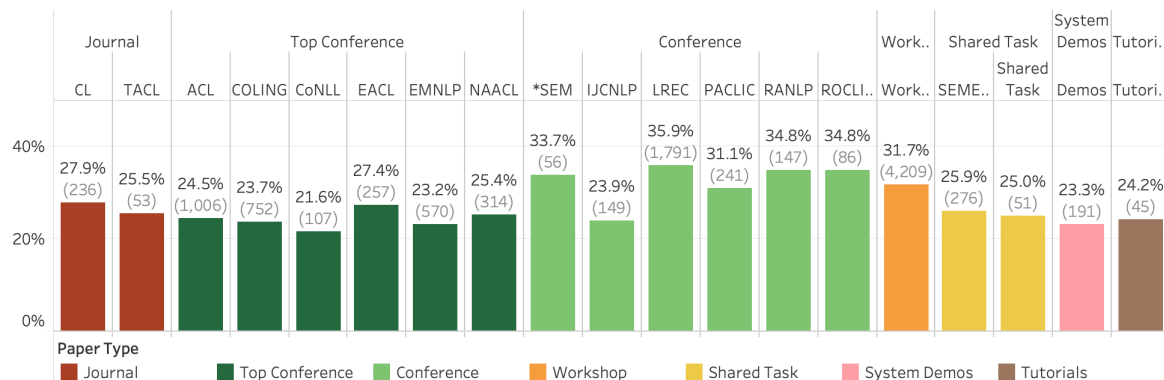


Figure 2: FFA percentage by venue and paper type. The number of FFA papers is shown in parenthesis.

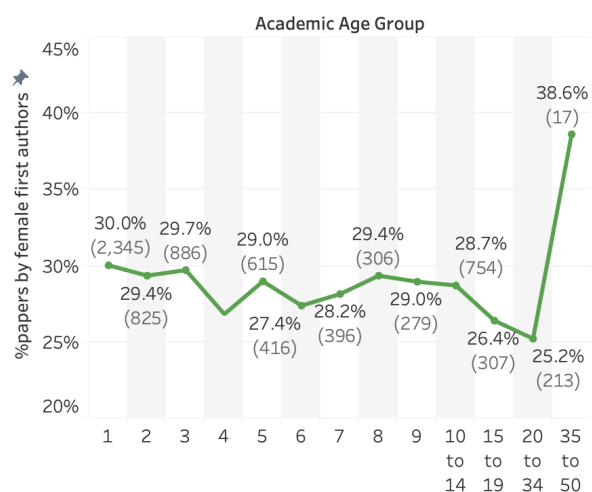


Figure 3: FFA percentage by academic age. The number of FFA papers is shown in parenthesis.

number of papers and the rank of a bigram by FFA%. This correlation was found to be only 0.16. This correlation is not statistically significant at $p < 0.01$ (two-sided p -value = 0.2). Experiments with lower thresholds (174 bigrams occurring in 50 or more papers and 1408 bigrams occurring in 10 or more papers) also resulted in very low and non-significant correlation numbers (0.11 and 0.03, respectively).

Discussion: Observe that FFA% varies substantially depending on the bigram. It is particularly low for title bigrams such as *dependency parsing*, *language models*, *finite state*, *context free*, and *neural models*; and markedly higher than average for *domain specific*, *semantic relations*, *dialogue system*, *spoken dialogue*, *document summarization*, and *language resources*. However, the rank correlation experiments show that there is *no* correlation between the popularity of an area (number of papers that have a bigram in the title) and the percentage of female first authors.

To obtain further insights, we also repeat some of the experiments described above for unigrams in paper titles. We found that FFA rates are relatively high in non-English European language research such as papers on Russian, Portuguese, French, and Italian. FFA rates are also relatively high for work on prosody, readability, discourse, dialogue, paraphrasing, and individual parts of speech such as adjectives and verbs. FFA rates are particularly low for papers on theoretical aspects of statistical modelling, and for terms such as *WMT*, *parsing*, *markov*, *recurrent*, and *discriminative*.

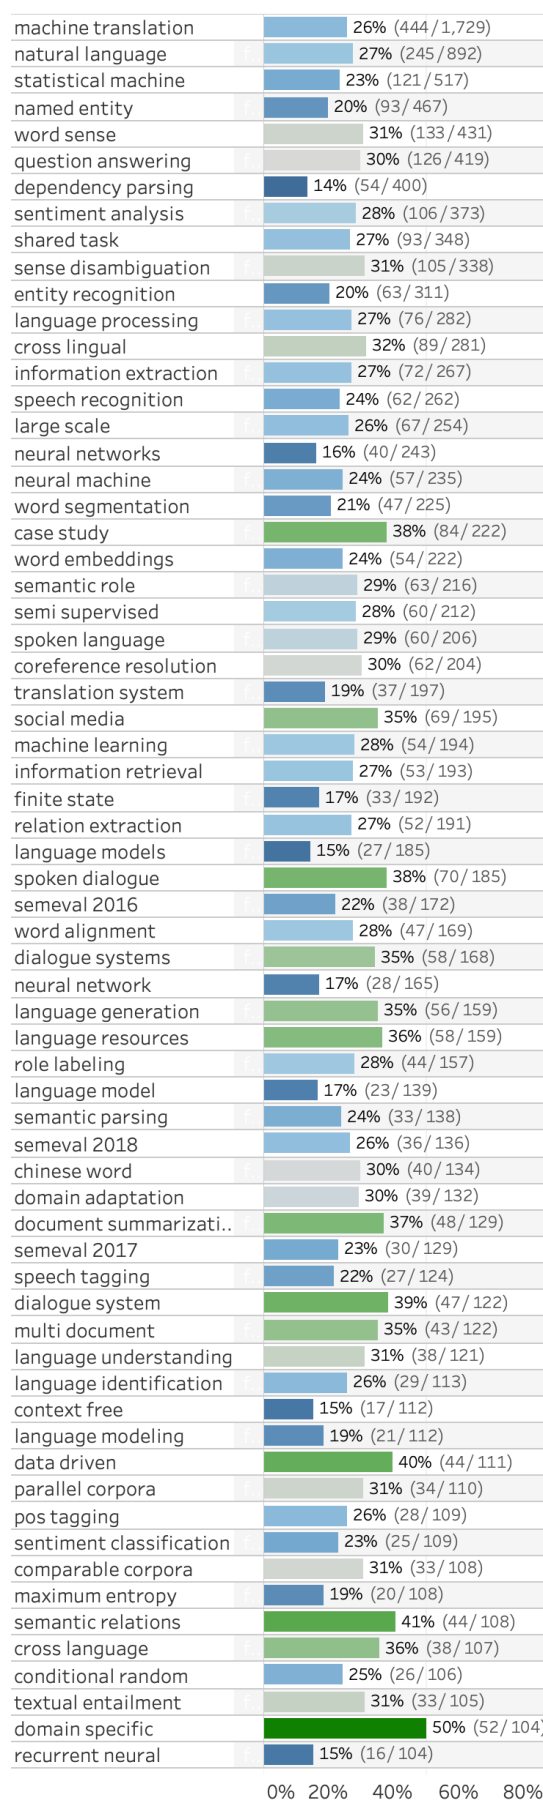


Figure 4: Top 66 bigrams in AA* titles and FFA% (30%: light grey; <30%: blue; >30%: green).

5 Gender Gap in Citations

Research articles can have impact in a number of ways—pushing the state of the art, answering crucial questions, finding practical solutions that directly help people, etc. However, individual measures of research impact are limited in scope—they measure only some kinds of contributions. The most commonly used metrics of research impact are derived from citations including: number of citations, average citations, h-index, and impact factor (Bornmann and Daniel, 2009). Despite their limitations, citation metrics have substantial impact on a researcher’s scientific career; often through a combination of funding, the ability to attract talented students and collaborators, job prospects, and other opportunities in the wider research community. Thus, disparities in citations (citation gaps) across demographic attributes such as gender, race, and location have direct real-world adverse implications. This often also results in the demoralization of researchers and marginalization of their work—thus negatively impacting the whole field.

Therefore, we examine gender disparities in citations in NLP. We use a subset of the 32,985 AA’ papers (§3.2) that were published from 1965 to 2016 for the analysis (to allow for at least 2.5 years for the papers to collect citations). There are 26,949 such papers.

Q5. How well cited are women and men?

A. For all three classes (females, males, and gender unknown), Figure 5 shows: a bar graph of number of papers, a bar graph of total citations received, and box and whisker plots for citations received by individuals. The whiskers are at a distance of 1.5 times the inter-quartile length. Number of citations pertaining to key points such as 25th percentile, median, and 75th percentile are indicated on the left of the corresponding horizontal bars.

Discussion: On average, female first author papers have received markedly fewer citations than male first author papers (37.6 compared to 50.4). The difference in median is smaller (11 compared to 13). The difference in the distributions of males and females is statistically significant (Kolmogorov–Smirnov test, $p < 0.01$).¹⁷ The large difference in averages and smaller difference in medians suggests that there are markedly more

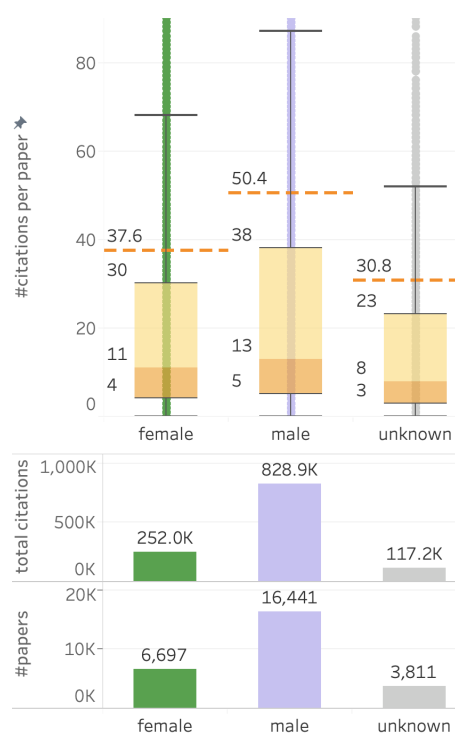


Figure 5: #papers, total citations, box plot of citations per paper: for female, male, gender-unknown first authors. The orange dashed lines mark averages.

very heavily cited male first-author papers than female first-author papers.

The differences in citations, or *citation gap*, across genders may itself vary: (1) by period of time; (2) due to confounding factors such as academic age and areas of research. We explore these next.

Q6. How has the citation gap across genders changed over the years?

A. Figure 6 (left side) shows the citation statistics across four time periods.

Discussion: Observe that female first authors have always been a minority in the history of ACL; however, on average, their papers from the early years (1965 to 1989) received a markedly higher number of citations than those of male first authors from the same period. We can see from the graph that this changed in the 1990s when male first-author papers obtained markedly more citations on average. The citation gap reduced considerably in the 2000s, and the 2010–2016 period saw a further reduction. It remains to be seen whether the citation gap for these 2010–2016 papers widens in the coming years.

It is also interesting to note that the gender-unknown category has almost bridged the gap with the male category in terms of average citations. Further, the proportion of the gender-unknown

¹⁷Kolmogorov–Smirnov (KS) test is a non-parametric test that can be applied to compare any two distributions without making assumptions about the nature of the distributions.

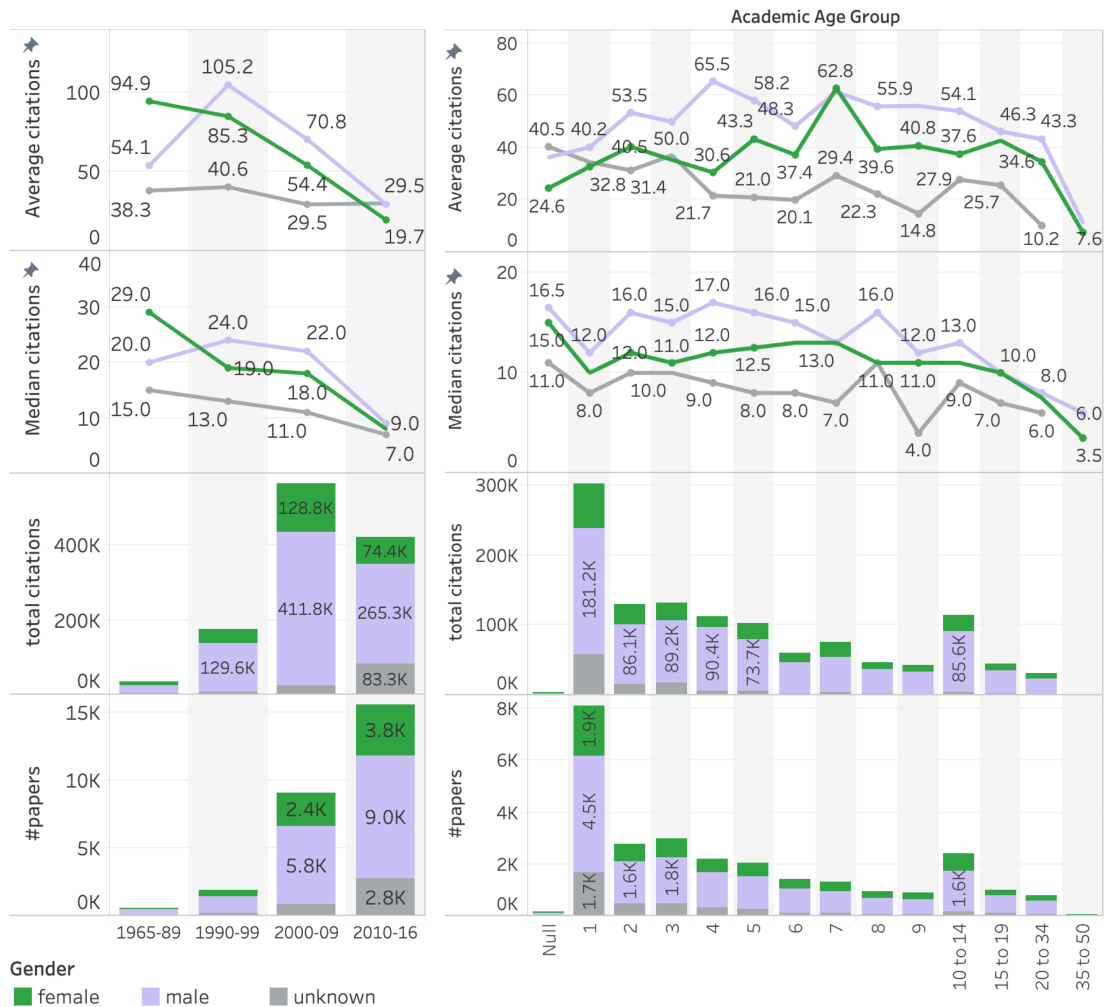


Figure 6: Citation gap across genders for papers: published in different time spans (left); by academic age (right).

authors has steadily increased over the years—arguably, an indication of better representation of authors from around the world in recent years.¹⁸

Q7. How have citations varied by gender and academic age? Is the citation gap a side effect of a greater proportion of new-to-NLP female first authors than new-to-NLP male first authors?

A. Figure 6 (right side) shows citation statistics broken down by gender and academic age.

Discussion: The graphs show that female first authors consistently receive fewer citations than male first authors across the spans of their academic age. (The gap is highest at academic age 4 and lowest at academic age 7.) Thus, the citation gap is likely due to factors beyond differences in average academic age between men and women.

¹⁸Our method is expected to have a lower coverage of names from outside North America and Europe because USSA and PUBMED databases historically have had fewer names from outside North America and Europe.

Q8. How prevalent is the citation gap across areas of research within NLP? Is the gap simply because more women work in areas that receive low numbers of citations (regardless of gender)?

A. On average, male first authors are cited more than female first authors in 54 of the 66 areas (82% of the areas) discussed earlier in Q4 and Figure 4. Female first authors are cited more in the sets of papers whose titles have: *word sense, sentiment analysis, information extraction, neural networks, neural network, semeval 2016, language model, document summarization, multi document, spoken dialogue, dialogue systems, and speech tagging.*

If women chose to work in areas that happen to attract less citations by virtue of the area, then we would not expect to see citation gaps in so many areas. Recall also that we already showed that FFA% is not correlated with rank of popularity of an area (Q4). Thus it is unlikely that the choice of area of research is behind the gender gap.

6 Limitations and Ethical Considerations

Q9. What are the limitations and ethical considerations involved with this work?

A. Data is often a representation of people (Zook et al., 2017). This is certainly the case here and we acknowledge that the use of such data has the potential to harm individuals. Further, while the methods used are not new, their use merits reflection.

Analysis focused on women and men leaves out non-binary people.¹⁹ Additionally, not disaggregating cis and trans people means that the statistics are largely reflective of the more populous cis class. We hope future work will obtain disaggregated information for various genders. However, careful attention must be paid as some gender classes might include too few NLP researchers to ensure anonymity even with aggregate-level analysis.

The use of female- and male-gender associated names to infer population level statistics for women and men excludes people that do not have such names and people from some cultures where names are not as strongly associated with gender.

Names are not immutable (they can be changed to indicate or not indicate gender) and people can choose to keep their birth name or change it (providing autonomy). However, changing names can be quite difficult. Also, names do not capture gender fluidity or contextual gender.

A more inclusive way of obtaining gender information is through self-reported surveys. However, challenges persist in terms of how to design effective and inclusive questionnaires (Bauer et al., 2017; Group, 2014). Further, even with self-report textboxes that give the respondent the primacy and autonomy to express gender, downstream research often ignores such data or combines information in ways beyond the control of the respondent.²⁰ Also, as is the case here, it is not easy to obtain self-reported historical information.

Social category detection can potentially lead to harm, for example, depriving people of opportunities simply because of their race or gender. However, one can also see the benefits of NLP techniques and social category detection in public health (e.g., developing targeted initiatives to improve health outcomes of vulnerable populations),

¹⁹Note that as per widely cited definitions, nonbinary people are considered transgender, but most transgender people are not non-binary. Also, trans people often use a name that is more associated with their gender identity.

²⁰<https://reallifemag.com/counting-the-countless/>

as well as in psychology and social science (e.g., to better understand the unique challenges of belonging to a social category).

A larger list of ethical considerations associated with the NLP Scholar project is available through the project webpage.²¹ Mihaljević et al. (2019) also discusses the ethical considerations in using author names to infer gender statistics in the Gender Gap in Science Project.²²

7 Conclusions

We analyzed the ACL Anthology to show that only ~30% have female authors, ~29% have female first authors, and ~25% have female last authors. Strikingly, even though some gains were made in the early years of NLP, overall FFA% has not improved since the mid 2000s. Even though there are some areas where FFA% is close to parity with male first authorship, most areas have a substantial gap in the numbers for male and female authorship. We found no correlation between popularity of research area and FFA%. We also showed how FFA% varied by paper type, venue, academic age, and area of research. We used citation counts extracted from Google Scholar to show that, on average, male first authors are cited markedly more than female first authors, even when controlling for experience and area of work. Thus, in NLP, gender gaps exist both in authorship and citations.

This paper did not explore the reasons behind the gender gaps. However, the inequities that impact the number of women pursuing scientific research (Roos, 2008; Foschi, 2004; Buchmann, 2009) and biases that impact citation patterns unfairly (Brouns, 2007; Feller, 2004; Gupta et al., 2005) are well-documented. These factors play a substantial role in creating the gender gap, as opposed to differences in innate ability or differences in quality of work produced by these two genders. If anything, past research has shown that self-selection in the face of inequities and adversity leads to more competitive, capable, and confident cohorts (Nekby et al., 2008; Hardies et al., 2013).

Acknowledgments

Many thanks to Rebecca Knowles, Ellen Riloff, Tara Small, Isar Nejadgholi, Dan Jurafsky, Rada Mihalcea, Isabelle Augenstein, Eric Joanis, Michael Strube, Shubhanshu Mishra, and Cyril

²¹<http://saifmohammad.com/WebPages/nlpscholar.html>

²²<https://gender-gap-in-science.org>

Goutte for the tremendously helpful discussions. Many thanks to Cassidy Rae Henry, Luca Soldaini, Su Lin Blodgett, Graeme Hirst, Brendan T. O'Connor, Lucía Santamaría, Lyle Ungar, Emma Manning, and Peter Turney for discussions on the ethical considerations involved with this work.

References

- Richard Alford. 1987. *Naming and identity: A cross-cultural study of personal naming practices*. Hraf Press.
- Jens Peter Andersen and Mathias Wullum Nielsen. 2018. Google Scholar and Web of Science: Examining gender differences in citation coverage across five scientific disciplines. *Journal of Informetrics*, 12(3):950–959.
- Ashton Anderson, Dan McFarland, and Dan Jurafsky. 2012. Towards a computational history of the ACL: 1980–2008. In *Proceedings of the Workshop on Re-discovering 50 Years of Discoveries*, pages 13–21.
- Herbert Barry III and Aylene S Harper. 2014. Unisex names for babies born in pennsylvania 1990–2010. *Names*, 62(1):13–22.
- Greta R Bauer, Jessica Braimoh, Ayden I Scheim, and Christoffer Dharma. 2017. Transgender-inclusive measures of sex/gender for population surveys: Mixed-methods evaluation and recommendations. *PLoS one*, 12(5):e0178043.
- Lutz Bornmann and Hans-Dieter Daniel. 2009. The state of h index research. *EMBO reports*, 10(1):2–6.
- Margo Brouns. 2007. The making of excellence: Gender bias in academia. In *Exzellenz in Wissenschaft und Forschung - neue Wege in der Gleichstellungspolitik*, pages 23–42. Wissenschaftsrat.
- Claudia Buchmann. 2009. Gender inequalities in the transition to college. *Teachers College Record*, 111(10):2320–2345.
- Catherine Connell. 2010. Doing, undoing, or redoing gender? Learning from the workplace experiences of transpeople. *Gender & Society*, 24(1):31–55.
- Helana Darwin. 2017. Doing gender beyond the binary: A virtual ethnography. *Symbolic Interaction*, 40(3):317–334.
- Michelle L Dion, Jane Sumner, and Sara McLaughlin Mitchell. 2018. Gendered citation patterns across political science and social science methodology fields. *Political Analysis*, 26(3):312–327.
- Jordi Duch, Xiao Han T Zeng, Marta Sales-Pardo, Filippo Radicchi, Shayna Otis, Teresa K Woodruff, and Luís A Nunes Amaral. 2012. The possible role of resource requirements and academic career-choice risk on gender differences in publication rate and impact. *PLoS one*, 7(12):e51332.
- Irwin Feller. 2004. Measurement of scientific performance and gender bias. In *Gender and Excellence in the Making*, pages 35–39. Luxembourg: Office for Official Publications of the European Communities.
- Marta Foschi. 2004. Blocking the use of gender-based double standards for competence. In *Gender and Excellence in the Making*, pages 51–55. Luxembourg: Office for Official Publications of the European Communities.
- Juan Miguel Gallego and Luis H Gutiérrez. 2018. An integrated analysis of the impact of gender diversity on innovation and productivity in manufacturing firms. Technical report, Inter-American Development Bank.
- Gita Ghiasi, Vincent Larivière, and Cassidy Sugimoto. 2016. Gender differences in synchronous and diachronous self-citations. In *21st International Conference on Science and Technology Indicators-STI 2016. Book of Proceedings*.
- GenIUSS Group. 2014. *Best practices for asking questions to identify transgender and other gender minority respondents on population-based surveys*. eScholarship, University of California.
- Namrata Gupta, Carol Kemelgor, Stefan Fuchs, and Henry Etzkowitz. 2005. Triple burden on women in science: A cross-cultural analysis. *Current science*, pages 1382–1386.
- Malin Håkanson. 2005. The impact of gender on citations: An analysis of college & research libraries, journal of academic librarianship, and library quarterly. *College & Research Libraries*, 66(4):312–323.
- Dalia S Hakura, Mumtaz Hussain, Monique Newiak, Vimal Thakoor, and Fan Yang. 2016. *Inequality, gender gaps and economic growth: Comparative evidence for sub-Saharan Africa*. International Monetary Fund.
- Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–13.
- Kris Hardies, Diane Breesch, and Joël Branson. 2013. Gender differences in overconfidence and risk taking: Do self-selection and socialization matter? *Economics Letters*, 118(3):442–444.
- Janet Shibley Hyde, Rebecca S Bigler, Daphna Joel, Charlotte Chucky Tate, and Sari M van Anders. 2019. The future of sex and gender in psychology: Five challenges to the gender binary. *American Psychologist*, 74(2):171.
- Suzanne J Kessler and Wendy McKenna. 1978. *Gender: An ethnomethodological approach*. IL: The University of Chicago Press.

- Molly M King, Carl T Bergstrom, Shelley J Correll, Jennifer Jacquet, and Jevin D West. 2017. Men set their own cites high: Gender and self-citation across fields and over time. *Socius*, 3:2378023117738903.
- Vincent Larivière, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R Sugimoto. 2013. Bibliometrics: Global gender disparities in science. *Nature News*, 504(7479):211.
- Stanley Lieberson, Susan Dumais, and Shyon Baumann. 2000. The instability of androgynous names: The symbolic maintenance of gender boundaries. *American Journal of Sociology*, 105(5):1249–1287.
- Linda L Lindsey. 2015. The sociology of gender theoretical perspectives and feminist frameworks. In *Gender roles*, pages 23–48. Routledge.
- The Linguistic Society of America LSA. 2017. The state of linguistics in higher education annual report 2017. Technical report, The Linguistic Society of America.
- Sangeeta Mehta, Karen EA Burns, Flavia R Machado, Alison E Fox-Robichaud, Deborah J Cook, Carolyn S Calfee, Lorraine B Ware, Ellen L Burnham, Niranjana Kisson, John C Marshall, et al. 2017. Gender parity in critical care medicine. *American journal of respiratory and critical care medicine*, 196(4):425–429.
- Helena Mihaljević, Marco Tullney, Lucía Santamaría, and Christian Steinfeldt. 2019. Reflections on gender analyses of bibliographic corpora. *Frontiers in Big Data*, 2:29.
- Helena Mihaljević-Brandt, Lucía Santamaría, and Marco Tullney. 2016. The effect of gender in the publication patterns in mathematics. *PLoS One*, 11(10):e0165367.
- Shubhanshu Mishra, Brent D Fegley, Jana Diesner, and Vetle I Torvik. 2018. Self-citation is the hallmark of productive authors, of any gender. *PLoS one*, 13(9):e0195773.
- Saif M. Mohammad. 2019. The state of NLP literature: A diachronic analysis of the ACL Anthology. *arXiv preprint arXiv:1911.03562*.
- Saif M. Mohammad. 2020a. Examining citations of natural language processing literature. In *Proceedings of the 2020 Annual Conference of the Association for Computational Linguistics*, Seattle, USA.
- Saif M. Mohammad. 2020b. NLP Scholar: A dataset for examining the state of NLP research. In *Proceedings of the 12th Language Resources and Evaluation Conference (LREC-2020)*, Marseille, France.
- Saif M. Mohammad. 2020c. NLP Scholar: An interactive visual explorer for Natural Language Processing literature. In *Proceedings of the 2020 Annual Conference of the Association for Computational Linguistics*, Seattle, USA.
- Lena Nekby, Peter Thoursie, and Lars Vahtrik. 2008. Gender and self-selection into a competitive environment: Are women more overconfident than men? *Economics Letters*, 100(3):405–407.
- Caroline Criado Perez. 2019. *Invisible women: Exposing data bias in a world designed for men*. Random House.
- Jane Pilcher. 2017. Names and “doing gender”: How forenames and surnames contribute to gender identities, difference, and inequalities. *Sex roles*, 77(11-12):812–822.
- Cristina Richards, Walter Pierre Bouman, and M Barker. 2017. *Non-binary genders*. London: Palgrave Macmillan.
- Patricia A Roos. 2008. Together but unequal: Combating gender inequity in the academy. *Journal of Workplace Rights*, 13(2):185–199.
- Natalie Schluter. 2018. [The glass ceiling in NLP](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2793–2798, Brussels, Belgium.
- Inger Skjelsboek and Dan Smith. 2001. *Gender, peace and conflict*. Sage.
- Brittany N Smith, Mamta Singh, and Vetle I Torvik. 2013. A search engine approach to estimating temporal changes in gender orientation of first names. In *Proceedings of the 13th ACM/IEEE-CS joint conference on Digital libraries*, pages 199–208. ACM.
- Vetle I Torvik and Neil R Smalheiser. 2009. Author name disambiguation in medline. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 3(3):1–29.
- Adam Vogel and Dan Jurafsky. 2012. He said, she said: Gender in the ACL Anthology. In *Proceedings of the Special Workshop on Rediscovering 50 Years of Discoveries*, pages 33–41, Jeju Island, Korea.
- World Economic Forum WEF. 2018. The global gender gap report 2018. Technical report, World Economic Forum, Geneva, Switzerland.
- Cassandra Willyard. 2011. Men: A growing minority. *GradPSYCH Magazine*, 9(1):40.
- Jonathan Woetzel et al. 2015. The power of parity: How advancing women’s equality can add \$12 trillion to global growth. Technical report, McKinsey Global Institute.
- Matthew Zook, Solon Barocas, Danah Boyd, Kate Crawford, Emily Keller, Seeta Pena Gangadharan, Alyssa Goodman, Rachelle Hollander, Barbara A. Koenig, Jacob Metcalf, Arvind Narayanan, Alondra Nelson, and Frank Pasquale. 2017. [Ten simple rules for responsible big data research](#). *PLOS Computational Biology*, 13(3):1–10.